

Challenges and approaches in bridging the biomimicry gap in biohybrid systems of fish and robots

Présentée le 23 novembre 2023

Faculté des sciences et techniques de l'ingénieur
Groupe de Systèmes Robotiques Mobiles
Programme doctoral en robotique, contrôle et systèmes intelligents

pour l'obtention du grade de Docteur ès Sciences

par

Vaios PAPASPYROS

Acceptée sur proposition du jury

Dr D. Gillet, président du jury
Prof. F. Mondada, directeur de thèse
Prof. S. Mintchev, rapporteur
Prof. D. Tarapore, rapporteur
Prof. A. Ijspeert, rapporteur

“Nothing in life is to be feared, it is only to be understood.
Now is the time to understand more, so that we may fear less.”
— Marie Curie

Acknowledgements

Reflecting on my PhD journey, I realize many aspects of my personal and professional life would not have been the same without the interactions and support of many people. Although I am not particularly famous for making long speeches, the least I can do is to dedicate a few words to each of them.

First and foremost, I would like to thank my family: my sister (Andriana), parents (Spyros & Giota), aunt (Athina & Maria), uncles (Nontas & Kostas), and — yes, of course — dog (Judy). Sadly, two of you, dad & Judy, are no longer with us. I miss you and wish I could share this part of my life with you as well. All I have achieved, to this day, would not have been possible without your (all of you) love and support for nearly 3 decades. Each of you have contributed, one way or another, to me being the person I am today. I deeply thank you for that. This thesis is dedicated to all of you, and perhaps a bit more to my sister. She was there for me when I needed it the most, and together we went through fun and rough patches of life that bonded us more than ever.

Certainly, this thesis would not have been the same without the contribution of my supervisor, Francesco Mondada. He supervised with great understanding, respect, and interest for science. He allowed me the freedom I needed to express my scientific curiosity, while at the same time being there to uphold the quality and novel direction of my work. Even during the “darker” times of this journey, he believed in me, trusted me, and was there to help me get back on track when needed. All in all, he taught me how to be a scientist (and that is no small feat), sometimes on purpose and others indirectly, due to his ethos and professional attitude. I am proud to close this chapter of my life being one of his PhD students. I am immensely thankful to him.

A special mention is also due to Guy Theraulaz, Clément Sire, and Ramón Escobedo from the Université de Toulouse – Paul Sabatier, with which I have collaborated from my very early PhD days. Thank you for welcoming me so warmly in Toulouse. Of course, I am also grateful for all the deep scientific discussions which produced very interesting research. To Guy Theraulaz, I am grateful for his contribution to my PhD, by means of hosting our robot experiments in his premises and for all the valuable insights into a field that was new to me, collective behavior. Clément Sire’s contribution to my thesis is also hard to put into a few lines, but he has taught me a great deal during our interactions and endless conversations on any imaginable topic, certainly fostering important scientific values in me. Last but not least, Ramón Escobedo,

my fellow research and aviation friend, has been a great influence and has provided a lot of support in various phases of the thesis. I am looking forward to visiting you all in Toulouse under a different context.

I would also like to thank the people I had the chance to meet and interact with in MOBOTS. I was embraced warmly from the first moment I stepped foot in the lab, and I have very fond memories over the last 5 years. You have all contributed to some of my best memories in life. Big thank you to all the MOBOTS & Learn people. I am making a special mention to those closest to me during this journey, in chronological order for fairness and to the best of my memory's ability: Special thanks goes to Daniel Burnier with whom I closely collaborated on numerous engineering challenges including the LureBot, which would have never seen the light of day if it was not for him. He taught me how to be a better engineer. I would also like to thank Claudia Marra, Norbert Crot, Frank & Evgenia Bonnet, Bertrand Collignon, and Eliot Ferragni, the first group of people I was introduced to at MOBOTS and made me feel at home. Christian Giang and Morgane Chevalier also deserve a very honorable mention. The only PhD students of MOBOTS at the time, we had a lot of fun collaborating on courses and I learned a lot from your experience. I am also very grateful to Laila el-Hamamsy, Matthieu Broisin, Rafael Barmak, Rob Mills, and Raphael Cherfan, who joined later and with whom I spent the most time. Although I am sad for not having spent even more time with you because of COVID (or overworking, or simply wearing noise-cancelling headphones), we since had a lot of fun together. I can not thank you enough for making my day better every time I was in the office space, or by organizing activities outside EPFL. The same applies to our more recent members, Kunal Massé, Cyril Monette and Jérôme Brender Guillaume. I am happy that can refer to each one of you as friends. Thank you all for listening to me complaining endlessly (mainly Laila, Rafael, Matthieu, Raphael, Dany, and Francesco), I am sure this will not be missed, but I believe I balanced it with equally good jokes. Those will be missed (?).

This would have been an incomplete section without thanking my friends outside EPFL, who helped me balance my life and take important decisions. Here we go, many thanks to my many — really a lot of — friends that I made in Lausanne, and which I will mention by first name only (the reason is twofold: 1. because Greek surnames are a tad too long, and 2. diplomacy, *i.e.*, if I forget someone, chances are their first name is still in the list): Sofia, Anastasia, Konstantinos, Iason, Irene, the Professor (Athanasios), Elena, Georgia, Giannis, Sevi, Theodora, Ioanna, Dimitra, and Dimitris. I could probably dedicate entire chapters about the ways you helped me through this journey, but I should refrain from doing so for brevity. I would, however, like to thank Sofia, Anastasia, and Konstantinos from the bottom of my heart, for being with me through so many good, bad, fun, sad, and serious experiences. This chapter of my life is only complete with you in it. Of course, it would also be impossible to not dedicate some words to Iason and Irene who opened up their home to me and made me feel part of it. Athanasios and I certainly overstayed the welcome, and all 4 of us have many stories to tell. All three of you, taught me how to deal with situations in more maturely (when I would eventually decide to listen to you).

Organized by current country of residence, special thanks also go to Konstantinos Chatzilygeroudis who I have consulted on academic matters prior and during the PhD thesis (and I will probably keep bugging in the future). Last but not least, I would like to thank my life-long friends, Gregory & Matina, Vangelis & Chrysa, and Andreas with whom I have shared the majority of my life events and always have my back.

I also owe a very special mention to (MEng, MSc, soon to be Dr.) Charalampia who put up with me during the last 2+ years of my thesis, which were incidentally the most stressful and demanding ones. She supported me despite my overworking habit — at multiple instances turning off my computer in non-democratic ways — and despite my intense character. She was there for everything, to listen, cook great meals, take me on trips or enforce vacation days, and was very patient throughout. She has acted as a mental anchor-point in my life, and I would not have completed this journey as happily, if at all, without her. The PhD title partially belongs to her (she will most likely tell you that herself, and it is true).

I am also grateful to the members of the PhD defense jury. It has been an honor to be reviewed by respected and accomplished experts in the field. Finally, I would like to thank the Swiss National Science Foundation (SNSF) and the EPFL, who gave me the opportunity and means to pursue the PhD thesis.

Lausanne, Sunday 15th October, 2023

V. P.

Abstract

As the field of ethology advances, especially over the past two decades, the role of animal-robot interaction tools has increasingly become essential. This importance arises from the need for controlled, repetitive, repeatable, and long-duration experiments, which not only relieve human researchers from tedious tasks but also offer novel insights into the rules that govern collective behavior. Such devices can infiltrate groups of animals (in this case, fish) and engage in closed-loop interactions, eliciting responses that range from biomimetic to modulated behavior. However, constructing harmonious biohybrid groups of animals and robots is an intricate task. Despite significant progress in the domain, many questions remain unanswered, necessitating further research and development in both robotics and collective behavior modelling.

This thesis delves into the intersection of collective behavior phenomena and robotics. It capitalizes on the advancements in electronics manufacturing, cutting-edge algorithms, and increased accessibility to computational power that have reshaped the field, resulting in a mixed society methodology that achieved unprecedented levels of biomimicry. Initially, a comprehensive exploration was conducted to understand how fish groups interact with artificial agents. We showed that models displaying active, bidirectional interactions lead to a higher probability of integrating the artificial agent into the fish group. This realization underlined the necessity of transitioning to more detailed and accurate models of interaction that can withstand comparisons to spontaneous fish interactions. Addressing these limitations, we carried out an extensive study on the key design factors that enhance the performance of social interaction models, also revealing the need for a rigorous spatio-temporal benchmarking metric set. This set ensures these models successfully generate realistic short- and long-term social dynamics. However, these models uncovered a secondary engineering problem. Transferring high-fidelity models back to reality demanded highly agile and responsive robotic equipment, a requirement unmet by the current state-of-the-art. Accordingly, within this thesis, we designed a novel framework, inclusive of an experimental setup, a mobile robot, and ancillary software (*e.g.*, for robot control, artificial intelligence and analytical behavioral models...) to overcome these limitations. This approach subsequently enabled us to make substantial strides towards bridging the “biomimicry gap” by transferring high-fidelity models from simulation to reality. The culmination of this thesis outlines our success in progressively bridging this gap, demonstrating unprecedented similarity between simulations, biohybrid, and spontaneous fish-only interaction experiments. By open-sourcing the entirety of the developed software and hardware tools, we aim to lay a solid foundation for future research in

the realm of robot-animal interaction.

In conclusion, this thesis significantly contributes to our understanding of the underlying rules that govern collective behavior. We hope it will pave the way for the creation of truly biohybrid groups and set the stage for future explorations into models of social interactions and their exploitation as robot controllers.

Keywords: Animal-robot interaction, ethorobotics, artificial intelligence, machine learning, deep learning, biohybrid systems, biomimetics, collective behavior, social interactions, complex systems modelling, real-time robot control, reality gap.

Résumé

Au fur et à mesure que le domaine de l'éthologie progresse, en particulier au cours des deux dernières décennies, le rôle des outils d'interaction entre les animaux et les robots est de plus en plus essentiel. Cette importance découle de la nécessité d'expériences contrôlées, répétitives, répétibles et de longue durée, qui non seulement soulagent les chercheurs humains des tâches fastidieuses, mais offrent également de nouvelles perspectives sur les règles qui régissent le comportement collectif. De tels dispositifs peuvent s'infiltrer dans des groupes d'animaux (dans ce cas, des poissons) et s'engager dans des interactions en boucle fermée, provoquant des réponses allant du comportement biomimétique à un comportement modulé. Cependant, la construction de groupes biohybrides harmonieux d'animaux et de robots est une tâche complexe. Malgré des progrès significatifs dans le domaine, de nombreuses questions restent sans réponse, nécessitant des recherches et un développement supplémentaires tant en robotique qu'en modélisation du comportement collectif.

Cette thèse explore l'intersection des phénomènes de comportement collectif et de la robotique. Elle s'appuie sur les avancées dans la fabrication d'électronique, les algorithmes de pointe, et l'accessibilité accrue à la puissance de calcul qui ont remodelé le domaine, aboutissant à une méthodologie de société mixte qui a atteint des niveaux sans précédent de biomimétisme. Dans un premier temps, une exploration approfondie a été menée pour comprendre comment les groupes de poissons interagissent avec des agents artificiels. Nous avons montré que les modèles présentant des interactions actives et bidirectionnelles conduisent à une probabilité plus élevée d'intégration de l'agent artificiel dans le groupe de poissons. Cette réalisation a souligné la nécessité de passer à des modèles d'interaction plus détaillés et précis qui peuvent résister aux comparaisons avec les interactions spontanées des poissons. Pour pallier ces limitations, nous avons mené une étude approfondie sur les facteurs de conception clés qui améliorent la performance des modèles d'interaction sociales, révélant également le besoin d'un ensemble rigoureux de mesures de référence spatio-temporelles. Cet ensemble assure que ces modèles génèrent avec succès des dynamiques sociales réalistes à court et à long terme. Cependant, ces modèles ont mis à jour un problème d'ingénierie secondaire. Transférer des modèles haute fidélité à la réalité exige un équipement robotique agile et réactif, une exigence non satisfaite par l'état de l'art actuel. Ainsi, dans cette thèse, nous avons conçu un nouveau cadre d'études, comprenant un dispositif expérimental, un robot mobile, et des logiciels auxiliaires (par exemple, pour le contrôle du robot, l'intelligence artificielle et les modèles comportementaux analytiques...) pour surmonter ces limitations. Cette approche nous a ensuite permis de faire des avancées substantielles vers la réduction

du “fossé du biomimétisme” en transférant des modèles de haute fidélité de la simulation à la réalité. L’aboutissement de cette thèse décrit notre succès à combler progressivement ce fossé, démontrant une similitude sans précédent entre les simulations, les expériences d’interaction biohybride et spontanée entre poissons uniquement. En mettant en open source l’intégralité des outils logiciels et matériels développés, nous visons à poser une base solide pour les futures recherches dans le domaine de l’interaction entre les robots et les animaux. En conclusion, cette thèse contribue de manière significative à notre compréhension des règles sous-jacentes qui régissent le comportement collectif. Nous espérons qu’elle ouvrira la voie à la création de véritables groupes biohybrides et posera les bases pour de futures explorations des modèles d’interactions sociales et de leur exploitation en tant que contrôleurs de robots.

Mots-clés : Interaction animal-robot, éthorobotique, intelligence artificielle, apprentissage automatique, apprentissage profond, systèmes biohybrides, biomimétisme, comportement collectif, interactions sociales, modélisation de systèmes complexes, contrôle robotique en temps réel, écart réalité-simulation.

Riassunto

Con l'avanzare del campo dell'etologia, soprattutto negli ultimi due decenni, il ruolo degli strumenti di interazione tra animali e robot è diventato sempre più essenziale. Questa importanza deriva dalla necessità di esperimenti controllati, ripetitivi, ripetibili e di lunga durata, che non solo alleviano i ricercatori umani dai compiti noiosi, ma offrono anche nuove intuizioni sulle regole che governano il comportamento collettivo. Tali dispositivi possono infiltrarsi in gruppi di animali (in questo caso, pesci) e impegnarsi in interazioni a ciclo chiuso, suscitando risposte che vanno dal comportamento biomimetico al comportamento modulato. Tuttavia, costruire gruppi bioibridi armoniosi di animali e robot è un compito complesso. Nonostante significativi progressi nel campo, molte domande rimangono senza risposta, richiedendo ulteriori ricerche e sviluppi sia in robotica che nella modellazione del comportamento collettivo.

Questa tesi esplora l'intersezione tra fenomeni di comportamento collettivo e robotica. Si avvale dei progressi nella produzione di elettronica, algoritmi all'avanguardia e l'aumentata accessibilità alla potenza di calcolo che hanno ridisegnato il campo, portando a una metodologia di società mista che ha raggiunto livelli senza precedenti di biomimetismo. Inizialmente, è stata condotta un'indagine esaustiva per capire come i gruppi di pesci interagiscono con gli agenti artificiali. Abbiamo dimostrato che i modelli che mostrano interazioni attive e bidirezionali portano a una maggiore probabilità di integrazione dell'agente artificiale nel gruppo di pesci. Questa constatazione ha sottolineato la necessità di passare a modelli di interazione più dettagliati e accurati che possano resistere ai confronti con le interazioni spontanee dei pesci. Per affrontare tali limitazioni, abbiamo condotto uno studio approfondito sui fattori di progettazione chiave che migliorano le prestazioni dei modelli di interazione sociali, rivelando anche la necessità di un rigoroso insieme di metriche di benchmarking spazio-temporali. Questo insieme garantisce che i modelli generino con successo dinamiche sociali realistiche a breve e lungo termine. Tuttavia, questi modelli hanno rivelato un secondo problema ingegneristico. Trasferire modelli ad alta fedeltà alla realtà richiedeva attrezzature robotiche agili e reattive, un requisito non soddisfatto dallo stato dell'arte attuale. Di conseguenza, in questa tesi, abbiamo progettato un nuovo framework, che comprende un setup sperimentale, un robot mobile e software ausiliario (ad es. per il controllo del robot, intelligenza artificiale e modelli comportamentali analitici...) per superare queste limitazioni. Questo approccio ci ha successivamente permesso di fare passi significativi verso il colmare il "divario biomimetico" trasferendo modelli ad alta fedeltà dalla simulazione alla realtà. Il culmine di questa tesi delinea il nostro successo nel colmare progressivamente questo divario, dimostrando una somiglianza senza precedenti tra simulazioni, esperimenti di interazione bioibrida e spontanea

solo tra pesci. Rendendo *open source* l'intero insieme degli strumenti software e hardware sviluppati, miriamo a gettare le basi solide per future ricerche nel campo dell'interazione tra robot e animali.

In conclusione, questa tesi contribuisce in modo significativo alla nostra comprensione delle regole che governano il comportamento collettivo. Speriamo che possa aprire la strada alla creazione di veri e propri gruppi bioibridi e preparare il terreno per future esplorazioni dei modelli di interazione sociale e del loro utilizzo come controllori di robot.

Parole chiave: Interazione animale-robot, etorobotica, intelligenza artificiale, apprendimento automatico, apprendimento profondo, sistemi bioibridi, biomimetica, comportamento collettivo, interazioni sociali, modellazione di sistemi complessi, controllo robot in tempo reale, divario realtà-simulazione.

Zusammenfassung

Mit dem Fortschreiten der Ethologie, insbesondere in den letzten zwei Jahrzehnten, hat die Rolle von Tier-Roboter-Interaktionstools zunehmend an Bedeutung gewonnen. Diese Bedeutung ergibt sich aus der Notwendigkeit von kontrollierten, wiederholten und langandauernden Experimenten, die nicht nur menschliche Forscher von mühsamen Aufgaben entlasten, sondern auch neue Einblicke in die Regeln bieten, die kollektives Verhalten bestimmen. Solche Geräte können sich in Tiergruppen (in diesem Fall Fische) einschleusen und in geschlossenen Interaktionsschleifen agieren, die von biomimetischem bis zu moduliertem Verhalten reichen. Dennoch ist die Bildung harmonischer biohybrider Gruppen aus Tieren und Robotern eine komplexe Aufgabe. Trotz erheblicher Fortschritte in diesem Bereich bleiben viele Fragen unbeantwortet, was weitere Forschung und Entwicklung sowohl in der Robotik als auch in der Modellierung kollektiven Verhaltens erfordert.

Diese Dissertation untersucht die Schnittstelle zwischen kollektiven Verhaltensphänomenen und Robotik. Sie nutzt die Fortschritte in der Elektronikfertigung, modernste Algorithmen und den erhöhten Zugang zu Rechenleistung, die das Feld geprägt haben, was zu einer gemischten Gesellschaftsmethodik führte, die beispiellose Biomimetik-Niveaus erreichte. Zunächst wurde eine umfassende Untersuchung durchgeführt, um zu verstehen, wie Fischgruppen mit künstlichen Agenten interagieren. Wir zeigten, dass Modelle, die aktive, bidirektionale Interaktionen aufweisen, eine höhere Wahrscheinlichkeit haben, den künstlichen Agenten in die Fischgruppe zu integrieren. Diese Erkenntnis unterstrich die Notwendigkeit des Übergangs zu detaillierteren und genaueren Interaktionsmodellen, die den spontanen Fischinteraktionen standhalten können. Um diese Einschränkungen zu beheben, führten wir eine umfangreiche Studie zu den Schlüsseldesignfaktoren durch, die die Leistung von sozialen Interaktionsmodellen verbessern, und zeigten auch die Notwendigkeit eines strengen raumzeitlichen Benchmarking-Maßstabs auf. Dieser stellt sicher, dass diese Modelle erfolgreich realistische soziale Dynamiken kurz- und langfristig erzeugen. Dennoch offenbarten diese Modelle ein sekundäres Ingenieurproblem. Die Übertragung von Hochleistungsmodellen zurück in die Realität erforderte hoch agile und reaktionsfähige robotische Ausrüstungen, eine Anforderung, die vom aktuellen Stand der Technik nicht erfüllt wurde. In dieser Dissertation entwickelten wir daher einen neuen Rahmen, einschließlich eines experimentellen Setups, eines mobilen Roboters und unterstützender Software (z. B. für Robotersteuerung, künstliche Intelligenz und analytische Verhaltensmodelle), um diese Einschränkungen zu überwinden. Dieser Ansatz ermöglichte es uns, bedeutende Fortschritte bei der Überbrückung der "Biomimetik-Lücke" zu machen, indem wir Hochleistungsmodelle von der Simulation in die Realität übertrugen. Der

Abschluss dieser Arbeit skizziert unseren Erfolg, diese Lücke schrittweise zu überbrücken und zeigt eine beispiellose Ähnlichkeit zwischen Simulationen, biohybriden und spontanen reinen Fischinteraktionsexperimenten. Durch die Open-Source-Bereitstellung aller entwickelten Software- und Hardwaretools wollen wir eine solide Grundlage für zukünftige Forschungen im Bereich der Tier-Roboter-Interaktion schaffen.

Abschließend trägt diese Dissertation erheblich zum Verständnis der zugrunde liegenden Regeln bei, die das kollektive Verhalten bestimmen. Wir hoffen, dass sie den Weg für die Schaffung von echten biohybriden Gruppen ebnet und die Bühne für zukünftige Untersuchungen von sozialen Interaktionsmodellen und deren Nutzung als Robotersteuerungen bereitet.

Schlüsselwörter: Tier-Roboter-Interaktion, künstliche Intelligenz, maschinelles Lernen, tiefes Lernen, biohybride Systeme, Biomimetik, kollektives Verhalten, soziale Interaktionen, komplexe Systemmodellierung, Echtzeit-Robotersteuerung.

Contents

Acknowledgements	i
Abstract (English/Français/Italiano/Deutsch)	v
List of Figures	xix
List of Tables	xxiii
I Motivation and thesis outline	1
1 Introduction	3
1.1 Collective animal behavior	4
1.2 Models of collective behavior	4
1.3 Robots in ethology	5
1.3.1 Behavioral robot controllers	7
1.3.2 The biomimicry gap	8
1.4 Fish as a model for understanding social dynamics	10
1.4.1 Burst-and-coast swimming	10
1.5 Biohybrid experiments: A temporal-technological duality perspective	11
2 Thesis outline	15
2.1 Thesis objectives	15
2.2 Summary of contributions and thesis organization	16
2.2.1 Part II	16
2.2.2 Part III	16
2.2.3 Part IV	17
2.2.4 Part V	17
II How to design models that encode animal behavior	19
3 Establishing a baseline of fish-robot dynamics	21
3.1 Introduction	21
3.2 Background	23
3.2.1 Animals	23
	xiii

3.2.2	Experimental setup	23
3.2.3	Robotic system for closed-loop zebrafish-robot interactions	23
3.2.4	Control and tracking software	24
3.3	Experimental procedure	26
3.4	Behavioral models	26
3.5	Data treatment	30
3.5.1	Data filtering	30
3.5.2	Data analysis	31
3.5.3	Statistical tests	34
3.6	Behavioral model assessment	34
3.6.1	Average angular distance	34
3.6.2	Collective U-turns	34
3.6.3	Transfer entropy	37
3.7	Conclusion	39

4 Designing and benchmarking social interaction models: from analytical to machine learning models, and from short to long timescale dynamics generation 43

4.1	Introduction	44
4.2	Experimental data	44
4.3	Quantification of individual and collective behavior in pairs of fish	45
4.4	Analytical and deep learning models of fish behavior	47
4.5	Assessing the generative capabilities of biomimetic social interaction models	54
4.5.1	Quantification of the instantaneous individual behavior	54
4.5.2	Quantification of the instantaneous collective behavior	56
4.5.3	Quantification of temporal correlations	57
4.6	Benchmarking the DLI against a similarly purposed neural network	58
4.6.1	Comparing the short- and long-term performance of DLI and D-LSTM	60
4.6.2	Quantification of instantaneous individual behavior	60
4.6.3	Quantification of instantaneous collective behavior	60
4.6.4	Quantification of time correlation	61
4.6.5	Comparing the short-term performance of DLI and D-LSTM	62
4.7	Validation of the DLI's scalability to other fish species	63
4.7.1	Quantification of the instantaneous individual behavior	63
4.7.2	Quantification of the instantaneous collective behavior	63
4.7.3	Quantification of temporal correlations	63
4.8	Assessing neural networks with no memory components	64
4.8.1	Quantification of the instantaneous individual behavior	66
4.8.2	Quantification of the instantaneous collective behavior	66
4.8.3	Quantification of temporal correlations	66
4.8.4	Complementary analyses	66
4.9	Conclusion	67

III	What makes for an efficient biohybrid interaction framework	69
5	Introducing a novel biohybrid interaction framework	71
5.1	Introduction	71
5.1.1	Chapter outline	72
5.2	Experimental Setup	73
5.3	LureBot	75
5.3.1	Mechanics	75
5.3.2	Electronics	77
5.4	Software architecture	78
5.5	Lures	82
5.5.1	Biomimetic (BM) Lure	83
5.5.2	Disc-shaped (DS) Lure	83
5.6	Trajectory generators	84
5.6.1	Circular trajectory (CT)	84
5.6.2	Eightfold rose trajectory (R8T)	84
5.6.3	Biomimetic interaction model (BIM)	85
5.7	Protocols and Data Treatment	85
5.7.1	Experimental Procedure	85
5.7.2	Data Filtering	86
5.8	Assessing the dynamics of biohybrid groups	87
5.8.1	Experiment № 1: Open-loop dynamics	88
5.8.2	Experiment № 2: Closed-loop dynamics	92
5.9	Conclusion	97
IV	Transferring social interaction models from simulation to reality	99
6	Evaluating the biomimicry gap in biohybrid pairs driven by a deep learning model	101
6.1	Introduction	101
6.2	Methods	103
6.2.1	Real-time tracking and robot control	103
6.2.2	Deep Learning Interaction model	103
6.2.3	Evaluating the short- and long-term interactions	104
6.3	Dynamics of pairs of agents	106
6.4	Results	106
6.4.1	Instantaneous individual quantities	107
6.4.2	Instantaneous collective quantities	109
6.4.3	Temporal correlation quantities	111
6.4.4	Complementary results for DLiv2 simulated pairs	113
6.5	Conclusion	113

7 Investigating the biomimicry gap with varying sizes of biohybrid groups driven by a high-fidelity model	117
7.1 Introduction	117
7.2 Methods	119
7.2.1 Multiagent simulations with the analytical burst-and-coast model	119
7.2.2 Real-time robot control with the analytical burst-and-coast model	120
7.2.3 Data processing	121
7.2.4 Experimental procedure	122
7.2.5 Evaluating the short- and long-term interactions	123
7.3 Results	124
7.3.1 Experiments with a single agent	124
7.3.2 Experiments with agent pairs	127
7.3.3 Experiments with groups of five agents (1 agent in perception field)	131
7.3.4 Experiments with groups of five agents (2 agents in perception field) . . .	137
7.4 Conclusion	141
V Conclusion	145
8 Contributions and research outlook	147
8.1 How can social interactions be modeled in an informative and expressive manner, with the purpose of constructing cohesive mixed (<i>i.e.</i> , biohybrid) groups of animals and robots?	149
8.1.1 (RQ1) How can we explore the bidirectionality of fish-robot interactions to measure how they facilitate the integration of a robot into fish groups?	149
8.1.2 (RQ2) How can we design and assess bidirectional models of high fidelity (<i>w.r.t.</i> , to observations from the spontaneous animal interactions)?	149
8.1.3 (RQ3) Which key elements make for robotic devices that can closely mimic fish motion and dynamics to subsequently transfer such models “back to reality”?	150
8.1.4 (RQ4) How can we measure, and what is the gap between real experiments and simulations with such models?	151
8.1.5 Technical limitations and future work	152
8.1.6 Modelling limitations and future work	153
8.1.7 Behavioral experiment, scalability limitations, and future work	153
8.2 Final words	155
A Appendix for Chapter 3 “Establishing a baseline of fish-robot dynamics”	157
B Appendix for Chapter 4 “Designing and benchmarking social interaction models: from analytical to machine learning models, and from short to long timescale dynamics generation”	165

C Appendix for Chapter 5 “Introducing a novel biohybrid interaction framework”	171
C.1 Scaling up from the FishBot v4.4 to the LureBot	174
D Appendix for Chapter 6 “Evaluating the biomimicry gap in biohybrid pairs driven by a deep learning model”	177
D.1 DLiv2 Supplementary figures (standard deviation)	177
D.2 DLiv2 Supplementary figures (confidence interval)	179
D.3 DLI biohybrid pair supplementary figures (confidence interval)	179
D.4 Supplementary Tables	183
E Appendix for Chapter 7 “Investigating the biomimicry gap with varying sizes of biohybrid groups driven by a high-fidelity model”	187
E.1 Experiments with a single agent	187
E.2 Experiments with agent pairs	188
E.3 Experiments with groups of five agents	190
E.3.1 Experiments with groups of five agents (1 agent in perception field) . . .	190
E.3.2 Experiments with groups of five agents (2 agents in perception field) . .	192
E.4 Tables for Chapter 7	194
Bibliography	214
Curriculum Vitae	215

List of Figures

1.1	Illustration of the sources of the biomimicry gap	8
1.2	Zebra <i>Danio rerio</i>	10
1.3	Rummy-nose tetra <i>Hemigrammus rhodostomus</i>	10
3.1	Experimental setup for fish-robot experiments.	24
3.2	FishBot and biomimetic zebrafish lure.	25
3.3	Closed-loop FishBot control.	29
3.4	Average angular distance metric.	35
3.5	Average number of collective U-turns per minute over all replicates.	36
3.6	U-turn initiation success rates.	37
3.7	Average transfer entropy (TE) for heading direction time-series.	38
4.1	Individual and collective variable visualization.	46
4.2	Deep Learning Interaction (DLI) network architecture.	51
4.3	Probability density functions (PDF) of observables characterizing individual behavior.	55
4.4	Probability density functions (PDF) of observables characterizing collective behavior.	56
4.5	Observables quantifying temporal correlations in the system.	57
4.6	Probability density functions (PDF) of all D-LSTM observables.	61
4.7	Short-scale trajectory comparison of the DLI and D-LSTM.	62
4.8	Probability density functions (PDF) of all zebrafish observables.	64
4.9	Probability density functions (PDF) of all observables (simulations produced by neural networks with no memory).	65
5.1	Behavioral Observation & Biohybrid Interaction framework (experimental setup and software) overview.	74
5.2	Description of the LureBot and its internal parts.	76
5.3	Lure construction process.	83
5.4	Open-loop trajectory generator patterns.	84
5.5	Short trajectories from the open-loop experiments.	88
5.6	Motion profile when a fish interacts with a passive DS or BM lure (CT).	89
5.7	Motion profile when a fish interacts with a passive DS or BM lure (R8T).	90
5.8	Single agent motion profile.	92

5.9	Short trajectories from the closed-loop experiments.	93
5.10	Motion profile for pairs of agents.	94
5.11	Motion profile for groups of 5 agents.	96
6.1	Closed-loop robot control with Deep Learning Interaction (DLI) model.	103
6.2	Individual and collective variables.	105
6.3	Instantaneous individual quantities.	107
6.4	Instantaneous collective quantities.	109
6.5	Temporal correlation quantities.	111
6.6	Quantification of biomimicry gap values between different experiments.	114
7.1	Closed-loop analytical burst-and-coast control of the robot.	121
7.2	Instantaneous individual quantities.	124
7.3	Temporal correlation quantities.	126
7.4	Instantaneous individual quantities.	127
7.5	Instantaneous collective quantities.	129
7.6	Temporal correlation quantities.	130
7.7	Instantaneous individual quantities.	132
7.8	Instantaneous collective quantities.	134
7.9	Temporal correlation quantities.	135
7.10	Instantaneous individual quantities.	137
7.11	Instantaneous collective quantities.	139
7.12	Temporal correlation quantities.	140
7.13	Quantification of biomimicry gap values between different experiments.	142
8.1	Biomimicry gap study plan.	148
A.1	Duration of consecutive movement towards a direction (across all 10 replicates).	160
A.2	Duration of consecutive movement towards a direction (across all 10 replicates).	160
A.3	Duration of consecutive movement towards a direction (across all 10 replicates).	161
A.4	Duration of consecutive movement towards a direction (across all 10 replicates).	161
A.5	Duration of consecutive movement towards a direction (across all 10 replicates).	162
A.6	Duration of consecutive movement towards a direction (across all 10 replicates).	162
A.7	Duration of consecutive movement towards a direction (across all 10 replicates).	163
A.8	Duration of consecutive movement towards a direction (across all 10 replicates).	163
B.1	DLI simulations without a rejection procedure in place.	168
C.1	LureBot <i>versus</i> FishBot v4.4 performance.	174
D.1	Instantaneous individual quantities (DLIv2).	177
D.2	Instantaneous collective quantities (DLIv2).	178
D.3	Temporal correlation quantities (DLIv2).	178
D.4	Instantaneous individual quantities (DLIv2).	179

D.5	Instantaneous collective quantities (DLiv2).	179
D.6	Temporal correlation quantities (DLiv2).	180
D.7	Instantaneous individual quantities.	180
D.8	Instantaneous collective quantities.	181
D.9	Temporal correlation quantities.	182
E.1	Instantaneous individual quantities.	187
E.2	Temporal correlation quantities.	188
E.3	Instantaneous individual quantities.	188
E.4	Instantaneous collective quantities.	189
E.5	Temporal correlation quantities.	189
E.6	Instantaneous individual quantities.	190
E.7	Instantaneous collective quantities.	191
E.8	Temporal correlation quantities.	191
E.9	Instantaneous individual quantities.	192
E.10	Instantaneous collective quantities.	193
E.11	Temporal correlation quantities.	193

List of Tables

1.1	Partial social interaction model & biohybrid framework literature overview. . .	12
A.1	Post hoc analysis for the average angular distance distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. . .	157
A.2	Post hoc analysis for the average angular distance distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. . .	157
A.3	Post hoc analysis for the U-turn initiation success rate distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. .	158
A.4	Post hoc analysis for the highest U-turn initiation success rate by any one individual distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion.	158
A.5	Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average outgoing TE including all agents)	158
A.6	Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average outgoing TE only for the robot’s contribution)	158
A.7	Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average outgoing TE only for the fish contribution)	159
A.8	Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average incoming TE including all agents – identical to outgoing case for all agents)	159
A.9	Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average incoming TE only for the robot’s contribution)	159
A.10	Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average outgoing TE only for the fish contribution)	159
B.1	Mean and standard deviation values for the experiment.	165
B.2	Mean and standard deviation values for the ABC model.	165

B.3	Mean and standard deviation values for the DLI model.	166
B.4	Mean and standard deviation values for D-LSTM model.	166
B.5	Mean and standard deviation values for the Euclidean distance between prediction and real trajectory (focal individual).	166
B.6	Mean and standard deviation values for the Euclidean distance between prediction and real trajectory (neighboring individual).	166
B.7	Implementation details of the DLI per layer. The 5 columns correspond to the increasing layer count, the type of layer, activation function, number of inputs, and number of outputs, respectively.	167
B.8	Mean and standard deviation values for <i>D. rerio</i> experiment.	167
B.9	Mean and standard deviation values for DLI model (<i>D. rerio</i>).	167
B.10	Neural network grid.	169
C.1	Hellinger distance for pairs of individuals (the LureBot follows a passive circular trajectory).	172
C.2	Hellinger distance for pairs of individuals (the LureBot follows a passive eightfold rose trajectory).	172
C.3	Hellinger distance for single individuals (the LureBot is commanded by the behavioral model).	172
C.4	Hellinger distance for pairs of individuals (the LureBot is commanded by the behavioral model).	173
D.1	Means and standard deviations of quantities.	183
D.2	Hellinger distances for pairs.	184
D.3	Statistical significance tests for the experiment, DLI-SP, and DLI-BP. We ran a Kruskal-Wallis test to measure whether the distributions differ significantly and a Tukey's honestly significant difference (T-HSD) test to test which distributions differ significantly.	185
D.4	Statistical significance tests for the experiment, DLI-SP, and DLIv2-SP. We ran a Kruskal-Wallis test to measure whether the distributions differ significantly and a Tukey's honestly significant difference (T-HSD) test to test which distributions differ significantly.	186
E.1	Means and standard deviations of quantities for single agents	194
E.2	Hellinger distances for single agents	194
E.3	Means and standard deviations of quantities for agent pairs.	195
E.4	Hellinger distances for agent pairs	196
E.5	Means and standard deviations of quantities for groups of five agents (ABC-BG, ABC-SG with $k = 1$)	197
E.6	Hellinger distances for groups of five agents (ABC-BG, ABC-SG with $k = 1$)	198
E.7	Means and standard deviations of quantities for groups of five agents (ABC-BG, ABC-SG with $k = 2$)	199
E.8	Hellinger distances for groups of five agents (ABC-BG, ABC-SG with $k = 2$)	200

Motivation and thesis outline **Part I**

1 Introduction

This section's content is adapted from the following article:

- Papaspyros, V., Bonnet, E., Collignon, B. and Mondada, F., 2019. Bidirectional interactions facilitate the integration of a robot into a shoal of zebrafish *Danio rerio*. *PLoS one*, 14(8), p.e0220559. <https://doi.org/10.1371/journal.pone.0220559> (reuse authorised under the CC 4.0 licence) – Papaspyros et al. (2019)
 - Vaios Papaspyros's contribution: Data curation, formal analysis, investigation, methodology, software, visualization, writing – original draft, writing – review & editing.
- Papaspyros, V., Escobedo, R., Alahi, A., Theraulaz, G., Sire, C. and Mondada, F., 2023. Predicting long-term collective animal behavior with deep learning. *bioRxiv*, pp.2023-02. <https://doi.org/10.1101/2023.02.15.528318> (reuse authorised under a non-exclusive distribution licence) – Papaspyros et al. (2023b)
 - Vaios Papaspyros's contribution: Conception, data curation, formal analysis, investigation, methodology, software, visualization, writing – original draft, writing – review & editing.
- Papaspyros, V., Burnier, D., Cherfan, R., Theraulaz, G., Sire, C. and Mondada, F., 2023. A biohybrid interaction framework for the integration of robots in animal societies. *IEEE Access*. <https://doi.org/10.1109/ACCESS.2023.3290960> (reuse authorised under the CC 4.0 licence) – Papaspyros et al. (2023a)
 - Vaios Papaspyros's contribution: Conception, data curation, formal analysis, investigation, methodology, software, hardware, visualization, writing – original draft, writing – review & editing.
- Papaspyros, V., Theraulaz, G., Sire, C. and Mondada, F., 2023. Quantifying the biomimicry gap in biohybrid systems. *arXiv*. <https://doi.org/10.48550/arXiv.2308.08978> (reuse authorised under the CC 4.0 licence) – Papaspyros et al. (2023c)
 - Vaios Papaspyros's contribution: Conception, data curation, formal analysis, investigation, methodology, software, visualization, writing – original draft, writing – review & editing.

1.1 Collective animal behavior

The complex spatio-temporal structures and interactions observed in groups of social animals have long fascinated and piqued the interest of scientists. How is it that from individual actions we observe the emergence of collective patterns of unique complexity Camazine et al. (2003); Theraulaz et al. (2003)? Or, similarly, how can we disentangle the collective dynamics to discover the individual primitives that led to a collective pattern Lukeman et al. (2010) on a global scale? These simple, yet difficult to answer, questions have triggered numerous studies on social animals throughout the years, some of which remain unanswered.

The resulting domain, that of collective behavior, has become a very active field of research, studying the fundamental mechanisms by which individuals coordinate their actions Sumpter (2010); Krause and Ruxton (2002); Ball (2011) and self-organize Camazine et al. (2003); Couzin et al. (2003). One of the most common forms of collective behavior can be observed in schools of fish and flocks of birds that have the ability to coordinate their movements to collectively escape predator attacks or improve their foraging efficiency Vicsek and Zafeiris (2012); Cavagna et al. (2018). This coordination at the group level mainly results from the social interactions between individuals. Important steps to understand these collective phenomena consist in characterizing these interactions and understanding the way individuals integrate interactions with other group members Deutsch et al. (2020); Herbert-Read (2016); Gautrais et al. (2012); Calovi et al. (2018).

1.2 Models of collective behavior

Notably, nowadays, new tracking techniques and tools for behavioral analysis have been developed that have greatly improved the quality of collective motion data Branson et al. (2009); Dell et al. (2014); Gallois and Candelier (2021); Anderson and Perona (2014); Pérez-Escudero et al. (2014); Romero-Ferrero et al. (2019); Walter and Couzin (2021). In particular, advances in computing have allowed the development of computationally demanding data-oriented model generation techniques Calovi et al. (2018); Escobedo et al. (2020); Jayles et al. (2020); Cazenille et al. (2019); Heras et al. (2019); Costa et al. (2020); Collignon et al. (2017); Aoki (1982); Bode et al. (2010); Reynolds (1987); Couzin et al. (2002); Lopez et al. (2012); Niwa (1996); Vicsek et al. (1995); Bertin et al. (2006); Chaté et al. (2008); Nagai et al. (2015); Gautrais et al. (2012); Zienkiewicz et al. (2015); Mwaffo et al. (2015) and the subsequent simulation of biological models Gilpin et al. (2020); Collignon et al. (2017). This has resulted in more realistic models that attempt to recover the social interactions that govern collective behaviors. Yet, the bottleneck with most of these approaches is that they rely on demanding and laborious mathematical work to obtain the interactions from experimental data, *e.g.*, by means of manually crafting functions that approximate the observed signals, and proceeding to computationally infer the best fit for the function parameters.

An alternative to such analytical models is to exploit machine learning (ML) techniques and let an algorithm learn the interactions directly from data. The know-how required to use these techniques is different from the one needed to design analytical models. Nevertheless, the structure of ML algorithms, *e.g.*, neural network, has an impact on the modeling performance, and requires extensive knowledge on the use and implementation of neural networks Mammadli et al. (2019). Furthermore, modelling a new species requires that we (potentially) retrain the network, but without changing its structure. The downside of this flexibility is that ML models are usually less explainable (“black box”). Yet, recent ML algorithms provide higher-level information mappable to more tangible formats, such as force maps, which show the strength and direction of behavioral changes experienced by an individual when interacting with other individuals in a moving group Heras et al. (2019); Costa et al. (2020). Despite their limited explainability, ML algorithms require only a few biological assumptions. They offer an almost hypothesis-free procedure Valletta et al. (2017) that can even outperform human experts in detecting subtle patterns Marques et al. (2018), making ML a very appealing complementary approach to analytical models.

For both analytical and ML models, several studies evaluate models over *short timescales* and through instantaneous quantities such as speed, acceleration, distance and angle to objects Cazenille et al. (2018b, 2019), or by measuring the error between predictions and ground truth Heras et al. (2019); Alahi et al. (2016); Kothari et al. (2021). Only more recently, long timescales have also been considered Jayles et al. (2020), that is, sufficiently long timescales to observe and verify the emergence and evolution of collective dynamics. Whereas the notion of long-term timescale may vary depending on the animal and type of dynamics under study, it is well established that collective behavior spans multiple timescales Ioannou and Laskowski (2023), thus, a model’s effectiveness in the short term does not guarantee similar performance over long timescales. Therefore, there are models that only exhibit accuracy in the short term. When these models are used for extended simulations (typically in the range of 1-3h for social interaction studies), a divergence from the actual observations of an animal’s inherent behavior may become apparent. The shortcomings of current social interaction models can be summarized as follows: 1) they are often employed to predict short-term quantities with minimal or no evaluation over long timescales, and 2) even when these models are used for extended simulations, the assessment often lacks rigor in terms of temporal quantities and focuses predominantly on instantaneous spatial ones.

1.3 Robots in ethology

While some knowledge can be obtained by observing and modelling the innate interactions within animal groups, in more recent years, scientists are presented with a new opportunity for studying collective behavior. That is, the use of robots and robotic devices to probe, control or mimic the behavior of animals. As a result of the rapid development of micro-electronic and embedded systems, scientists can now use those devices to stimulate the animals in such a way as to validate theoretical assumptions which were traditionally based on observation. In turn,

this allows for experiments with closed-loop interaction models stimulated in real-time with true feedback from the animal. To this day, these advancements, along with advancements in computer hardware and algorithms, have facilitated the development of smaller and more agile robotic systems that can be operated with improved perception systems and increasingly sophisticated motion models. In turn, this has allowed scientists to put these theoretical models to the test in real-life scenarios and with true feedback from the animals, in order to study their collective behavior. Some of them Pillot et al. (2010); Correll et al. (2008); Worm et al. (2018) relied on the use of teleoperated devices that produce signals (*e.g.*, visual, acoustic, electric) to attract or repel the animals, others rely on mobile robots that are not explicitly mimicking the animal under study (*e.g.*, it could be a sheepdog among sheep Strömbom et al. (2014)) and, in cases where visual stimulus is critical, some relied on mimetic lures, that is, on lures that mimic the shape, size, and appearance or behavior Faria et al. (2010); Bonnet et al. (2014); Kopman and Porfiri (2013); Halloy et al. (2007); Marras and Porfiri (2012). These studies demonstrated that artificial agents able to perceive and emit pertinent and adapted signals can influence and control self-organized choices by mixed groups of animals and robots Romano et al. (2019).

Over the years, the combination of these advances, along with a continued interest in deciphering the rules that govern collective behavior, has led to the design of a plethora of biohybrid systems, spanning from groups of fish Faria et al. (2010); Swain et al. (2011); Landgraf et al. (2013); Bonnet et al. (2014); Landgraf et al. (2016); Ruberto et al. (2016); Bartolini et al. (2016); Bonnet et al. (2018); Porfiri (2018); Romano and Stefanini (2021, 2022b,a), bees Landgraf et al. (2010); Stefanec et al. (2017); Griparić et al. (2017); Barmak et al. (2023), cockroaches Halloy et al. (2007), locusts Romano et al. (2020), rats Shi et al. (2014), and birds Jolly et al. (2016); Folkertsma et al. (2017); Gribovskiy et al. (2010, 2018); Simon et al. (2023). While these systems primarily serve as a means to examine animal behavior, they also offer a glimpse into potential strategies for preserving ecosystems, thereby contributing to environmental conservation efforts Ilgün et al. (2021); Schmickl et al. (2021).

In these systems, the investigation of collective responses within biohybrid groups typically involves closed-loop mechanisms that encompass environmental sensing (*e.g.*, tracking animal positions, temperature...), decision-making processes (*e.g.*, utilizing computational behavioral models), and targeted communication channels specific to the species under study. For example, bees may respond to air currents or hive temperature fluctuations Bonnet et al. (2012); Stefanec et al. (2017); Griparić et al. (2017); Barmak et al. (2023), while fish can be influenced by lures or visual stimuli Faria et al. (2010); Swain et al. (2011); Kopman et al. (2013); Landgraf et al. (2013); Bonnet et al. (2014); Landgraf et al. (2014, 2016); Bonnet et al. (2018); Porfiri (2018); Chemtob et al. (2020); Romano and Stefanini (2021, 2022b); Maxeiner et al. (2023).

1.3.1 Behavioral robot controllers

Although robots have proved to be a valuable asset in this context, they have also uncovered the need for adaptive behavioral models that can drive the robot's movement in an animal-like manner. From the first experiments of sheep herding Vaughan et al. (2000); Strömbom et al. (2014) to the more recent fish-robot interaction experiments Bonnet et al. (2018, 2019); Romano and Stefanini (2021, 2022a); Porfiri et al. (2019), scientists have been attempting to uncover the underlying dynamics that allow for modulating the animal behavior. This requires models capable of adapting to the collective dynamics (either to influence or mimic the decision-making process), despite the rapid or more gradual changes that might deem a non-adaptive model completely inadequate. This is indeed imperative to the development of truly autonomous animal-like robots, in the behavioral control sense, that can blend in the animal group for long periods of time by self-regulating their interactions with their social companions. However, the design of such generic models is particularly difficult due to 1) limited number of high-fidelity models that reproduce the underlying dynamics with high accuracy, 2) unaccounted for physical properties of robotic systems in the model, and 3) lack of adaptation to animals', here fish's, behavioral plasticity de Lourdes Ruiz-Gomez et al. (2008); Suriyampola et al. (2023).

A review of the literature of the animal-robot interaction domain, covers two general categories of control schemes for the robot(s): (1) open-loop, and (2) closed-loop control. The former consists of behavioral models implemented on a robot, that are inherently not adaptive, do not react or take into account the behavior of neighboring animals Bonnet et al. (2018); Ladu et al. (2015a); Phamduy et al. (2014); Polverino et al. (2013); Abaid et al. (2012a); Ladu et al. (2015b); Polverino et al. (2012); Spinello et al. (2013); Polverino and Porfiri (2013a); Ruberto et al. (2016); Bartolini et al. (2016); Kruusmaa et al. (2016); Abaid et al. (2013); Butail et al. (2013, 2014b) or they are based on replaying pre-recorded trajectories with respect only to the setup dimensions (*i.e.*, the robot does not explicitly model interactions with the setup boundaries, but will purposely avoid collisions). On the other end, closed-loop models are designed to provide means for the robot to perceive, locally or globally, its environment Faria et al. (2010); Swain et al. (2011); Landgraf et al. (2013); Cazenille et al. (2018c) (*e.g.*, fish, walls...). These models can be further classified as passive or active depending on their tendency to naively follow the collective decisions (*e.g.*, by mirroring their movements) or, in the second case, to take initiative and influence the shoal members given some criteria. Thus, creating a biohybrid group requires careful design and attention to animal-specific communication cues Mondada et al. (2013).

Despite a plethora of proposed behavioral robot controllers, similarly to experiments based on simulation, the majority makes use of simplified rules of interaction to command the robot. These rules often lack comprehensive validation, but are computationally efficient for use in real-time systems.

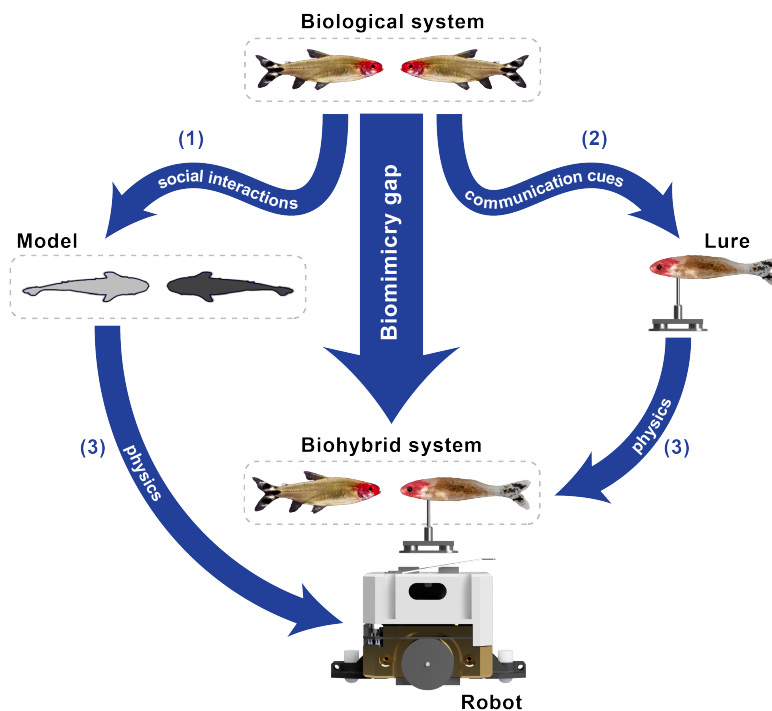


Figure 1.1: **Illustration of the sources of the biomimicry gap.** (1) The modeling phase may introduce a first source of discrepancy between the effect of social interactions on the swimming patterns in the model and the ones observed in real fish. (2) A second source of discrepancy between the visual appearance of the lure and that of a real fish might introduce imperfect communication cues and elicit unrealistic behavioral responses from neighboring organisms. (3) Finally, a third source of discrepancy between the characteristics of the movement produced by the model and its realization by the lure occurs when the numerical model is transferred to real-world scenarios due to the physics constraints that were not accounted for in the model. *H. rhodostomus* photo was taken by David Villa ScienceImage/CBI/CNRS, Toulouse.

1.3.2 The biomimicry gap

Nonetheless, akin to the well known reality gap observed when transferring simulated robot controllers to real-world applications Jakobi et al. (1995); Jakobi (1997); Mouret and Chatzilygeroudis (2017); Martinoli et al. (1999), biohybrid systems may exhibit a similar discrepancy, but more complex because of including biological entities interacting among them, with the environment and with the robot. This gap arises from the combination of the following elements (also depicted in Fig. 1.1):

1. Subtle behavioral patterns that social interaction models fail to capture. This causes a discrepancy in the decision-making process that effectively adheres to different (with varying degree) social dynamics and generates short- and long-term spatio-temporal patterns that may present differences from fish-only interactions in the same environ-

ment. We call this the *social interaction gap*;

2. the fidelity of the communication cues exhibited by the artificial device(s) toward the animal, which we call the *communication cue gap*. In this thesis' context, the level of biomimicry of artificial lures, which affects the way fish perceive and interact with it. This may not be relevant for some animal species, but is indeed shown to play a crucial role in fish-robot interactions Romano and Stefanini (2022a);
3. physics related to the operation of the robot in real life that were not accounted for in models. This refers to the discrepancies discussed in Jakobi et al. (1995); Jakobi (1997); Mouret and Chatzilygeroudis (2017), and in the context of modelling social interactions, especially, there is typically a complete lack of any physical modelling in social interaction models. We refer to this discrepancy as the *physics gap*.

We refer to the global resulting gap by coining the term “biomimicry gap”. This term encompasses all of the above aspects, illustrating the complexity of the phenomena that goes far behind the typical definition of the reality gap as it currently exists in the literature. In that respect, and to the best of our knowledge, the feasibility of bridging the social interaction reality gap by conducting long experiments in both simulated and real environments and comparing their outcomes has not been conclusively addressed. Throughout this thesis, we attempt to understand and implement ways to further bridge the biomimicry gap by advancing on all three of its stems.

The approach we followed in this thesis to bridge this gap is summarized in the 4-step methodology summarized below:

1. Measure the effect varying degrees of interactivity in models affects the collective dynamics. This step establishes a baseline for the biomimicry gap (discussed in Chapter 3).
2. Design and rigorously benchmark models in simulation to create high-fidelity models (discussed in Chapter 4).
3. Design and construct high-fidelity lures/replicas that are specially designed for the type of experiment (addressed in Chapter 5). For instance, mimicking fish behavior requires highly biomimetic lures of the same species, predation studies require good quality replicas of a fish's natural predator...).
4. Model the physics of the system in question. In this thesis, we do not directly answer this question, but we indirectly address it by using a physical system to replicate the dynamics that are validated in simulation (see Chapters 6 and 7). Then, we proceed to trace the potential sources of the gaps that arise. Intuitively, if we were to alleviate the effect of the imperfect models and communication cues, we would be able to directly measure the physics gap.

1.4 Fish as a model for understanding social dynamics

Many fish species exhibit social skills that lead to complex dynamics. Besides the social interaction studies, fish are an ideal candidate animal for laboratory-confined experiments, mainly due to their size and easy to replicate housing conditions. For this thesis work, we chose the zebrafish *Danio rerio* (wild-type, AB strain) Nasevicius and Ekker (2000) and the rummy-nose tetras *Hemigrammus rhodostomus* as our models.

The zebrafish is a freshwater fish species, very commonly found in aquariums. Notably, the zebrafish have 70% genetic homology with humans, which makes them ideal for drug development, cancer and clinical studies or for studies concerning the regenerative abilities. In terms of behavioral studies, the zebrafish have already served as a model for several studies on information propagation Crosato et al. (2018); Butail et al. (2016); Ruberto et al. (2016); Porfiri (2018) and animal-robot interaction Bonnet et al. (2014, 2018); Cazenille et al. (2017, 2018c); Ruberto et al. (2016); Porfiri (2018).



Figure 1.2: **Zebra *Danio rerio***.

Similarly, *H. rhodostomus* exhibit a strong social component, stronger than that of zebrafish Escobedo et al. (2020). Thus, they are an ideal fit to validate how social interaction models allow artificial systems (*e.g.*, mobile robot-lure systems) to blend into living fish groups with minimal to no effect on the collective dynamics.



Figure 1.3: **Rummy-nose tetra *Hemigrammus rhodostomus***. Photo by David Villa ScienceImage/CBI/CNRS, Toulouse.

1.4.1 Burst-and-coast swimming

H. rhodostomus or *D. rerio*, like many other fish species, move in a *burst-and-coast* manner, meaning that their swimming pattern consists of a sequence of abrupt accelerations each followed by a longer gliding period (see Fig. 4.1B), during which a fish moves more or less in a straight line (see Fig. 4.4C). The kicking instants observed in the curve of the speed can be interpreted as decision times when the fish potentially initiates a change of direction. In *H. rhodostomus*, the mean time interval between kicks and the typical kick length were experimentally found to be close to 0.5 s and 7 cm, respectively Calovi et al. (2018), and close

to 0.8 s and 8 cm for zebrafish. When confined in circular tanks, fish tend to swim close to the curved wall because their trajectory is made of quasi straight segments with limited variance of the heading angle between kicks, hence preventing the fish to escape from the tank wall (unless when a rare large heading angle change occurs) Calovi et al. (2018). When swimming in groups, *H. rhodostomus* and *D. rerio* tend to remain close to each other, especially when the number of fish in the tank is small. Naturally, the characteristics and variability of this motion and the emerging social interactions, is inherently linked with the experimental environment and conditions under which these motions are observed Li et al. (2021); Xue et al. (2023).

1.5 Biohybrid experiments: A temporal-technological duality perspective

As previously highlighted, biohybrid systems, comprising behavioral models and the robotic systems they command, have proved insightful for fish behavior studies, with a focus on various domains including biomimetics Landgraf et al. (2014); Cazenille et al. (2018c), predation Abaid et al. (2013), phenotypics Romano and Stefanini (2022a), and behavioral modulation Bonnet et al. (2018); Chemtob et al. (2020), among others. Nonetheless, only a limited number of studies have concentrated on evaluating the degree to which such systems gain acceptance and genuinely become indistinguishable from living animals (refer to Table 1.1) — a pivotal aspect that enables researchers to derive more confident and relevant conclusions regarding the origins of responses elicited by these systems. Specifically, it is pivotal to discern whether potential phenotypic and behavioral discrepancies between the robotic system and the study animal significantly widen the biomimicry gap, well before the effects of deliberately targeted stimuli are quantified. This issue casts biohybrid experiments in a nuanced light, challenging the degree to which a modern robotic system can furnish clear, unbiased insights into the decision-making processes of animals, specifically, fish.

Moreover, it has been observed that fish-robot systems generally engage in experiments over relatively brief durations, with a typical session spanning between 5-30 min and only seldom extending from 30 to 60 min. This approach may sufficiently serve the objectives of certain experiments, particularly those intended to investigate the instantaneous or short-term behaviors of fish to very specific stimuli Abaid et al. (2012a); Polverino and Porfiri (2013a); Polverino et al. (2013); Butail et al. (2014b,a); Phamduy et al. (2014); Ladu et al. (2015b,a); Donati et al. (2016), and it is plausible that associated studies can exclusively rely on short-term timescale metrics. Consequently, such experiments have traditionally employed either open-loop and fixed pattern models (refer to Table 1.1) or closed-loop models of low reactivity, primarily with the purpose of initiating a predetermined pattern of response subject to specific fish states.

Conversely, to explore the spontaneous movement and social interactions among fish, it is essential to center attention on closed-loop reactive or biomimetic models, which strive to more precisely emulate realistic behaviors. Within Table 1.1, we distinguish between reactive

Table 1.1: Partial social interaction model & biohybrid framework literature overview. The table lists notable scientific works published between 2010-2023. It also includes a taxonomy with respect to the number of individuals in the study, the type of patterns (*i.e.*, prerecorded trajectories or dynamically generated ones) exhibited during experiments, information on whether the experiment is in simulation or reality, the experiments' duration (excluding acclimatization periods), and the temporal timescale over which they are assessed.

Nº	Literature	Nº agents incl. robot(s)	Interaction type	Simulation	Reality	Interaction loop	Exp. time (min)	Timescale
1	Maxeiner et al. (2023)	2	Fixed	✗	✓	Open, Closed	10	Short
2	Romano and Stefanini (2022b)	2	Fixed	✗	✓	Open	15	Short
3	Romano and Stefanini (2021)	2	Fixed	✗	✓	Open	20	Short
4	Chemtob et al. (2020)	5, 6	Biomimetic ^a	✗	✓	Open, Closed	15	Short
5	Cazenille et al. (2019)	5	Biomimetic	✓	✗	-	30	Short
6	Calovi et al. (2018)	1, 2, 5	Biomimetic	✓	✗	-	n/a ^b	Short, Long ^c
7	Bonnet et al. (2018)	6	Fixed, Reactive	✗	✓	Open, Closed	30	Short
8	Cazenille et al. (2018a)	5	Biomimetic	✓	✗	-	30	Short
9	Cazenille et al. (2018b)	5	Biomimetic	✗	✓	Closed	30	Short
10	Cazenille et al. (2018c)	5	Random, Biomimetic	✗	✓	Open, Closed	30	Short
11	Cazenille et al. (2017)	5	Biomimetic	✓	✓	Closed	30	Short
12	Collignon et al. (2016)	1, 10	Biomimetic	✓	✗	-	60	Short
13	Ruberto et al. (2016)	2	Fixed	✗	✓	Open	10	Short
14	Bartolini et al. (2016)	2	Fixed	✗	✓	Open	15	Short
15	Kruusmaa et al. (2016)	≈ 3000	Fixed	✗	✓	Open	150-200 ^d	Short
16	Donati et al. (2016)	5	Fixed	✗	✓	Closed ^e	≈ 10	Short
17	Landgraf et al. (2016)	7	Fixed, Reactive, Biomimetic	✗	✓	Open, Closed	≈ 15	Short
18	Ladu et al. (2015a)	2	Fixed	✗	✓	Open	10	Short
19	Ladu et al. (2015b)	2	Fixed	✗	✓	Open	10 ^f	Short
20	Landgraf et al. (2014)	22	Reactive, Biomimetic	✗	✓	Closed	30	Short
21	Phamduy et al. (2014)	2	Fixed	✗	✓	Open	10	Short
22	Butail et al. (2014b)	2, 3	Fixed	✗	✓	Open	5	Short
23	Butail et al. (2014a)	2	Fixed	✗	✓	Open	5	Short
24	Polverino et al. (2013)	2	Fixed	✗	✓	Open	5 ^g	Short
25	Polverino and Porfiri (2013b)	2	Fixed	✗	✓	Open	5	Short
26	Spinello et al. (2013)	2	Fixed	✗	✓	Open	15	Short
27	Polverino and Porfiri (2013a)	4	Fixed	✗	✓	Open	5	Short
28	Butail et al. (2013)	4	Fixed	✗	✓	Open	5	Short
29	Abaid et al. (2013)	2 ^h	Fixed	✗	✓	Open	≈ 20	Short
30	Abaid et al. (2012a)	12	Fixed	✗	✓	Open	5	Short
31	Polverino et al. (2012)	1, 10, 11	Fixed	✗	✓	Open	40 ⁱ	Short
32	Faria et al. (2010)	2, 11	Fixed	✗	✓	Open	≈ 120 min ^j	Short, Long

^aThe robots follow a biomimetic model that is periodically interrupted and exchanged for a specific modulation strategy.

^bThe model is discrete, and simulations are conducted with respect to kicks that amount to ≈ 10 h of data.

^cThe evaluation does not include long-term timescale observables, but these observables are employed for the same model in Chapter 4.

^dThe robots act once every 15 min or 20 min depending on the experiment.

^eAn open-loop motion is triggered by a closed-loop detection procedure.

^fInitial acclimatization period of 20 min with caffeine concentration. 10 min acclimatization and 10 min experimentation (recorded segment) in another tank.

^gThere is a 62 min long acclimatization period pertinent to the experiments.

^hThere is also an external fright stimulus.

ⁱThe authors state that experiments are partitioned in 8 trials of 5 min.

^jThe robot operates only at 5, 30, 60, and 120 min intervals to perform a short, predefined, motion.

and biomimetic models based on the fidelity they demonstrate. To elaborate, a reactive model might indeed function in a closed-loop, displaying simplified behaviors of the fish group (for instance, following the group's centroid Landgraf et al. (2016), etc.), while a biomimetic model seeks to replicate fish decision-making with a high fidelity to the actual system. Regardless of the interaction loop (open or closed), collective dynamics in natural systems unfold across multiple timescales Ioannou and Laskowski (2023), the source of two pivotal challenges: 1) a necessity for systems to be validated across both short- and long-term timescales, and 2) determining the extent to which our task needs to be evaluated at the distant end of the long-term timescale.

However, regarding long-term timescale experiments, Table 1.1 reveals that from 2010 to 2023, only few researchers have explored and investigated social interactions in extended timescales, and even fewer have validated hypotheses in both simulation and reality. For instance, within the context of fish social interactions, only the authors in Calovi et al. (2018) have conducted prolonged simulations of their social interaction model, enabling the study of emergent long-term dynamics, but have only done so in later research on human interactions Jayles et al. (2020). In terms of real-world experiments, whereas we observe significant strides in implementing biomimetic models Cazenille et al. (2017); Landgraf et al. (2014, 2016), these studies also rely exclusively on instantaneous metrics. Nevertheless, a notable paradigm shift is observed for Cazenille et al. (2017), whose authors conducted the only study, to this date (late 2023), that offers insights into the social dynamics of fish in both simulated and real environments.

The absence of temporal assessments in the existing state-of-the-art necessitates a critical definition of the temporal horizon for current and future studies. Indeed, establishing this is a non-trivial task, contingent upon several factors including the animal's characteristics (size, housing configuration, etc.), the behavior(s) scientists aim to study, and the scale of the collective group. Nonetheless, establishing a lower bound serves as a more accessible starting point. For example, fresh-water fish species in the wild exhibit high assemblage variability even within short time periods Czeglédi et al. (2022, 2016); Olin and Malinen (2003), at small enough timescales of even mere hours. Whereas this may simply be the case of sampling bias during short period experiments, it still suggests a need for relatively long within-day experiments and multiple replicates. This also underscores that the conventional experimentation times, as listed in Table 1.1, may fall short of thoroughly observing the full collective dynamics of even relatively small groups of fish. Furthermore, the typical total number of replicates conducted in the studies listed in Table 1.1, usually translates to a cumulative experimentation time capped at approximately 4 hours. In contrast, more recent modeling approaches suggest that a substantial volume of social interaction data (potentially exceeding 10 hours) might be requisite to assure minimized per replicate noise Calovi et al. (2018). Therefore, future research on biohybrid interactions, particularly those aiming to replicate authentic social interactions should, at a minimum, establish a lower limit of 60 min, as well as, a significantly higher cumulative experimentation time (*e.g.*, greater than 8-10 h).

Next, identifying an upper bound proves somewhat less straightforward than determining a lower one. Guided by insights from Ioannou and Laskowski (2023), we might consider temporal timescales extending from the juvenile stage of fish through to adulthood, even within the realm of social interactions, acknowledging that the social behavior of fish indeed varies with age. This not only demands considerably extensive and intricate experimental plans, but also presupposes a more nuanced understanding of collective dynamics at shorter timescales before attempting to observe social behavior evolution across different age groups. This complexity often motivates the concentration of behavioral studies on adult fish, as evidenced by the majority of research endeavors listed in Table 1.1. Consequently, in behavioral studies focused on a single age group, one might assume that a few hours under a consistent experimental configuration would be adequate. Nonetheless, even within a single day, the circadian rhythm of fish can influence their behavior, for instance, in the periods shortly before or after feeding or during shifts in light intensity, and must be taken into account Boujard and Leatherland (1992). Thus, even within laboratory environments, we must carefully schedule experiments to avoid overlapping with feeding times or the periods early and late in the day (sunrise/sunset phases), typically resulting in a further constricted experimental window of 6-8 h. This could empirically be established as the upper bound for social interaction studies of this nature.

However, the experimental window for biohybrid systems is further constrained by the operational autonomy of existing robotic systems. While numerous experiments, outlined in Table 1.1, typically span 5-15 min, and even fewer extend between 20 and 30 min, the critical need for experiments on longer timescales has been somewhat overlooked and underplayed. Additionally, current biohybrid systems often lack the capability to operate over extended durations (elaborated further in Chapter 5), for instance, due to reliance on batteries Landgraf et al. (2016) or power delivery systems prone to high failure rates Bonnet et al. (2014), thereby substantially constraining the upper bound of experimental time. Furthermore, replicating realistic levels of social interactions using robotic systems not only necessitates considerable technological advancements to facilitate lengthier experiments, but also demands enhanced fidelity in fish-like motion, as previously discussed.

The task at hand, that is, to understand and potentially minimize the biomimicry gap poses a notable challenge, advocating for the development of high-fidelity biohybrid platforms that can conceal and mitigate artificially-induced bias, to genuinely observe fish responses on individual and collective scales, as well as, across both short and long timescales. Subsequent chapters detail our approach, which encompasses extended experiments (relative to the current state-of-the-art presented in Table 1.1), enabled by robotic systems with unprecedented agility and autonomy, and commanded by models closely mirroring the spontaneous social dynamics evident in fish schools.

2 Thesis outline

2.1 Thesis objectives

Deciphering collective behavior through the use of robotic devices has rapidly progressed since the early 2000s. However, as discussed in previous sections, there exist research directions that are yet to be explored and research questions (RQ) that remain unanswered. The main research we sought to answer is:

- How can we bridge the biomimicry gap with the purpose of constructing cohesive mixed (*i.e.*, biohybrid) groups of animals and robots that are indistinguishable from fish-only groups?

Naturally, this research question requires work spanning the domain of collective behavior and complex system modelling, and that of real-time robot design, control, and deployment. Ergo, our original research question encompasses the following sub-questions:

- (RQ1) How can we explore the bidirectionality of fish-robot interactions to measure how they facilitate the integration of a robot into fish groups?
- (RQ2) How can we design and assess bidirectional models of high fidelity (*w.r.t.*, to observations from the spontaneous animal interactions)?
- (RQ3) Which key elements make for robotic devices that can closely mimic fish motion and dynamics to subsequently transfer such models “back to reality”?
- (RQ4) How can we measure, and what is the gap between real experiments and simulations with such models?

2.2 Summary of contributions and thesis organization

2.2.1 Part II

This part of the thesis is dedicated to modelling strategies, as well as, the metrics that validate a model's fidelity with respect to spontaneous fish movements.

Chapter 3.

We study the bidirectionality of interactions between a robot and groups of fish. In this chapter, we establish a baseline understanding into how biomimetic models can lead to increased levels of acceptance in biohybrid groups. This research question is primarily aimed at establishing a baseline about the social interaction gap (depicted in Fig. 1.1).

Chapter 4.

We investigate how a deep learning algorithm can achieve state-of-the-art level predictions for social interactions. We expand on how models of behavior (deep learning or otherwise) should be stringently assessed in both short and long timescales and demonstrate that while models may predict well in the short term, this does not guarantee good predictions in the long term. We also conduct an ablation study to expand this methodology to multiple fish species. Following the work in Chapter 3, this chapter aims at reducing the effect of the social interaction model on the biomimicry gap (see also Fig. 1.1).

2.2.2 Part III

This part delves into the state-of-the-art laboratory equipment that enables us to perform biohybrid studies. We also present the equipment that allows us to bring the models from Part II to be transferred back to real-life in Chapters 6 and 7.

Chapter 5.

We redesign the modern small animal experimentation framework and introduce a novel experimental setup and robot for robot-animal interaction studies, which is aimed at reducing the physical gap observed when transferring models into real systems (see Fig. 1.1). Within this chapter, we also demonstrate how biomimetic characteristics are key to establishing strong communication cues with the animals (fish), thus further reducing the biomimicry gap owed to this component (depicted in Fig. 1.1).

2.2.3 Part IV

In this part, we assess the transferability of social interaction models into real-life through biohybrid groups. That is, we traverse the steps of Fig. 1.1, and attempt to quantify the biomimicry gap, also shedding some light into which of the sub-gaps contributes the biggest deviation from simulated and real-life fish groups.

Chapter 6.

We conduct experiments with a deep learning model for social interactions by commanding a robot interacting with live fish. We investigate the consistency of the robotic system and social interaction model across many hours of experiments. Finally, we perform a series of analyses to measure the “biomimicry gap” between the simulated model and its physical interpretation through the robot. We focus in quantifying the biomimicry gap observed in groups of two individuals.

Chapter 7.

We conduct experiments with an analytical model for social interactions by commanding a robot interacting with live fish. Yet again, we perform a series of analyses to measure the “biomimicry gap”, but this time we extend the experimentation to varying group sizes of fish. This not only further substantiates our methodology to quantify the gap, but also demonstrates how the gap is affected when changing the group size (and in turn, the complexity of dynamics).

2.2.4 Part V

Chapter 8.

We discuss on the sum of our findings, take note of the limitations of our and the domain’s current work, and expand on the directions of future work.

How to design models that encode animal behavior **Part II**

3 Establishing a baseline of fish-robot dynamics

This section's content is adapted from the following article:

- Papaspyros, V., Bonnet, E., Collignon, B. and Mondada, F., 2019. Bidirectional interactions facilitate the integration of a robot into a shoal of zebrafish *Danio rerio*. *PloS one*, 14(8), p.e0220559. <https://doi.org/10.1371/journal.pone.0220559> (reuse authorised under the CC 4.0 licence) – Papaspyros et al. (2019)
 - Vaios Papaspyros's contribution: Data curation, formal analysis, investigation, methodology, software, visualization, writing – original draft, writing – review & editing.
- The article's work was funded by the Swiss National Science Foundation project "Self-Adaptive Mixed Societies of Animals and Robots" (Grant No. 175731) and partially funded by the EU-ICT project ASSISIBf (Grant No. 601074).
- Ethics: The authorization for the experiments conducted in this research work was approved by the state ethical board of the Department of Consumer and Veterinary Affairs of the Canton de Vaud (SCAV) of Switzerland (authorization N° 2778).

3.1 Introduction

To well understand the full process that allows for creating a biohybrid society, we began by running an experiment involving all components of the system and investigated the various elements generating the biomimicry gap. We explored various types of models, that is, we created models aimed at producing different interactions and used a state-of-the-art robot and framework. The results are the baseline of this thesis and allow for starting the exploration on how to close the several levels of gaps to eventually generate a cohesive biohybrid society, where the robot is well integrated within it. Whereas thus far, it has been experimentally shown that fish decision-making processes can be manipulated with the use of artificial lures Faria et al. (2010); Swain et al. (2011); Landgraf et al. (2013); Bonnet et al. (2014); Landgraf et al. (2016); Ruberto et al. (2016); Bartolini et al. (2016); Bonnet et al. (2018); Porfiri (2018); Romano and Stefanini (2021, 2022b,a), there is limited insight into how the models commanding

the robot affect its successful integration to the fish group. Indeed, the degree to which a robot needs to exhibit social skills is not validated experimentally. That is, to investigate how interactive the robot needs to be, with respect to fish decisions, to strongly engage with and integrate into fish groups. The interactivity can be, in turn, interpreted as the robot's ability to communicate bidirectionally, and alternate between following passively and actively taking decisions. In many ways, this provides insight into the principles that should underlie social interaction models capable of bridging part of the biomimicry gap (social interaction gap, see Fig. 1.1). Currently, there are two main ways to command the robot.

A first strategy that addresses relies on open-loop models Phamduy et al. (2014); Polverino et al. (2013); Polverino and Porfiri (2013b); Abaid et al. (2012a); Butail et al. (2014a); Ladu et al. (2015a,b); Polverino et al. (2012); Spinello et al. (2013); Polverino and Porfiri (2013a); Ruberto et al. (2016); Bartolini et al. (2016); Kruusmaa et al. (2016); Abaid et al. (2013); Butail et al. (2013, 2014b) that do not actively react to or take into account the actions of the animals. In this case, it may be difficult to discern if the focal animal is reacting to external stimuli or to another agent considered as a shoal member. Therefore, they may often fail to provide a clear insight into the internal decision-making process and the natural information flow between the shoal members. A second strategy is based on closed-loop models Cazenille et al. (2018c); Faria et al. (2010); Swain et al. (2011); Kopman et al. (2013); Landgraf et al. (2013, 2016); Porfiri et al. (2019); Kim et al. (2018) attempting to achieve a conspecific status among individuals by engaging in mutual information exchange and could reveal the intrinsic decision-making mechanism of individuals.

Therefore, the following research question arises: can we discriminate between different behavioral rules by implementing them in a robot interacting with a shoal of zebrafish? And more importantly, how bidirectional do closed-loop models need to be in order to create harmonious biohybrid groups?

To investigate the extent to which the bidirectionality of models affect the collective behavior of fish, we devised the following experiment: (1) we make use of a circular corridor setup, a commonly used experimental arena for numerous studies on collective behavior Bonnet et al. (2018); Crosato et al. (2018); Abaid and Porfiri (2010); Lecheval et al. (2018); (2) we designed three behavioral models exhibiting three different dynamics, namely, a purely reactive model that explicitly follows the fish, an imposing direction model constantly attempting to dictate the collective swimming direction decision and a biomimetic model mimicking the decisions of zebrafish in a circular corridor environment; (3) we performed experiments with zebrafish-only and mixed groups of robots and zebrafish; (4) we analyzed the results using three inherently different approaches. We show that our robotic system was capable of participating in the collective decision-making and blending in the shoal without significantly perturbing its interaction dynamics. This, in turn, led to an improved integration state where the robot was not only accepted by the shoal, but it was also contributing to the collective decision and acting as a leader for the majority of the time. Thus, bidirectional interactions proved to be the key to forming cohesive biohybrid groups.

3.2 Background

3.2.1 Animals

For the experiments, 60 wild-type zebrafish *Danio rerio* with short fins were used (AB strain). The zebrafish were acquired from a pet shop and subsequently stored in a 60 L aquarium. The average length of the zebrafish used was approximately 4 cm. Water in the housing aquarium was kept at a temperature of 26 C°. The fish were fed once per day with commercial food between 16:00 and 18:00. Furthermore, the aquarium was enriched with plastic plants, Cladophora, gravel, rocks, and aquatic snails.

3.2.2 Experimental setup

The experimental arena pictured in Fig. 3.1 consists of a 10 cm wide circular corridor (from two circular walls: 1) an outer of 58 cm diameter, and 2) an inner of 38 cm diameter) placed in a $100 \times 100 \times 25 \text{ cm}^3$ glass tank, as in Bonnet et al. (2018). This setup presented the zebrafish with a binary choice for movement, that is, they could either move clockwise (CW) or counter-clockwise (CCW). In fact, this is a common setup for behavioral studies Abaid and Porfiri (2010); Jiang et al. (2017); Bonnet et al. (2018); Lecheval et al. (2018) because it allows for setting aside spatial complexities. Instead, this setup provides a symmetric arena that enables researchers to analyze multiple instances of the same behavioral traits (e.g., U-turns Crosato et al. (2018) where the fish will perform a direction change greater than or equal to 180°) and quantify their consistency across different types of behavioral models.

The bottom part of the experimental tank was covered with a Teflon plate to allow for smoother motion of the robotic fish lure (see Sec. 3.2.3) and avoid any stimuli produced by reflections or by the mobile robot moving below the setup. Furthermore, the setup was confined behind white sheets to isolate the fish from the rest of the room, while also maintaining a consistent lighting environment. A uniform luminosity for the room was provided by four 110-W fluorescent lamps placed at each of the four sides of the tank.

3.2.3 Robotic system for closed-loop zebrafish-robot interactions

For the zebrafish-robot interaction experiments, we used one miniature wheeled robot, the FishBot Bonnet et al. (2014, 2012) (see Fig. 3.2B). The robot was placed between two conductive plates located below the experimental setup (see Fig. 3.1) and was powered using brushes that were constantly in contact with them. This configuration allowed the robot to operate for long periods of time and powered the motors that were, in turn, capable of achieving the necessary speed and acceleration in order to quickly adapt to the rapid spatial displacements of zebrafish. The FishBot was additionally equipped with a Bluetooth chip that allowed it to wirelessly communicate with a computer that was providing the necessary motor commands.

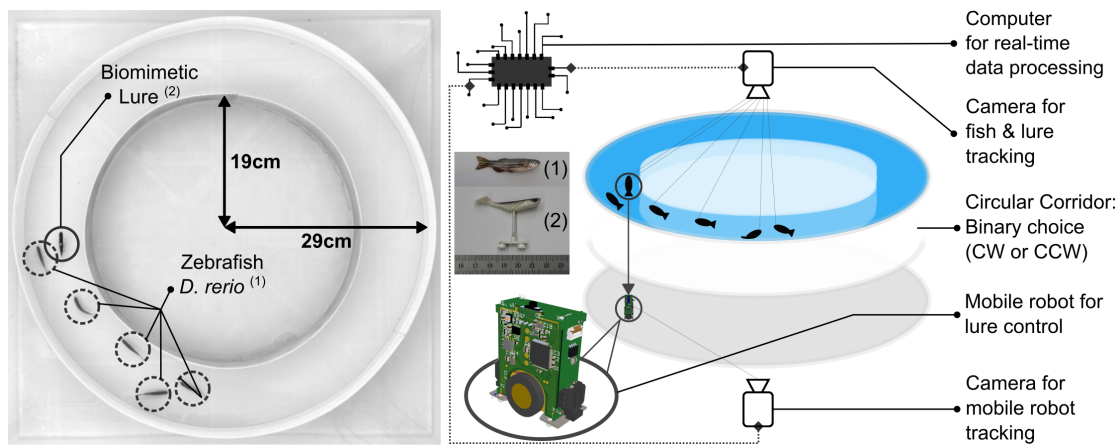


Figure 3.1: **Experimental setup.** (left) Top view depicting the setup's dimensions, *i.e.*, inner ring radius of 19 cm and outer of 29 cm. The dotted circles indicate the positions of zebrafish while the full circle indicates the position of the biomimetic lure, and (right) breakdown view of the setup depicting individual components that are necessary for closed-loop interaction. The mobile robot (FishBot) is moving below the tank and drives a biomimetic lure inside the tank through a magnetic coupling. The top and bottom mounted cameras capture frames at a rate of 15 Hz and transmit the information to a computer. The computer will then fuse the information to determine the positions and heading of fish and robot(s) alike. In gray, we denote the conductive plates that are used to power the FishBot.

A soft biomimetic lure of approximately 4.5 cm length (see Fig. 3.2A) was selected to physically interact with the animals. This lure was designed to mimic the morphology of fish and passively beat its tail during its underwater motion. As described in Bonnet et al. (2014) this specific lure achieved strong acceptance in groups of zebrafish. Subsequently, it was mounted on a carbon stick at a height of 3 cm to ensure that it was visible by neighboring fish. An iron plate located at the bottom of this stick carried two magnets that allowed for a magnetic coupling (similarly to Bonnet et al. (2014); Faria et al. (2010); Swain et al. (2011); Landgraf et al. (2013, 2016), see Fig. 3.2D) to the robot located below the setup.

3.2.4 Control and tracking software

In order to close the interaction loop between the fish and the robot, we made use of the Control and tracking for multi-agent animal-robots groups (CATS) framework Bonnet et al. (2017). CATS continuously monitored the positions of the robot and the animals through image frames obtained by the two cameras located above and below the setup (see Fig. 3.1). In particular, the overhead camera was set to simultaneously stream video in two resolutions; (1) a 1040×1040 stream that was recorded and used for the analysis of the experiment, and (2) a 512×512 stream that was used in CATS for the detection of the fish and/or lure in real-time. The image frames from the camera located at the bottom of the setup were directly fed to CATS for processing.




Figure 3.2: **FishBot and biomimetic lure.** (A) The biomimetic lure (approximately 4.5 cm long) fixed on a white carbon stick, (B) the FishBot on its side (approximately 5.5 cm long), (C) relative size of the lure, FishBot and FishBot cover (from left to right), and (D) assembled robotic system; the white cover protects critical parts of the FishBot, and the lure magnetically coupled to the FishBot.

More specifically, the agents' positions were determined by feeding the image frames of the overhead camera to a corner detection method Shi et al. (1994) implemented using the OpenCV Bradski (2000) library; while the camera located below the setup was used to localize the mobile robot, which was equipped with 6 light-emitting diodes of blue color, using a blob detector. Subsequently, CATS fused information from both cameras to distinguish the artificial lure from the living individuals. We note that, both cameras operated at a rate of 15 frames per second.

Once CATS had finished determining the positions, the resulting spatial information (2D position and heading direction) became available in the control layer of CATS. This layer is responsible for the higher level control procedure of the robot (*i.e.*, deciding which is the next desired state for the robot). The behavioral models presented in the following section were implemented within the control layer of CATS and output higher level commands such as desired velocity, position, and orientation. Those commands were then fed to a micro-controller unit where a proportional-integral-derivative controller (PID) translated the higher level commands to motor commands, similarly to Bonnet et al. (2012).

Additionally, to the online control procedure that was devised for the fish-robot interaction experiments, the videos of each experiment underwent post-processing using the idTracker software Pérez-Escudero et al. (2014) to extract the trajectories of each agent for each experi-

ment. This time-consuming and computationally expensive process is capable of recreating the trajectories of each identified agent (6 agents for 30 minute long experiments) with on average 95% accuracy, correcting any mistakes made in trajectories due to crossings that occurred naturally throughout the experiment.

3.3 Experimental procedure

For the duration of the experiments, we maintained a constant height of 6 cm of water in the setup. Under these conditions, the fish were not additionally stressed, and their movement was on average constrained to a specific height, thus, reducing spatial complexities on the z-axis. Prior to placing the zebrafish in the setup, the water temperature was brought to 26 C°. Thereafter, a shoal of zebrafish was randomly caught from the rearing tanks with a fishnet and placed in the experimental setup. After a 5-minute acclimatization period during which the FishBot remained stationary, we started and recorded the experiment for 30 minutes. No individual was used twice in the same day.

We conducted 10 experiments with shoals of six zebrafish and no FishBot to observe the baseline behavior of the fish (hereby referred to as “fish-only” experiments) when no artificial stimuli were provided. Then we conducted a total of 3×10 experiments with five zebrafish and one FishBot with three different behavioral models (described in the following section) for the FishBot. Each model was tested in random order to account for the fish getting accustomed to a specific behavior exhibited by the robot.

3.4 Behavioral models

Follower Model (FM). We designed a closed-loop following model, the “Follower Model”, where the robot was simply instructed to head towards the point in space that was on average most dense in terms of fish occupancy (*i.e.*, the robot is given an explicit Cartesian goal position that corresponds to this point). FM is a purely reactive, passive model in the sense that it does not actively model or embed interaction in its design and instead reacts only to the position of the fish by always following it.

Despotic Model (DM). We also designed an open-loop model, which is an adaptation of the approach described in Bonnet et al. (2018) that uses only one robot which was instructed to perform a CW movement throughout the experiment, that is, the robot does take spatial information into account. Contrary to FM, this model is always attempting to impose the collective movement direction decision, thus we call it “Despotic Model”.

Feedback-Initiative Model (FIM). Finally, we implemented a closed-loop parametric behavioral model similar to Abaid and Porfiri (2010). In Abaid and Porfiri (2010), the authors

described a model that operates in a one-dimensional decision space (*i.e.*, CW or CCW movement). First, the circular corridor arena is divided into 40 equal cells of 9-degree arc length each. Then, the focal individual will take a directional decision according to the perceived heading directions of its neighbors and a probability to “disobey” this collective decision. More specifically, the next heading direction of an individual is given by the following expression:

$$h(\text{fish}_j, t+1) = \frac{h(\text{fish}_j, t) + \sum_{i=0, i \neq j}^{N_p} h(\text{fish}_i, t)}{\left| h(\text{fish}_j, t) + \sum_{i=0, i \neq j}^{N_p} h(\text{fish}_i, t) \right|}, \quad (3.1)$$

where $h(\text{fish}_j, t) \in \{-1, 1\}$ the heading direction at time t and N_p the subset of the set of all individuals, N_A , that are in the perceptual range of the focal individual. The perceptual range is defined as the set $F \subseteq N_A$ of the individuals that are within p_r cells in the forward direction of the focal individual. At every timestep and after computing the new heading direction, the focal individual is given a probability $1 - P_{\text{obey}}$ to choose the opposite direction to the one computed through the above interaction metric. We refer to this opposition to the collective decision of swimming direction as the initiative of the focal individual.

In order to investigate the degree to which our robot can influence the directional decisions taken by the group, we propose a variant of this model where the focal individual attempts to closely mimic zebrafish interaction in similar arenas [Abaid and Porfiri \(2010\)](#); [Jiang et al. \(2017\)](#); [Bonnet et al. \(2018\)](#) while at the same time it intelligently embeds initiative in its decision-making according to the feedback perceived by the neighboring social companions. Therefore, we call this model “Feedback-Initiative Model”. Similarly to [Abaid and Porfiri \(2010\)](#), we discretized the circular corridor in cells and controlled the robot’s direction of movement in the one-dimensional space of CW or CCW movement. More specifically, we separated the setup in 40 cells, each of which corresponds to approximately 4 cm of arc length, the average length of a single zebrafish. Essentially, increasing the number of cells, *i.e.*, decreasing the arc length per cell, would allow for more detailed separation of the fish in terms of cell occupancy, but would be subject to the noise produced by CATS (see Sec. 3.2.4). This discretization process reduced the locomotion control dimensions to 1D space, and the model need only to output a simple instruction at every timestep: move one cell CW or CCW. The model made a prediction about the best candidate direction at time $t+1$, every 0.25 seconds (*i.e.*, the controller timestep is equal to 0.25 seconds). In the context of this chapter, the best candidate swimming direction was considered to be the one that has the highest probability to elicit a collective U-turn (*i.e.*, a switch of the swimming direction for the majority of individuals).

Our source of inspiration for those interactions was derived by the innate behavior of zebrafish in similar setups and can be summarized in the following two key features for a focal individual:

1) an innate tendency to align with the shoal, and 2) a tendency to perform a U-turn when few or no agents are in the field of view of the focal, initiating a direction change that might propagate throughout the shoal Bumann and Krause (1993); Collignon et al. (2017); Lecheval et al. (2018); Jiang et al. (2017). The former is formalized as a weighted sum of the direction of a focal individual's neighbors and is defined as the dot product of the heading (-1 or 1) and an exponential function that can rapidly increase or decrease the impact of a neighbor according to its position, as follows:

$$h(\text{fish}_j, t + 1) = \sum_{i=0, i \neq j}^{N_A} h(\text{fish}_i, t) \cdot e^{\alpha \cdot p(\text{fish}_i, t)} \quad (3.2)$$

where $p(\text{fish}_j, t) \in [0, |N_C|] \subseteq \mathbb{Z}$ the position of the fish j at time t , N_C the set of cells that correspond to the discretization process and $\alpha \in \mathbb{R}$ a regulatory parameter to set the slope.

To allow for more flexibility in FIM we used two separate parameters α_f and α_b for fish that are in leading and following positions respectively (see Algo. 1) and defined the sets F and B of fish in the forward and backward positions respectively. Intuitively, the difference between the parameters α_f and α_b can express biologically observed behaviors such as that fish within the immediate field of view of a focal individual have more influence on it Bumann and Krause (1993); Collignon et al. (2017); Crosato et al. (2018); Lecheval et al. (2018), while the followers can still be perceived due to the water flow Dijkgraaf (1963) and might have less, but significant influence. The parameters α_f and α_b were manually tuned for this specific configuration and are given the values -0.2 and -0.5, respectively. Despite the sum's ability (see Eq. 3.2) to phase out the influence of perceived agents over distance, we explicitly limited the robot's knowledge within 15 cells ($p_f = 7$ cells forward, $p_b = 7$ cells backwards and current cell occupied; or 135° of perceptual range for the choice of 40 cells). Indeed, for very low values of α (see Eq. 3.2), a conspecific could have been perceived in front and behind the focal fish due to the circular design of the setup. The limits were empirically selected to emulate the effect that the two circular walls have on the ability of individuals to see neighbors due to the setup's curvature. Finally, the focal individual's tendency to perform a U-turn and even disobey the collective decision of the shoal concerning the direction of movement was modeled as a probability P_{obey} . The probability P_{obey} is dependent on the amount of fish in the forward direction to account for an individual's intuition to not wander too far from the fish school or to simply initiate a random direction change. More specifically, this probability is regulated by two constant parameters (see Algo. 1): 1) $a_{\text{influence}} = 4$, which allows one to increase or decrease the amount of influence of forward individuals concerning the obedience, 2) a constant upper bound $\tau = 0.95$, value which we estimated through past fish-only experiments. A subset of the parameter space that could be used in different scenarios include:

- $\alpha_f = \alpha_b = -\text{inf}$, $\tau = 1 \Rightarrow P_{obey} = 1$, will produce purely following behavior.

- $\alpha_f = \alpha_b = \tau = 0 \Rightarrow P_{obey} = 0$, will produce a behavior that always contradicts the collective.
- $a_{influence} = 0$, will retain a constant and neighbor-independent P_{obey} .
- Setting the perceptual range to zero (*i.e.*, $Fish' = \emptyset$) and $\tau = 1$, will produce an imposing direction behavior (the initial direction will be followed throughout the experiment).

We designed this model to be parametric and include stochastic elements of decision-making. The parametric design allows for modification of the model to comply with different scenarios of interaction or species of fish (*e.g.*, the robot could be instructed to emphasize on following by changing a few parameters), either prior to deployment or during an experiment, while the stochasticity serves as a way to promote initiative in the model. In Abaid and Porfiri (2010) the authors described a model where the focal individual's next direction will be with high certainty decided by the average swimming direction of the neighboring individuals. Conversely, here, the goal was to elicit a different effect from the fish and influence them to change their swimming direction.

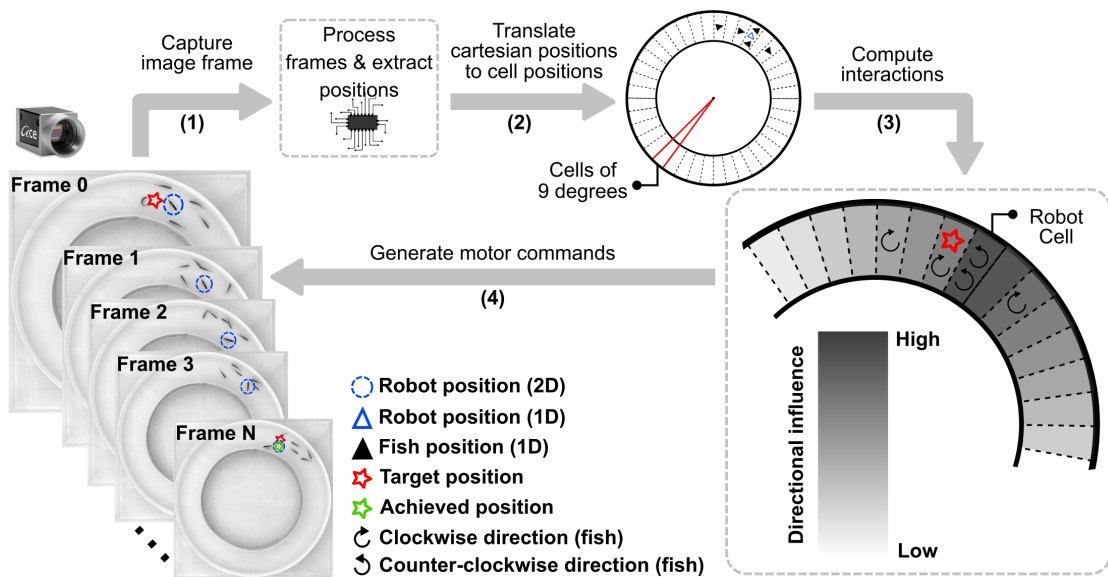


Figure 3.3: **Closed-loop robot control.** For each system cycle: 1) a high resolution image frame is captured by the overhead camera, 2) the frame sent to a high performance computer, where it is processed to determine the positions, velocities, and headings of each individual, 3) the extracted positions are discretized and each individual is placed in its corresponding cell, and 4) the discretized positions and headings of each individual are forwarded to the FIM, which in turn, weighs the heading direction of the neighboring individuals and produces a desired position (red star) for the next timestep. After a few timesteps, an approximation of the target position will be achieved (green star) and the process is repeated.

Algorithm 1 Pseudocode summarizing the Feedback-Initiative model (FIM). Let two sets F and B contain the fish in forward and backward cells. Then, the decision of a focal individual in time $t + 1$ will be determined by the sum of two exponential functions weighted by a_f and a_b for the forward and backward fish, respectively. Finally, the focal individual reverses its predicted direction with probability P_{obey} , determined according to a lower bound τ and an obedience coefficient $a_{\text{influence}}$.

procedure STIMULATE(Fish')

Split Fish' to:

$$C = \{\text{fish in focal cell}\} \subseteq \text{Fish}'$$

$$F = \{\text{fish in } p_f \text{ forward cells}\} \subseteq \text{Fish}'$$

$$B = \{\text{fish in } p_b \text{ backward cells}\} \subseteq \text{Fish}'$$

$$s = \sum_{i=0, \text{fish}_i \in F}^{|F|} h(\text{fish}_i, t) \cdot e^{\alpha_f p(\text{fish}_i, t)} + \sum_{j=0, \text{fish}_j \in C \cup B}^{|C \cup B|} h(\text{fish}_j, t) \cdot e^{\alpha_b p(\text{fish}_j, t)}$$

$$h' = \begin{cases} s & , s \neq 0 \\ h(\text{focal}, t) & , \text{otherwise} \end{cases}$$

$$P_{\text{obey}} = \begin{cases} \tau * \left(1 - (|F \cup C| + 1)^{-\alpha_{\text{influence}}}\right) & , \text{Fish}' \neq \emptyset \\ \tau & , \text{otherwise} \end{cases}$$

with probability $1 - P_{\text{obey}}$ reverse h'

return h'

procedure MOVE($h(\text{focal}, t + 1)$)

$$p(\text{focal}, t + 1) = p(\text{focal}, t) + h(\text{focal}, t + 1)$$

procedure FEEDBACKINITIATIVEMODEL

$\forall \text{fish} \in \text{Fish}$, where $|\text{Fish}| = |N_A|$:

Initialize **position** $\in [0, |N_C|)$

Initialize **heading** $\in \{\text{CW} = -1, \text{CCW} = 1\}$

while stopping criteria not met **do**

$\forall \text{fish} \in \text{Fish}$:

Fish' = {fish within the perceptual range of the focal fish} $\subseteq \text{Fish} \cup \emptyset$

$h(\text{focal}, t + 1) = \text{STIMULATE}(\text{Fish}')$

MOVE($h(\text{focal}, t + 1)$)

3.5 Data treatment

3.5.1 Data filtering

The 15 frames per second capture rate made it possible to detect even minor fluctuations in the displacement of an individual. On one hand, this rate is useful for tracking fast moving objects or animals, but, in the case of zebrafish that move with an average speed below 20-

25 cm/second in this setup, it might induce noise due to temporary loss of the position or the image processing algorithm reporting minor differences in the position of an individual at every timestep. Therefore, throughout the following experiments we filtered the data reported in two ways: (1) 3 frames (0.2 seconds of interaction) were averaged to calculate the centroidal position, heading, or velocity of each agent and (2) the behavioral models presented in following sections discretized the positions in bins, the number of which was selected to further filter the measurements where necessary.

3.5.2 Data analysis

In this section, we introduce a set of metrics based on spatial, directional and information theoretic measures, as well as the statistical methods followed to evaluate and compare the behavioral models. We note that, all the raw trajectory data are available at <https://doi.org/10.5281/zenodo.6783303>.

Average angular distance

In collective behavior, denser groups often suggest a more cohesive, aligned and organized movement Buhl et al. (2006); Jiang et al. (2017), thus, density-based measures have been widely used for inferring the interaction rules within a group of animals Lukeman et al. (2010); Hemelrijk and Hildenbrandt (2008). Here, we computed a similar measure, by calculating the average angular distance between all pairs of agents. The angular distance between two fish is defined by the angle $\theta_{ij}(t)$ at time t , where i, j are two individuals and $\theta_{ij}(t) \in [0, \pi]$ is the angle between i and j with respect to the origin point of the setup (center of both rings). We note that, $\theta_{ij}(t)$ refers to the acute angle between the two individuals (*i.e.*, we only evaluate the angular proximity). The average angular distance was computed as the average of all the pairwise angular distances and is summarized in the following expression:

$$\text{averageAngularDistance}(t) = \frac{1}{N_A(N_A - 1)} \sum_{i=1}^{N_A} \sum_{j=1, i \neq j}^{N_A} \theta_{ij}(t) \quad (3.3)$$

Collective U-turns

Although the average angular distance provides useful topological information concerning the closeness of the group and thus its cohesive and synchronized movement, it would be incomplete without a complementary metric concerning the interactions within it. Here, we captured these interactions in the number of collective direction changes performed (*e.g.*, from CW to CCW or vice versa), which in this binary choice scenario we defined as collective U-turns. The U-turn in schools of fish has attracted attention Crosato et al. (2018); Lecheval et al. (2018); Jiang et al. (2017) as it provides insight on how information is propagated among

individuals. Consequently, the effect that each behavioral model has on the occurrence of collective U-turns is representative of its ability to mimic, modulate or perturb the collective decision-making.

To calculate the number of collective U-turn events, we first defined the polarization of a zebrafish shoal in this context, as follows:

$$\text{pol}(t) = \frac{1}{N_A} \sum_{i=0}^{N_A} h(\text{fish}_i, t) \quad (3.4)$$

A collective U-turn occurs when the polarization of the shoal switches from one direction to another ($\text{pol}(t) \cdot \text{pol}(t-1) < 0$; *i.e.*, we did not take into account transitions from CW/CCW to 0). Complementary figures concerning the duration of consecutive movement before a collective U-turn occurs are available in Appendix A section (Figs. A.1, A.2, A.3, A.4, A.5, A.6, A.7, A.8).

Transfer entropy

To complete the aforementioned metrics, we also employed an information theoretic measure based on the Shannon entropy Shannon and Weaver (1949), called transfer entropy (TE) Schreiber (2000); Vicente et al. (2011). Recent studies in collective behavior Crosato et al. (2018); Butail et al. (2014a); Porfiri (2018); Mwaffo et al. (2017); Porfiri and Marín (2017); Ruberto et al. (2018) have been increasingly using TE to provide insight on the mutual interactions of individuals over time or with time delays Wibral et al. (2013). Here, we adopted the notation of TE with embedded delay, as defined in Takens (1981). More specifically, given two time-series X and Y , TE measures the amount of information provided by the source X about the target Y and is defined as follows:

$$T_{X \rightarrow Y} = \sum p(y_{n+1}, y_n^{(k)}, x_n^{(l)}) \log \frac{p(y_{n+1} | y_n^{(k)}, x_n^{(l)})}{p(y_{n+1} | x_n^{(k)})} \quad (3.5)$$

where l and k are the history lengths for the two time-series:

$$x_n^{(l)} = \{x_n, x_{n-\tau_k}, x_{n-2\tau_k}, \dots, x_{n-(k-1)\cdot\tau_k}\} \quad (3.6)$$

$$y_n^{(l)} = \{y_n, y_{n-\tau_l}, y_{n-2\tau_l}, \dots, y_{n-(l-1)\cdot\tau_l}\} \quad (3.7)$$

and τ_k, τ_l the time delay for the source and destination signal, respectively.

In the context of this chapter, we defined the time-series X , Y to be the heading direction of two separate individuals over time. More specifically, we represented the direction of each agent in a discrete signal with values -1 (CW) or 1 (CCW), by sampling the direction of each individual every 0.2 seconds. For this computation, a more detailed trajectory is required, and thus we used a discretization with a cell count of $N_C = 160$ (*i.e.*, 1 cm per cell; for this procedure, we used the positions extracted from idTracker). In case an individual had not moved during this period, we assumed that its heading has remained the same as in the previous timestep. Furthermore, considering that the influence of one individual to another will be delayed in time, we shifted the source and target time-series by a factor τ_k and τ_l , respectively. The intuition behind this measure is that given the direction Y_n of an individual, we gain information about the next direction X_n of another individual. It is rather obvious that the direction change of a single individual will not propagate instantly and, thus, there exist the parameters τ_k and τ_l that express this delay.

However, choosing the latter parameters is a non-trivial task, as the values need to be meaningful with respect to the experiment in question and at the same time expressive enough to allow for observing potential differences in the fish-robot experiments. To that end, we adopted the same technique of the authors in Crosato et al. (2018), that is, we run a simple search algorithm to find the parameters that maximize the average TE for the fish-only experiments. To reduce the size of the search space, we only considered values of $k \in [1, 15]$ and $\tau_k \in [1, 15]$ (*i.e.*, up to three seconds of signal length and delay). We explicitly set the target delay to $\tau_l = 1$ and the length to $l = 1$ (*i.e.*, 0.2 seconds) as we are interested in the effects of the source signal. We note that, the robot-fish experiments were considered for this optimization step to account for the bias that was introduced due to the use of the lure and the models.

Subsequently, we calculated all the pairwise TE values and computed the mean TE across all individuals during one experiment. To do so, we used the JIDT Lizier (2014) framework to calculate the TE with the optimized parameters $k = 4$ and $\tau_k = 1$ (*i.e.*, $k = 4$ corresponds to 0.8 seconds of history and $\tau_k = 1$ corresponds to a delay of 0.2 seconds). In fact, the time of 0.8 seconds is found to be approximately equal to the average decision time of zebrafish Escobedo et al. (2020), whereas a delay of 0.2 seconds is empirically found to correspond to the reaction time of zebrafish Roberts et al. (2019). The optimized parameters are empirically found to correspond to the time that is necessary for an individual to perform a U-turn and fully propagate it to the shoal.

We computed two separate mean TE values: one for the outgoing and one for the incoming amount of information exchanged. For each case, we sum the resulting TE for the trajectories of all the fish-fish or fish-robot pairs. Intuitively, the metric expresses the average information flow direction (incoming or outgoing) when the robot is used. For each behavioral model and each of the outgoing and incoming cases, we computed three different quantities: (1) the overall TE for all individuals (2) the TE related to the robot alone (outgoing and incoming) and (3) the average TE of all fish (*i.e.*, excluding the robot's contribution).

3.5.3 Statistical tests

To further validate the interpretation of the resulting data, we performed a Kruskal–Wallis (KW) test followed by a post hoc analysis using Tukey’s honest significant difference (T-HSD), for each measure presented in the following section. The Kruskal–Wallis is chosen due to the fact that the variance of the data-sets in question differs depending on the experiment type.

3.6 Behavioral model assessment

3.6.1 Average angular distance

First, we compiled the average angular distance between the shoal members across the entire observation period for each experiment of the four tested conditions (Fig. 3.4). For the fish-only condition (*i.e.*, where no artificial stimuli are present in the setup), we observed an average angular distance of $35.04 \pm 12.62^\circ$. Compared to this baseline measure, all models showed a higher inter-individual angular distance (FM: $64.72 \pm 15.95^\circ$, DM: $59.31 \pm 11.87^\circ$, FIM: $51.04 \pm 10.68^\circ$), with these distributions differing significantly from each other (KW test, $p < 0.0001$, $\chi^2 = 22.45$). A more detailed comparison revealed that both the FM and DM significantly differ from the fish-only (T-HSD post hoc test, $p < 0.0001$ and $p < 0.001$, respectively) while the FIM differs significantly from the FM and DM (T-HSD post hoc test, $p \approx 0.22$ and $p \approx 0.53$, respectively) but not from the fish-only distribution (T-HSD post hoc test, $p \approx 0.072$). Complementary statistics are available in Table A.1 of Appendix A. Thus, while FIM still did not perform as well as the control experiments ($16^\circ \approx 8$ cm arc distance), its ability to mimic the collective decision-making allowed the robot to maintain the cohesion of the shoal with on average 14° better than FM and 9° better than DM. Moreover, the results observed are consistent over time as shown in Fig. 3.4B depicting the average angular distance for every minute of the experiment.

The mean performance and amount of variance in the FM model is indicative of its deficiency when it with regard to its ability to be accepted and integrated with the shoal. Furthermore, such result is in direct contradiction with FM’s explicit goal, which was to head towards the densest point of the shoal and thus promote a more cohesive behavior, and could suggest that its movement patterns were too aggressive to be accepted by the shoal and contribute to its operation. DM, on the other hand, performed on average worse than the fish-only but exhibits similar variance and seemed to perturb the shoaling behavior less. Finally, FIM was the most consistent over time, which could be indicative of an overall better acceptance by the shoal.

3.6.2 Collective U-turns

The collective U-turns performed per minute (see Fig. 3.5). (1) fish-only had median of 13.05 turns and a mean of 12.38 ± 3.71 , (2) FM had median of 7.25 turns and a mean of 8.27 ± 2.30 , (3) DM had median of 6.97 turns and a mean of 7.91 ± 1.60 and (4) FIM had median of 12.38

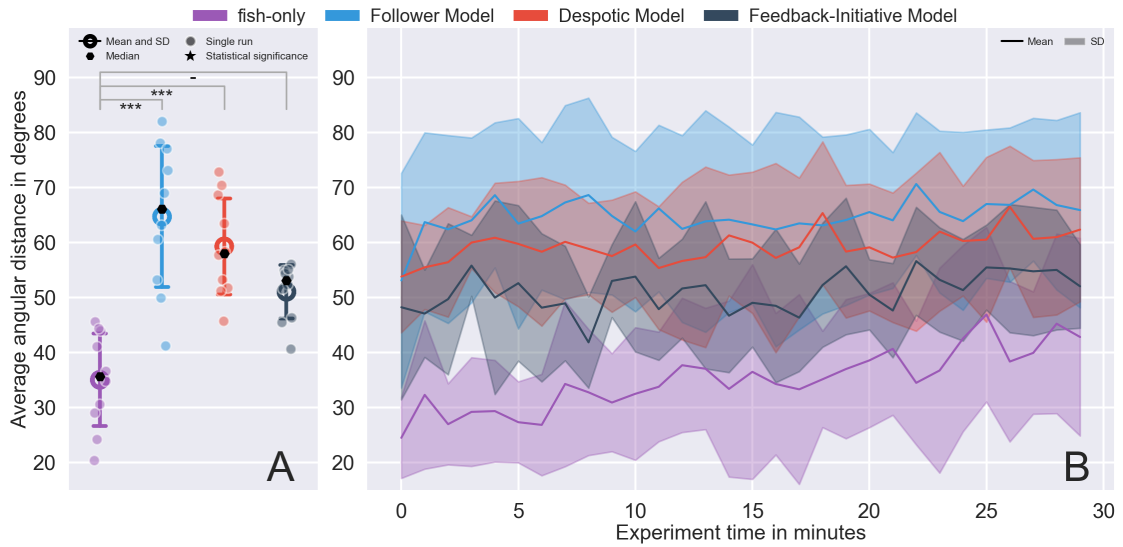


Figure 3.4: (A) **Average angular distance in degrees over ten runs** and (B) **average angular distance in 1-minute timesteps over all replicates**. Annotations of the statistical significance (Kruskal–Wallis test followed by Tukey’s honest significant difference criterion post hoc analysis) are marked with a dash or stars. The dash corresponds $p > 0.05$, a single star to $p < 0.05$, two stars to $p < 0.01$, three stars to $p < 0.001$, and four stars to $p < 0.0001$.

turns and a mean of 11.91 ± 2.77 . These distributions differ significantly from each other (KW test, $p < 0.001$ and $\chi^2 = 16.63$). An additional post hoc T-HSD analysis showed that: FM versus fish-only had significantly different mean rank ($p < 0.05$); DM versus fish-only also showed significantly different mean rank ($p < 0.001$); FIM versus fish-only showed no significant difference ($p > 0.99$) but FIM versus FM ($p < 0.05$) and DM ($p < 0.01$) showed a significant difference (a detailed table of the post hoc analysis is available in Table A.2 of Appendix A).

These results showed that FM’s and DM’s poor performance in terms of angular distance translated in poor performance in terms of collective U-turns. While this was to be expected for the DM that was instructing the robot to move CW, FM once again appeared to disrupt the collective dynamics of the shoal. More specifically, DM’s low number of U-turns demonstrated its ability to influence the collective decision-making rather than to participate in it (the shoal moves CW $\approx 65\%$ of the time, similar to Bonnet et al. (2018)). Conversely, there was no significant difference between FIM and the fish-only experiments regarding the collective U-turns. This indicates that FIM had strong biomimetic capabilities due to its design, that explicitly embedded the ability to follow but also initiate direction changes.

In addition to the collective U-turns, we also investigated the success rate of the robot to initiate a collective U-turn. In Fig. 3.6A we depict the percentage of successful U-turns that were owed to the robot’s motion, in Fig. 3.6B we depict the highest percentage of successful U-turns exhibited by any one individual taking part in the experiment, and in Fig. 3.6C we depict the percentage of the robot that was the most influential individual. We note that, in the case of the fish-only experiments we chose one random individual and we excluded the DM experiments




Figure 3.5: **Average number of collective U-turns per minute over all replicates.** Annotations of the statistical significance (Kruskal–Wallis test followed by Tukey’s honest significant difference criterion post hoc analysis) are marked with a dash or stars. The dash corresponds $p > 0.05$, a single star to $p < 0.05$, two stars to $p < 0.01$, three stars to $p < 0.001$, and four stars to $p < 0.0001$.

since the robot would never perform a U-turn. Intuitively, the above measurements can provide an estimate of the leadership characteristics of each model compared to the innate behavior of the zebrafish. Moreover, the distributions depicted in Fig. 3.6A and Fig. 3.6B provide, once again, insight on the degree to which the robot might have been perturbing or naturally interacting with the living individuals.

We performed KW test for the successful U-turn initiation distribution and obtained the values of $p < 0.05$ and $\chi^2 = 7.65$. A follow-up T-HSD post hoc test revealed that fish-only does not differ significantly from the FM ($p > 0.5$), while it did indeed differ significantly from FIM ($p < 0.5$). FM also differs significantly from FIM ($p < 0.05$). While Fig. 3.6B alone does not provide a lot of additional information (the distributions do not differ significantly KW test, $p > 0.05$), in combination with Fig. 3.6A we notice that the robot’s U-turn initiation success rate was very similar to the distribution for the fish-only individuals with the highest success rate. In Fig. 3.6C we quantified the latter in terms of the percentage that the robot acted as the leading individual and found that in FM experiments this corresponds to 20% and in FIM experiments to 70%. While FIM was clearly more successful in initiating a U-turn, FM’s success rate was greater than what could be expected by a following model. This is, was fact, directly linked to the densest point alternation (see Sec. 3.4), which could very well have triggered a U-turn for the robot if the densest centroid appeared in the reverse direction. Overall, this provides evidence that the FIM was capable of producing patterns that did not perturb the collective and simultaneously allowed it to have a leadership role with high probability.

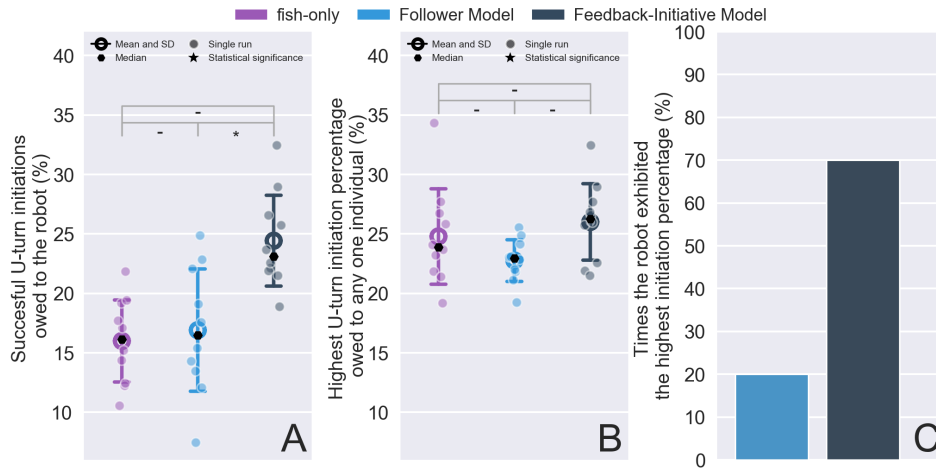


Figure 3.6: **U-turn initiation success rates.** (A) **successful U-turns that were initiated by the robot (or a random individual in the case of fish-only),** (B) **highest success rate in an experiment owed to any one individual, and** (C) **percentage of experiments in which the robot had the highest success rate.** Annotations of the statistical significance (Kruskal–Wallis test followed by Tukey’s honest significant difference criterion post hoc analysis) are marked with a dash or stars. The dash corresponds $p > 0.05$, a single star to $p < 0.05$, two stars to $p < 0.01$, three stars to $p < 0.001$, and four stars to $p < 0.0001$.

3.6.3 Transfer entropy

We complete this section by evaluating the information propagation capabilities of each model by resorting to information theory and more specifically to the use of TE (see section Sec. 3.5.2). In Fig. 3.7, we measured the TE for all shoal members (fish and robots) quantifying the average influence that individuals exerted on (outgoing TE) or received from (incoming TE) the others.

First, we analyzed the average outgoing entropy for all agents (Fig. 3.7A). Again, the distributions significantly differ from each other across the different treatments (KW test, $p < 0.00001$, and $\chi^2 = 26.78$). The fish-only condition, that showed the highest TE values, differs significantly from the FM (T-HSD post hoc test, $p < 0.01$) and DM (T-HSD post hoc test, $p < 0.0001$) but not from the FIM (T-HSD post hoc test, $p \approx 0.59$).

The lower performance observed in the mixed groups could be partly attributed to the robot’s slow response to stimuli or the models’ lack of locomotive aspects that might play an important role in good integration (*e.g.*, biomimetic locomotion patterns). Conversely, it is important to note that the amount of outgoing directional information exchanged did not significantly differ in the case of FIM versus fish-only which in turn implies that the robotic lure had a considerable impact on the shoal. In fact, FIM stood out compared to the rest of the models in terms of distribution similarity, therefore, we conclude that its biomimetic decision-making was indeed important when it came to propagating information within the shoal.

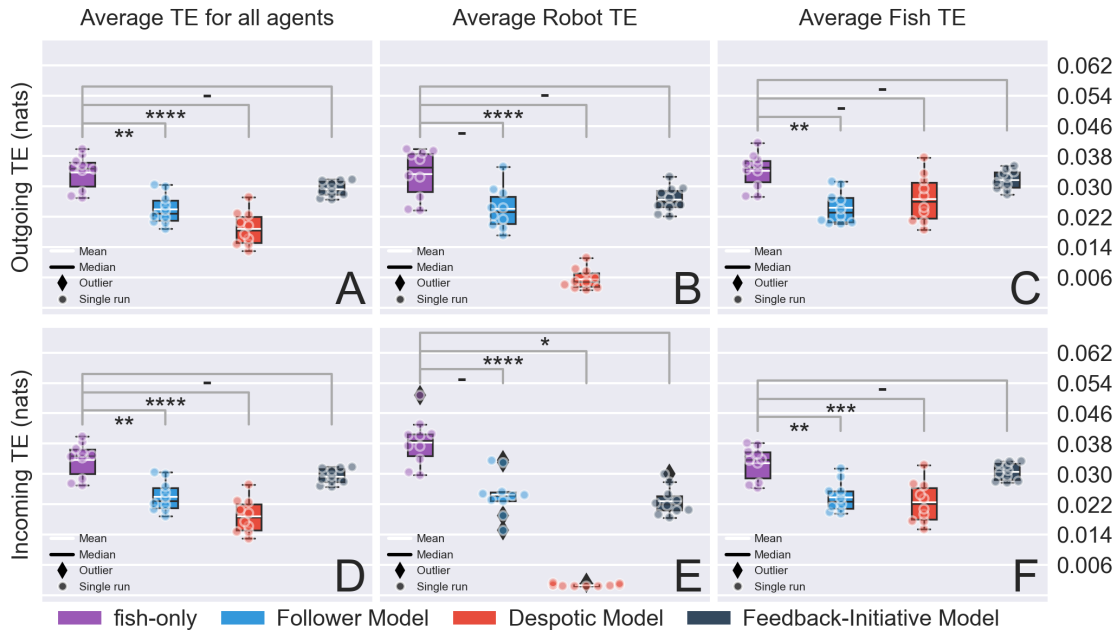


Figure 3.7: **Average transfer entropy (TE) for heading direction time-series.** The first row (*i.e.*, Fig. A, B, C) corresponds to the average outgoing TE of the shoal (*i.e.*, amount of TE from focal towards other individuals) and the second (*i.e.*, D, E, F) corresponds to the average incoming TE (*i.e.*, amount of TE from other individuals towards the focal). A & D: TE for the mixed group (all individuals are considered), B & E: TE only for the robotic agent (for fish-only experiments a random fish replaces the robot) and C & F: TE only for living individuals (*i.e.*, the robot is excluded in the computation and for fish-only experiments a random fish and excluded from the analysis). Annotations of the statistical significance (Kruskal–Wallis test followed by Tukey’s honest significant difference criterion post hoc analysis) are marked with a dash or stars. The dash corresponds $p > 0.05$, a single star to $p < 0.05$, two stars to $p < 0.01$, three stars to $p < 0.001$, and four stars to $p < 0.0001$. The complete pairwise comparisons can be found in the Supporting information Tables A.5, A.6, A.7, A.8, A.9, A.10.

To highlight the role played by the robot in the shoal dynamics, we separated the average outgoing TE of the robot (Fig. 3.7B) and the average outgoing TE of the fish (Fig. 3.7C). The intuition behind this threefold separation (average TE for all agents, average robot TE, average fish TE) is summarized as follows: (1) an overall estimate of how each model affected directional information transfer in the shoal, (2) a quantification of the robot’s interaction with the fish and (3) an evaluation of the perturbations in information transfer between the fish due to the presence of the robot. We note that, for the fish-only box-plots of Fig. 3.7B and 3.7C, we chose one fish at random, since no robot is used.

For the robot TE case in Fig. 3.7B, the distributions differ significantly (KW test, $p < 0.00001$ and $\chi^2 = 27.19$). A multiple comparison of the distributions showed that: fish-only did not differ significantly from FM (T-HSD post hoc test, $p \approx 0.11$) and FIM (T-HSD post hoc test, $p \approx 0.54$) but differed significantly from DM (T-HSD post hoc test, $p < 0.0001$).

Then, we computed the average outgoing TE exchange only among the fish for the different conditions (Fig. 3.7C). The KW test showed that the distributions differ significantly with each other (KW test, $p < 0.05$). The multiple comparisons post hoc test revealed that: fish-only differ significantly from FM (T-HSD post hoc test, $p < 0.01$), but not from DM (T-HSD post hoc test, $p > 0.05$) and FIM (T-HSD post hoc test, $p \approx 0.93$). In these cases, FIM performed closely to the control experiments. We believe that this could be linked to the degree of acceptance of the robot by the society. More specifically, if the robot does not perturb the directional information propagation, or ideally contributes to it, it might have higher chances to be accepted as a conspecific. In that respect, FIM seems to be the better model out of the ones we tested.

Similarly, we computed the average incoming TE for all individuals Fig. 3.7D, only the robot Fig. 3.7E, and only the fish for all conditions Fig. 3.7F. For the incoming TE for all agents, the results were identical to the one obtained for the average outgoing TE shown in Fig. 3.7A as the amount of information exchanged is preserved within the system but distributed differently among the fish and the robot.

Concerning the average incoming TE of the robot (see Fig. 3.7E), the distributions were significantly different, as observed for the average outgoing TE (KW test, $p < 0.0001$, $\chi^2 = 32.21$). From the complementary T-HSD post hoc analysis we obtained the values $p \approx 0.07$, $p < 0.0001$, $p < 0.05$ for fish-only versus FM, fish-only versus DM and fish-only versus FIM, respectively. In this case, the FIM under-performed compared to FM. However, this could be expected as the FM, a purely reactive model, was constantly instructing the robot to follow the fish while the FIM could lead the robot to take an initiative that contradicted the behavior of the fish.

For the average incoming TE of only the fish (Fig. 3.7F), we also observed a significant effect of the conditions on the average TE (KW test, $p < 0.001$). However, contrarily to the average outgoing TE, the multiple comparisons showed that fish-only significantly differs from the FM (T-HSD post hoc test, $p < 0.01$) and DM (T-HSD post hoc test, $p < 0.001$) but not from the FIM (T-HSD post hoc test, $p \approx 0.89$). These results confirmed that the robot controlled by the FIM did not impede the transfer of information between the fish.

3.7 Conclusion

Testing theoretical hypotheses in realistic conditions is an imperative step towards understanding the collective dynamics of natural systems. However, generating specific patterns that are valuable to validate those hypotheses requires sophisticated physical systems. In the case of animal studies, and specifically the study of zebrafish's group interactions, such systems must blend well enough in the shoal as to allow for natural and life-like interaction dynamics to emerge. Thus, apart from visual biomimetic cues, a robotic device ought to behave as close to the living creature as possible. In turn, this raises questions on the necessity of complex behavioral models in order for an artificial agent to socially interact with a high degree of integration in the group. Here, we showed that a model, that has been simplified to be implemented on a physical system, allowed the robot to establish an increased degree of

life-like interactions with a shoal of fish through a bidirectional communication scheme.

While one could assume that a simple following model that instructs a robotic agent to move towards the fish would succeed in “infiltrating” the shoal, we showed that our such a Follower model (FM), failed to do so. Intuitively, a lure that is not attempting to interact with the living agents is less appealing to them and at times is completely disregarded (in the collective decision-making sense). In fact, the results of the previous section indicate that this passive control scheme perturbed the behavior of the fish, as we observed greater mean inter-individual distance and at the same time fewer U-turns performed on the global scale. This was also validated by visual inspection of the corresponding experiment recordings. To trace the source of this failure, we used a TE metric to estimate the amount of information that is exchanged when this model is active on the robot. While such a metric can not be used to safely draw causal conclusions alone, the results indicate that there was a significantly different trend in the information flow for this model (see Fig 3.7) that could explain the dissimilarity on the global scale (*i.e.*, inter-individual distance, U-turns and successful U-turn initiation).

Conversely, the Despotic Model (DM), tests the response of the fish when the scenario is inverted, that is, when the robot is disregarding their decisions concerning the direction of movement. More specifically, we aimed to test two extreme cases and observed the responses for each one. As shown in Sec. 3.6.3, the DM also fails to capture the interest of the living individuals for long periods of time. However, the Feedback-Initiative model (FIM) managed to exhibit patterns that proved to lead to similar dynamics on a global scale. Especially the results depicted in Fig. 3.5 concerning the collective U-turns, showed that the living individuals interact with the robot and between them in a marginally different manner to that of groups of fish (only). Similarly, the TE measurements indicated that the robot managed to establish stronger communication channels with the fish that could, in turn, explain the similarity of the resulting U-turn distributions. Notably, We noticed that the robot had a leading role (*i.e.*, a direction change of the robot was likely to propagate to the remaining group members) for the majority of the experiments and had a very similar influence to the most influential fish individual of the fish-only experiments.

This highlights the importance of encoding bidirectional (*i.e.*, closed-loop) behaviors in robot controllers to create biohybrid groups. While this notion is further studied in following chapters, the results of this section establish the foundation for the work conducted in 2 directions: 1) here we tested a simplified 1D model of interaction, but our findings indicated that bridging the gap between this model and the real behavior of fish requires more sophisticated models of social interaction, and 2) the robotic system (composed of the mobile robot and biomimetic lure) must also exhibit realistic traits (visual and in terms of motion; physics and communication cue gap in Fig. 1.1) to be truly accepted and form a biohybrid group.

The simplified models we have used in this study inevitably raise questions about the robot’s ability to blend seamlessly with the fish group and meaningfully interact with its neighbors.

This ambiguity is a common thread in much of the existing literature on fish-robot interactions. In essence, there is a noticeable lack of evidence pointing to the specific components of the biohybrid system that might cause significant discrepancies, leading to interaction dynamics that substantially deviate from those observed in natural settings. This indicates a pressing need for closer examination of the three factors contributing to the biomimicry gap (as outlined in Fig. 1.1) with the aim of creating biohybrid groups that more closely emulate natural fish populations. For instance, in this and previous studies employing the same system Bonnet et al. (2018, 2016), the lure used is a commercially available fishing lure which, despite superficial resemblances to the zebrafish, possesses markedly different physiological properties (*e.g.*, see the tail in Fig. 3.2A). Likewise, as discussed in subsequent chapters (*e.g.*, Chapter 5 and Appendix C.1), the FishBot v4.4 exhibits limitations in generating diverse dynamics across various environments, further widening the physics gap. Keeping these limitations in mind, in subsequent chapters, we first systematically address the social interaction gap, and then we follow up with ways to mitigate the physics and communication cue gap. We then return to biohybrid experiments to quantitatively evaluate our success in bridging the biomimicry gap.

4 Designing and benchmarking social interaction models: from analytical to machine learning models, and from short to long timescale dynamics generation

This section's content is adapted from the following article (Under review/Preprint available):

- Papaspyros, V., Escobedo, R., Alahi, A., Theraulaz, G., Sire, C. and Mondada, F., 2023. Predicting long-term collective animal behavior with deep learning. bioRxiv, pp.2023-02. <https://doi.org/10.1101/2023.02.15.528318> (reuse authorised under a non-exclusive distribution licence) – Papaspyros et al. (2023b)
 - Vaios Papaspyros's contribution: Conception, data curation, formal analysis, investigation, methodology, software, visualization, writing – original draft, writing – review & editing.
- The article's work was funded by the Swiss National Science Foundation project "Self-Adaptive Mixed Societies of Animals and Robots" (Grant No. 175731). Our collaborators from the Université Toulouse – Paul Sabatier were funded by the French National Research Agency (ANR-20-CE45-0006-01).^a
- Ethics: The experiments conducted with *H. rhodostomus* were approved by the Ethics Committee for Animal Experimentation of the Toulouse Research Federation in Biology no. 1 and comply with the European legislation for animal welfare. The experiments conducted with *D. rerio* were approved by the state ethical board of the Department of Consumer and Veterinary Affairs of the Canton de Vaud (SCAV) of Switzerland (authorization no. 2778).

^aAll the code concerning the data pre-processing, neural networks, and plot scripts are publicly available in <https://doi.org/10.5281/zenodo.7634912>. Experimental and generated data are available in <https://doi.org/10.5281/zenodo.7634687>.

4.1 Introduction

As discussed in Chapter 1, and after conducting a baseline study in Chapter 3, in this chapter we address the issue of designing and benchmarking high-fidelity models of social interaction. As shown in the previous chapter, models are more likely to command robots that integrate to animal groups when they are efficient in their bidirectional information exchange. Intuitively, the more accurate the interaction models that are subsequently used to command a robot, the smaller the gap to be traversed. In the remaining of this chapter, we validate this intuition with experimental data.

We demonstrate that ML models can generate realistic synthetic data with minimal biological assumptions, and that they allow to accelerate and generalize the process of collective behavior modeling. More specifically, we present a social interaction model composed of a deep neural network that captures both the short- and long-term dynamics observed in schooling fish. We apply our approach to pairs of rummy-nose tetra (*Hemigrammus rhodostomus*) swimming in a circular tank, and show that it can also be applied to fish species with similar burst-and-coast swimming (zebrafish; *Danio rerio*). Our ML model is benchmarked against the state-of-the-art analytical model for this species Calovi et al. (2014), showing that it performs as well as the latter, even for very subtle quantities measured in the experiments. Moreover, we also introduce a systematic methodology to stringently test the results of an analytical or ML model against experiment, at different timescales, and in the context of animal collective motion.

This chapter serves as a stepping stone to closing the biomimicry gap and generating accurate robot controllers, by significantly reducing the social interaction gap (outlined in Fig. 1.1). Naturally, fully bridging this gap is an iterative process requiring thorough validation. As noted in previous paragraphs, in this chapter, we also introduce a set of metrics to measure the extent to which a model can bridge the biomimicry gap between simulation and real-life observation. Following chapters further delve into how this is translated to a biohybrid group composed of fish and a robot.

4.2 Experimental data

The trajectory data used in this chapter were originally published in Calovi et al. (2018) for *Hemigrammus rhodostomus* swimming either alone or in pairs in a 50 cm diameter circular tank Calovi et al. (2018). This species is characterized by a burst-and-coast swimming mode, where the fish perform a succession of sudden short acceleration periods (kick), each followed by a longer gliding period almost in a straight line. The instant of the kicks, when heading changes take place, are assimilated to decision instants.

The dataset corresponds to 40 hours of video recordings at 25 Hz. Fish are tracked with idTracker Pérez-Escudero et al. (2014), an image analysis software which extracts the 2D trajectories of all individuals. Occasionally, the tracking algorithm is temporarily unable to report positions accurately. This can be due to small fluctuations in lighting conditions,

fish standing still or moving at very low speed, fish swimming very close to the surface, to the border, or to each other. These instances are corrected using several filtering processes. Since our analyses focus on social interactions, we remove the periods during which fish are inactive. Fish body length (BL) is of about 3.5 cm, and the intervals of time during which fish velocity is less than 1 BL/s are removed. Large leaps in fish trajectories during which fish move by more than 1.5 BL \approx 5.25 cm between two consecutive frames, meaning that fish move at almost 65 cm/s, are also identified and removed, as they result from tracking errors. Finally, a linear interpolation procedure fills missing points. The final dataset used in this work represents approximately 10 hours of trajectory data for pairs of *H. rhodostomus*. For this work, trajectories of the original dataset are resampled with a timestep of $\Delta t = 0.12$ s instead of the original 0.04 s provided by the camera, and data points have been converted from pixel space to a normalized $[-1, 1]$ range to facilitate the training of the networks. The subsampling rate was chosen carefully to reduce the random noise between subsequent camera frames at the very short timescale of 0.04 s, while maintaining a sufficiently small timestep to study and model the social interactions.

4.3 Quantification of individual and collective behavior in pairs of fish

We use a set of observables to quantify how close the results of the models are from the measures obtained in the experiments Calovi et al. (2018); Escobedo et al. (2020); Jayles et al. (2020). These observables constitute a challenging benchmark when designing and testing a model. In the case of deep learning techniques, those observables also serve as means to partially explain what the algorithm has learned. In both cases, the observables constitute a stringent validation test.

Let us first define the temporal variables characterizing the individual and collective behavior of the fish. Fig. 4.1A shows two fish swimming in a circular tank of radius $R = 25$ cm. The position vector of a fish i at time t is given by its Cartesian coordinates $\vec{u}_i(t) = (x^i(t), y^i(t))$ in the system of reference, centered at the center of the tank $C(0, 0)$. The components of the velocity vector $\vec{v}_i(t) = (v_x^i(t), v_y^i(t))$ are given by $v_x^i(t) = (u_x^i(t) - u_x^i(t - \Delta t)) / \Delta t$ (similar expression for $v_y^i(t)$). The heading angle of the fish is assumed to indicate its direction of motion and is therefore given by the angle that the velocity vector forms with the horizontal, $\phi_i(t) = \text{ATAN2}(v_y^i(t), v_x^i(t))$.

The motion of a given fish i is then described using the three following instantaneous variables: the speed, $V_i(t) = \|\vec{v}_i(t)\|$, the distance of the fish to the wall, $r_w^i(t) = R - \|\vec{u}_i(t)\|$, and the angle of incidence of the fish to the wall, $\theta_w^i(t)$, defined by the angle formed by the velocity vector and the normal to the wall: $\theta_w^i(t) = \phi_i(t) - \text{ATAN2}(y^i(t), x^i(t))$, see Fig. 4.1A.

When there are two fish i and j in the tank, their relative motion is characterized by means of three variables: the distance between fish, $d_{ij}(t) = \|\vec{u}_i(t) - \vec{u}_j(t)\|$, the difference between

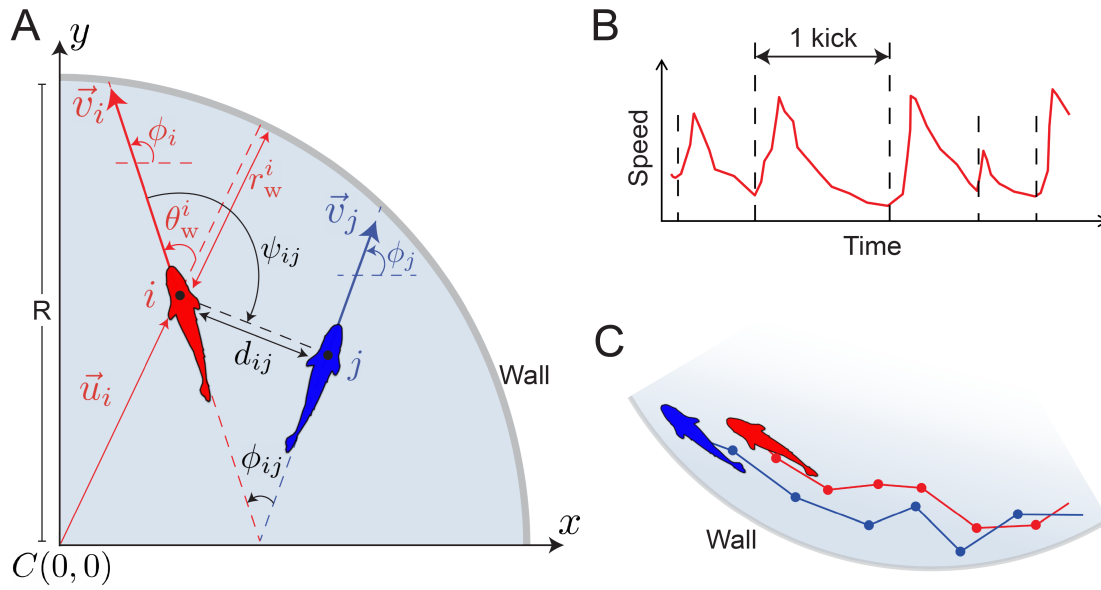


Figure 4.1: **A.** Individual and collective variables characterizing the instantaneous state of an individual (focal fish in red) and its pairwise relation with a neighbor (blue): distance to the wall $r_w^i(t)$, angle of incidence to the wall $\theta_w^i(t)$, heading angle $\phi_i(t)$, distance between individuals $d_{ij}(t)$, difference of heading angles $\phi_{ij}(t)$, and angle of perception $\psi_{ij}(t)$. Positive angles (curved arrows) are defined in the anti-clockwise direction, starting from the positive semi-axis of abscissas. The radius of the circular setup is $R = 25$ cm. For visualization purposes, the size of fish is not to scale with the tank. **B.** Typical profile of the fish speed, $V(t)$, showing the typical sequence of kicks (abrupt accelerations followed by longer gliding phases). **C.** Trajectories of two fish close to the wall due to their burst-and-coast swimming mode. The dots in the trajectories denote the instants of the kicks, where fish decision-making is assumed to take place.

their heading angles, $\phi_{ij}(t) = \phi_j(t) - \phi_i(t)$, which measures the degree of alignment between both fish, and the angle of view, $\psi_{ij}(t)$, which is the angle with which fish i perceives fish j , and which is generally independent of $\psi_{ji}(t)$. See Fig. 4.1A for the graphical representation of these quantities. The angle of perception of the fish also allows us to define the notion of *geometrical leadership* for two fish: fish i is the *geometrical leader* (and therefore, j is the *geometrical follower*) if $|\psi_{ij}(t)| > |\psi_{ji}(t)|$, meaning that i has to turn by a larger angle to face j than the angle that j has to turn to face i . In practice, these definitions of the geometrical leader and follower provide a precise and intuitive characterization of a fish being ahead of the other. Note that being the leader or the follower is an instantaneous state that can change from one kick to the other.

These 6 quantities $V_i(t)$, $r_w^i(t)$, $\theta_w^i(t)$, $d_{ij}(t)$, $\phi_{ij}(t)$, and $\psi_{ij}(t)$ being defined, the measure of their probability distribution functions (PDF) constitutes a set of observables probing the individual and collective instantaneous fish dynamics in a fine-grained and precise manner. The PDF of $V_i(t)$, $r_w^i(t)$, $\theta_w^i(t)$ probe the behavior of a focal fish sampled over the observed dynamics, and are hence called *instantaneous individual observables*. The PDF of $d_{ij}(t)$,

$\phi_{ij}(t)$, and $\psi_{ij}(t)$ characterize the correlations between 2 fish *at the same time* t and are hence called *instantaneous collective observables*. These 3 collective observables can be easily generalized to a group of arbitrary size $N > 2$, by considering i and j as pairs of nearest neighbors, or pairs of second-nearest neighbors (or even farther neighbors), or even averaging them over all pairs in the group (then probing the size, the polarization, and the anisotropy of the group). Ultimately, comparing experimental results and model predictions for these individual and collective observables constitute a stringent test of a model.

Moreover, to characterize the *temporal correlations* arising in the dynamics, we make use of 3 additional observables involving quantities measured *at two different times*, for a given focal fish Jayles et al. (2020): the mean-squared displacement $C_X(t)$, the velocity autocorrelation $C_V(t)$, and, especially challenging, the autocorrelation of the angle of incidence to the wall $C_{\theta_w}(t)$, defined respectively by

$$C_X(t) = \left\langle [\vec{u}_i(t+t') - \vec{u}_i(t')]^2 \right\rangle, \quad (4.1)$$

$$C_V(t) = \left\langle \vec{v}_i(t+t') \cdot \vec{v}_i(t') \right\rangle, \quad (4.2)$$

$$C_{\theta_w}(t) = \left\langle \cos \left[\theta_w^i(t+t') - \theta_w^i(t') \right] \right\rangle, \quad (4.3)$$

where $\langle w(t) \rangle$ is the average of a variable $w(t)$ over all reference times t' (assumption of a stationary dynamics, where correlations between two times depend solely on their time separation), over all focal fish, and over all experimental runs. Note that, although $C_V(t)$ can be shown to be the second derivative of $C_X(t)$, both quantities are measured independently. In principle, these correlation observables can also be generalized to probe the (collective) time correlations between the two different fish (or between nearest neighbors in a group of $N > 2$ individuals). For instance, one could consider $C_V(t) = \langle \vec{v}_i(t+t') \cdot \vec{v}_j(t') \rangle$, where the average is now over nearest neighbor pairs. However, in the present study, we will limit ourselves to the study of the 3 (individual) correlation functions listed in Eqs. 4.1-4.3.

4.4 Analytical and deep learning models of fish behavior

The *H. rhodostomus*' motion is characterized by a burst-and-coast swimming pattern (see 1.4.1). In the following section, we leverage aspects of this swimming pattern to “inform” our models and allow them to generate more realistic social interactions in simulation. Since the fish typically swim in close proximity, the social interactions between them reflect the combined tendency to align with and follow their neighbors while at the same time maintaining a safe distance with the wall. At a given kicking instant, only a few neighbors (one or two) have a relevant influence on the behavior of a fish Lei et al. (2020). The decision-making of fish displaying a burst-and-coast swimming mode can thus be reproduced by considering only pairwise interactions. Naturally, if one considers pairs of fish, like here, it suffices to consider the relative state of the neighboring fish (relative position and velocity) and the effect of the distance and the relative orientation to the wall Calovi et al. (2018); Escobedo et al. (2020).

Analytical Burst-and-Coast model

The Analytical Burst-and-Coast model (hereafter called ABC model) quantitatively reproduces the interaction dynamics of *H. rhodostomus* swimming alone or in pairs under the hypothesis that fish decision-making times correspond exactly to their kicking times, that is, the new direction of movement, the duration, and the length of the kick are decided precisely at the end of the previous kick Calovi et al. (2018).

Given a pair of agents i and j at a respective state (\vec{u}_j^n, ϕ_j^n) and (\vec{u}_i^n, ϕ_i^n) at time t^n , the state of agent i at the next instant of time t_i^{n+1} is given by

$$t_i^{n+1} = t_i^n + \tau_i^n, \quad (4.4)$$

$$\phi_i^{n+1} = \phi_i^n + \delta\phi_i^n, \quad (4.5)$$

$$\vec{u}_i^{n+1} = \vec{u}_i^n + l_i^n \vec{e}(\phi_i^{n+1}), \quad (4.6)$$

where $\vec{e}(\phi_i^{n+1})$ is the unitary vector pointing in the heading direction ϕ_i^{n+1} , τ_i^n and l_i^n are the duration and length of the n -th of agent i kick, and $\delta\phi_i^n$ is the heading change of agent i . During the gliding phase, the kick length is defined $l_i^n = v_i^n \tau_0 (1 - \exp[-\frac{\tau_i^n}{\tau_0}])$, where the speed is empirically found to decrease exponentially with a dissipation time $\tau_0 = 0.8$ s. That is, v_i^n , l_i^n , and t_i^n are generated from bell-shaped probability density functions that are in agreement with experimental data from *H. rhodostomus* Calovi et al. (2018).

The heading angle change $\delta\phi_i^n$ is the result of three effects: the effect of the wall, the effect induced by the social interactions with the other fish (repulsion/attraction and alignment), and the natural spontaneous fluctuations of fish motion (cognitive noise). The social influence depends only on the relative state of both agents, determined by the triplet $(d_{ij}, \psi_{ij}, \phi_{ij})$. The derivation of the shape and intensity of the functions involved in $\delta\phi_i^n$ is based on physical principles of symmetry of angular functions and a data-driven reconstruction procedure detailed in Calovi et al. (2018) for the case of *H. rhodostomus* and in Escobedo et al. (2020) for the general case of animal groups. Thus, Eq. 4.5 can be further expanded as:

$$\delta\phi_i^n = \delta\phi_{R,i}^n(r_{w,i}^n) + \delta\phi_{w,i}^n(r_{w,i}^n, \theta_{w,i}^n) + \delta\phi_{Att,ij}^n(d_{ij}^n, \psi_{ij}^n, \phi_{ij}^n) + \delta\phi_{Ali,ij}^n(d_{ij}^n, \psi_{ij}^n, \phi_{ij}^n) \quad (4.7)$$

where $\delta\phi_{R,i}^n(r_{w,i}^n) = \gamma_R g$, with g a Gaussian random variable of zero mean and unit variance and γ_R is the heading direction fluctuation strength which is found to be of the order of 0.35 rads for single and pairs of fish, but may vary for larger groups or different species. Then, $\delta\phi_{w,i}^n(r_{w,i}^n, \theta_{w,i}^n) = \gamma_{w,i} \sin(\theta_{w,i}^n) f_{w,i}^n(r_{w,i}^n)$ is the repulsive force exerted by the wall on the fish i at time t^n , with $f_{w,i}^n(r_{w,i}^n) \in [0, 1]$ a regulatory parameter to adjust the repulsive effect of the wall such that $f_{w,i}^n(r_{w,i}^n) \rightarrow 0$ when the individual i is far from the wall and $f_{w,i}^n(0) = 1$, and γ_w the repulsion intensity. The expected attraction and alignment are defined as the product of

three functions as follows:

$$\delta\phi_{\text{Att},ij}^n(d_{ij}^n, \psi_{ij}^n, \phi_{ij}^n) = F_{\text{Att}}^n(d_{ij}^n)O_{\text{Att}}^n(\psi_{ij}^n)E_{\text{Att}}^n(\phi_{ij}^n) \quad (4.8)$$

$$\delta\phi_{\text{Ali},ij}^n(d_{ij}^n, \psi_{ij}^n, \phi_{ij}^n) = F_{\text{Ali}}^n(d_{ij}^n)O_{\text{Ali}}^n(\psi_{ij}^n)E_{\text{Ali}}^n(\phi_{ij}^n) \quad (4.9)$$

where O are odd functions (*e.g.*, as the focal fish turns by equal angles but of the opposite sign), and E functions are even. These functions have been computationally fit on *H. rhodostomus* data obtained during spontaneous movements in a similarly-sized, circular, arena. Here, the parametrization of these functions closely follows the original definition in Calovi et al. (2018), where:

$$F_{\text{Att}}^n(d_{ij}^n) = \frac{d_{ij}^n - 3}{3(1 + (d_{ij}^n/20)^2)} \quad (4.10)$$

$$O_{\text{Att}}^n(\psi_{ij}^n) = \sin(\psi_{ij}^n)(1 - 0.33 \cos(\psi_{ij}^n)) \quad (4.11)$$

$$E_{\text{Att}}^n(\phi_{ij}^n) = 1 \quad (4.12)$$

$$F_{\text{Ali}}^n(d_{ij}^n) = \exp[-(d_{ij}^n/20)^2] \quad (4.13)$$

$$O_{\text{Ali}}^n(\psi_{ij}^n) = \sin(\phi_{ij}^n)(1 + 0.33 \cos(2\phi_{ij}^n)) \quad (4.14)$$

$$E_{\text{Ali}}^n(\phi_{ij}^n) = 1 + 0.6 \cos(\psi_{ij}^n) - 0.32 \cos(2\psi_{ij}^n) \quad (4.15)$$

Starting from the initial condition (\vec{u}_i^0, ϕ_i^0) of fish i , the length, and the duration of its next kick, l_i^0 and τ_i^0 , are sampled from the experimental distributions obtained in Calovi et al. (2018). Then, the timeline t_i^1 of fish i is updated with Eq. 4.4, the heading angle of the next kick ϕ_i^1 is calculated with Eq. 4.5, and the position of the fish at the end of the kick \vec{u}_i^1 is obtained with Eq. 4.6. As kicks of different fish are asynchronous, the next kick can be performed by any of the two fish. Each fish has thus its own timeline, but is subject, at each of its kicks, to the evolution of the other fish along its own kicks.

The ABC model is, therefore, a discrete model that generates kick events instead of continuous time positions. To directly compare with the DLI model presented in the next section, which is a continuous time model, we re-sampled the trajectories made of kick events produced by the ABC model and build continuous time trajectories with a timestep of size $\Delta t = 0.12$ s. We produced trajectories that add up to a total of approximately 16.6 hours of duration, that is, 500,000 timesteps. To do so, we exploit the viscous dynamics (due to water grad) for times $0 \leq t \leq t_i^n$ to derive the following continuous time expression:

$$\tilde{u}_i(t_i^n + t) = \tilde{u}_i^n + l_i^n \frac{1 - \exp[-\frac{t}{\tau_0}]}{1 - \exp[-\frac{\tau_i^n}{\tau_0}]} \tilde{e}(\phi_i^{n+1}) \Rightarrow \quad (4.16)$$

$$\Rightarrow \tilde{u}_i(t_i^n + \tau_i^n) = \tilde{u}_i(t_i^{n+1}) = \tilde{u}_i^n + l_i^n \tilde{e}(\phi_i^{n+1}) = \tilde{u}_i^{n+1}, \quad (4.17)$$

Deep Learning Interaction model

The Deep Learning Interaction model (hereafter called DLI model) consists of an Artificial Neural Network (ANN) which is fed with a set of variables characterizing the motion of *H. rhodostomus* and which provides the necessary information to reproduce the social interactions of these fish by estimating their motion along timestep of length $\Delta t = 0.12$ s. At time t , the DLI is designed to take sequences of states as input to capture the short- and long-term dynamics. Then, it generates predictions for the acceleration components of the fish at the following timestep $t + \Delta t$.

For the DLI model, the state of an agent i at time t is defined by

$$\mathbf{s}_i(t) = (\tilde{u}_i(t), \tilde{v}_i(t), r_w^i(t)) \in \mathbb{R}^5. \quad (4.18)$$

The state of an agent includes redundant information: in a fixed geometry, r_w^i can be deduced from \tilde{u}_i , and \tilde{v}_i^n from the input sequence $\tilde{u}_i^{n-4}, \dots, \tilde{u}_i^n$. This redundancy is intended to facilitate the training process of the neural network.

The system's state $\mathbf{S}(t)$ is then defined as the combination of both agent states, in addition to their inter-individual distance $d_{ij}(t)$ (also a redundant variable):

$$\mathbf{S}(t) = (\mathbf{s}_i(t), \mathbf{s}_j(t), d_{ij}(t)) \in \mathbb{R}^{11}. \quad (4.19)$$

Fig. 4.2 shows the structure of the ANN, consisting of 7 layers: two Long-Short Term Memory (LSTM) layers Hochreiter and Schmidhuber (1997), and 5 fully connected (Dense) layers.

The first LSTM layer consists of 256 neurons and is located at the input of the ANN, where it receives the sequence of the 5 last states of the system, *i.e.*, a matrix of dimension 5×11 : $(\mathbf{S}(t-4), \dots, \mathbf{S}(t))$. This history length of 4 timesteps (0.48 s) is borrowed from the biology of the fish: as already mentioned, the time it takes for a fish to display its characteristic behavior, a kick, is 0.5 s Calovi et al. (2018), therefore, we input the current state plus the states that correspond to the average duration of a kick. The output of the first LSTM is then gradually reduced in dimension by two successive dense layers, and then scaled up again with a second LSTM, whose configuration is also based on a history of 5 states. Then, two other dense layers are used to reduce the dimension of the output of the second LSTM, and a last dense layer is applied to provide the final output of the ANN. More details about the configuration of the ANN are given in Table B.7.

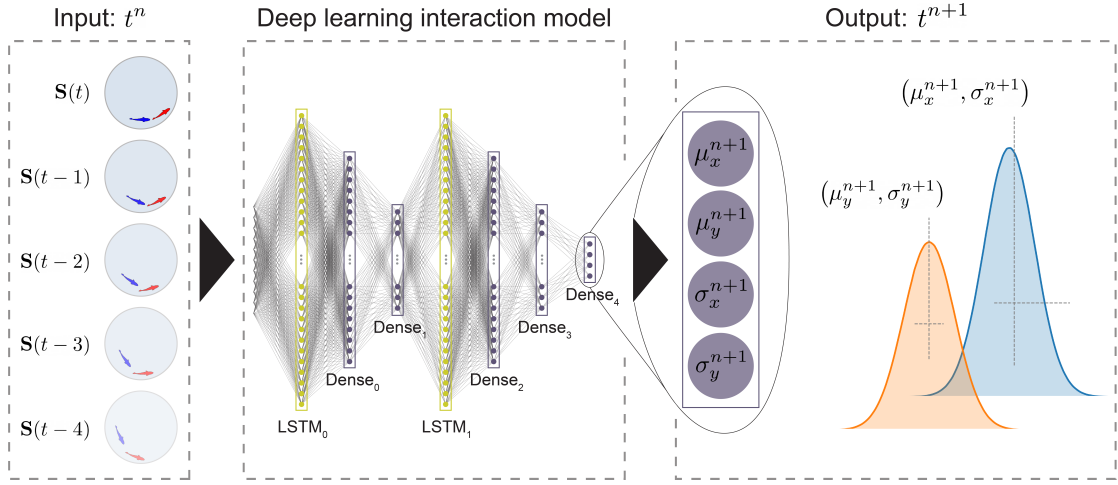


Figure 4.2: Structure of the Artificial Neural Network (ANN) used in the DLI model. From left to right: *Input* of the ANN: the 5 last states, $(\mathbf{S}(t-4), \dots, \mathbf{S}(t))$ at time t . Where $\mathbf{S}(t) = (\mathbf{s}_i(t), \mathbf{s}_j(t), d_{ij}(t)) \in \mathbb{R}^{11}$ and each state is parametrized as $\mathbf{s}_i(t) = (\bar{u}_i(t), \bar{v}_i(t), r_w^i(t)) \in \mathbb{R}^5$; the 7 layers (two LSTM layers and 5 Dense Layers) capturing the social dynamics; *Output*: the two pairs of values (μ_x, σ_x) and (μ_y, σ_y) corresponding respectively to the mean and standard deviation of the probability distribution function (assumed to be Gaussian) of each component a_x and a_y of the instantaneous acceleration vector \vec{a} at time $t+1$, constituting the prediction of the DLI model.

The output of the ANN consists of two pairs of values, (μ_x, σ_x) and (μ_y, σ_y) , corresponding to the expected value and standard deviation of the x and y components of the predicted acceleration, which are assumed to be Gaussian distributed Chua et al. (2018), as actually found for *H. rhodostomus* Calovi et al. (2018). Hence, the predicted acceleration of the agent, $\vec{a} = (a_x, a_y)$, can be written

$$a_x = \mu_x + \sigma_x g_x, \quad a_y = \mu_y + \sigma_y g_y, \quad (4.20)$$

where g_x and g_y are independent standard Gaussian random variables drawn from $\mathcal{N}(0, 1)$. Then, the velocity vector of the agent i at the time t^{n+1} is given by

$$\vec{v}_i^{n+1} = \vec{v}_i^n + \Delta t \vec{a}_i^n, \quad (4.21)$$

and the position of the agent is updated according to

$$\vec{u}_i^{n+1} = \vec{u}_i^n + \Delta t \vec{v}_i^{n+1}. \quad (4.22)$$

We note that in the DLI model, the predicted variance of the acceleration accounts for the fish intrinsic spontaneous behavior exhibited during their decision process (cognitive noise), and hence translates the fact that 2 real (or modeled) fish will not act the same if put twice in the same given state characterized by Eq. 4.19.

The *prediction* of the ANN for at time t^{n+1} is thus a vector of dimension 1×4 that can be written as $(\vec{\mu}_{\text{pred}}^{n+1}, \vec{\sigma}_{\text{pred}}^{n+1})$, where

$$\vec{\mu}_{\text{pred}}^{n+1} = (\mu_x^{n+1}, \mu_y^{n+1}) \quad \text{and} \quad \vec{\sigma}_{\text{pred}}^{n+1} = (\sigma_x^{n+1}, \sigma_y^{n+1}). \quad (4.23)$$

The ANN is then trained to approach the *real/observed* values $\vec{\mu}_{\text{real}}^{n+1}$ by means of the Adaptive Moment Estimation Optimizer (Adam) with a time-decaying learning rate $\lambda = 10^{-4}$ and a negative log-likelihood loss function ℓ defined in terms of the prediction error $\vec{\epsilon}_{n+1} = \vec{\mu}_{\text{pred}}^{n+1} - \vec{\mu}_{\text{real}}^{n+1}$ and the standard deviations as follows Kingma and Ba (2014):

$$\ell(\vec{\epsilon}_{n+1}, \vec{\sigma}^{n+1}) = \frac{1}{2} \sum_{n=1}^{N_h} [\vec{\epsilon}_{n+1}]^T \mathcal{C}^{-1}(\vec{\sigma}^{n+1}) \vec{\epsilon}_{n+1} + \frac{N_h}{2} \log \|\mathcal{C}(\vec{\sigma}^{n+1})\|, \quad (4.24)$$

where N_h is the number of timesteps in the history of the input of the ANN (here $N_h = 5$) and \mathcal{C} is a diagonal covariance matrix with the values of $\vec{\sigma}_{\text{pred}}^{n+1}$ in the diagonal and zeroes elsewhere.

The training of the ANN is carried out with a subset of the experimental dataset. More specifically, the training process is given a budget of 45 epochs with a batch size of 512 samples on a dataset that was split 80%, 15%, and 5% for training, validation, and test, respectively. Then, the DLI is used to produce trajectories of 500,000 timesteps of size $\Delta t = 0.12$ s, as done with the ABC model. At the beginning of the simulation, each agent is given a copy of the DLI model and both agents are initialized with a random 5 timesteps long trajectory. At each timestep t^n , the state vector $\mathbf{S}(t^n)$ is built and introduced in the network, which provides the estimated instantaneous acceleration distributions at time t^{n+1} . Then, the acceleration is evaluated according to Eq. 4.20, and the next positions of the agents \vec{u}_i^{n+1} and \vec{u}_j^{n+1} are obtained from the equations of motion, Eqs. 4.21, 4.22.

Designing the DLI model

Designing and selecting an appropriate ANN structure to model a system is for the most part non-trivial and requires either an extensive search through automatic methods (*e.g.*, neuro-evolution Martín et al. (2007); Mwaura and Keedwell (2015); Sekaj et al. (2019)) or an exhaustive number of empirical attempts for very specific applications Cazenille et al. (2019); Heras et al. (2019); Costa et al. (2020). Here, we followed a hybrid approach consisting in empirically designing an ANN based on biological insight and automatically searching for its optimal structure by bootstrapping the search. Once we established this initial model, we performed an automated search for similar neural networks using the same input and output for different combinations of *i*) the number of layers, *ii*) the size of the layers, and *iii*) the activation functions (*i.e.*, transfer functions tasked with mapping the inputs of a neuron to a single weighted output value passed to the next layer). The search included a total of 96 neural network structures (see Table B), out of which the ANN shown above is the best performing.

Three notable categories of networks were considered: *i*) non-probabilistic networks that only generate $(\mu_x^{n+1}, \mu_y^{n+1})$ (and hence, not explicitly including the cognitive noise), *ii*) probabilistic networks that do not have memory cells (hence, missing the fact that fish are gliding passively on a timescale of order 0.5 s), and *iii*) probabilistic networks that implement memory thanks to LSTM layers. Non-probabilistic networks (*i*) provide the mean value of the components of the acceleration for the next timestep with high accuracy, but miss the essential variability that is intrinsic to the spontaneous behavior of fish and which allows for the emergence of social interactions. Probabilistic networks without memory (*ii*) are able to partly capture this intrinsic variability, but do not fully capture the non-linear nature of the problem (see Fig. 4.9 and Video 4 of Appendix B). Finally, probabilistic networks with memory (*iii*) performed generally well, and we found that the structure used in the DLI model consistently provides the best results for the number of epochs set for training and for the ANNs considered by the automatic search.

Our search approach revealed the existence of two crucial ingredients that must be considered in the model, both accounting for biological characteristics of fish behavior observed experimentally. First, the neural network must be fed with information covering the typical timescale along which relevant changes take place in the behavior of the fish. Since real fish kicks last 0.5 s on average, the NN needs information about the fish behavior over time intervals of at least this duration (that is, 4 to 5 timesteps of 0.12 s). However, we found that using longer vector lengths (up to 10 timesteps) for the case of *H. rhodostomus* does not lead to any significant improvement in the results, while considerably increasing the training time. Second, the output of the network must contain a sufficiently wide diversity of predictions so that the agents reproduce the high variability of responses that fish display when behaving spontaneously and reacting to external stimuli.

ANNs without memory tend to make too similar predictions, and agents do not initiate the typical direction changes that are observed in the experiments. A possible solution could be to add some phenomenological noise to the predictions of the NN. However, this would result in an unrealistic behavior, albeit an improvement over not adding noise at all. For example, when a fish swims close to the wall, it does not have the same liberty to turn toward or away from the wall, which would not be captured by a too crude implementation of the fish cognitive noise. Our approach accounts for this behavioral uncertainty for each state (position, velocity, distance to the neighbor and to the wall) and for both degrees of freedom during the training phase of the ANN, being therefore able to capture these complex behavioral patterns. The performance of the two variants is depicted in Fig. 4.9.

Rejection procedures

Notably, none of these models is explicitly encoding the existence of the circular wall bounds ($R = 25$ cm). On one hand, the ABC exploits the distance to the wall in order to compute its influence on the individual's decision-making by means of repulsion, but does not explicitly guarantee that the individual will always choose to stay within the bounds. On the other hand,

the DLI’s black-box nature does not provide a lot of evidence with respect to the underlying dynamics that it learned, however, the DLI is explicitly given the distance to the wall for the focal and neighboring individual, implicitly guiding it to learn this constraint. Nevertheless, as the primary goal is to study the social interactions within the walls, we implement a rejection procedure pertinent to each model, to guarantee the pair does not escape.

For the ABC, upon generating a prediction that will lead an individual outside the wall, the authors of Calovi et al. (2018) opted to repeat the stochastic decision-making loop until a valid prediction is made, or a threshold is reached. In the latter case, the individual is explicitly positioned in a position parallel to the wall before repeating the decision-making loop once more. Similarly, if sampling the Gaussian distributions that the DLI generates yields an invalid trajectory, then the sampling process is repeated until a valid prediction is made. We have no empirical data from the DLI to suggest that this behavior occurs frequently, or that the resampling has to be thresholded, similarly to the ABC, for computational efficiency reasons. In fact, when the rejection procedure is lifted, the DLI pairs tend to stay within the wall confines for the majority of the simulations we ran, indicating that it indeed encoded the interaction with the wall (see Fig. B.1). Conversely, despite the use of relevant information, *i.e.*, distance to the wall, the ABC is not designed or likely to operate well without the rejection procedure.

4.5 Assessing the generative capabilities of biomimetic social interaction models

When fish swim in a circular tank (here, of radius $R = 25$ cm), they interact with each other and with the tank wall. The resulting collective dynamics can be finely characterized by exploiting the 9 observables introduced and described in the Materials and Methods Section. As explained there, these observables probe 1) the instantaneous individual behavior, 2) the instantaneous collective behavior, and 3) the temporal correlations of the dynamics.

Hereafter, three 16-hour trajectory datasets are analyzed: the first one corresponds to pairs of *H. rhodostomus* in our experiment, the second one to the Analytical Burst-and-Coast model (ABC), and the third one to the Deep Learning Interaction model (DLI). Video 1 of Appendix B shows typical trajectories for these three conditions. The aim of this section is to quantitatively validate the qualitative agreement observed in this video.

4.5.1 Quantification of the instantaneous individual behavior

The individual fish behavior is characterized by three observables: the probability distribution function (PDF) of the speed V , of the distance to the wall r_w , and of the angle of incidence to the wall θ_w . When swimming in pairs, fish tend to adopt a typical speed of about 7 cm/s (see the peak of the PDF in Fig. 4.3A), but can also produce high speeds up to 25–30 m/s. Both fish remain close to the wall of the tank (a consequence of the fish burst-and-coast

swimming mode Calovi et al. (2018)), the leader being closer to the wall (typically, at about 0.5 BL) than the follower (at about 1.2 BL; see Fig. 4.3B). This feature is due to the follower fish trying to catch up with the leader fish by taking a shortcut while taking the turn. Moreover, fish spend most of the time almost parallel to the wall: see the peaks of both PDFs at $\theta_w \approx \pm 90^\circ$ in Fig. 4.3C. A slight asymmetry is observed in the PDF of θ_w , showing that, in the experiments, fish have turned more frequently in the counter-clockwise direction. Values of the mean and the standard deviation of the PDFs presented in this section are given in Tables B.1, B.2, and B.3.




Figure 4.3: Probability density functions (PDF) of observables characterizing individual behavior: **A** Speed V , **B** distance to the wall r_w , and **C** angle of incidence to the wall θ_w . Black lines: experimental fish data. Blue lines: agents of the Analytical Burst-and-Coast model (ABC). Red lines: agents of the Deep Learning Interaction model (DLI). Dashed lines: geometrical leader; dotted lines: geometrical follower.

Both ABC and DLI models produce agents that move at the same mean speed as fish in the experiments, and Fig. 4.3A shows that the speed PDF for both models are in excellent agreement with the one observed in real fish. Moreover, the agents of the ABC model are as close to the wall and as parallel to it as fish are. The PDF of the ABC leader is in good agreement with that of the fish leader (Fig. 4.3B). However, the PDF for the ABC follower has a peak at about the same distance to the wall as that of the leader, while the corresponding peaks are more separated for real fish. Yet, the PDF for the ABC follower is broader than for the leader, showing that the ABC follower tends to be farther from the wall than the leader, as observed for real fish. For the DLI model, the peaks of both leader and follower PDFs are at about the same position as for real fish, although their height is smaller than for fish, meaning that DLI-agents tend to explore more frequently the interior of the tank (observe the thicker tails of the PDF of r_w for the DLI model in Fig. 4.3B). Alignment with the wall is also well reproduced by both models (Fig. 4.3C), including the asymmetry in the direction of rotation around the tank: their peak at $\theta_w > 0$ is higher than the one at $\theta_w < 0$. As already seen in the PDF of r_w , DLI-agents visit more often the interior of the tank, and are hence less aligned with the wall than the real fish and ABC agents. Note that the tendency of DLI-agents to rotate more frequently in the counterclockwise direction is learned from the training set, while this asymmetry has to be explicitly implemented in the ABC model, by introducing an asymmetric term in the analytical expression of the wall repulsion function. A closer look at Fig. 4.3C

shows that fish actually follow the wall with a most likely angle of incidence $|\theta_w|$ that is slightly smaller than 90° , a feature resulting from the burst-and-coast swimming mode inside a tank with positive curvature: fish are found more often going toward the wall than escaping it.

4.5.2 Quantification of the instantaneous collective behavior

H. rhodostomus is a social species, and Fig. 4.4A shows that the two fish remain most of the time close to each other, with the PDF of their distance d_{ij} presenting a peak around $d_{ij} \approx 7 \text{ cm} \approx 2 \text{ BL}$ (mean and standard deviation of the PDFs presented in this section are given in Tables B.1, B.2, and B.3). The fish have a strong tendency to align with each other, as shown in Fig. 4.4B, with the PDF of their relative heading ϕ_{ij} being sharply peaked at 0° . In addition, the PDF of the viewing angle ψ_{ij} reveals that the fish are swimming one behind the other rather than side-by-side. This is illustrated in Fig. 4.4C by the sharp difference in the PDF of the viewing angle for the leader and the follower. The PDF of ψ_{leader} is peaked around $\pm 160^\circ$, meaning that the follower fish is almost right behind the leader fish, but slightly shifted to the right or left. A slight left-right asymmetry in the PDF of the viewing angles is also visible, the follower being more frequently on the left side of the leader, a consequence of the fact that the fish in the experiment follow the wall by turning more often counterclockwise (Fig. 4.3C), with the follower swimming farther from the wall than the leader (Fig. 4.3B).




Figure 4.4: Probability density functions (PDF) of observables characterizing collective behavior: **A** Distance between individuals d_{ij} , **B** difference in heading angles ϕ_{ij} , and **C** angle of perception of the geometrical leader and follower ψ_{ij} . Black lines: experimental fish data. Blue lines: agents of the Analytical Burst-and-Coast model (ABC). Red lines: agents of the Deep Learning Interaction model (DLI). Dashed lines : geometrical leader; dotted lines: geometrical follower (in **C**).

All these features are well reproduced by both models, with only some small quantitative deviations. The ABC model reproduces almost perfectly the experimental PDF of the distance between fish, whereas the PDF for the DLI model is only slightly wider and presents slightly more weight at very small distance than found for real fish or in the ABC model (Fig. 4.4A). The DLI model is in turn better than the ABC model at reproducing the PDF quantifying the alignment of the fish, the latter producing more weight near 0° than for real fish (Fig. 4.4B).

Both models fail at reproducing the very small increase in the PDF at $\phi_{ij} \approx \pm 180^\circ$, which corresponds to sudden U-turns that real fish sometime perform. The PDF of the viewing angles for the leader and the follower (Fig. 4.4C) are also fairly reproduced by both models, including the slight left-right asymmetry observed in real fish, although the peak in the PDF at $\psi_{\text{follower}} = 0^\circ$ (and to a lesser extent at $\psi_{\text{leader}} \approx -160^\circ$) is not quite as sharp as in the experiment.

4.5.3 Quantification of temporal correlations

Fig. 4.5 shows the 3 observables defined in Eqs. 4.1-4.3 and probing the emerging temporal correlations in the system: the mean squared displacement $C_X(t)$, the velocity autocorrelation $C_V(t)$, and the autocorrelation of the angle of incidence to the wall $C_{\theta_w}(t)$, as function of the time difference t between observations. The figure reveals that both models fail to fully reproduce quantitatively these very non-trivial observables, which indeed constitute the most challenging benchmark characterizing the correlations emerging from the fish behavior.




Figure 4.5: Observables quantifying temporal correlations in the system. **A** Mean squared displacement $C_X(t)$, **B** Velocity temporal autocorrelation $C_V(t)$, **C** Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Black lines: experimental fish data. Blue lines: agents of the Analytical Burst-and-Coast model (ABC). Red lines: agents of the Deep Learning Interaction model (DLI). Dashed lines: geometrical leader; dotted lines: geometrical follower; full lines: average over the 2 fish or agents.

Fish data present 3 distinct regimes: a quasi-ballistic regime at short timescale ($t \lesssim 1.5$ s) where $C_X(t) \approx \langle v^2 \rangle t^2$, followed by a second short diffusive regime ($1.5 \text{ s} \lesssim t \lesssim 5$ s) where $C_X(t) \approx Dt$, which is limited by the finite size of the tank, ultimately leading to a third regime of saturation ($t > 5$ s) characterized by slowly damped oscillations since fish are guided by the wall (Fig. 4.5A). Accordingly, the velocity correlation function starts from $C_V(t=0) = \langle v^2 \rangle$ at short time and also presents damped oscillations (Fig. 4.5B). The negative minima of the oscillations in $C_V(t)$ correspond to times when the focal fish is essentially at a position diametrically opposite to its position at the reference time $t=0$, its velocity then being almost opposite to that at $t=0$. Similarly, positive maxima correspond to times when the fish returns to almost the same position it had at $t=0$, with a similar velocity, guided by the tank wall. Of course, these oscillations are damped as correlations are progressively lost, and the velocity correlation function $C_V(t)$ ultimately vanishes at large time $t \gg 20$ s, due to the actual stochastic nature of

the trajectories at this timescale (possible U-turns, or the fish randomly crossing the tank). Note that $C_X(t)$ is markedly different for the leader and follower fish, with a higher saturation value for the leader, which swims closer to the wall, as mentioned above.

The ABC model is able to fairly reproduce the short and intermediate regimes for $C_X(t)$ (Fig. 4.5A), as well as the position of its first peak, reached only slightly later than for fish (1 s after). The ABC model also reproduces the experimental saturation value of $C_X(t)$ averaged over the two fish. As for the DLI model, its predictions are only slightly worse than that of the ABC model, due to the fact that the DLI agents are moving a bit farther to the wall compared to ABC agents and real fish. Yet, both models equally fail at producing more than one oscillation, and the correlations are damped faster compared to the experiment.

As for the velocity autocorrelation $C_V(t)$ (Fig. 4.5B), the ABC model reproduces almost perfectly the short and intermediate regimes and the position of the first negative minimum (hence, up to $t = 6$ s), while the DLI model underestimates the depth of this first minimum. But again, both models fail at reproducing the persistence of the correlations, producing a too fast damping of the oscillations (an effect slightly stronger in the DLI model).

Both models struggle at reproducing the correlation function $C_{\theta_w}(t)$ of the angle of incidence to the wall (Fig. 4.5C), where the fish curve first sharply decreases up to $t = 6$ s and then remains close to $C_{\theta_w} \approx 0.2$. The ABC model is clearly unable to reproduce both the decreasing range (clearly diverging before $t = 2$ s) and the correct saturation value (never falling below $C_{\theta_w} \approx 0.6$). As for the DLI model, it produces a slightly sharper decay of $C_{\theta_w}(t)$ than for real fish, up to $t \approx 6$ s, but fails to reproduce the non-negligible remaining persistence of the correlation observed in fish for $t > 7$ s, with $C_{\theta_w}(t)$ in the DLI model decaying rapidly to zero. In fact, both models fail to reproduce the experimental $C_{\theta_w}(t)$ for opposite reasons. The ABC model exhibits a too high persistence of the correlations of θ_w compared to real fish, presumably because real fish indeed often follow the wall but can also produce sharp U-turns, as observed in Fig. 4.3C. On the other hand, the failure of the DLI model in reproducing $C_{\theta_w}(t)$ stems from the fact that DLI agents move farther from the wall and cross through the tank more often than real fish and ABC agents (see the discussion of Fig. 4.3B above), hence leading to a too fast, and ultimately total, loss of correlation for θ_w .

4.6 Benchmarking the DLI against a similarly purposed neural network

Along with ABC, we also adapted a neural network from Kothari et al. (2021), that was initially intended for human trajectory forecasting in crowded spaces, for pair-wise fish interactions. More specifically, we used a D-MLP-ConC-LSTM as described in Kothari et al. (2021), which we refer to as D-LSTM for brevity. The models presented in Kothari et al. (2021) were introduced in the context of forecasting human trajectories in arbitrary scenes of humans walking. We found that there are many similarities in the approach and goals of human trajectory fore-

casting research works and the goals of our work, *i.e.*, to model the interactions between fish. Therefore, we chose to adapt and use this algorithm to obtain a baseline of the performance that our neural network, presented in the following subsection, achieves. We opted to not intervene with the core structure of the neural network and use the framework provided by the authors of Alahi et al. (2016); Kothari et al. (2021)¹, and parametrize it to generate accurate trajectory estimates. An example of generate trajectory simulation for the MLI can be found in S14 Video.

The model's name is partly owed to its directional pooling layer, where the relative velocity of neighboring individuals is combined with the focal individual's to estimate its future trajectory. To obtain predictions, we trained the network with the same input data, only adapted to a different format compatible with the TrajNet++ framework Kothari et al. (2021). Furthermore, we opted to convert the inputs to meters, instead of the arbitrary $[-1, 1]$ scale, to allow for more intuitive parameter selection. The TrajNet++ framework provides a series of different tools for different tasks, including pre-processing of datasets. Before training the D-LSTM, we performed a pre-processing methodology provided by TrajNet++ that categorizes the data in 4 categories depending on the trajectory type. Namely, static (Type I), linear (Type II), non-linear Type III) and non-interacting (Type IV) trajectories. Here, we parametrized the percentage of the trajectories that correspond to those types and are subsequently used for training, with values 80%, 100%, 100%, and 100%, respectively. Then, the TrajNet++ tool split the dataset in 80-15-5% manner for the training, validation and test datasets. We also set the frame rate option for the pre-processing tool to 25 frames per second, which corresponds to the timescale chosen for ABC, DLI, and the experiments. Finally, we parametrized the group distance threshold to 0.5 m, that is, the diameter of the setup to always consider both individuals in the dataset.

Then, similarly to DLI, the parameters of the model were learned by minimizing a negative log-likelihood loss using the Adam optimizer Kingma and Ba (2014) with an initial learning rate of $\lambda = 0.0001$. A step learning rate scheduler was used to reduce the learning rate with every 10 epochs. The network was given a budget of 80 epochs and a batch size of 8. With respect to the directional pooling layer and the underlying grid, we selected a grid of 32 cells with a cell side size of 0.005 m (see details in Kothari et al. (2021)).

The network itself is structured as follows; 1) the input state, rectangular position, of the focal individual is passed to a Multi-layered perceptron (MLP) structure. We opted to maintain the original input information of the D-LSTM, therefore, contrary to the DLI, the D-LSTM is missing 3 input variables, namely, distance to the wall for the focal and the neighboring individual, and inter-individual distance. 2) the directional information (relative velocity and position) of the top- k (in the context of this work $k = 1$) neighbors are concatenated, and 3) the concatenation is passed to an LSTM layer. Both the MLP and LSTM layers consist of 256 neurons. The network was given observations of length 5, similarly to DLI, and asked to predict the mean and standard deviation of the components of acceleration of 3 future time

¹<https://github.com/vita-epfl/trajnetplusplusbaselines>

steps (duration of $\Delta t = 0.36 s$), instead of 1 that was required from DLI, for two reasons; 1) to put pressure during training on the network to learn more accurately, and 2) to compare its longer horizon forecasting capabilities against DLI. In D-LSTM's simulations, we considered only the prediction corresponding to $t + \Delta t$ and followed the same simulation logic of DLI (see the Materials and Methods section of the main text). That means, a new trajectory was generated at every time step ($\Delta t = 0.12 s$). It is indeed common, in such cases, that only one time step is used for simulation and the remaining part of the trajectory is used to understand the estimated intent of the agent.

4.6.1 Comparing the short- and long-term performance of DLI and D-LSTM

4.6.2 Quantification of instantaneous individual behavior

D-LSTM produces agents that do not swim at the same speed as fish in the experiments. The location of the peak of the PDF of V (see Fig. 4.6A) is indeed in good agreement with the experiments, but the distribution decays much faster than in the experiments and DLI and goes to zero at approximately $25 cm/s$. Fig. 4.6B shows that the PDFs of the D-LSTM leader and follower are not in good agreement with the experiments or DLI and both agents swim very close to the wall most of the time. Furthermore, agents tend to swim at a farther distance from the wall more often than in the experiments (see the dashed and dotted lines in Fig. 4.6B between $6 - 25 cm$) and the DLI. Similarly, Fig. 4.6C shows that alignment with the wall is not well reproduced by the D-LSTM model. The locations of the peaks are well reproduced with respect to the experiments, including an angle of incidence θ_w smaller than 90° . However, the D-LSTM agents swim parallel to the wall considerably less than in the experiments, see the PDF for values of $\theta_w < -90^\circ$ and $\theta_w > 90^\circ$.

4.6.3 Quantification of instantaneous collective behavior

The PDF of the inter-individual distance d_{ij} is quantitatively different from what is measured for the experiments. Notably, the location of the peak is located approximately $1 cm$ more to the left (i.e., closer to the neighboring agent) than in the experiments (see Fig. 4.6). Furthermore, the D-LSTM agents swim much closer to each other and almost never swim more than $15 cm$ away from each other, contrary to the experiments. Conversely, the D-LSTM reproduces the alignment PDF ϕ_{ij} very well, albeit D-LSTM agents tend to be unaligned more often than in the experiments and DLI. Similarly, to ABC and DLI, the D-LSTM fails to reproduce the small increase in the PDF at $\Delta phi \approx 180^\circ$. Viewing angles are very well reproduced by the D-LSTM. In fact, the PDF is very similar to those of ABC and DLI, although the D-LSTM follower agent (see dotted lines in Fig. 4.6) tends to swim less often parallel to the leader agent.

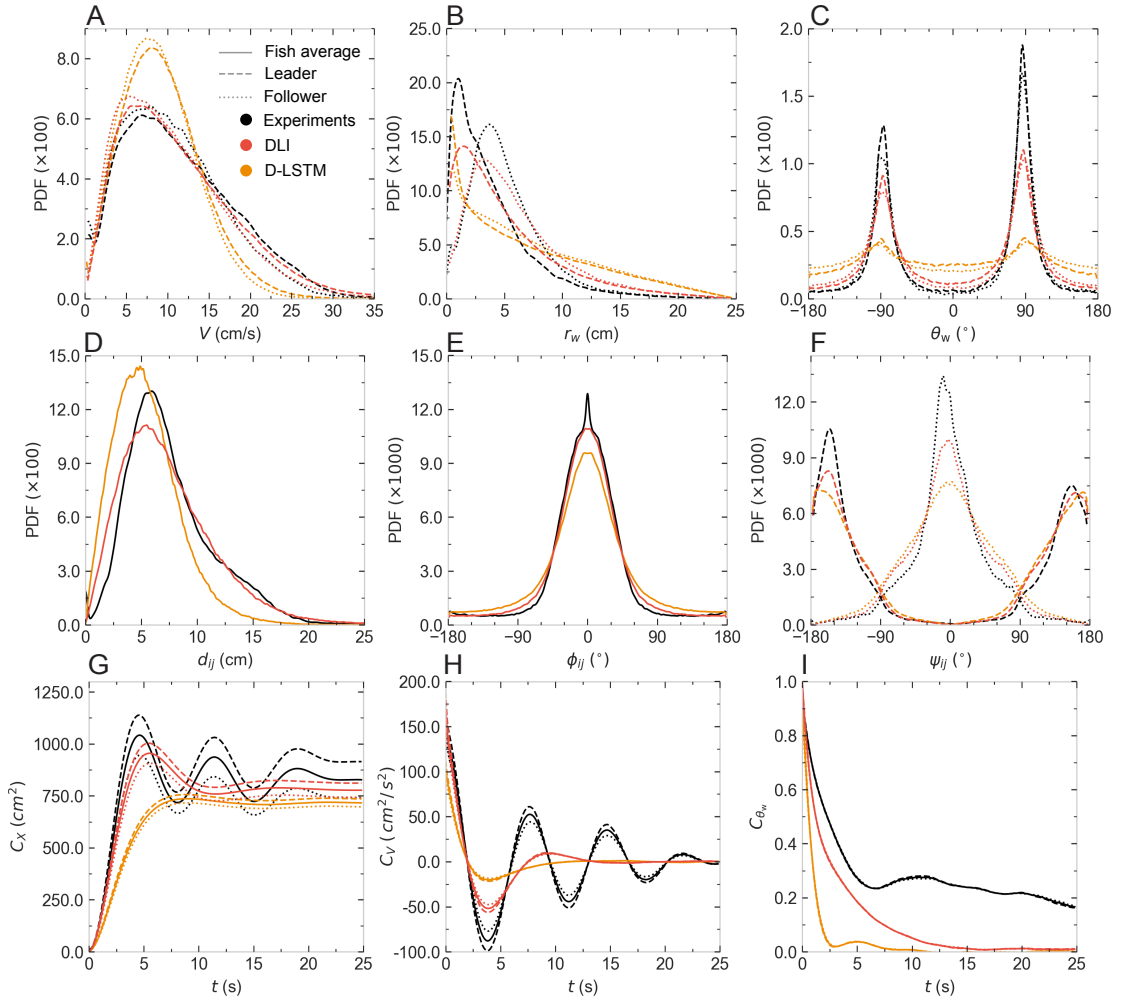


Figure 4.6: Probability density functions (PDF) of all observables. **A** Speed V , **B** distance to the wall r_w , **C** angle of incidence to the wall θ_w , **D** Distance between individuals d , **E** difference in heading angles ϕ_{ij} , **F** angle of perception of the geometrical leader and follower ψ , **G** Mean squared displacement (*i.e.*, $C_X(t) = \langle [\vec{u}(t+t') - \vec{u}(t')]^2 \rangle$), **H** Velocity temporal correlations (*i.e.*, $C_V(t) = \langle \vec{v}(t+t') \cdot \vec{v}(t') \rangle$), and **I** Temporal correlations of the heading of a fish relative to the wall (*i.e.*, $C_{\theta_w}(t) = \langle \cos[\theta_w(t+t') - \theta_w(t')] \rangle$). Black lines: experimental data of real fish. Orange lines: agents of D-LSTM. Red lines: agents of DLI (Deep learning interaction model). Dashed lines correspond to the leader, and dotted lines to the follower.

4.6.4 Quantification of time correlation

Fig. 4.6G shows that the D-LSTM cannot well reproduce the short or intermediate regimes of $C_X(t)$. The curve saturates later than the fish in the experiments (2 s more) being considerably farther distance from the wall compared to fish or ABC and DLI. Similarly to fish, ABC and DLI, the leader and follower have different curves for $C_X(t)$ and $C_V(t)$ (Fig. 4.6G and H, respectively), although in D-LSTM the effect is less pronounced than the other two interaction models. The oscillations of $C_V(t)$ in D-LSTM are sharply damped and disappear at $t = 5$ s. Similarly to

ABC and DLI, the oscillations in $C_{\theta_w}(t)$ are not in good agreement with fish data and fail to reproduce the correlation of fish for $t > 2.5 s$.

4.6.5 Comparing the short-term performance of DLI and D-LSTM

Trajectory forecasting algorithms provide a longer trajectory horizon than $t + \Delta t$ in an attempt to capture the intent of the agents. Here, the ANNs' goal is to reproduce similar movement trajectories as fish.

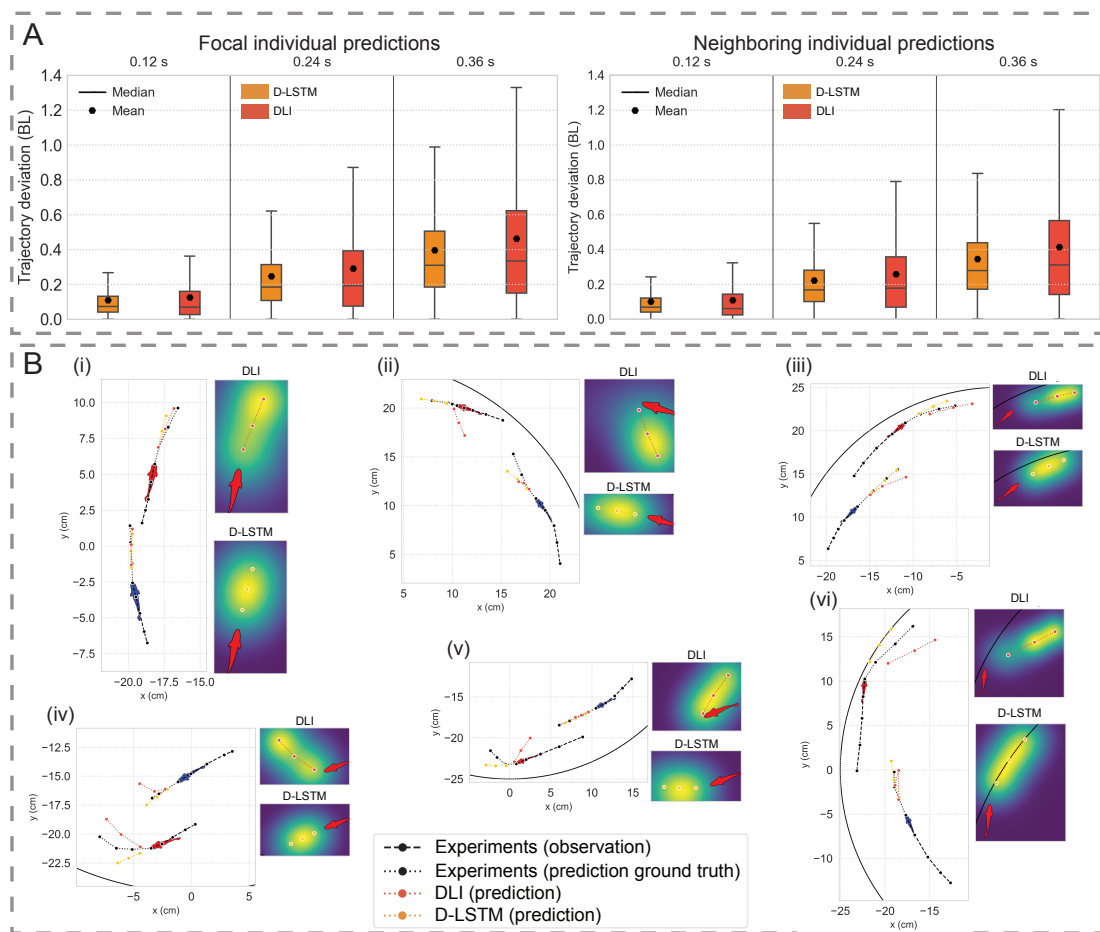


Figure 4.7: **A** Mean squared error between the generated position and the actual position of the focal individual (left) and the neighbor (right), and **B** examples of generated trajectories for DLI and D-LSTM (left) and uncertainty of the models (right) for all panels (i)-(vi), where certainty is depicted with hues of blue to yellow, for low to high certainty areas, respectively.

In Fig. 4.7, we depict a box-plot of the performance of DLI and D-LSTM with respect to their mean squared error measured against the fish in the experiments. Both models have achieved comparable performance, as the D-LSTM performs systematically better by marginal difference. As expected, the models demonstrated an increasing error (see Fig. 4.7A) for the more distant horizon (*i.e.*, at 0.36 s). In Fig. 4.7B(i-vi), we depict a few trajectory examples along

with their uncertainty maps (on the right-hand side of each example). Fig. 4.7A show that the D-LSTM makes more confident decisions, that is, the generated trajectories exhibit less variability. After visual qualitative inspection of approximately 3,500 trajectories, such as the ones of Fig. 4.7B, we did not observe any differences or particularities for either model.

4.7 Validation of the DLI's scalability to other fish species

The premise of ML algorithms is that they can more easily scale up to solve similar tasks. In this subsection, we put this to the test. That is, to validate the DLI's performance with a different fish species. Therefore, we conducted experiments of *D. rerio* pairs, similarly to the *H. rhodostomus*. Then, we trained the DLI as described in previous sections of this chapter.

4.7.1 Quantification of the instantaneous individual behavior

Fig. 4.8A is very well reproduced by the DLI, capturing the peak of the PDF correctly. Marginal difference is noticed in velocities between 15 and 30 cm/s. The DLI is also reproducing well the PDF of the leader and follower distance to the wall r_w (Fig. 4.8B), with marginal differences producing a wider PDF for the DLI. In Fig. 4.8C, the PDFs of the angle of incidence to the wall θ_w show good agreement, with marginal differences between angles of -90° and 90° .

4.7.2 Quantification of the instantaneous collective behavior

In terms of the collective behavior, the DLI is marginally worse when compared to real fish for the interindividual distance d_{ij} of agents (see Fig. 4.8D). Similarly, the DLI is not fully capable of recovering the PDF heading angles exhibiting more occurrences in the tails of the PDF ($\phi_{ij} < -50^\circ$ and $\phi_{ij} > 50^\circ$). However, when DLI leader and follower agents swim, they are able to reproduce the PDF of viewing angle ψ_{ij} with good agreement to the original fish experiments.

4.7.3 Quantification of temporal correlations

Contrary to *H. rhodostomus*, *D. rerio* data present only 2 distinct regimes: quasi-ballistic regime at short timescale ($t \lesssim 2$ s) where $C_X(t) \approx \langle v^2 \rangle t^2$, followed by a saturation regime ($t > 5$ s) characterized by slowly damped oscillations (Fig. 4.8G). Similarly, the velocity correlation function starts from $C_V(t) \approx \langle v^2 \rangle$ at short time and presents oscillations (Fig. 4.8H). Oscillations are eventually damped at a time greater than $t \gg 20$ s. Contrary to *H. rhodostomus*, leader and follower showcase approximately the same saturation values. DLI agents reproduce the reproduce well the oscillations of $C_X(t)$ up to $t \approx 2.5$ s. Contrary to the real experiments, the DLI PDF is quickly damped after this time. The same is true for $C_V(t)$, that is also quickly damped after $t \approx 2.5$ s. However, the $C_{\theta_w}(t)$ PDF of the DLI model is in better agreement with the experiments, but the two curves deviate in amplitude at $t \approx 2$ s.

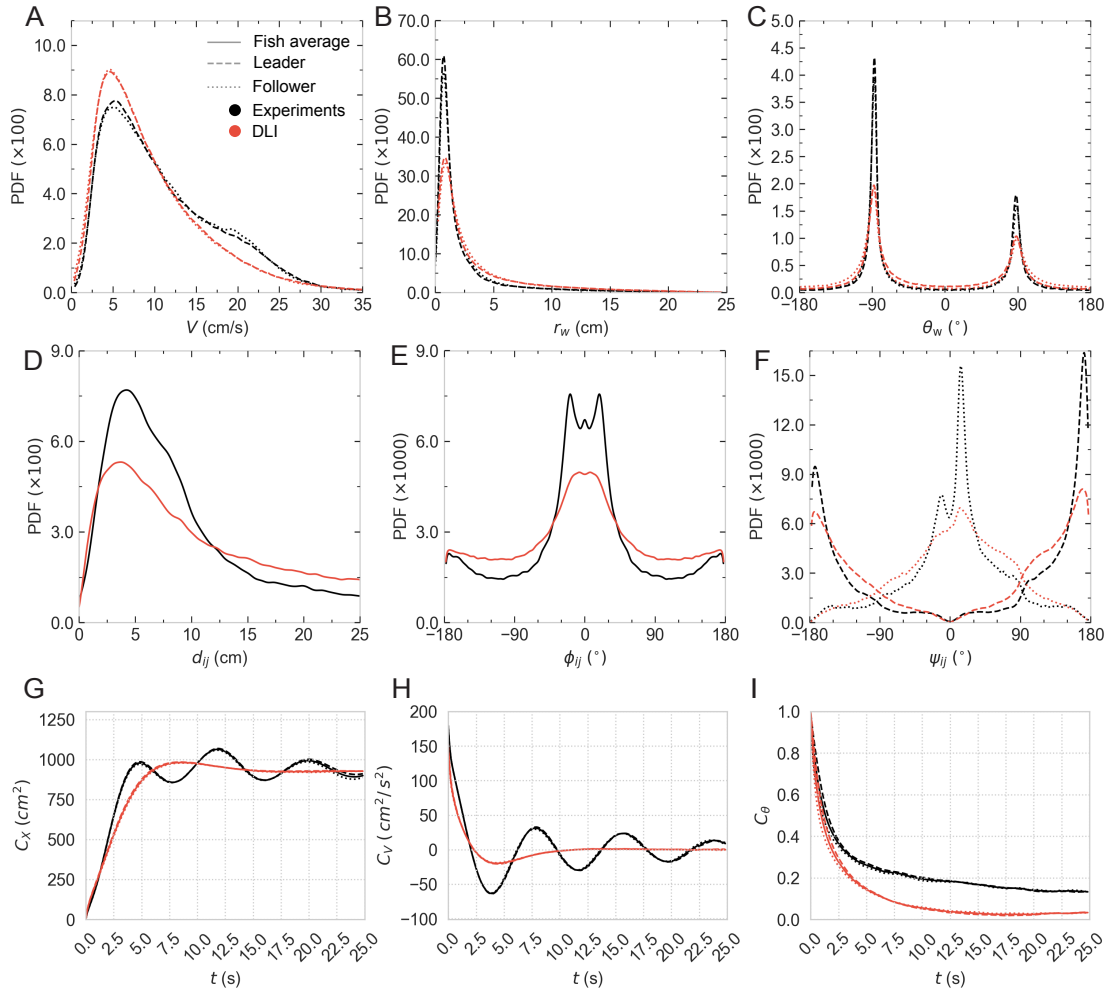


Figure 4.8: Probability density functions (PDF) of all observables. **A** Speed V , **B** distance to the wall r_w , **C** angle of incidence to the wall θ_w , **D** Distance between individuals d , **E** difference in heading angles ϕ_{ij} , **F** angle of perception of the geometrical leader and follower ψ , **G** Mean squared displacement (*i.e.*, $C_X(t) = \langle [\vec{u}(t+t') - \vec{u}(t')]^2 \rangle$), **H** Velocity temporal correlations (*i.e.*, $C_V(t) = \langle \vec{v}(t+t') \cdot \vec{v}(t') \rangle$), and **I** Temporal correlations of the heading of a fish relative to the wall (*i.e.*, $C_{\theta_w}(t) = \langle \cos[\theta_w(t+t') - \theta_w(t')] \rangle$). Black lines: experimental data for zebrafish. Red lines: agents of DLI. Dashed lines correspond to the leader, and dotted lines to the follower.

4.8 Assessing neural networks with no memory components

One of the main alternatives investigated during the automated search, described in previous sections of this chapter, is a Multi-layered Perceptron Interaction (MLI) version of the DLI. That is, we maintained the general shape of the structure of 7 layers, but all of them are dense layers. That also means that the input is not a sequence of states, but only the last state at time t^n . For reference, we also tested various architectures where a sequence of states was concatenated and provided at the input of the network, but we present MLI for two reasons: 1)

concatenating past states at the input would provide some sort of memory to the ANN which is not the comparison we set out to make, and 2) the ANNs with sequence of states at the input (and no memory) did not give better results than MLI. Naturally, the lack of LSTM layers that consist of memory cells, means that the MLI is composed of significantly less free parameters than the DLI (approximately 1/6 th). An example of generate trajectory simulation for the MLI can be found in S16 Video.

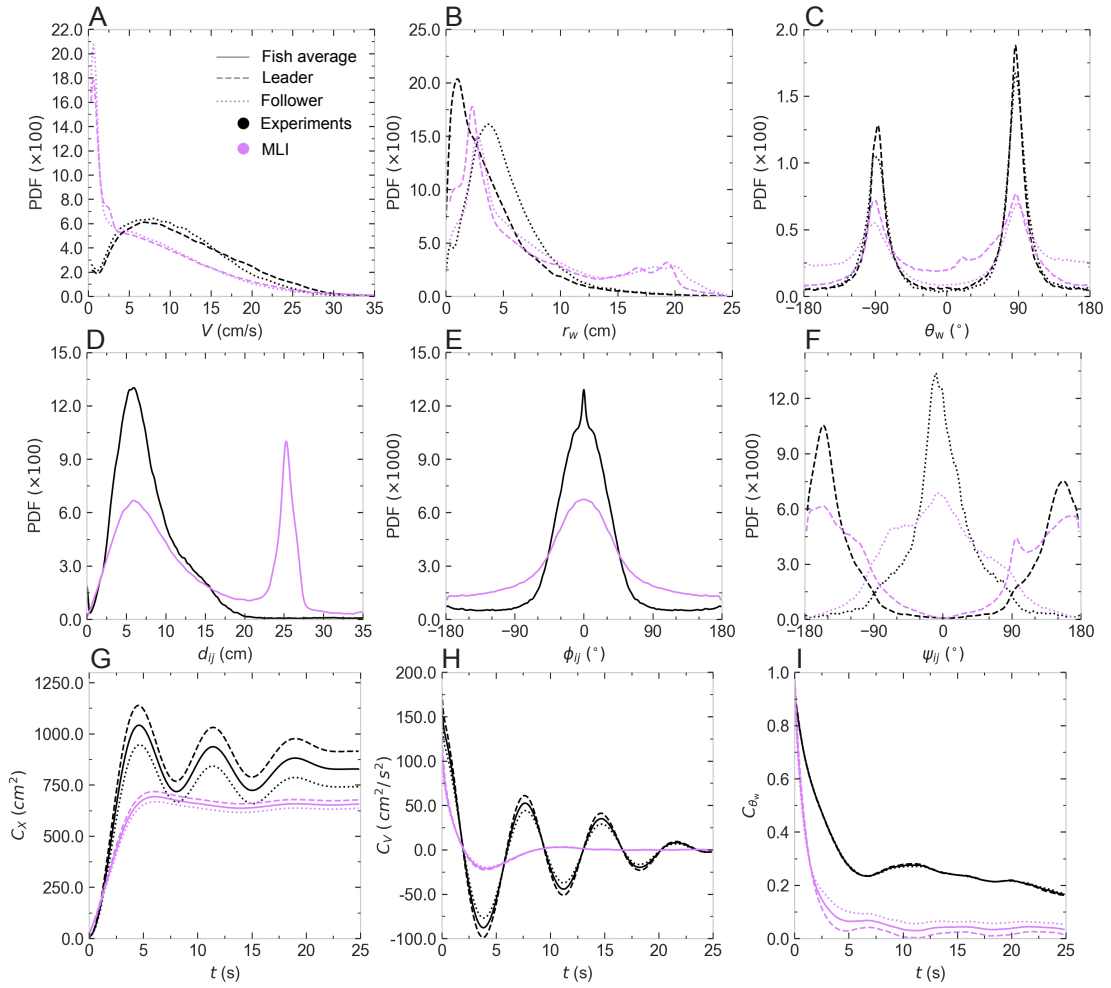


Figure 4.9: Probability density functions (PDF) of all observables. **A** Speed V , **B** distance to the wall r_w , **C** angle of incidence to the wall θ_w , **D** Distance between individuals d , **E** difference in heading angles ϕ_{ij} , **F** angle of perception of the geometrical leader and follower ψ , **G** Mean squared displacement (*i.e.*, $C_X(t) = \langle [\vec{u}(t+t') - \vec{u}(t')]^2 \rangle$), **H** Velocity temporal correlations (*i.e.*, $C_V(t) = \langle \vec{v}(t+t') \cdot \vec{v}(t') \rangle$), and **I** Temporal correlations of the heading of a fish relative to the wall (*i.e.*, $C_{\theta_w}(t) = \langle \cos[\theta_w(t+t') - \theta_w(t')] \rangle$). Black lines: experimental data of real fish. Orange lines: agents of D-LSTM. Red lines: agents of MLI (Multi-layered Perceptron Interaction model). Dashed lines correspond to the leader, and dotted lines to the follower.

4.8.1 Quantification of the instantaneous individual behavior

Fig. 4.9A shows that the PDF of V is considerably different from the experiments. More specifically, the location of the peak is located at a much lower velocity ($\approx 1 \text{ cm/s}$). Similarly, the PDF of the MLI leader and follower are not peaked at the same location as the experiments. However, the PDF of the MLI follower is peaked at a greater distance to the wall than the leader MLI agent, in agreement with what is measured in the experiments. Despite this, the PDF of the leader and follower also demonstrate smaller peaks around a distance of $d_w = 19 \text{ cm}$. The alignment with the wall is also not very well reproduced by MLI (Fig. 4.9), although the asymmetry in the direction of rotation is captured to a small extent. Furthermore, whereas the PDFs of both the leader and follower MLI agent capture the slightly smaller angle of incidence θ_w when $\theta_w > 0$, the angle of incidence for $\theta_w < 0$ is approximately equal to 90° .

4.8.2 Quantification of the instantaneous collective behavior

In Fig. 4.9D the inter-individual distance d_{ij} PDFs of MLI are not in good agreement with the experiments. Although the location of the peak is located at the same value as in the experiments, there is a second peak at $d_{ij} \approx 26 \text{ cm}$ which is not present in fish data. Similarly, Fig. 4.9E shows that MLI fails in reproducing the PDF of alignment ϕ_{ij} , although the peak location is in good agreement with the experiments. However, MLI agents swim considerably less amount of time aligned than the experiments. The viewing angle (Fig. 4.9F) leader and follower PDF peaks are in good agreement with the experiments, but the asymmetry present in the fish data is not captured. Furthermore, both PDFs of the MLI are wider around the peaks than the experiments.

4.8.3 Quantification of temporal correlations

The MLI is not able to accurately reproduce the oscillations of $C_X(t)$ similarly to experiments. It saturates approximately 1 s later than the fish in the experiments, while considerably farther distance from the wall. Similarly to fish, ABC and DLI, the MLI leader and follower have different curves for $C_X(t)$ and $C_V(t)$ (Fig. 4.9G and H, respectively). Fig. 4.9H shows that MLI is also unable to reproduce the oscillations of $C_V(t)$. The curve deviates from $t = 0$ and reaches the first oscillation at time $t = 5 \text{ s}$, similarly to the experiments, but is quickly damped after. $C_{\theta_w}(t)$ are not in good agreement with fish data and fail to reproduce the correlation of fish for $t > 2.5 \text{ s}$. Notably, the leader and follower curves of the MLI are considerably different, in contrary to the fish data.

4.8.4 Complementary analyses

H. rhodostomus Lei et al. (2020), like many other group-living species Cavagna et al. (2018), effectively only interact with a few influential neighbors, at a given time. Thus, for a given agent in a group of $N > 2$ agents, the DLI for *H. rhodostomus* should only retain the influence of

typically the two agents leading to the highest acceleration Wang et al. (2022); Lei et al. (2020), as predicted by the DLI model. Video 3 of Appendix B illustrates this procedure for $N = 5$ agents, resulting in a cohesive and aligned group, in qualitative agreement with experimental observation Lei et al. (2020).

4.9 Conclusion

Studying social interactions in animal groups is crucial to understand how complex collective behaviors emerge from individuals' decision-making processes. Very recently, such interactions have been extensively investigated in the context of collective motion by exploiting classical computational modeling Calovi et al. (2018); Escobedo et al. (2020); Jayles et al. (2020) and automated machine learning-based methods Heras et al. (2019); Costa et al. (2020). Although ML algorithms have been shown to provide insight into the interactions of hundreds of individuals at short timescales Heras et al. (2019); Costa et al. (2020), their ability to reproduce the complex dynamics in animal groups at long timescales has not yet been assessed.

In this chapter, we have presented a deep learning interaction model (DLI) which reproduces the behavior of fish swimming in pairs. We have also introduced the appropriate tools for its validation when compared to experimental results and when confronted with the state-of-the-art analytical model (ABC). In fact, our study establishes a systematic methodology to assess the long-term predictive power of a model (analytical or ML), by introducing a set of fine observables probing the individual and collective behavior of model agents, as well as the subtle correlations emerging in the system. These observables, which can be straightforwardly extended to groups of $N > 2$ agents, provide an extremely stringent test for any model aimed at producing realistic long-term trajectories mimicking that of actual animal groups. In particular, we consider that the usual validation of an ML model at a short timescale should be complemented by the type of long timescale analysis that we propose here, in order to fully assess its performance.

The DLI model closely reproduces the dynamics of real fish at both the individual (speed, distance to the wall, angle of incidence to the wall) and collective (distance between individuals, relative heading angle, angle of perception) levels during long simulations corresponding to more than 16 hours of fish swimming in a tank, hence successfully generating life-like interactions between agents. When compared to experiment, the ABC model and the DLI model essentially performs equally well. Notably, the DLI model better captures the most likely distance of the leader and follower from the wall. However, the DLI model is less accurate in reproducing the temporal correlations quantified by the mean-squared displacement and the velocity autocorrelation. Yet, both ABC and DLI models fail at capturing the temporal correlations of the angle of incidence to the wall, but for very different reasons.

Our study demonstrates two advantages of ML techniques: 1) they can drastically accelerate the generation of new models (as illustrated here for zebrafish), and 2) with minimal expertise in biology or modeling. This is especially useful in robotics, where models often act as

behavioral controllers (*i.e.*, trajectory generators) that guide the robot(s). Although there already exist many bio-hybrid experiments in the literature, most of them rely on simplified models for behavioral modulation Bonnet et al. (2016, 2018) (see also Chapter 3), few of them exploit realistic models (analytical or ML) Cazenille et al. (2018c,b), and, to our knowledge, none of them are tested in the long term in simulations or real-life. In this context, ML has the potential to benefit multidisciplinary studies, provided such techniques are thoroughly validated in simulations.

However, accelerating the production of collective behavior models with ML comes at a cost. Indeed, the DLI is a black-box model, and although it captures the subtle impact of social interactions between individuals, it is impossible to retrieve the interaction functions themselves. Some approaches partially address this issue by providing insight into how the network operates for specific sets of inputs Heras et al. (2019); Costa et al. (2020). Yet, they still do not offer explicit interaction functions. On the other hand, analytical models supplemented by a procedure to reconstruct social interactions Calovi et al. (2018); Escobedo et al. (2020) provide a concise and explicit description of the system in question. Moreover, varying the parameters of such models allows for investigating their relative impact on the dynamics, and to make predictions for various sets of these parameters Wang et al. (2022). This is not feasible with ML models, unless they are retrained or specifically structured to allow it. Finally, whereas the DLI was shown to also perform well for the zebrafish, this required completely retraining the network on zebrafish data. However, given that both *H. rhodostomus* and zebrafish swim in a burst-and-coast manner, in future work could investigate only partially retraining the network (*e.g.*, the last few layers that intuitively decode the social dynamics captured in the model) by exploiting transfer learning techniques Zhuang et al. (2020); Weiss et al. (2016).

In summary, this chapter shows that DLI-like models may now be considered as firm candidates to shed light on groundbreaking problems such as how social interactions take place and affect collective behavior in living groups. Yet, we have emphasized that social interaction models should be precisely tested at both short *and* long timescales. Future work could include the design of ANNs that provide additional information about the learned dynamics, possibly by exploiting symbolic regression algorithms Quade et al. (2016); Chen et al. (2019). Furthermore, in Chapter 6 we thoroughly investigate how the DLI model can be transferred back to real-life through a robotic system to create biohybrid groups that exhibit life-like social interactions.

What makes for an efficient biohybrid interaction framework **Part III**

5 Introducing a novel biohybrid interaction framework

This section's content is adapted from the following article:

- Papaspyros, V., Burnier, D., Cherfan, R., Theraulaz, G., Sire, C. and Mondada, F., 2023. A biohybrid interaction framework for the integration of robots in animal societies. IEEE Access. <https://doi.org/10.1109/ACCESS.2023.3290960> (reuse authorised under the CC 4.0 licence) – Papaspyros et al. (2023a)
 - Vaïos Papaspyros's contribution: Conception, data curation, formal analysis, investigation, methodology, software, hardware, visualization, writing – original draft, writing – review & editing.
- The article's work was funded by the Swiss National Science Foundation project "Self-Adaptive Mixed Societies of Animals and Robots" (Grant No. 175731). Our collaborators from the Université Toulouse – Paul Sabatier were funded by the French National Research Agency (ANR-20-CE45-0006-01).
- Ethics: All *H. rhodostomus* experiments were conducted at the Centre de Recherches sur la Cognition Animale, Centre de Biologie Intégrative, CNRS, Université de Toulouse – Paul Sabatier. Experiments were approved by the local ethical committee for experimental animals and were performed in an approved fish facility (A3155501) under permit APAFIS#27303-2020090219529069 v8 in agreement with the French legislation.

5.1 Introduction

Nearly a decade after the initial iteration of the FishBot v4.4 Bonnet et al. (2012, 2014) and other similar platforms Swain et al. (2011); Landgraf et al. (2013, 2016), advancements in algorithms and precision construction methodologies have enabled significant improvement. This progress also extends to the ancillary hardware and software that supports the robot's operation, such as cameras and communication protocols. In this chapter, we present a novel Behavioral Observation and Biohybrid Interaction (BOBI) framework, comprising: 1) an

experimental setup; 2) a wheeled mobile robot, the LureBot; 3) a lure-building methodology, primarily designed to study fish interactions, as well as small animals or insects in general; 4) an open source distributed codebase utilizing Robot Operating System (ROS) 1 (Melodic Morenia Stanford Artificial Intelligence Laboratory et al. (2018)) packages to manage various aspects of the setup and robot.

This new experimental framework facilitates and considerably extends continuous experiments (exceeding 3 h), at robot speeds that can reproduce or even surpass the motion profiles of multiple fish species (*e.g.*, Rummy-nose Tetra – *Hemigrammus rhodostomus* – and zebrafish – *Danio rerio*), with potential for applications to other small animal species. Additionally, BOBI incorporates significant upgrades over the FishBot v4.4 Bonnet et al. (2012, 2014) and the Control and Control And Tracking Software (CATS) Bonnet et al. (2017) (see Sec. C.1 for a performance comparison between the FishBot v4.4 and the LureBot), in terms of real-time robot control, individual detection and identification algorithms, and hardware components such as cameras for capturing more precise data. In addition, in the context of the present work where the LureBot is interacting with rummy-nose tetra fish (*Hemigrammus rhodostomus*), the LureBot is commanded by a data-driven model which was shown in simulations to faithfully reproduce the collective behavior of this species Calovi et al. (2018); Lei et al. (2020).

Furthermore, we have designed the new experimental platform to be mobile, compact, and user-friendly for non-engineering personnel, a feature that, to the best of our knowledge, has not been addressed in similar works. Finally, we demonstrate two sets of control experiments with groups of *H. rhodostomus* and a lure manipulated by the LureBot. We validate its operational stability and confirm its acceptance by the fish. This new framework allowed us to reduce the physics gap (see Fig. 1.1) between simulation and preliminary real-world experiments, the communication gap (see Fig. 1.1) by introducing a new biomimetic lure construction methodology, and acted as a first step towards the extended biohybrid experiments presented in Chapters 6 and 7, that allowed us to measure the biomimicry gap.

5.1.1 Chapter outline

The remainder of the chapter is organized as follows:

- in Sec. 5.2, we describe the design aspects of the experimental setup that includes the fish tank, recording equipment, and the robot;
- in Sec. 5.3, we introduce our new robot, the LureBot, and provide a description of its mechanical and electronics design;
- in Sec. 5.4, we detail our software architecture that allows the interplay of the experimental setup and robot;
- in Sec. 5.5, we present our lure design methodology;

- in Sec. 5.6, we describe 3 trajectory generation methodologies that we used to control the robot, and in particular, to test its acceptance by the fish;
- in Sec. 5.7.1, we detail the handling conditions of the fish and provide a step-by-step procedure that we have followed during our experiments;
- in Sec. 5.7.2, we describe some data-related post-processing operations;
- in Sec. 5.8, we present the results of our experiments. We first focus on assessing the acceptance of the robot by the fish by exploiting open-loop trajectory generators with a biomimetic lure and a disc-shaped lure. Then, we quantify the behavior of a single fish and of the robot alone and of a fish swimming with the robot. Finally, we present preliminary results for a group of 4 fish interacting with the LureBot;
- in Sec. 5.9, we highlight the main contributions of this chapter's work and also discuss some of its limitations and potential areas for future exploration.

5.2 Experimental Setup

The proposed behavioral setup, depicted in Fig. 5.1, measures $810\text{ mm} \times 810\text{ mm} \times 1810\text{ mm}$ ($w \times d \times h$), which is smaller than the previous one Bonnet et al. (2017, 2012), which measured $1000\text{ mm} \times 1000\text{ mm} \times 2100\text{ mm}$ ($w \times d \times h$). All aluminum beams are indented in the middle to allow for mounting external attachments. Therefore, all the necessary equipment (*e.g.*, lights, cameras...) can be directly attached to it for use and storage. Conversely, in Bonnet et al. (2017, 2012) some equipment (*e.g.*, lights) was permanently fixed to the experimental room walls. Consequently, it was particularly difficult to move it and proved rather restricting for maintenance and experimental procedures alike, and it was almost impossible to use the setup in different experimental facilities, even if those were located in the same building. The proposed setup stands on omnidirectional wheels (equipped with brakes) attached to each corner of the frame, which allows the experimenter to move with great ease, even when the full equipment is mounted on its frame.

In the current setup configuration, a 10 mm thick glass tank of $720\text{ mm} \times 720\text{ mm} \times 140\text{ mm}$ ($w \times d \times h$) is fixed on the outer frame, and is used to bound the experimental area. The tank's bottom surface is fitted with a white polytetrafluoroethylene layer (PTFE) due to its small friction coefficient, which allows for smooth motion of the lure (see Sec. 5.5.1). Inside the tank, we use the same circular shaped arena that was used in our previous setup (refer to Chapter 3). An additional inox plate is glued on the outer bottom part of the tank and connected to the positive pole of an external power supply. A second plate is placed below the glass tank, on four spring-loaded supports (see Fig. 5.1) that are attached to the frame. This space is where the wheeled robot(s) operate. The plate consists of a perforated inox layer glued on transparent plexiglass. The holes in the inox layer allow for light to pass through, and the plexiglass serves two purposes: 1) maintain the inox layer flat, and 2) diffuse the light passing through the perforated surface. Our previous setup entirely lacked those two characteristics. The




Figure 5.1: **Experimental setup and software overview.** (a) Basler acA4024-29um monochrome camera mounted at the center top of the metal frame. (b) Color webcam (ELP-USBFHD01M-L170) equipped with a fisheye lens and mounted at the center bottom of the metal frame. (c) Sample frame from the top camera depicting 5 agents (including a lure). (d) Sample frame from the bottom camera depicting the two LED lights of the robot, namely, red and green located at the front and back of the robot, respectively. (e) Cross-section view of the experimental setup depicting the robot operating below the tank and the lure moving inside it. (f) Robot detector module. Extracts the robot positions and headings from the robot's LED lights, using images from the bottom camera. (g) Individual detector routine. No identities assigned at this point. (h) Trajectory assignment module. Given the tracked positions of robots (bottom camera) and individuals (top camera), and using past observations, this module assigns unique trajectory IDs to individuals. (i) Behavioral control module. High-level decision-making process that generates target vectors (positions or velocities) for the robot to follow. (j) High-level robot control module. The module consists of regulation routines that take instructions from decision-making models and translate them into motor commands. (k) Robot interface module. The communication interface that exchanges information with the robot(s). (l) On-board (low-level) robot control module. (m) Data logging module. The module collects and stores the outputs of all other routines for analysis and debugging purposes. (n) Sample frame depicting the fused information from the top and bottom cameras after executing the trajectory assignment routine. The red marker signifies that the individual is the artificial robot; blue markers correspond to all other individuals; the orange marker represents the robot's desired position; and the two green circular markers are the two most influential neighbors of the lure (see Lei et al. (2020) for more details).

distance of this plate to the tank bottom is regulated by the mounting location of the springs on the frame and support elements (gauge blocks) in between. The adjustable springs and support elements also serve as means to fine-tune the magnetic coupling strength between the lure and robot, by adjusting the distance between them. Additionally, the springs allow the experimenter to easily add or remove robots by momentarily compressing the springs, instead of the external crank lift system used in our previous setup.

We attach two cameras of different specifications in two aluminum beams at the top and bottom of the outer frame, for each camera respectively (see Figs. 5.1a, b). The top camera, which is used to monitor the contents of the glass tank, is a monochrome high resolution (4024 pixels \times 3036 pixels) Basler acA4024-29uc fitted with a low distortion lens. The camera is capable of delivering frames at 30 Hz, twice the rate of the camera used in Bonnet et al. (2017, 2012), with added resolution. The bottom camera is the same full high-definition (1920 pixels \times 1080 pixels) color webcam (ELP-USBFHD01M-L170) from Bonnet et al. (2017, 2012). It is capable of retrieving frames at 30 Hz and is equipped with a fisheye lens that can keep the second plate in frame even when it is mounted at a low position (close to the camera in Fig. 5.2e). The bottom part of the frame, which includes the plate where robots operate, is enclosed by Medium-density fiberboard (MDF) boards. This eliminates environmental light sources and allows for easier detection of the LED lights located at the bottom part of the robot's chassis.

5.3 LureBot

5.3.1 Mechanics

A full depiction of the LureBot's mechanical design¹ is presented in Fig. 5.2, and can be generally summarized as a two-part design that consists of its chassis and cover. The robot's outer dimensions are 49 mm \times 50 mm \times 42.7 mm ($w \times d \times h$), not including the magnets (Fig. 5.2a) and ball casters (Fig. 5.2i) or 49 mm \times 78 mm \times 46.7 mm ($w \times d \times h$), otherwise. For reference, the FishBot v4.4's outer dimensions (without its magnets) are 22 mm \times 43 mm \times 67 mm ($w \times d \times h$) Bonnet et al. (2014).

The LureBot's chassis is roughly equal to the bottom half of its height and houses two motors symmetrically placed in its center. We explicitly designed it in a differential drive configuration to allow for rapid direction changes, including rotations in-place thanks to the motor placement symmetry, much like the U-turn movements exhibited in fish schools Crosato et al. (2018). We use two independent Faulhaber AM1524-0450 stepper motors (Fig. 5.2h), that are capable of producing speeds up to 100 cm/s and accelerations of 175 cm/s², which are more than adequate to express the motion profile of fish species like *H. rhodostomus* or *Danio rerio* (refer to Chapter 4). The motors are directly attached to an aluminum frame (Fig. 5.2f), a material we explicitly chose to dissipate the large amounts of heat (up to 90 °C) generated

¹Designs available at <https://doi.org/10.5281/zenodo.7796299>




Figure 5.2: Description of the robot and its internal parts. (a) Conductive metal link which is used during operation to supply power (positive pole) to the robot. An identical link exists on the bottom part of the robot (negative pole). (b) Two rectangular neodymium magnets measuring are housed on the top robot cover; half of each magnet is below the cover surface and half above. (c) A thin metal post is used to penetrate the lure and maintain it at a fixed height. (d) At the base of the metal post, there is an additional metal plate with two neodymium magnets attached below it. (e) A low friction Poly(methyl methacrylate) (PMMA) plate houses the magnets and metal post. (f) An aluminum part covers the robot base's middle area. (g) Two brass parts cover the robot base's sides. (h) Two rubber wheels are directly driven by two stepper motors. (i) Two ball casters support the robot on the back and front. (j) Two RGB LED lights, on the back and front of the robot, are attached to and driven by the main board. (k) Three infrared (proximity; IR) sensors and (l) two supercapacitors are soldered on the (m) custom electronics circuit board. (n) Arduino Nano 33 IoT single-board.

by the motors during long experiments. Furthermore, to facilitate the transfer of heat, we apply thermal paste between surfaces that come into contact with the motor's outer shell (e.g., the aluminum frame). We attach two brass parts (Fig. 5.2g) to the sides of the aluminum frame, primarily to add weight and improve traction during operation. As a result, the LureBot weighs 300 g and is significantly heavier than the FishBot v4.4's 80 g. A recessed area on each brass component accommodates two rubber wheels (Fig. 5.2h), each one directly attached to a motor.

In comparison, the FishBot v4.4 was unstable in high-speed movements (primarily angular), and for the majority of our past experiments, its maximum speed was limited to 20 – 25 cm/s

Bonnet et al. (2018) (refer to Chapter 3). Hence, it was unable to reproduce the entire spectrum of speeds observed in fish schools. Increasing its speed required that we: 1) increase the torques (*i.e.*, using new motors), and 2) decrease the operating distance of the robot to the top plate to ensure sufficient magnetic coupling in faster movements. However, the latter reduces traction in the wheels, and while this could be partially solved by adding weight, the motors on the FishBot v4.4 did not produce enough torque to support this choice.

On the lower part of the LureBot's chassis, there exist two ball casters extending the LureBot's length by 18 mm, that serve a dual purpose, namely: 1) they support the robot during its movement, and 2) have recessions made of light diffusive material for the LED lights (Fig. 5.2j) facilitating its detection by the bottom camera. Finally, we place a conductive link, depicted in (Fig. 5.2a) below and along the center axis of the robot, to limit the effect of uneven friction forces when the link is touching the conductive plate located below the robot.

The second part of the LureBot's mechanical design, the cover, is designed to house two neodymium magnets (Nd-Fe-B) measuring at $8\text{ mm} \times 8\text{ mm} \times 4\text{ mm}$ ($w \times d \times h$) with a magnetic strength of 14.7 N (1.5 kg), depicted in Fig. 5.2b, used to magnetically move the lure within the tank (Fig. 5.2b). The magnets protrude by half their size (*i.e.*, 4 mm) above the cover to allow for easier replacement if necessary, while remaining firmly fixed in a symmetric arrangement during operation. Additionally, we attach a conductive link along the center axis of the cover (Fig. 5.2a), similarly to the bottom one, that makes contact with the inox plate located below the glass tank. The cover is also used to enclose and protect the robot's electronics (discussed in the following section). Similarly to the FishBot v4.4, and unlike battery powered robots Swain et al. (2011); Landgraf et al. (2013, 2016), the two conductive links allow for continuous access to a power source during operation, *i.e.*, the robot can operate normally for extended periods of time ($\gg 3\text{ h}$).

5.3.2 Electronics

The LureBot is driven by two main components, the Arduino Nano 33 IoT (see Fig. 5.2n) connected on a custom circuit board (see Fig. 5.2m). We use the Arduino as an endpoint (see Fig. 5.1l) that communicates with the high-level interface (see Fig. 5.1k) and receives motor commands. We use those commands, *i.e.*, the desired motor speeds, to time-schedule the step signals of the two Faulhaber motors at 100 Hz. The resulting signals are propagated to the custom circuit board that eventually powers and drives the two Faulhaber motors. The motors have a nominal voltage of 2 V, but we power them at 5.3 V to ensure that they have sufficient current to achieve the desired holding torques. However, this means that the motors produce more heat than intended. In addition to the heat dissipation measures described in the previous section, we use a Pulse-Width modulation (PWM) signal operating at a 70-30% ratio at 1 A and 0.2 A, respectively, to limit the amount of time the motors spend on high currents, and effectively reduce the heat they produce. Therefore, the ratio choice is a tradeoff between producing sufficient torques and low temperatures, and we extensively tested the

robot to obtain one that yields a good balance. The circuit board also houses and powers two LED lights (see Fig. 5.2j) and three TCRT1000 infrared (IR) sensors (see Fig. 5.2k). Finally, we add two supercapacitors (charging at 3 V with 10 F capacity, each), on two sides of the board (see Fig. 5.2l). Their role is to ensure that temporary loss of power (*e.g.*, conductive link not making good contact with the plate or voltage drops) does not cause a complete power down or lead to dimmed LED lights, which are critical for the robot(s)' detection and control.

5.4 Software architecture

The experimental setup and robot/lure is operated by the Behavioral Observation and Biohybrid Interaction (BOBI) software framework, for which, modules are implemented as ROS 1 packages². We use ROS because of its: 1) widespread use and support in the robotics community; 2) package-based architecture that allows modules to be easily replaced, maintained, adapted, or extended; 3) distributed architecture that allows programs (nodes) to be run on different machines to balance the computational load; 4) support of many common programming languages, *e.g.*, C++ and Python, which makes it easier to alter high-level operational aspects (*e.g.*, the behavioral model generating trajectories) for people with limited knowledge or interest in systems programming.

In comparison, our previous Control And Tracking Software (CATS) Bonnet et al. (2017), was packaged as a single C++ software suite, with explicit dependencies to external libraries (*e.g.*, GStreamer 0.1 Team (2016)). These choices shortened its life span and proved rather limiting: 1) for experimenters with little programming know-how, as it required good knowledge of C++ programming; 2) for maintenance, as GStreamer has since moved to a newer version that is not backwards compatible; 3) for software and hardware upgrades, because newer operating systems present dependency conflicts with CATS' external libraries which have been considerably updated and often do not support newer hardware (*e.g.*, GStreamer 0.1 with USB Basler cameras).

As depicted in Fig. 5.1, BOBI is currently comprised of 4 main packages for vision, control, logging tasks, and interfacing with the robot (orange, green, blue, and red color, respectively). Additionally, there are 2 supporting packages that hold ROS message and service definitions, *i.e.*, the definitions of communications that allow the system to distribute information in real-time, and, finally, simulation models for the LureBot. In the following subsections, we detail the implementation of the three operationally fundamental modules, namely, vision, control, and robot interface. **Video 1** of Appendix C includes a depiction of a single operational cycle within the framework.

²Code available at <https://doi.org/10.5281/zenodo.7796357>

Vision

Similarly to our previous work Bonnet et al. (2017), we utilize two cameras installed at the top and bottom of the experimental setup (see Figs. 5.1a, b and Sec. 5.2) to track the individuals (including the lure, if any) and robot, respectively. The cameras operate at their maximum rate of 30 Hz, twice the rate used in our previous setups Bonnet et al. (2017) (see also Chapter 3), which allows for tracking even the very rapid and fine movements of fish or other rapidly moving animals.

Both camera streams are downsampled at source from their original resolution to 512 pixels \times 512 pixels and 640 pixels \times 480 pixels for the top and bottom camera, respectively, to reduce the computational load. Operating the cameras at their highest resolution is computationally demanding and not required for fish individuals, but the extra resolution could prove useful for smaller animals, *e.g.*, ants. At the beginning of an operational cycle, we obtain two frames, one from each of camera stream. The two frames are subject to distortion caused by each camera's lens, therefore, in the modules depicted in Fig. 5.1f and g, we first undistort them, and optionally apply a mask to retain user-defined regions of interest (*i.e.*, to further reduce the computational load and assist the detection). Then, we run two routines for each frame in parallel.

For the bottom camera frame, we apply a color threshold twice, once to isolate the light blob produced by the front LED and once for the light blob produced by the back LED. The two LED lights have been purposely given different colors to indicate the heading of the robot instantly, unlike the single LED color used in Bonnet et al. (2017) which requires that the heading is inferred from the robot's movement.

For the top camera frame, we use the first 500 frames to train a Mixture of Gaussians background subtractor variant (namely, MOG2) KaewTraKulPong and Bowden (2002) with a learning rate $\lambda = 0.05$ using the implementation provided by OpenCV Bradski (2000). Once the training is complete, we instead start by subtracting the background from the frame to remove static objects (*e.g.* the tank's walls) from following operations. Then, we run a blob detection algorithm to extract regions of interest that could potentially be identified as individuals. We directly remove small blobs with size much smaller than the animals being tracked. We use the remaining blobs to create a masked frame which contains black pixels outside their boundaries and the original pixel intensities inside, substantially reducing the computational load of subsequent operations. We apply a corner detection algorithm, namely, the Shi-Tomasi Shi et al. (1994), on the masked frame to extract candidate pixel positions of the individuals' heads. These two pieces of information, that is, the coordinates of the head and the centroid of the blob, give a good estimate of the individuals' instantaneous heading directions and positions at any time, without the use of past frames.

Finally, we convert all coordinates from their pixel values to their actual coordinates in meters. To achieve this, we designed an auto-calibration routine that uses the robot and a lure to traverse a grid of points, while recording the positions of the robot and lure separately. After

the grid is traversed, we use the two resulting matrices to solve a Perspective- n -Point system Marchand et al. (2015) and to retrieve the roto-translation matrices that allow for converting coordinates between the two cameras' different coordinate systems. This allows for precisely fusing information from the top and bottom cameras to identify which individual is indeed the lure, an important feature for the computation of behavioral models. In Fig. 5.1n, we depict an example of what this result looks like, where the individual in red color has been identified to be the lure.

In our experiments with up to 25 individuals, the described processes were completed in approximately half the time it takes for a new frame to arrive (*i.e.*, at ≈ 60 Hz), using a computer equipped with an AMD Ryzen Threadripper 2970wx 24-core processor and an NVIDIA GeForce RTX 2080 Ti graphics card. We take advantage of the remaining time to apply a trajectory assignment routine (see Fig. 5.1h). More specifically, within the routine we apply operations to account and correct for missing individuals in frames (*i.e.*, in case of overlapping or very small distances between individuals – which is particularly common in fish groups). More importantly, we attempt to maintain unique identifiers (IDs) for each individual across different frames in real-time. Whereas there exist many solutions to solve this problem offline Pérez-Escudero et al. (2014); Romero-Ferrero et al. (2019); Gallois and Candelier (2021), those methodologies are not fit for real-time systems, as they depend on past and future trajectory points to discern between individuals. Instead, real-time systems, like CATS Bonnet et al. (2017), often bypass this problem by using behavioral models that, by design, do not require trajectory information or work under the assumption that errors are small. In BOBI, we formulate the ID assignment as a combinatorial optimization problem and apply a Hungarian method Kuhn (1955) to solve it. That is, we solve the minimization problem that follows in real-time:

$$\min_P \text{Trace}(P \cdot C), \quad (5.1)$$

where P is a permutation matrix and C a cost matrix of size $m \times n$. Assuming that there is no overlapping between individuals of two subsequent frames (our algorithm is able to automatically correct overlapping coordinates in some cases), then $m = n$. Otherwise, the minimization problem (5.1) is solved for the cost matrix $C_{m \times n}$ and the missing individual coordinates are directly copied from past observations. In some cases, if prior observations are also uncertain, our algorithm allows the system to return a vector of coordinates smaller than the actual number of individuals, but still guarantees the smallest possible assignment cost. We experimentally parametrize and define the cost matrix as

$$C_{m,n} = \begin{pmatrix} f(1,1) & f(1,2) & \cdots & f(1,n) \\ f(2,1) & f(2,2) & \cdots & f(2,n) \\ \vdots & \vdots & \ddots & \vdots \\ f(m,1) & f(m,2) & \cdots & f(m,n) \end{pmatrix} \quad (5.2)$$

with

$$f(i, j) = a \cdot \|\mathbf{p}_i(t) - \mathbf{p}_j(t-1)\| + b \cdot |\phi_i(t) - \phi_j(t-1)|, \quad (5.3)$$

and where $\mathbf{p}_i(t)$, $\mathbf{p}_j(t-1)$ are the Cartesian position vector of the individuals with ID i and j , at time t and $t-1$, respectively. Similarly, ϕ denotes the heading direction of individuals, and a , b are weighing parameters. For fish experiments, we found that $a = 0.95$ and $b = 0.08$ yields good ID assignments results.

Control

We semantically separate the robot control into two types of controllers, namely, behavioral and motion controllers. The former often consist of decision-making models or any general purpose trajectory generator, while the latter translate the desired trajectories into motor velocities.

Our motion control routine, depicted in Fig. 5.1j, comprises a pair of proportional-integral-derivative (PID) controllers to regulate the linear and angular speed of the robot. The output of these controllers is subsequently converted into two motor speed commands, one for each motor. Specifically, the PID variant we utilized incorporates a priori knowledge of the desired speed, since most of our behavioral models generate decisions in the form of velocities.

The mathematical expression for the speed PID controller used in our study can be expressed as follows:

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \frac{de(t)}{dt} + K_a V_a, \quad (5.4)$$

where $u(t)$ is the new motor speed, K_p , K_i , K_d , and K_a correspond to the proportional, integral, derivative gains, and an optional a priori speed gain, respectively. The e function denotes the error signal, which represents the difference between the desired and actual speed, while V_a refers to the desired or a priori speed that can be explicitly specified in the behavioral control module.

For the linear speed PID, we use (5.4), and define the error function as $e_{\text{linear}}(t) = \|\mathbf{p}(t) - \mathbf{p}(t-1)\|$, *i.e.*, the Euclidean distance between the current and goal position. Similarly, for the angular speed PID, we use the error function $e_{\text{angular}}(t) = (\phi(t) - \phi(t-1)) \bmod \pi$. We typically find that when using the additional a priori component, the robot has smoother acceleration and deceleration phases, and simultaneously accounts for the loss of speed (due to the phase smoothing) with the proportional gain, and effectively maintains the desired average speed V_a over short durations of time.

We denote V' and ω' the linear and angular speed generated by the two PIDs, respectively. Then, the left and right motor speeds for the differential drive, denoted as V_l and V_r , respectively, are computed as follows:

$$V_l = \frac{2V' - \omega'L}{2}, \text{ and } V_r = \frac{2V' + \omega'L}{2} \quad (5.5)$$

where $L = 0.0451$ m the distance between the center point of the two wheels.

Robot Interface

We currently offer support for 3 robots in BOBI, namely, the LureBot, the FishBot v4.4 used in our previous studies Bonnet et al. (2019, 2017) and Chapter 3 and the Thymio II Mondada et al. (2017). In the remaining of this section, we focus on the LureBot, our latest hardware addition that is presented here for the first time.

To interface with the LureBot we use the BLE protocol, which offers a publish-subscribe communication model³. More specifically, the high-level interfacing routine, depicted in Fig. 5.1k, is responsible for three tasks: 1) to send the motor commands to the corresponding BLE characteristic; 2) to react to notifications from the LureBot which include information about the robot state (*e.g.*, current motor speeds, IR sensor values, etc.), and to communicate those back to the motion control module (see Fig. 5.1j); 3) to verify the stability of the communication link. In fact, the latter is a safety procedure we devised on both ends of the communication, the robot, and the computational machine running the high-level routine, and can be summarized as an abrupt motor stop in case there has not been any incoming message for more than 0.5 s, or one of the two ends is detecting largely desynchronized communication. In the last part of **Video 1** of Appendix C, we illustrate how the LureBot operates in a closed loop and ultimately interacts with the fish.

5.5 Lures

In our previous studies Chapter 3, Bonnet et al. (2018, 2016); Cazenille et al. (2015), we made use of commercially available fishing lures, with a body length of 4.0 cm. They were factory-painted, and their tail was specifically designed to oscillate passively to attract attention. However, this presented two problems: 1) small-sized lures (length < 4 cm) are hard to find in commercial stores, and more importantly, 2) they rarely replicate the color patterns of actual fish in detail, especially those that are primarily lab animals, like rummy-nose tetra or zebrafish. In this section, we present the methodology we developed to construct two lures, a visually biomimetic and a non-biomimetic disc-shaped, with equipment that is readily available in most laboratories.

³Robot code (low level; see Fig. 5.1l) available at: <https://doi.org/10.5281/zenodo.7802052>




Figure 5.3: **Lure construction process.** (a) Picture of an actual *Hemigrammus rhodostomus* (photo by David Villa ScienceImage/CBI/CNRS, Toulouse). (b) Mold (cross-section view) constructed from a 3D scan of an actual *H. rhodostomus*. PLASTIBAITS® Low Shore liquid plastic mixed with a silver pearl color pigment is inserted from the top. (c) Picture of the lure without paint. (d) Picture of the final hand-painted biomimetic lure after the application of protective coating. (e) Picture of a disc-shaped black lure.

5.5.1 Biomimetic (BM) Lure

To construct a biomimetic (BM) lure replicating the color patterns of *H. rhodostomus*, we first obtained a high quality 3D model of an actual *H. rhodostomus*. We built the replicas out of plastic, namely, PLASTIBAITS® Low Shore, which is rated at 22 Shore hardness. We empirically found that this hardness index provides enough flexibility in-water to allow for sufficient deformation when in contact with objects (e.g., the tank wall or neighboring fish), and for passive tail movements due to water flow. The plastic, which is initially in liquid form, contains a hardener that is activated at high temperature, approximately 150°. Therefore, we built the molds out of high-temperature tolerance resin material (see Fig. 5.3b). Then, we replicated the color patterns of *H. rhodostomus* in a 4-step process: 1) we introduced a small amount of silver pearl color pigment in the liquid plastic before pouring it into the mold; 2) we let the lures dry completely (usually a few hours) (see Fig. 5.3c); 3) we hand-painted the lures with consumer-level, alcohol-based markers and let them dry; 4) we coated the painted lures with varnish to protect the color. The final result is the lure depicted in Fig. 5.3d.

5.5.2 Disc-shaped (DS) Lure

For validation purposes, we also built a disc-shaped (DS) lure (see Fig. 5.3e) with the same process as described in the previous section. The DS lure was painted matte black using the same alcohol-based markers used for the biomimetic lure. Additionally, we molded the disc-shaped lure to have the same exact volume as the biomimetic, such that the only difference between them lies in their color and shape.

5.6 Trajectory generators

In this section, we describe the three trajectory generators that we have implemented on the robot to validate its capabilities. In the first two cases, the robot is passive and does not react to the fish, and it follows a predetermined circular trajectory or an eightfold rose trajectory. We specifically designed these non-biomimetic trajectories to study the extent to which actual fish interact with the two different lures, BM and DS. In the third case, the robot decisions are commanded by means of a state-of-the-art behavioral model Calovi et al. (2018) that generates realistic trajectories for *H. rhodostomus*.




Figure 5.4: **Open-loop trajectory generator patterns.** (a) Circular (clockwise) trajectory. (b) Eightfold rose. The trajectory is constructed by rotating an eight-like pattern by a step of 45° (see also Fig. 5.7a for the full trajectory).

5.6.1 Circular trajectory (CT)

The robot performs a clockwise circular trajectory of radius $R' = 22.5$ cm (see Fig. 5.4a; **Video 1** of Appendix C). During this movement, the robot uses the PID controller summarized in (5.4) to track circular trajectory waypoints. Individual waypoints are generated every 0.2 s and advance the goal position by 0.1 rad, *i.e.*, which corresponds to a theoretical constant speed movement of $V = 11$ cm/s, similar to the average speed exhibited by *H. rhodostomus* (refer Chapter 3). In this case, the robot follows a trajectory similar to what an actual single fish would do, by staying very close to the tank wall. This model is open-loop, in the sense that the robot does not react to the movements of the fish or to its relative position to the wall.

5.6.2 Eightfold rose trajectory (R8T)

We have designed a second open-loop trajectory generator, where the robot is instructed to perform a trajectory with an eightfold symmetry, similar to a so-called rose in mathematics (see Fig. 5.4b; **Video 1** of Appendix C). Contrary to the previous case of a circular trajectory, the robot now spends long periods in the innermost area of the experimental tank. In fact, the interest of this trajectory lies in the fact that it allows us to explore whether a fish is willing to follow a DS or BM lure into areas of the tank that it would not normally visit.

5.6.3 Biomimetic interaction model (BIM)

We have also implemented a closed-loop behavioral model describing the social interactions exhibited by *H. rhodostomus*, as defined in Calovi et al. (2018). This model has been shown to reproduce the actual fish interaction dynamics with great accuracy, in numerical simulations Calovi et al. (2018), Chapter 4 and by implementing the model in groups of CUBOID robots Lei et al. (2020) (these CUBOID robots were moving on a disc in open air and were not interacting with any fish). Here, we implement the model in the LureBot to conduct preliminary experiments with fish interacting with a lure (biomimetic or disc-shaped) propelled by the robot.

The model is constructed through a computational methodology Escobedo et al. (2020) for fish species that perform a burst-and-coast swimming mode Calovi et al. (2018), like rummy-nose tetra (*H. rhodostomus*) and zebrafish (*Danio rerio*). Indeed, their motion consists in a succession of short and sudden acceleration phases (kick or burst), each followed by a longer deceleration period, almost in a straight-line (gliding or coasting phase). The updated heading and speed of a focal agent are decided during the kick phase and depend on its heading and distance with respect to the tank wall, and on its relative heading, relative position and distance to neighboring agents Calovi et al. (2018); Lei et al. (2020). We exploit the tracking features presented in previous sections to obtain real-time information and plan the robot's motion at time $t + \Delta t$, where $\Delta t = 1/30$ s, the smallest possible duration between two image frames. During the kick phase, the robot computes its target position and rapidly accelerates towards it. Subsequently, during the gliding phase, the robot attempts to maintain a straight-line movement with the speed profile computed at kick time. We also use the a priori component of the PID in (5.4) to approximate the speed profile that the model computes.

5.7 Protocols and Data Treatment

5.7.1 Experimental Procedure

Hemigrammus rhodostomus (rummy-nose tetras) were purchased from Amazonie Labège in Toulouse, France. Fish were kept in 16 L aquariums on a 12/12 hour, dark/light photoperiod, at 27°C with a 30 min dimming period between phases and were fed *ad libitum* with fish flakes. The average body length of the fish used in these experiments is 35 mm. A trained technician feeds and verifies the housing conditions every day between 8:30 am and 9:30 am.

All behavioral experiments presented in the following sections were conducted within a circular arena of radius $R = 25$ cm. We filled the experimental arena in such a way that the water level inside the circular arena was approximately 5 cm. The water is supplied by the same water filtering system used for housing the fish. Therefore, it has the same salinity and conductivity conditions. Similarly, the temperature inside the tank is maintained at an average temperature of 27°, to match the one in the rearing tanks. The behavioral setup presented in previous sections is housed in a separate room to the one containing the rearing tanks.

Experiments were performed with fish that have been fed and only during the fully lit periods, between 9:00 am and 20:30 pm.

For all experiments, we follow the same procedure, summarized below:

1. Fish are randomly selected from a rearing tank that has not been used during the previous day, and placed in the circular arena.
2. They are subsequently given an acclimatization period of 15 minutes. During this period, if a fish is exhibiting a high-stress level (belly up; floating without moving at all...), we return it to its rearing tank and randomly pick another fish.
3. Each fish or fish group is left to interact with a lure for 1 hour and is then returned to its corresponding tank. For biohybrid experiments we adapt our procedure as follows:
 - (a) we first introduce the BM lure and conduct the experiment for 1 hour;
 - (b) we remove the BM lure and let the fish rest for 15 minutes;
 - (c) we introduce the DS lure and conduct another 1-hour experiment with the same fish.
4. We return all fish to their rearing tank.

This procedure ensures that no fish is tested twice in the same day, nor two days in a row, to avoid the fish getting accustomed to the lure or to specific patterns it exhibits.

5.7.2 Data Filtering

We use `idtracker.ai` (v4) Romero-Ferrero et al. (2019) with a high-resolution recording (1500×1500 pixels) of the experiment to track the movements of all agents and extract their 2D trajectories offline. The software reports an average tracking success rate greater than 99.5%, a result that is further validated by manual inspection. Additionally, we run a post-processing procedure that checks for and corrects any remaining instances where agents' identities are misclassified or missing (e.g., when `idtracker.ai` cannot detect all individuals at a specific video frame). Analyses of later sections focus on the interaction of the lure with the fish, therefore, long periods during which the fish are barely moving are removed. More specifically, given the 35 mm average body length of *H. rhodostomus*, we remove intervals during which the fish or the lure speed is less than 1 BL/s. This procedure also removes instances where the magnetic coupling between the robot and lure is lost, and the lure is stationary. The percentage of the experiment that corresponds to those periods is further discussed hereafter, and is referred to as the "inactivity percentage", for brevity. Finally, all trajectories are resampled with a timestep of $\Delta t = 0.1$ s instead of the original 1/30 s. The new timestep is carefully selected to reduce the random noise introduced between subsequent frames and the dimensionality of the dataset, but is still small enough to study the interactions of fish and lure.

5.8 Assessing the dynamics of biohybrid groups

We validate our robotic system software and hardware by conducting experiments first focusing on the impact of the choice of lure (DS or BM), and then, studying the behavior of the robot, in particular, when interacting with 1 or 4 actual fish⁴. Furthermore, we consider 5 observables to evaluate the robot's ability to move similarly to a fish and validate the importance of the visually biomimetic lure:

1. a qualitative comparison of the agents' trajectories by means of density heatmaps indicating the areas of the circular experimental tank visited by the robot and the fish during a given experiment;
2. the probability density function (PDF) of an agent's speed V ;
3. the PDF of an agent's acceleration α ;
4. for $N > 1$ agents, the probability density function of the group's interindividual distance:

$$d = \sqrt{\frac{1}{N-1} \sum_{j \neq i} \|\mathbf{p}_i - \mathbf{p}_j\|^2}, \quad (5.6)$$

where i is the focal agent and the index j runs over its $N - 1$ neighbors;

5. for $N > 1$ agents, the inactivity percentage, *i.e.*, the percentage of time when the focal agent is barely moving, as detailed in the Sec. 5.7.2.

Moreover, in order to quantify the (dis)similarity between the PDF produced by the fish and the DS/BM lures, we have also considered the Hellinger distance between these PDF (for more details, including the detailed definition, see Appendix C; Tables C.1 to C.4 there).

In the first set of experiments, we study the impact on the behavior of a fish of using a biomimetic (BM) lure compared to a disc-shaped (DS) lure. We test each lure in 2 separate open-loop cases: 1) the lure is performing a circular trajectory (CT; see Sec. 5.6); 2) the lure is performing an eightfold rose trajectory (R8T; see Sec. 5.6). Hence, in this first series of experiments, the robot/lure is passive and does not react to the fish.

The second set of experiments consists of 3 closed-loop cases, where the robot is now commanded by the biomimetic interaction model (BIM): 1) without any fish in the setup, and generating realistic velocity commands when the robot is interacting with the tank wall alone (a baseline to validate that the robot can reproduce the basic motion profile of a single *H. rhodostomus*); 2) in presence of a single *H. rhodostomus* fish, and for both the BM and DS lure, and comparing the results to the spontaneous motion of a pair of *H. rhodostomus*; 3) with 4 *H. rhodostomus*, and comparing the results to the spontaneous motion of 5 actual *H. rhodostomus*. **Video 1** of Appendix C shows side-by-side comparisons of all experiments.

⁴Data available at <https://doi.org/10.5281/zenodo.7796158>

5.8.1 Experiment N° 1: Open-loop dynamics




Figure 5.5: **Short trajectories from the open-loop experiments.** (a) Disc-shaped (DS) lure swimming in a circular motion with a *H. rhodostomus* in the tank. (b) Biomimetic (BM) lure swimming in a circular motion with the same *H. rhodostomus* in the tank. (c) DS lure swimming in an eightfold rose motion with a *H. rhodostomus* in the tank. (d) BM lure swimming in an eightfold rose motion with the same *H. rhodostomus* in the tank. On all trajectories, a dot is shown every second, and the X marker represents the start of a trajectory. In both open-loop experiments, the fish responds more faithfully to the passive motion of the BM lure (see also Figs. 5.6 and 5.7 for a more quantitative assessment).

Circular trajectory

Fig. 5.5a and b display a short excerpt of the circular passive trajectory of the DS and BM lures, along with that of the fish interacting with the lure (see also **Video 1** of Appendix C). More quantitatively, Fig. 5.6 presents the results corresponding to a 1-hour-long circular trajectory of the robot and the resulting trajectory of a single fish interacting with the DS lure (insets a-d) or the BM lure (insets e-h). Fig. 5.6i shows that when the fish swims with the DS lure in the tank, it maintains a typical distance of 46 cm from the lure, almost the maximum possible distance from the robot's predefined trajectory of radius $R' = 22.5$ cm. Remarkably, when the DS lure is replaced by the BM lure, the same fish maintains a much smaller typical distance of 3 cm. The DM lure is hence unable to capture the fish's attention. In fact, in the presence of the DS lure, the fish remains inactive for long periods, with an inactivity percentage (see Sec. 5.7.2) of 71.1% (see Fig. 5.6j). However, in the presence of the BM lure, the fish is only inactive for 6.4% of the experiment's duration.

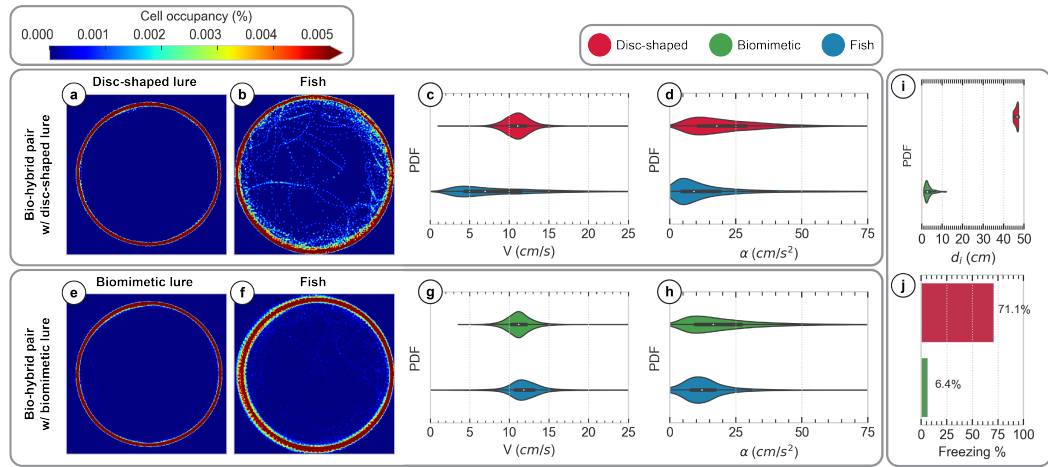


Figure 5.6: Motion profile when a fish interacts with a passive DS or BM lure (CT). (a) 1-hour-long robot circular trajectory (CT) with the disc-shaped (DS) lure. (b) 1-hour-long *H. rhodostomus* spontaneous trajectory swimming with the DS. (c) Probability density function (PDF) of the DS's (red color) and fish's (blue color) speed V . (d) PDF of the DS's (red color) and fish's (blue color) acceleration α . (e) 1-hour-long robot circular motion trajectory with the biomimetic (BM) lure. (f) 1-hour-long *H. rhodostomus* spontaneous trajectory swimming with the BM. (g) PDF of the BM's (green color) and fish's (blue color) speed V . (h) PDF of the BM's (green color) and fish's (blue color) acceleration α . (i) PDF of the pair's interindividual distance (d ; BM and fish in green color, DS and fish in red color). (j) Inactivity percentage (BM and fish in green color, DS and fish in red color). In all PDFs, the white dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile.

The robot typically moves at 11.3 cm/s (see the red and green PDF peak and median value that match in Fig. 5.6c, g), regardless of the lure choice, which matches the intended speed for this motion (see Sec. 5.6). When the fish swims with the DS lure, the speed of the fish is typically 4 cm/s, much lower than that of the lure, and rarely reaches more than 10 cm/s (see the PDF of the fish speed in Fig. 5.6c). However, when the same fish swims with the BM lure, its typical speed is now 11.4 cm/s (see PDF peak in Fig. 5.6g), very similar to the typical speed of the lure, and the speed PDF of the fish and the lure are also fairly similar (see Fig. 5.6g). Note that the speed PDF of the DS and BM lures display a ± 2 cm/s fluctuation around the median speed, which naturally occurs when the robot slips or deviates from the circular trajectory momentarily and the PID controller (see Sec. 5.4) tries to compensate for it.

In terms of acceleration, the robot produces typical values of 17.5 cm/s² (see the peak of the red and green PDF in Fig. 5.6d, h), regardless of the lure choice. The PDF of the robot's acceleration is relatively narrow, as the robot is producing higher accelerations for compensatory movements alone. When the fish swims with the DS lure, its typical acceleration is 9 cm/s² (see the PDF peak in Fig. 5.6d), but when the same fish swims with the BM lure, the typical acceleration is 3 cm/s² higher (see the PDF peak in Fig. 5.6h). In the DS case, we recorded more instances where the fish had a low acceleration than in the BM case.

We have also computed the Hellinger distance between the PDF of speed and acceleration of the fish and that of the DS and BM lures. The Hellinger distance is a quantifier of the (dis)similarity of 2 PDF and is precisely introduced in the Appendix C. The results of Table C.1 confirm that the PDF of speed and acceleration for the BM lure are in much better agreement with the corresponding PDF for the fish, than for the DS lure.

Eightfold rose trajectory

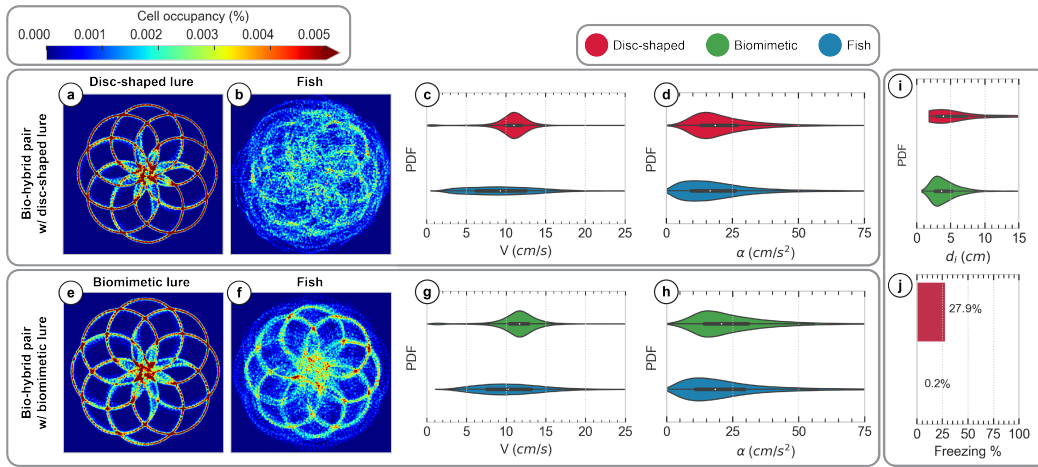


Figure 5.7: **Motion profile when a fish interacts with a passive DS or BM lure (R8T).** (a) 1-hour-long robot eightfold rose trajectory (R8T) with the disc-shaped (DS) lure. (b) 1-hour-long *H. rhodostomus* spontaneous trajectory swimming with the DS. (c) Probability density function (PDF) of the DS's (red color) and fish's (blue color) speed V . (d) Probability density function of the DS's (red color) and fish's (blue color) acceleration α . (e) 1-hour-long robot eightfold rose trajectory with the biomimetic (BM) lure. (f) 1-hour-long *H. rhodostomus* spontaneous trajectory swimming with the BM. (g) Probability density function of the BM's (green color) and fish's (blue color) speed V . (h) Probability density function of the BM's (green color) and fish's (blue color) acceleration α . (i) Probability density function of the pair's interindividual distance (d_i ; BM and fish in green color, DS and fish in red color). (j) Inactivity percentage (BM and fish in green color, DS and fish in red color). In all PDFs, the white dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The fish swimming with the DS and BM lures is actually the same fish (see Sec. 5.7.1).

Fig. 5.5c and d display a short excerpt of the eightfold rose passive trajectory of the DS and BM lures, along with that of the fish interacting with the lure (see also **Video 1** of Appendix C). More quantitatively, a 1-hour-long eightfold rose trajectory of the robot results in the pattern shown in Fig. 5.7a, e. When the fish swims with the DS lure, it qualitatively appears to follow the lure, but not enough to reproduce the clear pattern performed by the robot (see Fig. 5.7b). However, when the BM lure is in the tank with the same fish, the fish trajectory reproduces the eightfold rose pattern with a much better contrast than with the DS lure (compare Fig. 5.7b and f).

Fig. 5.7i shows the PDF of the interindividual distance for the pair. In the DS lure case, the pair typically moves at a distance of 4 cm, but also at greater distances, up to 15 cm. Notably, the PDF shows that the pair never swims closer than 3 cm (see the head of the red PDF in Fig. 5.7i). The fish typically swims at a distance of 3 cm from the BM lure, and can get as close as 1 cm (see the green PDF head and peak in Fig. 5.7i). In addition, the tail of the PDF of the distance between the fish and the lure is notably thinner than in the case of the DS lure. When the fish swims with the DS lure, it is inactive 27.9% of the experiment's duration. On the other hand, when the same fish swims with the BM lure, it practically does not stop moving (inactive 0.2% of the time; see Fig. 5.7j).

The robot's PDF of speed for both lures are marginally different, with typical speed of 11.3 cm/s and 11.7 cm/s for the DS and BM lure, respectively (see the peaks of the speed PDFs in Fig. 5.7c, g). The median speed for the two lure variants coincides with their corresponding peak speed. For both lures, the speed PDF of the fish is markedly narrower than that of the robot. In the DS case, the fish moves at a typical speed of 9.5 cm/s (see Fig. 5.7c). When swimming with the BM lure, the same fish moves slightly faster (see Fig. 5.7g), with a typical speed of 10 cm/s, and with a narrower speed PDF than in the DS case.

The acceleration PDF of the robot for the two lures are marginally different, with peak values of 15 cm/s² and 17 cm/s² for the DS and BM lure, respectively. This is due to more frequent slipping and trajectory deviation instances in the BM's case, which are caused by the forces applied by the water on its larger footprint. The fish accelerated in an almost identical manner in both cases with a typical acceleration of 12 cm/s² (see the peaks of the PDF in Fig. 5.7d, h), although, in the BM case, the fish acceleration PDF presents a slightly thicker tail, also observed in the BM lure's PDF.

Again, the results of Table C.2 in the Appendix C for the Hellinger distance confirm that the PDF of speed and acceleration for the BM lure are in better agreement with the corresponding PDF for the fish, than for the DS lure. However, Table C.2 also points to the fact the PDF of the speed for the BM lure does not reproduce that of the fish as well as for circular trajectories (see Table C.1).

In conclusion, the results of our first series of experiments indicate that a rummy-nose tetra fish interacts much strongly with our biomimetic lure than with the disc-shaped lure. For a circular trajectory, the DS lure is totally unable to capture the attention of the fish, whereas the fish faithfully follows the BM lure, staying at close distance. For the eightfold rose trajectory, the fish follows the DS lure, but without precisely matching the trajectory of the lure. When swimming with the BM lure, the fish remains closer to the lure, and reproduces much more faithfully the complex trajectory of the lure. Finally, for both types of trajectory, the fish is much more active when interacting with the BM lure than with the DS lure.

5.8.2 Experiment № 2: Closed-loop dynamics

Single agent interacting only with the tank wall

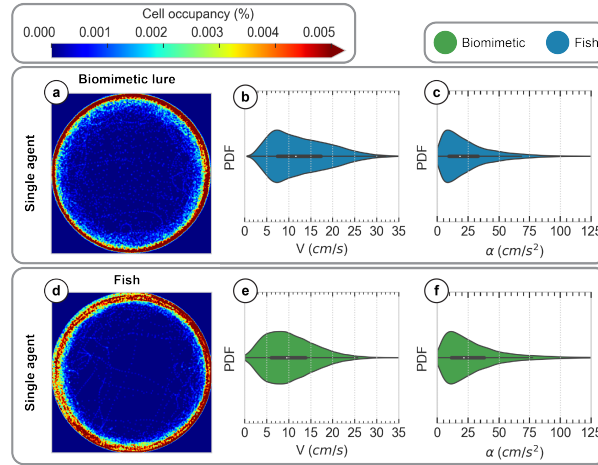


Figure 5.8: **Single agent motion profile.** (a) 1-hour-long robot trajectory generated by the Biomimetic Interaction Motion (BIM) model. (b) Probability density function (PDF) of the robot speed V . (c) Probability density function of the robot acceleration α . (d) 1-hour-long *H. rhodostomus* spontaneous trajectory. (e) Probability density function of fish speed V . (f) Probability density function of fish acceleration α . In all PDFs, the white dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile.

Fig. 5.8a and Fig. 5.8d depict the 1-hour-long trajectory of the robot and fish, respectively (see also **Video 1** of Appendix C). The fish and the robot (driven by the BIM) both stay close to the tank wall – a consequence of the burst-and-coast dynamics Calovi et al. (2018) – although the robot is on average slightly closer to the wall than the fish.

The fish typically moves at a speed of 7 cm/s (see the peak of the PDF in Fig. 5.8b), but also exhibits faster movements of up to 30 – 35 cm/s. On the lower end, the fish does not move below 2 cm/s very often. The robot moves at a marginally different typical speed of 7.5 cm/s (see the peak of the PDF in Fig. 5.8e) and produces a narrower speed PDF tail in Fig. 5.8e), and maximum speeds in the range 25 – 30 cm/s, but also moves more often with a speed between 0 and 2 cm/s. The differences at the head and tail of the two PDFs are also reflected in the difference between the robot’s median speed of 9.6 cm/s and the fish’s median speed of 11.6 cm/s.

The robot and the fish move with a typical acceleration of 9.5 – 10 cm/s² (see the peaks of the PDFs in Fig. 5.8c, f). Overall, the PDF of the acceleration of the robot and the fish are in excellent agreement (Fig. 5.8c, f), although the PDF for the robot displays a slightly fatter tail, which results in a slightly higher median acceleration for the robot (21.7 cm/s²) than for the fish (18.5 cm/s²).




Figure 5.9: **Short trajectories from the closed-loop experiments.** (a) Trajectories recorded from the spontaneous interactions of a *H. rhodostomus* pair in the tank. (b) Disc-shaped (DS) lure commanded by the behavioral model and interacting with a *H. rhodostomus* (Fish 0). (c) Biomimetic (BM) lure commanded by the behavioral model and interacting with a *H. rhodostomus* (Fish 0). Note that fish 0 is the same across the 3 experiments. On all trajectories, a dot is shown every second, and the X marker represents the start of a trajectory. Again, the fish responds more faithfully to the motion of the BM lure than to the DS lure (see also Fig. 5.10 for a more quantitative assessment).

Finally, the results of Table C.3 in the Appendix C for the Hellinger distance confirm that the PDF of speed and acceleration for the BM lure are in good agreement with the corresponding PDF for the fish.

Biohybrid group of 2 agents

Fig. 5.9a-c respectively displays a short excerpt of the trajectories of 2 fish, of a fish interacting with the DS lure, and of the same fish interacting with the BM lure (see also **Video 1** of Appendix C). More quantitatively, in Fig. 5.10, we show the full trajectories and the corresponding PDF for pairs of *H. rhodostomus* fish (insets a-d), the DS lure and a fish (insets e-h), and the BM lure and a fish (insets i-l). We use one fish (with ID 0) across all 3 cases to obtain a direct comparison. In both the biohybrid pair interaction experiments, the trajectories follow a similar trend, where the pair usually stays close to the wall, and only rarely swims away from it. This is in line with what is observed in actual groups of *H. rhodostomus* (see Fig. 5.10a, b). However, in the case of the DS, the actual fish tends to swim markedly closer to the wall than the fish-only pair and the biohybrid pair with the BM lure.

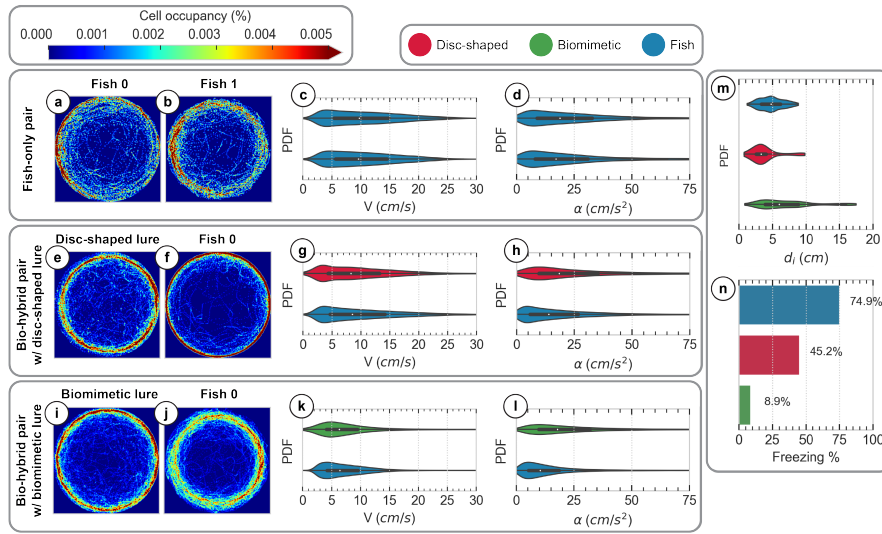


Figure 5.10: **Motion profile for pairs of agents.** (a) and (b) 1-hour-long spontaneous trajectories for a pair of *H. rhodostomus* fish. (c) Probability density functions (PDF) of the speed V for the *H. rhodostomus* pair. (d) Probability density functions of the acceleration α for the *H. rhodostomus* pair. (e) 1-hour-long robot (using the biomimetic interaction model) trajectory with the disc-shaped (DS) lure. (f) 1-hour-long *H. rhodostomus* spontaneous trajectory swimming with the DS. (g) Probability density function of the speed V for the DS (red color) and fish (blue color). (h) Probability density function of the acceleration α for the DS (red color) and fish (blue color). (i) 1-hour-long robot (using the biomimetic interaction model) trajectory with the biomimetic (BM) lure. (j) 1-hour-long *H. rhodostomus* spontaneous trajectory swimming with the BM. (k) Probability density function of the speed V for the BM (green color) and fish (blue color). (l) Probability density function of the acceleration α for the BM (green color) and fish (blue color). (m) Probability density function of the pair's interindividual distance (d_i ; BM and fish in green color, DS and fish in red color, *H. rhodostomus* pair in blue color). (n) Inactivity percentage (BM and fish in green color, DS and fish in red color, *H. rhodostomus* pair in blue color). In all PDFs, the white dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. In all three experiments, fish 0 is actually the same fish (see Sec. 5.7.1).

The PDF of the speed of both *H. rhodostomus* agents are almost identical, and shows that they typically swim at a speed of 4 cm/s, but at times reach higher speeds of order 27 – 30 cm/s (see the peaks and tails of the speed PDFs in Fig. 5.10c). The typical swimming speed for the biohybrid group with the DS lure, is 3.5 cm/s (see the peaks of the speed PDFs in Fig. 5.10g), with the two agents having an almost identical distribution of speed. The width of the speed PDF in the DS case is also marginally different from the ones from the fish-only case. In the biohybrid group with the BM lure, the speed PDF of the two agents are once again marginally different, and their typical speed is 4 cm/s and 4.8 cm/s for the fish and lure, respectively (see the peaks of the PDFs in Fig. 5.10k). Furthermore, the two speed PDFs show that the agents swim generally slower, almost always below 19 cm/s. Similar results are obtained for the PDFs of acceleration. Indeed, the two fish acceleration PDFs are almost identical, with

typical accelerations of 8 cm/s^2 (see the peaks of the PDFs in Fig. 5.10d) and at times as high as 75 cm/s^2 (see the tails of the PDFs in Fig. 5.10d). In the biohybrid group with the DS lure, the robot's typical acceleration is 9 cm/s^2 , while the fish typical acceleration is 4 cm/s^2 (see the peaks of the PDFs in Fig. 5.10h). The two agents do not often accelerate with high values ($> 50 \text{ cm/s}^2$). For the biohybrid group with the BM lure, the acceleration profile is marginally different from the DS case (see the PDF of acceleration in Fig. 5.10l).

The pair of *H. rhodostomus* tend to swim at a distance of 5 cm from each other (see the peak and median of the PDF of interindividual distance in Fig. 5.10m). In the DS case, the pair swims on average at a smaller distance of 3 cm, and sometimes as far as 10 cm. When the same fish swims with the BM lure, their typical distance is 4 cm, but as shown in Fig. 5.10m, the PDF of interindividual distances is wider than in the other two cases, meaning that the fish swim more often far from each other, as far as 17 cm. However, when the robot is following trajectories generated by BIM, it is by design attracted to its neighbor(s), regardless of its movement. Therefore, the interindividual distance is effectively coupled with the inactivity percentages shown in Fig. 5.10n. As shown in Fig. 5.10n, the fish-only pair does not move for 74.9% of the experiment's duration. When the DS lure replaces one of the agents, then, the pair is inactive less often, only 45.2% of the time. However, the lowest inactivity percentage out of the three cases is reported when the same fish is swimming with the BM lure. Then, the pair is inactive for 8.9% of the experiment's duration, which potentially explains the wider PDF of interindividual distance in Fig. 5.10m.

The results of Table C.4 in the Appendix C for the Hellinger distance show that the PDF of speed and acceleration for the BM and DS lures are in comparable and fair agreement with the corresponding PDF for the fish (Hellinger distances $\lesssim 0.2$). However and as noted above, for both lures, the PDF of the distance to the fish is markedly different from the PDF of the distance between 2 fish, resulting in Hellinger distance values of 0.455 and 0.411 for the DS and BM lures, respectively.

Finally, in all our experiments presented up to now, we find that the biomimetic lure has a stimulating effect on the fish, resulting in the fish having a higher activity than when it interacts with the disc-shaped lure, or even when it interacts with a conspecific (see **Video 1** of Appendix C).

Biohybrid group of 5 agents

A qualitative assessment of the density heatmaps of Fig. 5.11a and d reveals that the biohybrid group (the robot and 4 fish) and the 5-fish group both tend to move close to the wall, although the radial dispersion of the biohybrid group appears to be larger. Yet, the robot radial dispersion is, in fact, very similar to that of fish 0 and 1 in the fish-only group. Furthermore, in both cases, we observe that the robot and the fish tend to adopt a similar trajectory radius throughout the experiment (in particular, see the atypical low radial dispersion of fish 3 in Fig. 5.11, or the smallest trajectory radius observed for fish 3 in Fig. 5.11d).

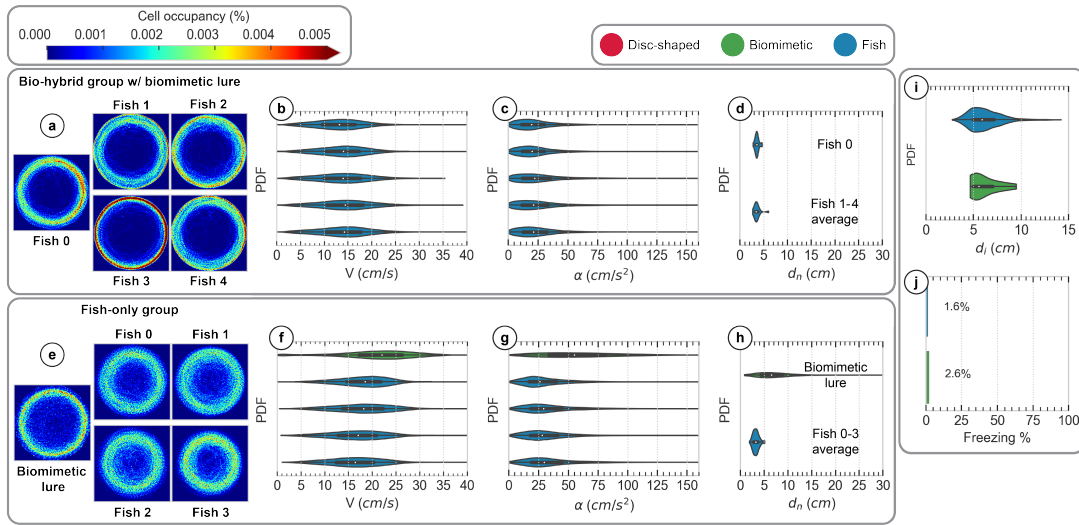


Figure 5.11: Motion profile for groups of 5 agents. (a) 1-hour-long spontaneous trajectories of 5 *H. rhodostomus*. (b) Probability density functions (PDF) of the speed V for each agent in the fish-only group. (c) Probability density functions of the acceleration α for each agent in the fish-only group. (d) 1-hour-long fish and robot (using the biomimetic interaction model; with the biomimetic lure) trajectories. (e) Probability density functions of the speed V for each agent in the biohybrid group (robot in green color and fish in blue color). (f) Probability density functions of the acceleration α for each agent in the biohybrid group (robot in green color and fish in blue color). (g) Probability density function of the pair's interindividual distance (d ; biohybrid group in green color, and fish-only group in blue color). (h) Inactivity percentage (biohybrid group in green color, and fish-only group in blue color). In all PDFs, the white dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile.

The group of 5 fish swims at synchronized speeds, as shown by their speed PDFs in Fig. 5.11b. All fish typically swim at a speed of 14 – 15 cm/s and not faster than 25 cm/s (see the peak and tail of the PDF of speed, respectively). In the biohybrid group, all agents swim typically faster than the 5 fish-only group, with typical speeds in the range 17 – 19 cm/s for the fish and 23 cm/s for the robot (see the peaks of the corresponding speed PDFs in Fig. 5.11e). They also reach higher speeds than the fish-only group, up to 30 cm/s for the fish and 35 cm/s for the robot. In other words, the robot commanded by the BIM swims faster, and the fish adjust their speed profile accordingly.

Fig. 5.11c shows that agents in the 5 fish-only group have almost identical acceleration PDFs, with a typical acceleration of 16 cm/s², and less frequent high values up to 60 cm/s²). In the biohybrid group case, the PDF of acceleration for the robot is wider, with a typical value of 35 cm/s², but also producing higher acceleration, up to 90 cm/s² (see the peak and tail of the PDF of acceleration for the robot in Fig. 5.11f). The fish in the biohybrid group also tend to produce higher accelerations, in the range 20 – 25 cm/s², but rarely going above 75 cm/s².

The group of 5 *H. rhodostomus* typically swims in close formation with a typical interindividual distance $d \sim 5$ cm (see the PDF of interindividual distance in Fig. 5.11g). The group does not often swim with a distance greater than 9 cm or smaller than 2.5 cm (see the tail and head of the PDF of interindividual distance). Despite the speed and acceleration differences, the biohybrid group also swims in close formation, with a typical $d \sim 6$ cm (see the peak of the PDF of interindividual distance in Fig. 5.11g). Similarly to the fish-only group, the interindividual distance does not get larger than 9 cm, but is remaining above 4 cm (see the tail and head of the PDF of interindividual distance; Fig. 5.11g). The median interindividual distance is very similar in both cases (see the white points in Fig. 5.11g). Both groups of agents remained very active throughout the experiment (see Fig. 5.11h), with an inactivity percentage of 2.6% and 1.6% for the biohybrid and fish-only group, respectively.

Despite the higher radial dispersion observed in the biohybrid group, and the higher typical robot speed compared to that of its 4 companion fish, the most important conclusion of this experiment is that the biohybrid group remains essentially as compact as the fish-only group. In particular, not only is the robot able to participate in the collective dynamics, but its presence does not lead to the dislocation of the group (see **Video 1** of Appendix C). This, along with our experiments with one fish interacting with a passive biomimetic lure (see Sec. 5.8.1) or in closed-loop with the BIM robot (see the previous section), shows that the robot is fairly well accepted by the fish, which respond to its presence in a way consistent with their response to the presence of one or four conspecifics.

5.9 Conclusion

The complexity of biohybrid systems often necessitates experimentation with small groups of animals, simplified behavioral models, and limited durations. The latter constraint is particularly limiting when studying long-term interactions that emerge within fish groups (see Chapter 3).

In this chapter, we have introduced a versatile framework for behavioral experiments and open-sourced its hardware and software components. Our new and improved tools enable the execution of real-time, highly precise biohybrid animal experiments while maintaining a minimal physical footprint in the experimental space. This aspect, to our knowledge, has not been explicitly addressed in similar studies, but we believe it significantly enhances both the experimenter's experience and laboratory conditions. Furthermore, we present a new robotic system, the LureBot, which demonstrates increased agility, precision, and extended operation duration compared to our former system, the FishBot v4.4. Our findings show that, with the robot's augmented torque and new perception and control algorithms, it can effectively interact with up to four *H. rhodostomus*, a very demanding task for any miniature robot. Additionally, we showcase our procedures for constructing biomimetic lures, emphasizing that, in many contexts, such as for fish groups, their visual appearance plays a critical role in facilitating interactions.

More specifically, using this system, we show that fish exhibit a strong preference for biomimetic lures over non-biomimetic ones. Once this has been established, we deploy our robots with a closed-loop behavioral model to interact with groups of fish. This data-driven behavioral model exploits the reconstructed social interactions of *H. rhodostomus* Calovi et al. (2018) and has been shown in simulations to faithfully reproduce the collective behavior of groups of 1, 2, and 5 fish Calovi et al. (2018); Lei et al. (2020). First, we demonstrate that the LureBot can express the entire motion profile spectrum of the species under study (e.g., *H. rhodostomus*), when interacting with the tank wall alone. Next, we present the results of experiments involving one fish and a lure, showing that the fish reacts to the robot similarly as when swimming with another fish. Notably, the biohybrid experiment consistently exhibits high activity, indicating that the robot can continuously engage in interactions. This finding holds significant potential for experiments requiring the collection of large, reproducible datasets. We also conduct experiments with larger groups, consisting of one lure and four fish, and show that the biohybrid group behaves comparably to fish-only experiments. However, in the five-agent experiments, the robot moves faster than individual fish in either the biohybrid group or fish-only experiments. Despite this, group cohesion is maintained, suggesting that although the biohybrid does not perfectly reproduce fish-only experiments, the fish in the group do not seem adversely affected, unlike when non-biomimetic models are used (as shown in Chapter 4).

The research work presented in this chapter not only contributes to the development of more efficient and sophisticated biohybrid systems, but also paves the way for future studies to explore the intricacies of animal behavior and social interactions. The open-source nature of the framework should encourage further advancements and adaptations within the field, thereby fostering a deeper understanding of the mechanisms underlying the collective behavior of various animal species. Furthermore, this chapter's work contributes towards our goal of closing the biomimicry gap, and more specifically in reducing the contribution of the physics and communication cue gap as outlined in Fig. 1.1. That is, our new framework allows for long real-time control experiments that produce marginally different dynamics than those observed in experiments, and the biomimetic lure constructed as part of this investigation is shown to be accepted by fish and facilitates the interactions. The effect of the robot on the social dynamics is further explored in Chapters 6 and 7.

Transferring social interaction models from simulation to reality **Part IV**

6 Evaluating the biomimicry gap in biohybrid pairs driven by a deep learning model

This section's content is adapted from the following article:

- Papaspyros, V., Theraulaz, G., Sire, C. and Mondada, F., 2023. Quantifying the biomimicry gap in biohybrid systems. arXiv. <https://doi.org/10.48550/arXiv.2308.08978> (reuse authorised under the CC 4.0 licence) – Papaspyros et al. (2023c)
 - Vaios Papaspyros's contribution: Conception, data curation, formal analysis, investigation, methodology, software, visualization, writing – original draft, writing – review & editing.
- The article's work was funded by the Swiss National Science Foundation project "Self-Adaptive Mixed Societies of Animals and Robots" (Grant No. 175731). Our collaborators from the Université Toulouse – Paul Sabatier were funded by the French National Research Agency (ANR-20-CE45-0006-01).
- Ethics: All *H. rhodostomus* experiments were conducted at the Centre de Recherches sur la Cognition Animale, Centre de Biologie Intégrative, CNRS, Université de Toulouse – Paul Sabatier. Experiments were approved by the local ethical committee for experimental animals and were performed in an approved fish facility (A3155501) under permit APAFIS#27303-2020090219529069 v8 in agreement with the French legislation.

6.1 Introduction

In previous chapters, we have devised and implemented both software and hardware tools with the objective of minimizing the social interaction gap (Chapter 4), the physics gap (Chapter 5), and the communication cue gap (Chapter 5). The integration of these efforts now allows us, in this chapter, to explore the biomimicry gap as a comprehensive entity. We examine this concept with a focus on pairs of agents, be they pairs of real *H. rhodostomus*, simulated *H. rhodostomus*, or a robot and one *H. rhodostomus*.

Notably, a few flocking models for fish behavior, analytical or machine learning, have been evaluated in extended simulations to study long-term emergent collective behavior Collignon et al. (2016); Calovi et al. (2018); however, these models have not been tested and validated in biohybrid groups. Conversely, numerous models have been deployed on robotic devices without being tested in simulations Abaid and Porfiri (2010); Bonnet et al. (2016); Cazenille et al. (2017, 2018c,b, 2019). Furthermore, the majority of these studies involve robot experiments lasting no more than 30 minutes, with the resulting interactions typically assessed solely in the short term. Consequently, none of these models have been stringently benchmarked in both short- and long-term timescales within both simulation and fish-robot biohybrid experiments. In fact, previous results indicate that certain models may yield satisfactory biomimetic outcomes in the short term while failing to reproduce emergent dynamics accurately on longer time scales (refer to Chapter 4).

Moreover, the transfer of computer models of social interactions into robot controllers that operate in real situations involving animals is not simple and can generate a discrepancy with their numerical simulation, akin to the reality gap observed when transferring simulated robot controllers to real-world applications Jakobi et al. (1995); Jakobi (1997); Mouret and Chatzilygeroudis (2017). However, this gap is only but one out of three primary sources of discrepancies in biohybrid systems (depicted in Fig. 1.1): 1) subtle behavioral patterns that social interaction models fail to capture, 2) physics related to the operation of the robot in real life that were not accounted for, and 3) the extent of biomimicry exhibited by artificial lures Romano and Stefanini (2022a); Chapter 5. We refer to the cumulative effect of these gaps with the term “biomimicry gap”. Therefore, the biomimicry gap is an inherent aspect of the multifaceted, cross-domain process of creating biohybrid groups composed of animals and robots. To the best of our knowledge, the feasibility of bridging this biomimicry gap — achieved by conducting extended experiments in both simulated and real-world environments, and comparing their results — has yet to be conclusively and rigorously validated across all these levels in a single approach.

In this chapter, we investigate this notion by employing the (pretrained) machine learning model presented in Chapter 4. We implement this model on a robotic system, the LureBot (see Chapter 5), and execute approximately 11 hours of multiple pair experiments wherein a biomimetic lure interacts with a single *H. rhodostomus*. This allows us to measure the behavioral differences between actual and simulated pairs *H. rhodostomus*, as well as, pairs of 1 biomimetic lure and 1 *H. rhodostomus*. In turn, this yields the first end-to-end approach aimed at quantifying and reducing the biomimicry gap, and is presented in the following sections.

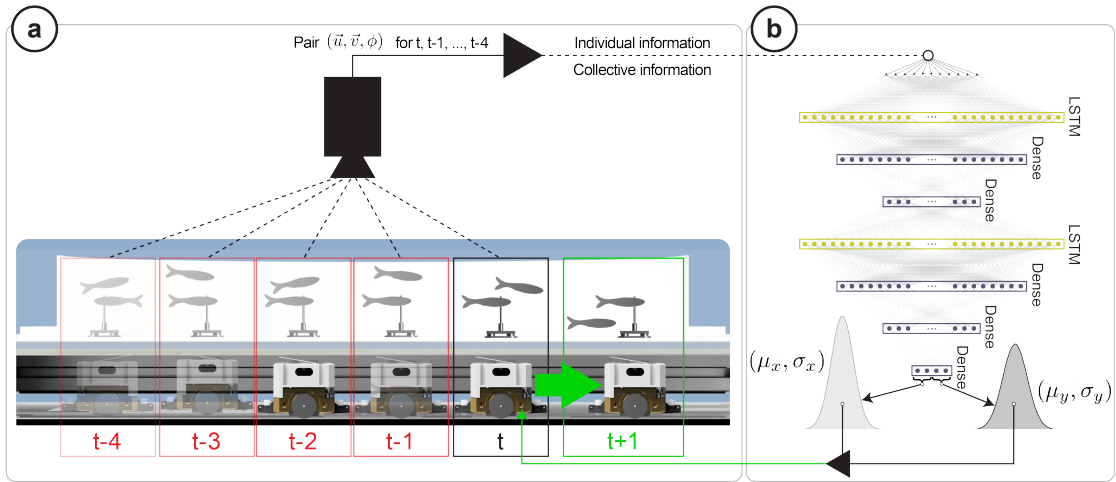


Figure 6.1: **Closed-loop robot control with Deep Learning Interaction (DLI) model.** (a) We use the top setup camera to track all agents (fish and/or lure) in real-time, and store unique trajectories for each agent. A 5×11 vector of individual and collective states, spanning 5 timesteps is fed to the DLI. (b) The DLI outputs two acceleration distributions, one for each Cartesian component. Then, the accelerations are converted to desired speed and position for $t + 1$ and communicated to the robot.

6.2 Methods

6.2.1 Real-time tracking and robot control

Experiments were performed with the Behavioral Observation and Biohybrid Interaction (BOBI) framework introduced in Chapter 5, including the LureBot, to propel a *H. rhodostomus* lure. We exploit BOBI to track multiple agents (here, only 2 are used) in real time, while maintaining unique IDs for each agent's trajectory. These agent-specific sequences of spatial displacement can be exploited by a behavioral model (see Section 6.2.2) to compute real-time individual and collective quantities concerning the biohybrid group, and close the loop of interaction by adapting the robot's behavior with instructions on future movements.

Here, we use the Proportional-Integral-Derivative (PID) controller as defined in BOBI (see Chapter 5), that incorporates a priori velocity information provided by the behavioral model. The PID combines the linear and angular errors between the LureBot's current and desired position, as well as the model's predicted velocity profile, to smoothly displace the robot.

6.2.2 Deep Learning Interaction model

We use a pretrained version of the Deep Learning Interaction (DLI) model, introduced in Chapter 4, to generate real-time goal positions for the LureBot. For readability purposes, this section summarizes the key elements of the DLI. For the full description, refer to Chapter 4.

The DLI consists of 7 layers (see Fig. 6.1b; Chapter 4): 1st and 4th are LSTM layers Hochreiter and Schmidhuber (1997); the remaining are densely connected layers; ReLU activations are used for all layers except for the last one which is linear. For a single agent i , the state at time t is defined as a 1×5 vector $\mathbf{x}_i(t)$:

$$\mathbf{s}_i(t) = (\vec{u}_i(t), \vec{v}_i(t), r_w^i(t)) \in \mathbb{R}^5, \quad (6.1)$$

where $\vec{u}_i(t)$, $\vec{v}_i(t)$ are the 2D position and velocity, respectively, and $r_w^i(t)$ the distance of the individual i from the wall at time t . Then, the pairwise state at time t is summarized in the following 1×11 vector:

$$\mathbf{S}_{ij}(t) = \left(\underbrace{\mathbf{s}_i(t)}_{\substack{\text{individual (focal) \\ information}}, \underbrace{\mathbf{s}_j(t)}_{\substack{\text{individual (neighbor) \\ information}}, \underbrace{d_{ij}(t)}_{\substack{\text{collective} \\ \text{information}}} \right) \in \mathbb{R}^{11}, \quad (6.2)$$

with i the focal individual for which we generate trajectory predictions, j its neighbor, and d_{ij} their interindividual distance. In real-time, we feed the DLI with a 5×11 sequence $(\mathbf{S}(t-4), \dots, \mathbf{S}(t))$ of the pair-wise states (see Fig. 6.1a), where we make sure that i (focal individual) corresponds to the LureBot.

Subsequently, the DLI model outputs the expected acceleration mean and standard deviation value, (μ_x, σ_x) and (μ_y, σ_y) , of the Cartesian components x and y . Assuming a Gaussian distribution for the acceleration Chua et al. (2018), we sample this distribution to produce acceleration predictions $\vec{a} = (a_x, a_y)$ and use the following motion equations to generate velocity commands and the goal position of the LureBot at time $t+1$:

$$\vec{v}_i(t+1) = \vec{v}_i(t) + \Delta t \vec{a}, \quad (6.3)$$

$$\vec{u}_i(t+1) = \vec{u}_i(t) + \Delta t \vec{v}_i(t+1), \quad (6.4)$$

where $\Delta t = 0.12$ s, a choice made with respect to the data filtering procedure applied on the raw data to generate an intermediate training dataset for the DLI. The 2D velocity commands, defined in (6.4), and goal position, defined in (6.4), are given to the BOBI's PID, and eventually translated to motor commands (see Section 6.2.1; see Fig. 6.1a,b).

In Chapter 4, this approach was validated in long simulations and is shown to be capable of reproducing the social dynamics of *H. rhodostomus* pairs faithfully with respect to experiments. In the following sections, we test the extent to which the DLI can produce faithful interactions when deployed on a physical robot-fish group instead of a simulated group.

6.2.3 Evaluating the short- and long-term interactions

To characterize the short- and long-term interactions of the agent pairs, we employ the same observables introduced in Chapter 4. In short, we use 3 observables that correspond to instantaneous quantities at the *individual* level, for which we measure their probability




Figure 6.2: **Individual and collective variables.** For the focal agent i (light gray) we define the individual quantities: \vec{u}_i its Cartesian position, \vec{v}_i its instantaneous velocity, r_w^i its distance to the wall, ϕ_i agent i 's heading angle. We also define the collective quantities from i 's perspective when its neighbor j (dark gray) swims in the tank: interindividual distance d_{ij} . $R = 25$ cm corresponds to the radius of the circular setup, ϕ_{ij} , the heading difference between both fish, and ψ_{ij} , the angle with which fish j is perceived by fish i . Note that, for visualization purposes, the size of agents is not to scale.

density function (PDF): the speed, V of an individual; its distance to the wall, r_w ; and its heading angle relative to the normal to the wall, θ_w . Then, 3 additional observables probe the instantaneous *collective* dynamics: the distance d_{ij} between the pair of individuals; the difference $|\phi_{ij}|$ between the heading directions of the two individuals; and the angle ψ_{ij} at which an individual perceives its neighbor. Finally, we consider 3 *temporal correlation functions* that probe the social dynamics at a very fine level Jayles et al. (2020), and which are generally particularly difficult to reproduce:

$$C_X(t) = \left\langle [\vec{u}_i(t+t') - \vec{u}_i(t')]^2 \right\rangle, \quad (6.5)$$

$$C_V(t) = \langle \vec{v}_i(t+t') \cdot \vec{v}_i(t') \rangle, \quad (6.6)$$

$$C_{\theta_w}(t) = \left\langle \cos \left[\theta_w^i(t+t') - \theta_w^i(t') \right] \right\rangle. \quad (6.7)$$

C_X is the mean-squared displacement, C_V the velocity autocorrelation, and C_{θ_w} the autocorrelation of the angle of incidence to the wall. In general, we denote $C_q(t) = \langle q(t+t')q(t') \rangle$ as the average of the quantity $q(t)q(t+t')$ over the reference times t' , over individuals, and over different experiments. Assuming the stationarity of the system, the temporal correlation function $C_q(t)$ only depends on the time difference between observations, and is often noted $C_q(t) = \langle q(t)q(0) \rangle$ (implicitly implying an average over the reference time $t' = 0$).

6.3 Dynamics of pairs of agents

In this chapter, we focus on the social dynamics that arise from pairwise interactions in three different conditions. First, we consider ≈ 11 h of experiments involving pairs of *H. rhodostomus*, to characterize and quantify the spontaneous social interactions when no artificial devices are present in the tank.

Secondly, we consider ≈ 16 h of effective trajectories for DLI simulated pairs (DLI-SP) from Chapter 4, as a baseline to the robot's underlying model in ideal conditions. This DLI model was originally trained in on a different series of experimental data obtained in Calovi et al. (2018) for the same species (*H. rhodostomus*), but in different conditions (different tank but of same radius $R = 25$ cm, lightning conditions...). We will also mention the results obtained after retraining the DLI model with the fish data considered in the present chapter's work, which we will refer to as the DLIv2-SP (see Appendix D and in particular Table D.1 and Figs. D.1, D.2, and D.3).

Finally, we conducted ≈ 11 h of experiments where the LureBot propels a biomimetic lure moving inside the circular arena, which is interacting in closed-loop with an actual *H. rhodostomus*. For brevity, in the following analysis of the results, we will simply refer to the LureBot and the lure attached to it as the LureBot. The LureBot is given a pre-trained copy of the DLI model of, which is queried in real time to generate biomimetic trajectories (see Section 6.2.2). We refer to these data as DLI biohybrid pairs (DLI-BP). We did not perform experiments with the LureBot trained with the DLIv2 model.

We explicitly designed a protocol which did not allow the use of the same fish in an experiment for at least 48 h after their first test, to avoid potential conditioning effects when the fish interact with the lure. The fish housing conditions and experiments have been approved by the local ethical committee and are described in detail in Chapter 5.

6.4 Results

This section reports the detailed comparison between the three test cases: 1) (fish-only) experiments with pairs of *H. rhodostomus*; 2) DLI simulated pairs (DLI-SP); and 3) DLI biohybrid pairs (DLI-BP), that consist of the LureBot interacting with a *H. rhodostomus*. In addition, at the end of this section, we will briefly present results for DLIv2 simulated pairs (DLIv2-SP).

The comparison between the different test cases exploits the observables described in Section 6.2.3. For all quantities (PDF and correlation functions), we have computed the statistical and sample to sample standard error by using a bootstrap method. In addition, for each PDF, we report the mean and standard deviation (SD) in Table D.1, as well as their standard error that we will omit to mention in the hereafter analysis of the results, for readability purposes (except when their value is relevant to the discussion). Moreover, in order to compare the PDF for a given quantity between two given test cases, we compute the Hellinger distance (see Ap-

pendix C). In general, a Hellinger distance $H(F|G) \lesssim 0.1$ points to a good agreement between both PDF, $0.1 \lesssim H(F|G) \lesssim 0.2$ points to a fair similarity between them, while $H(F|G) \gtrsim 0.2$ indicates that the two distributions are significantly dissimilar.

6.4.1 Instantaneous individual quantities

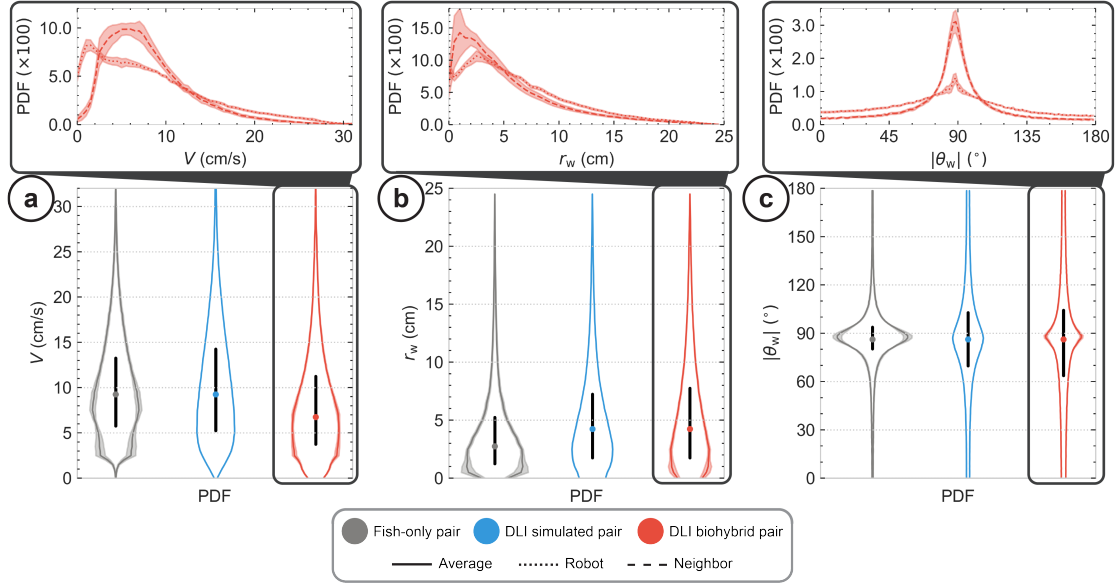


Figure 6.3: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the DLI simulated pairs, and the DLI biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the DLI biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

Fig. 6.3a shows the speed PDF for the three cases we considered. Fish pairs swim at a mean speed of 10.5 cm/s, associated to a standard deviation (SD) of 5.7 cm/s (see Table D.1). DLI simulated pairs produce a rather similar speed PDF (Hellinger distance $H = 0.09$; see Table D.2), albeit slightly wider (SD of 7.0 cm/s), with a nearly identical mean of 11.1 cm/s. For biohybrid pairs, the fish and the LureBot have a very similar mean speed (identical within error bars; see Table D.1), which is 20 % smaller than in the fish-only experiments, although the SD is similar to that of the fish experiments, resulting in a Hellinger distance of $H = 0.18$.

In Fig. 6.3b, we plot the PDF of the distance to the wall, r_w , for each case. Fish pairs swim very close to the wall, with a mean distance of 4.4 cm and a SD of 3.9 cm, both comparable to the fish typical body length (~ 3.5 cm). This is a consequence of the burst-and-coast swimming mode exhibited by *H. rhodostomus*, as shown in Calovi et al. (2018). Indeed, the motion of

this species is characterized by a succession of sudden acceleration periods (“kicks” or bursts of typical duration 0.1 s), each followed by a longer gliding period of typical duration 0.5 s, during which the fish moves in a quasi straight line. Because of the rather narrow distribution of heading changes between kicks, even observed when a fish is far from the wall Calovi et al. (2018), the fish is unable to escape the concave boundaries of the wall, except when rare large heading changes occur. The mean distance to the wall is 5.7 cm for the DLI simulated pairs, and the associated PDF compared to that for fish experiments is $H = 0.13$, showing that the DLI model captures reasonably well the tendency of the fish to move close to the wall. For biohybrid pairs, we found that the fish swims farther to the wall than in fish-only experiments, with a mean distance of 5.5 cm. In this case, the LureBot is even farther to the wall, at a mean distance of 6.6 cm, which likely also causes the fish to swim farther to the wall than in fish-only experiments.

Finally, in Fig. 6.3c, we plot the PDF of the absolute value of the heading angle relative to the normal to the wall, $|\theta_w|$. As a consequence of the agents (fish, DLI model, or LureBot) moving close to the wall, we naturally find that the mean of $|\theta_w|$ is very close to, but slightly below 90° (see Table D.1), a difference which is statistically significant. Indeed, as already reported in the experiments of Calovi et al. (2018), the agents spend slightly more time heading toward the wall ($|\theta_w| < 90^\circ$) than moving away from it ($|\theta_w| > 90^\circ$). The PDF for the three considered cases are symmetric around their mean, but we find that the fish experiments lead to the narrowest distribution, with a SD of 22° , compared to a SD of 35° for the DLI simulated pairs, and a SD of 33° and 43° for the fish and the LureBot in a biohybrid pair. The values of these SD are naturally correlated with the mean distance of the agent to the wall: the farther the agent, the larger are the fluctuations (SD) of its heading angle relative to the wall. Note that although the SD is larger for biohybrid pairs than for DLI simulated pairs (and for fish pairs), the intensity of the peak near $|\theta_w| = 90^\circ$ is larger for biohybrid pairs, which will have further consequences when we will address the temporal correlations of θ_w (see Section 6.4.3).

In summary, DLI-BP achieve fair agreement with the experimental results for all quantities. Concurrently, DLI-BP and DLI-SP show smaller dissimilarity, indicating that the transposition of the simulated model into the robot was successful leading a good overall precision (see Table D.2). We also observe that, in some cases (*e.g.*, $|\theta_w|$), the fish’s behavior guides the DLI-powered robot to approximate the experimental dynamics better, either due to unaccounted for dynamics in the model or due to an adapted fish behavior caused by the robot’s presence. Nonetheless, the observables test the DLI’s performance at a very fine level, especially in the case of DLI-BP, where the physical aspect is also impeding the precise reproduction of the social dynamics, either due to the imperfect (with respect to fish) motion of the robot or the varying degree the robotic system’s acceptance by the fish.

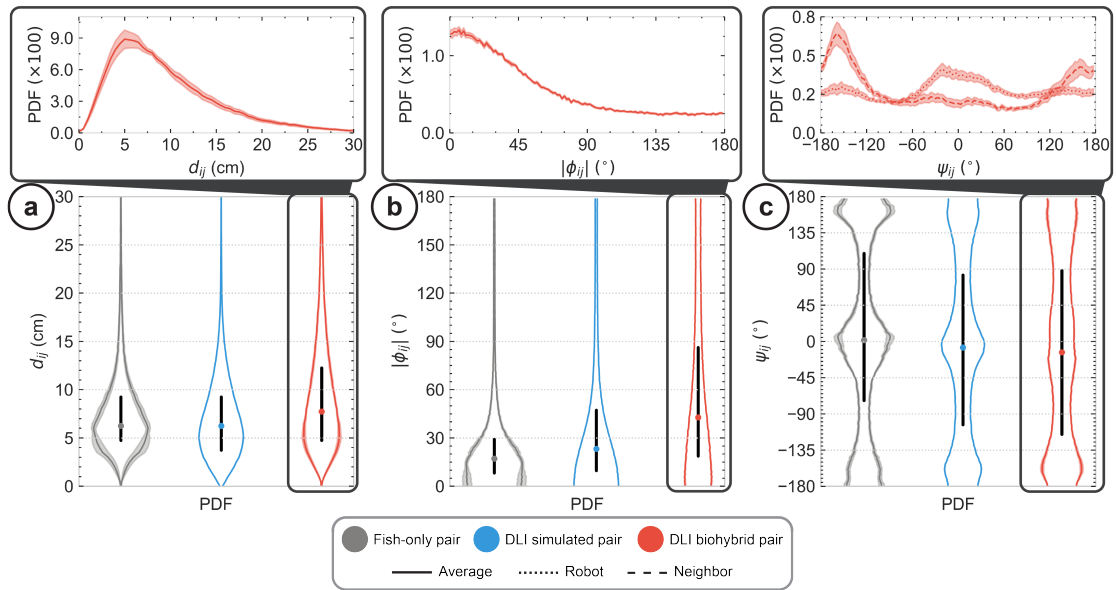


Figure 6.4: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles ϕ_{ij} probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, DLI simulated pairs and DLI biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the DLI biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

6.4.2 Instantaneous collective quantities

H. rhodostomus have a natural tendency to swim in close proximity to each other. In our experiments, fish pairs typically maintain a median interindividual distance d_{ij} of less than two body lengths (see Fig. 6.4a), with a mean distance of 8.05 ± 0.71 cm and a SD of 5.1 cm (see Table D.1). The dynamics of DLI simulated pairs results in a very similar PDF ($H = 0.16$), with a mean of 7.43 ± 0.03 cm, which is within one standard error (for the fish experiments) from the mean obtained for fish. As for the biohybrid pair, it is less bound than pairs of fish or DLI, with a mean distance between the fish and the LureBot of 9.96 ± 0.48 cm. The distribution is also slightly wider, with a SD of 6.3 cm. In fact, although the peak of the interindividual distance PDF is located at a similar value as for fish or DLI pairs (5 – 6 cm in the three cases), the biohybrid pairs are more often separated by a distance larger than 15 cm.

H. rhodostomus is a social species, often found to form well aligned schools. In fact, their pairwise alignment interaction was quantitatively measured in Calovi et al. (2018), showing that this interaction remains strong up to three body lengths, well within the typical distance between fish. In Fig. 6.4b, to quantify the alignment within pairs of agents, we plot the distribution of the absolute value of the difference between the heading angles of the two

agents, $|\phi_{ij}|$ (see the graphical definition in Fig. 6.2). The mean heading difference observed in fish experiments is 27° , with a rather narrow PDF associated with a SD of 30° , confirming the good level of alignment between the two fish. The DLI simulated pairs are not as aligned as fish pairs, with a larger mean and SD equal to 38° , although the Hellinger distance between the two PDF ($H = 0.14$) remains satisfactory. The corresponding PDF for biohybrid pairs exhibits the largest disagreement with the fish experiments of all the PDF presented here ($H = 0.30$). Indeed, despite also being peaked at $|\phi_{ij}| = 0$, the PDF has a non-negligible weight for $|\phi_{ij}| > 90^\circ$, resulting in a much larger mean of 59° and a SD of 48° . This wider PDF is a consequence of the fact that the fish and the LureBot, despite remaining close to each other on average, have a much higher probability than fish pairs to be at a distance above the range of the alignment interaction. Moreover, when the fish and the LureBot are far apart and attempt to get closer, they have a high chance to be actually anti-aligned during this process, hence the significant weight of the PDF near $|\phi_{ij}| = 180^\circ$.

Finally, Fig. 6.4c shows the PDF of the angle of perception ψ_{ij} , defined in Fig. 6.2. For pairs of fish, the PDF presents clear peaks at $\psi_{ij} = 0^\circ$ and near $|\psi_{ij}| = 180^\circ$. This indicates that the well aligned fish are following each other rather than swimming side by side. For DLI simulated pairs, the same pattern is observed but with slightly less pronounced peaks, although the Hellinger distance of $H = 0.05$ confirms the excellent agreement between both PDF. As for the biohybrid pair, the PDF averaged over the fish and the LureBot again presents the same peaks as before, but even less pronounced. Again, the less sharp peaks are a consequence of the fact that the biohybrid pairs stand farther from the wall than fish pairs, and above all, of the fact that their distance has a higher probability to be large enough so that their angle of perception ψ_{ij} becomes uncorrelated. The lesser alignment of the biohybrid pairs (see above) originates from the same causes, and in turn also results in a more homogeneous distribution of the angle of perception. However, the apparent reasonable agreement with the PDF for the fish-only and DLI-SP pairs masks the difference between the PDF for the fish and for the LureBot shown in the top inset of Fig. 6.4c. There, we observe that the peak near $\psi_{ij} = 0^\circ$ is dominated by the contribution of the fish, showing that the fish more often follows the LureBot than the converse. In addition, we find that the PDF for the fish is also peaked slightly above $\psi_{ij} = -180^\circ$, while the PDF for the LureBot has a corresponding peak slightly below $\psi_{ij} = +180^\circ$. By periodicity of 360° , these two peaks are obviously located at almost the same angle, but this slight angular shift translates to the fact that the fish is, on average, slightly closer to the wall than the LureBot, as noted in Section 6.4.1.

The instantaneous collective quantities demonstrate that despite the dissimilarities measured in the individual behavior of both DLI-SP and DLI-BP with respect to the fish-only experiment, the collective dynamics are fairly reproduced. Furthermore, the DLI is transferred in a physical system with good agreement compared to its simulated version, and the living agent responds positively. However, the angular control of the robot is arguably less precise, which contributes to the general deviation from the experimental angle-related distributions.

6.4.3 Temporal correlation quantities

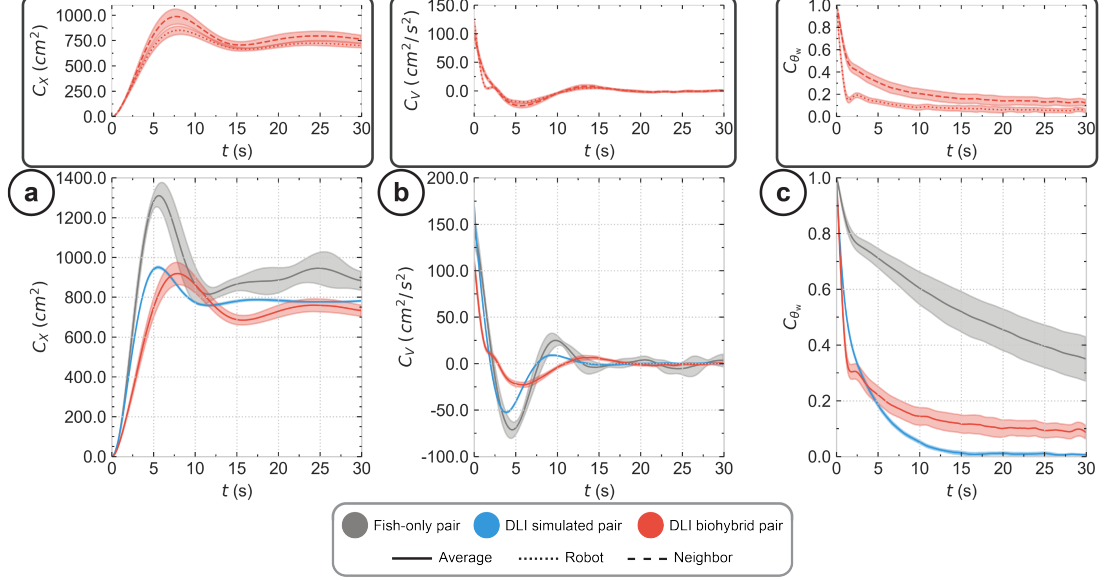


Figure 6.5: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, DLI simulated pairs and DLI biohybrid pairs, respectively. Dotted, dashed and solid lines indicate the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

In Fig. 6.5, we plot the three observables used to quantify the temporal correlations that emerge in the system during the long-term dynamics, which are defined in Section 6.2.3.

Fig. 6.5a shows the mean square displacement of the agents, $C_X(t)$, in the three considered cases. After a rapid growth, $C_X(t)$ presents a peak and an ultimate decay to a mean level equal to twice the mean square of the distance to the center of the tank. Indeed, for large time difference, the positions at time t' and $t + t'$ become uncorrelated, and we obtain

$$\begin{aligned}
 C_X(t) &= \left\langle [\vec{u}_i(t+t') - \vec{u}_i(t')]^2 \right\rangle \\
 &\underset{t \rightarrow +\infty}{\approx} \langle \vec{u}_i^2(t+t') \rangle + \langle \vec{u}_i^2(t') \rangle \\
 &= 2\langle \vec{u}_i^2(t') \rangle,
 \end{aligned} \tag{6.8}$$

which becomes time-independent due to the stationarity of the dynamics. Although $C_X(t)$ has the same qualitative form in the three cases, one observes differences in the position and height of the peak and in the asymptotic value. The latter is explained by the fact that the closer the agents are to the wall, the larger is the mean square of their distance to the center of the tank, $\langle \vec{u}_i^2(t') \rangle$. Indeed, we have found, in Section 6.4.1, that fish pairs swim closest to the wall, while biohybrid pairs are the farthest, which is consistent with the asymptotic behavior of $C_X(t)$ observed in Fig. 6.5a. Furthermore, the top inset of Fig. 6.5a for the biohybrid

pairs shows that $C_X(t)$ for the fish is systematically larger than for the LureBot, which is also consistent with the fact that the fish swims slightly closer to the wall than the LureBot. As for the position of the peaks in Fig. 6.5a, it roughly corresponds to the time for the corresponding agent to travel half of the tank perimeter. This time is directly correlated with the mean speed of the agent. In Section 6.4.1, we found that the fish pairs and DLI simulated pairs had essentially the same mean speed, which explains the agreement between the position of the corresponding peaks in $C_X(t)$. However, we also found that the biohybrid pairs were 20 % slower, which explains the fact that the peak in their $C_X(t)$ is reached at a later time than for fish and DLI pairs.

Fig. 6.5b shows the velocity autocorrelation, $C_V(t)$, in the three considered cases, which vanishes for t large enough, when the velocity at time $t + t'$ becomes uncorrelated with that at time t' . It can be formally shown that $C_V(t) = d^2 C_X(t) / dt^2$ (although this relation is only approximate, when the 2 quantities are observed independently over a finite sampling time), so that the interpretation of the shape of $C_V(t)$ results from the analysis that we have presented above for $C_X(t)$. In particular, the peaks of the first two oscillations in $C_V(t)$ roughly correspond to the two inflection points just before and after the main peak in $C_X(t)$. In addition, $C_V(t = 0)$ is the mean square velocity, and we indeed observe an agreement between its value for fish and DLI pairs, while the slower biohybrid pairs result in a lower initial value of $C_V(t = 0)$ in this case.

Finally, the (most subtle) temporal correlation function of the heading of an agent relative to the wall, $C_{\theta_w}(t) = \langle \cos [\theta_w^i(t + t') - \theta_w^i(t')] \rangle$, is shown in Fig. 6.5c. For very large time t , $C_{\theta_w}(t)$ must obviously decay, but we observe that for fish pairs, we still have $C_{\theta_w}(t = 30 \text{ s}) \approx 0.35$, indicating strong correlations. For DLI simulated pairs, we find that $C_{\theta_w}(t)$ vanishes very rapidly ($C_{\theta_w}(t = 15 \text{ s}) \approx 0$). Finally, for biohybrid pairs, we still observe some weak remnant correlations at $t = 30 \text{ s}$, with $C_{\theta_w}(t = 30 \text{ s}) \approx 0.1$ (although the correlation is dominated by the contribution of the fish, as shown in the top inset of Fig. 6.5c). Here, the decay rate of $C_{\theta_w}(t)$ is strongly related to the sharpness of the peak near $\theta_w = 90^\circ$ in the PDF of θ_w (see Fig. 6.3c and Section 6.4.1). Indeed, a sharp peak suggests that it can take a long time to explore values of θ_w far from 90° , leading to a slower decay of $C_{\theta_w}(t)$. Accordingly, we indeed found that the least sharp peak in the PDF of θ_w is observed for DLI simulated pairs, resulting in the fastest decay of $C_{\theta_w}(t)$ in this case.

Both the DLI-SP and DLI-BP fail to precisely reproduce the correlation function $C_{\theta_w}(t)$, producing a very similar sharp decay compared to the one of real fish. This is derived by the DLI's tendency to frequently produce trajectories farther from the wall than what observed in the experiment. Despite an overall dissimilarity between the experiment and the DLI derivatives, the DLI-BP remains fairly faithful to the DLI-SP. Yet again, this indicates that the DLI is missing some aspects of the social dynamics before being deployed on the robot, but the robot performs reasonably well in reproducing its underlying model.

6.4.4 Complementary results for DLiv2 simulated pairs

In addition to the Deep Learning Interaction (DLI) pretrained network utilized in the previous sections, we have also considered an updated version, the DLiv2. This version was retrained on data gathered from the present fish-only experiments under new lighting conditions, concurrently to the robot experiments presented in this chapter, so that retraining was only feasible after their completion. However, it provided us with the opportunity to test the scalability and predictive performance of the pretrained DLI with new input samples, which, while not fundamentally different, originated from altered social dynamics. For this purpose, we conducted extensive simulations with the DLiv2, and found that their results are in excellent agreement with the present fish-only pair experiments (see Tables D.1 and D.2 for further details) for the individual (see Fig. D.1) and collective (see Fig. D.2) observables, and for the temporal correlation functions (see Fig. D.3). The performance of the simulated DLiv2 model present a significant improvement compared to that of the pretrained DLI model, and one could expect that the LureBot commanded by the DLiv2 model would lead to better results than for the LureBot commanded by the pretrained DLI model. Yet, our point here is that the pretrained DLI model, in different experimental conditions, can still interact with a fish in a similar way as a fish would do.

6.5 Conclusion

Despite the abundance of studies on fish-robot interactions, to our knowledge, no prior research has drawn comparisons between the social interaction dynamics produced by fish-only, biohybrid, and simulated groups. This comparison also introduces an intriguing issue: while the reality gap in robotics Jakobi et al. (1995); Jakobi (1997); Mouret and Chatzilygeroudis (2017) typically pertains to the transferability of robot controllers from simulation to real-life conditions, a parallel can be drawn for biohybrid social interactions, termed the biomimicry gap. Addressing this gap is complicated by: 1) subtle behavioral patterns that researchers or machines fail to consider during the modelling process, 2) the inherently imperfect biomimetic properties of artificial lures and devices, and 3) the absence of physics in most models. Constructing biohybrid systems with minimal or, ideally, no biomimicry gap, thus making them indistinguishable from pure animal groups, could open doors to groundbreaking research in biological systems. This would enable, for example, the introduction of controlled, localized perturbations to accurately gauge an animal's reaction. Such endeavors require that any non-biomimetic effects of the robot be stringently assessed and alleviated. It is therefore crucial to ensure that models do not simply overfit experimental data, but genuinely translate to real-world scenarios, through robotic systems that can reproduce these models as expected.

Alas, while substantial progress has been made in the intersection of behavioral modeling and robotics hardware which has long been promoted as the key to unraveling and understanding the mechanisms that underlie collective behavior in animal groups, this assertion has not been convincingly demonstrated in the literature. Especially in terms of models, most biohybrid

systems are designed to modulate collective decisions between a limited set of spatial choices (which often resolves to a binary choice). Many of these systems rely on simplified passive (**open-loop**) Phamduy et al. (2014); Polverino and Porfiri (2013b); Abaid et al. (2012a); Ladu et al. (2015b); Spinello et al. (2013); Ruberto et al. (2016); Bartolini et al. (2016); Kruusmaa et al. (2016); Butail et al. (2014b) or reactive (**closed-loop**) Cazenille et al. (2018c,a); Faria et al. (2010); Swain et al. (2011); Landgraf et al. (2013, 2016); Porfiri et al. (2019); Kim et al. (2018) models, with only a handful utilizing biomimetic models. Even fewer biomimetic models have been successfully tested in biohybrid groups Cazenille et al. (2018c,a) to emulate real-life dynamics of fish groups. Simultaneously, to our knowledge, no end-to-end machine learning (ML) model has been examined in this context, despite the booming field of ML. Assessing a model's fidelity is particularly challenging in the case of ML models, which are often black-box (*i.e.*, not easily explainable). Moreover, our understanding of how well robotic devices can integrate and interact with living animal groups primarily relies on metrics that do not account for model limitations. Similarly, many models evaluated in simulations, rely on metrics that do not account for how such models can scale to real life Cazenille et al. (2015, 2017, 2018a). Finally, the role of biomimetic lures and agile robotic devices in such studies is usually underplayed. That is, there is insufficient evidence to support the claim that biomimetic systems can bridge the biomimicry gap.

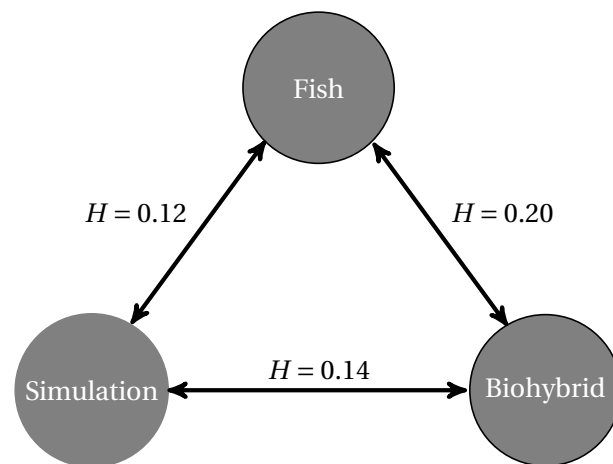


Figure 6.6: Quantification of biomimicry gap (root mean squared Hellinger distance of all observables) values between different experiments. Fish data compared to simulations yield a root mean squared Hellinger distance value of $H = 0.12$; simulations compared to biohybrid groups yield a mean squared distance $H = 0.14$; and Biohybrid groups have a mean squared distance $H = 0.20$ when compared to fish-only data.

In this chapter, through the precise and comparative quantification of collective behavior in pairs of agents (fish-only pairs, DLI simulated pairs and DLI biohybrid pairs), we demonstrate that our biomimetic lure and robot system (introduced in Chapter 5), combined with the DLL, are capable of bridging a substantial part of the biomimicry gap. More specifically, our study reveals that the overall gap between actual pairs of *H. rhodostomus* and simulation is smaller (root mean squared Hellinger distance of 0.12 of all observables) than that of the simulation

and the biohybrid experiment (root mean squared Hellinger distance of 0.14 of all observables), or the biohybrid pairs compared to fish-only pairs (root mean squared Hellinger distance of 0.20 of all observables), a direct result of the sum of the discrepancies indicated. In essence, our DLI model is successful in generating realistic social interactions (see also Chapter 4), our robotic system is very capable of replicating its instructions with small discrepancies, but the transferred model results in greater discrepancies and the gap widens compared to the simulation (see Fig. 6.6). Nonetheless, the biohybrid pair is not well aligned compared to fish groups, which is confirmed when we compute the largest Hellinger distance value out of the quantities considered (Hellinger distance between angle of incidence PDF is $H = 0.30$), and is the largest contributor to widening the social interaction discrepancies (*i.e.*, the largest contributor, out of all observables, to increasing the Hellinger distance). These discrepancies are consequently observed for the correlation function of the angle of incidence to the wall. It is worth noting that even in fish-only experiment comparisons, the Hellinger distance will probably exceed zero (albeit it is expected to be close to it), owing to the inherent behavioral variability fish display between experiments. This suggests that the primary goal for robotic systems might be to significantly reduce the gap, rather than entirely eliminating it.

Therefore, despite the positive results highlighted in our study, we demonstrated that further closing the biomimicry gap requires work on minimizing all three discrepancy sources (depicted in Fig. 1.1). First, it requires that we refine our modelling approaches (*e.g.*, by repeating the experiments with the DLIv2). Second, the physics-related discrepancies, primarily attributable to the transposition of the model into the robot, remains relatively small, but also requires measurable improvement in the robotic system's operation to fully bridge the gap. Finally, discrepancies in the communication cues pose a considerable challenge in terms of evaluation and could only be fully measured in the absence of the other two sources of discrepancies. These discrepancies might correlate with the way social interactions are faithfully modeled, under the assumption that the real fish's responses change due to the artificial lure, potentially influencing the interaction dynamics it is exhibiting in unexpected ways (*i.e.*, not frequently or at all observed in spontaneous interactions). Whereas it remains indeterminate as to what constitutes the ultimate goal in reducing the biomimicry gap, our study introduces a methodology and the resulting baseline values that produced realistic interactions. Furthermore, the correlation between improving the Hellinger distance and reducing the gap is not yet known, but future studies can leverage the values presented in this study as a baseline comparison score. To that end, along with the between-sample noise that we measured with our bootstrap approach, future work could include independently repeating the entirety of the fish-only experiments to provide insight into the behavioral variability of *H. rhodostomus*. In turn, this could provide better understanding into a sensible lower bound for the average Hellinger distance value, that indicates excellent agreement between any two datasets.

We believe our study may be the first of many that emphasizes on both simulated and biomimetic biohybrid experiments in a single end-to-end approach. An approach that, as demonstrated, allows us to gain more insight into which aspects led to behavioral differences in the biohybrid experiments and, therefore, need improvement. In turn, this contributes to: 1)

establishing a better experimentation pipeline to investigate the varying sources of gaps (*e.g.*, physical robot limitations, social interaction gaps in models, and eventually discrepancies in the communication cues exploited to evoke responses), and 2) drawing more definitive and insightful behavioral conclusions without the unrealistic effects introduced by the robotic system and social interaction models. Future studies could benefit by the advances presented in this chapter, by reporting the biomimicry gap score (in the form of the quantities presented). We hope that future studies may standardize and utilize a common methodology to evaluate the fidelity of biohybrid systems with respect to natural ones.

7 Investigating the biomimicry gap with varying sizes of biohybrid groups driven by a high-fidelity model

This section's content is adapted from the following article:

- This is partial and preliminary work that will undergo more analysis and will be included in a future publication. The results of this work have been generated in collaboration with Guy Theraulaz, Clément Sire, and Francesco Mondada.
 - Vaios Papaspyros's contribution: Conception, data curation, formal analysis, investigation, methodology, software, visualization, writing..
- The article's work was funded by the Swiss National Science Foundation project "Self-Adaptive Mixed Societies of Animals and Robots" (Grant No. 175731). Our collaborators from the Université Toulouse – Paul Sabatier were funded by the French National Research Agency (ANR-20-CE45-0006-01).
- Ethics: All *H. rhodostomus* experiments were conducted at the Centre de Recherches sur la Cognition Animale, Centre de Biologie Intégrative, CNRS, Université de Toulouse – Paul Sabatier. Experiments were approved by the local ethical committee for experimental animals and were performed in an approved fish facility (A3155501) under permit APAFIS#27303-2020090219529069 v8 in agreement with the French legislation.

7.1 Introduction

As discussed in previous chapters, fish schooling behavior is a captivating example of collective intelligence and emergent patterns in nature. Uncovering the underlying mechanisms governing these behaviors could provide valuable insights into various applications, including robotics, environmental monitoring, and even human social dynamics. Yet, despite a plethora of biohybrid studies, only a handful rely on the real time use of highly biomimetic models (as presented in Chapter 6), and, to our knowledge, none provide insight into the biomimicry gap between underlying models that control robots and the dynamics created in biohybrid experiments.

In Chapter 6, we showed that a deep learning model of social interactions can be transposed back to real life, albeit imperfectly. More importantly, our methodology introduced measurable observables to quantify the distance from its simulated counterpart, and provided insight as to which biomimicry gap components (see Fig. 1.1) contribute to discrepancies. In this chapter, we conduct a similar study, where we scale this experiment to varying sizes of biohybrid groups (of one, two, and five individuals), and we additionally investigate, for the first time in a biohybrid system, the role of multiagent perception in fish groups Lei et al. (2020) on reducing the social interaction gap (see Fig. 1.1). To do so, we employ our robotic system (see Chapter 5), wherein a robot propels a biomimetic lure realistically, and is designed to interact with a varying number of live fish. More specifically, we conduct an ablation study on how the analytical burst-and-coast model (introduced in Calovi et al. (2018) and presented in Chapter 4), can be transposed back to real life when: 1) the LureBot interacts with the wall of the tank alone, 2) the LureBot interacts with one *H. rhodostomus*, and 3) the LureBot interacts with four *H. rhodostomus*. This process affords us the opportunity to partially unravel the structure of social dynamics within biohybrid groups, and facilitates an evaluation of the model's and lure's abilities to reproduce the natural social interactions inherent to the group. Concurrently, it sheds light on two of the three facets of the biomimicry gap defined in Chapter 1 (depicted in Fig. 1.1). Specifically, we probe: 1) whether any subtle behavioral patterns elude the behavioral model steering the robot (social interaction gap), and 2) if part of this gap can be attributed to the robot's physical properties (physics gap).

Intriguingly, the novel biohybrid system introduced in Chapter 5, when coupled with the compelling evidence garnered in Chapter 6, elucidates that pairing with a high-fidelity social interaction model can markedly diminish the biomimicry gap. This, in turn, permits the derivation of more robust, minimally biased and more confident conclusions regarding the inherent behaviors of fish. Oftentimes, artificial systems have illuminated biological insights in non-biomimetic social interaction settings. Examples include artificial stimuli that span from phenotypically divergent lures Romano and Stefanini (2021) and physiologically distinct ones Bartolini et al. (2016); Polverino and Porfiri (2013a); Kruusmaa et al. (2016), to lures mimicking predators Abaid et al. (2013), and other visual stimuli Abaid et al. (2012b). Conversely, as contended in prior chapters, we assert that deriving confident biological conclusions relies on our comprehension and mitigation of the biomimicry gap. Specifically, it relies on biohybrid systems capable of exhibiting profound similarity to biological groups and minimal discrepancies originating from the artificial instruments. Consequently, studying the spontaneous responses of living individuals to deliberately targeted robotic stimuli is afforded augmented confidence concerning biological insight.

To that end, in this chapter, we explore a biological question, newly possible to validate in reality through our biohybrid system, concerning the impact of multiagent stimuli on a focal individual and the resultant emerging effects on collective-level dynamics. Although the effects of multiagent stimuli have been previously probed in the context of fish social interaction within simulations and robot-only swarms Lei et al. (2020), to our knowledge, no similar experiment has ever been conducted utilizing a biohybrid system. Consequently, in

subsequent sections, we delve into this concept, enabling the focal individual (*i.e.*, the artificial lure) to perceive: 1) the single most influential neighbor (*i.e.*, with the highest social influence force, see Calovi et al. (2018) and Chapter 4), and 2) the two most influential neighbors.

7.2 Methods

7.2.1 Multiagent simulations with the analytical burst-and-coast model

The analytical burst-and-coast (ABC) model simulations are initially configured as indicated in Chapter 4, regardless of the number of individuals tested. The parameters of the model are, however, adapted on a group-size basis as follows:

1. We conduct a series of simulations until the generated dynamics are closely fitting the experimental data, and in order to provide a good initial parameter set for the biohybrid system.
2. These parameters are then tested with the real-time system and are further hand-tuned to reduce notable discrepancies introduced by the biohybrid system's hardware and software components.
3. Finally, the hand-picked parameters are used in long simulations that are compared to the biohybrid results. Although using the same parameters prioritizes good performance for the biohybrid system, it may lead to reduced performance in simulation. However, this choice allows for a direct quantification of the biomimicry gap, one of the primary objectives of this study.

The model parameters affected during this process typically correspond to the fish-wall (γ_w) and fish-fish (γ_{Att}) attraction strength, as well as the fish-fish alignment (γ_{Ali}) preference. Furthermore, in the real-time system, we reduce the average kick duration from 0.52 s to 0.2 s to account for the robot's inherently slower response time compared to that of fish. The latter adaptation is not applied to simulations, as there are no mechanical aspects that require for this compensatory choice. As per the model's definition (see Calovi et al. (2018) and the corresponding section in Chapter 4), during a focal individual's decision-making process, we compute the social influence of each individual on the focal, as an effect of the latter's heading direction for its next motion. The effect is defined as the sum $\delta\phi_{Att}(d, \psi, \Delta\phi) + \delta\phi_{Ali}(d, \psi, \Delta\phi)$, that is, the effect of the fish-fish attraction and alignment strength, respectively. Typically, for fish groups of $N > 2$ individuals, the default definition of the model sums the resulting influence of all focal-neighbor pairs, before applying the cumulative effect on the focal individual's heading direction.

As previously mentioned, investigating a focal individual's behavior with limited perceptive abilities, requires that instead of considering the cumulative effect of all neighbors, we only

consider a subgroup subject to a selection criterion. In the following sections, this criterion refers to the most influential neighbor, and corresponds to the individual(s) with the largest value $\delta\phi_{\text{Att}}(d, \psi, \Delta\phi) + \delta\phi_{\text{Ali}}(d, \psi, \Delta\phi)$. In practice, during the decision-making process for the robotic system, we first sort the individuals in descending order of their influence, and test two cases where the focal (robotic) individual considers the top $k = 1$ and $k = 2$ neighbors, similarly to Lei et al. (2020).

7.2.2 Real-time robot control with the analytical burst-and-coast model

To generate biomimetic trajectories for the robot to follow, we use the analytical burst-and-coast model implementation described in the previous subsection and Chapter 4. Notably, we alter the model such as to generate predictions for time $t + 1$ in continuous time, as opposed to the (discrete) kick-based logic of the model. At a first level, the ABC remains the same, that is, we generate kicks (*i.e.*, peak velocity, duration...), then we sample the kick every with a 30 Hz rate to obtain and send the desired motor commands to the robot. In fact, we make use of the a-priori PID as defined in Chapter 5, where we tuned a PID controller to act on top of an a-priori known average velocity. We found that this controller was able to reproduce the velocity profile better than the traditional PID. Once the kick duration has elapsed, we compute a new kick and repeat the above process. We note that, the robot itself does not truly perform a kick. Whereas the PID controller produces kick-like speed profiles, those are, in fact, not highly faithful to the ones observed by fish. The latter is indeed within the LureBot's capabilities, but at the time we conducted the experiments the low-level controller that can perform this motion was not available.

Fig. 7.1 summarizes the closed-loop system we employed for the experiments presented in following sections. First, our cameras capture the latest frame (see Fig. 7.1a). Then, once the frame is available, BOBI (see Chapter 5) extracts the position, velocity, and direction of each individual (see Fig. 7.1b), information which is used to compute individual and collective quantities. This information, is then fed to the analytical burst-and-coast model which stochastically computes a kick to be performed (see Fig. 7.1c). As discussed already, in order to control the robot, we convert this discrete time information to continuous time segments by sampling the kick (see Fig. 7.1d bottom) every $\Delta t = 1/3$ s to produce its corresponding speed transitions. These speed transitions combined with the desired heading direction are used to command the robot at a rate of 30 Hz until $t = t_i^{n+1}$. The robot then magnetically propels a biomimetic lure that is placed within the tank. For simplicity and brevity, we refer to the lure and the LureBot as the robot.

Notably, for group sizes $N > 2$, we tested the two perception cases where the most influential neighbor is the one whose effect will contribute the largest heading variation, relative to the focal individual (as described in the previous subsection). Additionally, the robotic agent's desired speed is selected in a probabilistic manner to either: 1) retain its previous speed, 2) randomly draw it from the speed distribution as measured in groups of fish that interact

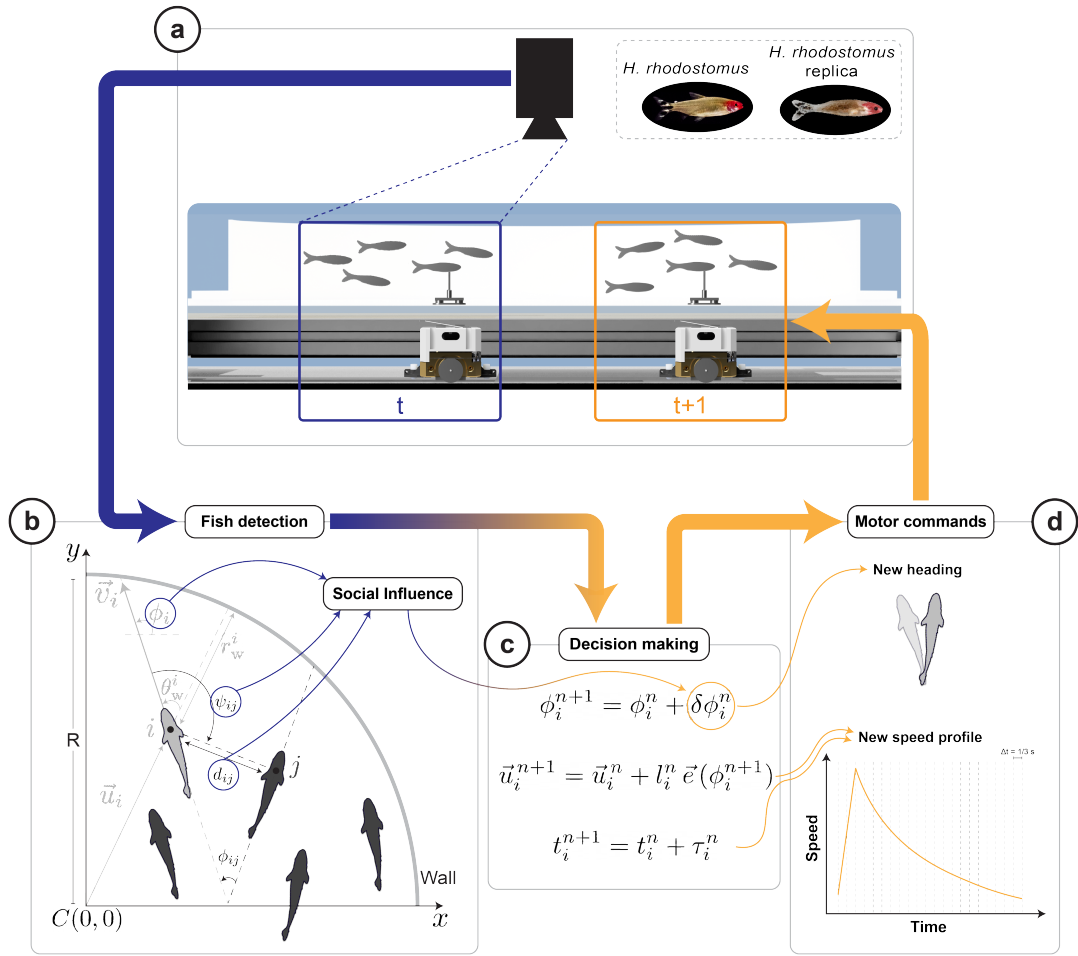


Figure 7.1: **Closed-loop analytical burst-and-coast control of the robot.** (a) the camera system captures a frame at time t , (b) the frame is used to extract the position, velocity, and direction of all individuals, (c) the resulting information is used by the analytical burst-and-coast (ABC) model to generate trajectory predictions, (d) the trajectory is adapted and sampled for continuous time control of the robot, and motor commands are sent to the robot.

spontaneously, or 3) copy the average speed of its neighbors. This probability, along the kick generation parameters, may vary depending on the number of fish in the tank, and is empirically set for the different group sizes. All experiments are conducted in the same experimental setup presented in Chapter 5, that is, the fish swim in a circular arena of radius $R = 25$ cm.

7.2.3 Data processing

We follow a similar data treatment as in Chapter 5. First, we employ idtracker.ai (v4) Romero-Ferrero et al. (2019) in conjunction with high-resolution recordings (1500×1500 pixels) from the experiment to monitor all agent movements and extract their 2D trajectories offline. The

software reports an average tracking success rate exceeding 99.5%, which is further validated through manual checks. We also incorporate a post-processing step that rectifies any instances of misclassified or missing agent identities (*e.g.*, when idtracker.ai fails to identify all individuals in a particular video frame). Subsequent analyses are centered on the lure's interaction with the fish, hence, lengthy periods with minimal fish movement are excluded. Considering the average body length of *H. rhodostomus* is 35 mm, we discard segments where the fish or lure's speed is under 1 BL/s. This procedure also eliminates occasions where the magnetic linkage between the robot and lure breaks, leaving the lure stationary. Ultimately, all trajectories undergo resampling with a timestep of $\Delta t = 0.1$ s, replacing the original 1/30 s. This adjusted timestep is deliberately chosen to mitigate the random noise arising between consecutive frames and minimize the dataset's dimensionality, while preserving enough detail to scrutinize the interactions between the fish and lure.

7.2.4 Experimental procedure

The *Hemigrammus rhodostomus* (rummy-nose tetras) used in our experiments were acquired from Amazonie Labège located in Toulouse, France. The fish were accommodated in 16 L tanks, maintained at a temperature of 27°C, and subjected to a 12/12 hour dark/light photoperiod with a 30-minute transition period. Fish were fed freely with fish flakes. A trained technician ensured proper feeding and living conditions each day between 8:30 am and 9:30 am. The average length of the fish used in our studies is 35 mm.

The subsequent behavioral experiments were carried out in a circular enclosure with a radius of $R = 25$ cm (the same setup from Chapters 5 and 6). The water level inside the enclosure was kept around 5 cm. The water used in the enclosure came from the same filtering system as the fish tanks, thus ensuring similar salinity and conductivity conditions. The temperature was maintained at an average of 27°C, similar to the fish tanks. The experimental setup was housed in a separate room from the fish tanks. All experiments were conducted during fully lit periods, between 9:00 am and 20:30 pm, with fed fish.

Our experimental process is summarized as follows:

1. Fish were randomly chosen from a tank not used on the previous day and transferred to the circular arena;
2. we then allowed the fish a 15-minute acclimation period. Fish showing significant stress (*e.g.*, remaining motionless or floating belly-up) were returned to their tank, and another fish was selected at random;
3. each fish or fish group was given an hour (1 h) to interact with the robot, after which they were returned to their respective tank;
4. This process is repeated until we collect approximately 11 h of data for each case.

This procedure ensured that no fish was subjected to tests two days consecutively or twice in a single day. This precaution helped prevent the fish from becoming accustomed to the lure or any specific patterns it displayed.

7.2.5 Evaluating the short- and long-term interactions

To evaluate the dynamics produced by fish-only, simulated and biohybrid groups, we follow the same procedure introduced in Chapter 6 and Jayles et al. (2020). That is, we exploit the 9 observables considered in Chapter 4 (see Fig. 6.2), with adaptation for groups $N > 2$, as described below.

The first 3 observables correspond to instantaneous quantities at the *individual* level, for which we measure their probability density function (PDF): the speed, V of an individual; its distance to the wall, r_w ; and its heading angle relative to the normal to the wall, θ_w . 3 additional observables probe the instantaneous collective dynamics, although as denoted below, those 3 observables differ for the case of $N = 2$:

- the distance d_{ij} between the pair of individuals (or between nearest-neighbor pairs in a group of more than 2 individuals);
- the difference ϕ_{ij} between the heading directions of the two individuals;
- and the angle ψ_{ij} at which an individual perceives its neighbor,

and for $N = 5$:

- cohesion, *i.e.*, the average of the average distance between pairs of individuals

$$P(C_i) = \sqrt{\frac{1}{N-1} \sum_i^N (\vec{r}_i - \vec{r}_j)^2}; \quad (7.1)$$

- the average of the average nearest neighbor distances for each fish in the group $P(d_i)$;
- and the average polarization of the group $P(P_i)$

$$P(P_i) = \frac{1}{N-1} \sum_i^N (\cos(\phi_i - \phi_j)), \quad (7.2)$$

where $i \in [1, N]$ and we plot the quantities $P(\cdot)$ two times, 1) for all $i \in [1, N]$, and a second for all $i \in [1, N]$ with $i \neq \text{ROBOT_ID}$.

Finally, we consider 3 *temporal correlation functions* that probe the social dynamics at a very fine level Jayles et al. (2020), and which are generally particularly difficult to reproduce. C_X is

the mean-squared displacement, C_V the velocity autocorrelation, and C_{θ_w} the autocorrelation of the angle of incidence to the wall. In general, we denote $C_q(t) = \langle q(t+t')q(t') \rangle$ as the average of the quantity $q(t)q(t+t')$ over the reference times t' , over individuals, and over different experiments. By the assumption of stationarity of the system, the temporal correlation function $C_q(t)$ only depends on the time difference between observations, and is often noted $C_q(t) = \langle q(t)q(0) \rangle$ (implicitly implying an average over the reference time 0).

For all quantities (PDF and correlation functions), we have computed the statistical and sample to sample standard error by using a bootstrap method. Additionally, we measure the Hellinger distance between PDFs to quantify the (dis)similarity between the different quantities and cases (see Appendix C).

7.3 Results

7.3.1 Experiments with a single agent

Instantaneous individual observables

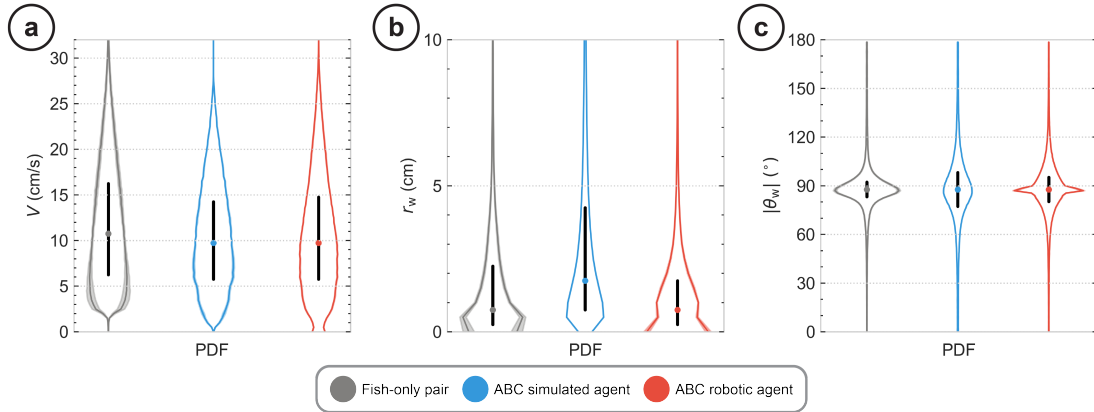


Figure 7.2: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Absolute angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the ABC simulated agents, and the ABC robotic agents, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The shaded areas correspond to the standard deviation.

In Fig. 7.2a we plot the three cases considered in this chapter. Single fish typically swim at a mean speed of 12.16 cm/s, with a standard deviation (SD) of 6.68 cm/s (see Table E.1). ABC simulated agents produce a fairly similar distribution of velocity (Hellinger distance $H = 0.20$; see Table E.2), although narrower (SD of 5.77 cm/s) and with a smaller mean of 10.96 cm/s. The robot, exhibits a very similar means speed with ABC-SP ($H = 0.13$) and are in very good agreement with single fish ($H = 0.18$).

Fig. 7.2b shows the PDF of the distance to the wall, r_w , for the three cases. Firstly, fish agents swim very close to the wall with a mean distance of 2.03 cm and SD of 2.49 cm, less than one body length (≈ 3.5 cm). Similarly to what was discussed in Chapter 6, this is a consequence of the burst-and-coast swimming mode of *H. rhodostomus* Calovi et al. (2018). Due to this movement, the fish is unable to escape the concave boundaries of the wall. ABC simulated agents exhibit a greater distance to the wall, with a mean value 3.72 cm and an SD of 3.99 cm, which differs from that of living fish with $H = 0.31$. Conversely, the robot typically produces smaller distances to the wall than those of fish-only and the simulated agent. More specifically, the robot swims at a mean distance of 1.73 cm with an SD of 2.04 cm from the wall, producing similar distributions compared to fish with $H = 0.15$. Notably, the robot behavior is rather different when compared to simulated agents, with their in-between difference measured at $H = 0.39$.

Finally, Fig. 7.2c shows the PDF of the absolute value of the heading angle relative to the normal to the wall, $|\theta_w|$. As described in Chapters 4 and 6, due to the swimming proximity of the agents (fish, simulated, or robotic) to the wall, we measure the means to be slightly smaller than 90° (see Table E.1). Out of the three cases, fish agents exhibit the narrowest distribution with SD of 18.16° , with a close second being the robot with an SD of 18.39° , and the simulated agents producing SD of 25.12° . The SD values are correlated with the mean distance to the wall, meaning that larger distances from the wall would lead to larger SD for the angle of incidence to the wall. In fact, the robot is in very good agreement with single fish distributions ($H = 0.11$), better than the agreement with simulated agents ($H = 0.18$).

The robot experiments demonstrated very good agreement with fish-only experiments, while exhibiting a more significant difference from their simulated counterpart (see Table E.2). The differences between simulation and real-world experiment are most likely owed to physics, the PID controller inducing some noise in the motion, as well as the delays propagating into the real-time system and causing delayed responses (albeit, by a small amount).

Temporal correlation functions

Fig. 7.3 depicts three observables probing the emergent temporal correlations, as defined in Sec. 4.3.

In Fig. 7.3a, we plot the mean squared displacement of the agents, $C_X(t)$, for all three cases. The $C_X(t)$ for single fish agents presents a rapid growth towards its peak at ≈ 5 s, after which point, there is a small decay to a mean level equal to twice the mean square of the distance to the center 403 of the tank. Although all three curves present a very similar growth up to ≈ 3 s, they differ in their peak height as well as their asymptotic value. Generally, the closest an agent swims to the wall, the larger the mean square of their distance to the wall. This is in line with the higher asymptotic value of the robot, which swims on average closer to the wall than both the other agents (fish-only and simulated). Conversely, the simulated agent has the lowest asymptotic value due to its, on average, larger distance to the wall. Focusing

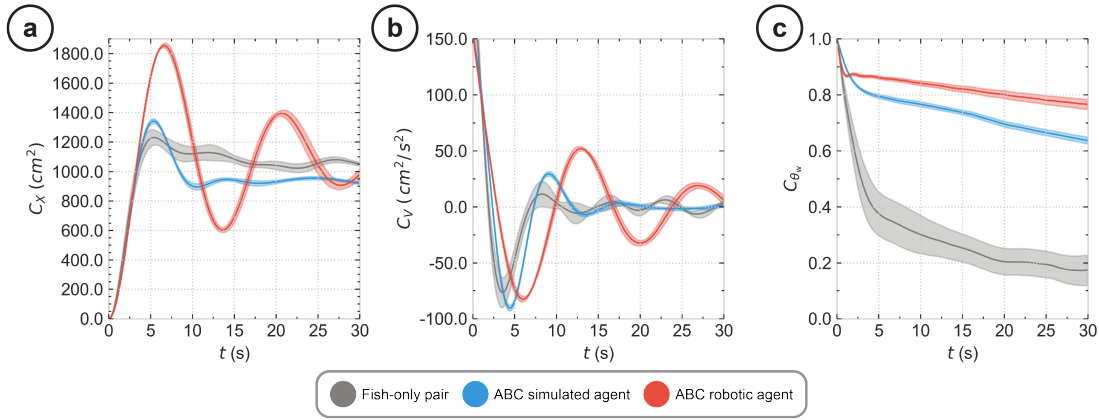


Figure 7.3: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, ABC simulated agents and ABC robotic agents, respectively. Solid lines indicate the average agent distributions. The shaded areas correspond to the standard deviation.

on the position of the peaks, which corresponds to the time for an agent to travel half of the tank perimeter, *i.e.*, correlated to its velocity, we also note differences among the curves. The robot exhibits a peak later than the other two cases, caused by its smaller average speed than that of fish. Finally, the robot curve exhibits oscillations for a longer amount of time, a direct outcome of its very close movement to the wall.

In Fig. 7.3b, we plot the velocity autocorrelation, $C_V(t)$, which vanishes when the velocity of time $t + t'$ becomes uncorrelated with that of time t' , for large values of t . We note that, as discussed in Chapter 6, $C_V(t)$ and $C_X(t)$ are independently observed over a finite sampling time, despite the fact that $C_V(t)$ can be approximated using $C_X(t)$. The peaks of the oscillations of $C_V(t)$ correspond to the inflection points before and after the main peak in $C_X(t)$. For $C_V(t = 0)$, we obtain the mean square velocity, which shows, as discussed in the previous section, the initial velocity of the robot is smaller than that of fish.

Finally, in Fig. 7.3c, we plot the temporal correlation function of the heading of an agent relative to the wall, $C_{\theta_w}(t) = \langle \cos[\theta_w^i(t + t') - \theta_w^i(t')] \rangle$. In all three cases, we expect $C_{\theta_w}(t)$ to decay, at a rate which is strongly related to the sharpness of the peak in Fig. 7.2c. Indeed, we observe that the robot presents strong correlations even after long time, with $C_{\theta_w}(t = 30 \text{ s}) \approx 0.78$, related to the very sharp peak it produced for $|\theta_w|$, which suggests that it may take a long time to explore values of θ_w far from 90° . The slowest decay is presented by the simulated agent, albeit there the agent still presents strong correlations for a long time $C_{\theta_w}(t = 30 \text{ s}) \approx 0.62$. Lastly, the fish decay fast, but also maintain strong correlations for a long time $C_{\theta_w}(t = 30 \text{ s}) \approx 0.18$.

Both the simulated agent and the robotic fail to precisely reproduce the correlation function $C_{\theta_w}(t)$, producing a similar but not as sharp decay compared to the one of the real fish. For the robot, this is owed to its tendency to swim very close to the wall, while simulated agents

swam farther but maintain an angle of incidence of almost 90° throughout the experiments. Similarly to our findings in Chapter 6, the ABC seems to not precisely reproduce the underlying correlations, but the robot was adequately able to follow the dynamics commanded by the ABC model (see the relative difference of $C_{\theta_w}(t)$ between robot and simulation).

7.3.2 Experiments with agent pairs

Instantaneous individual observables

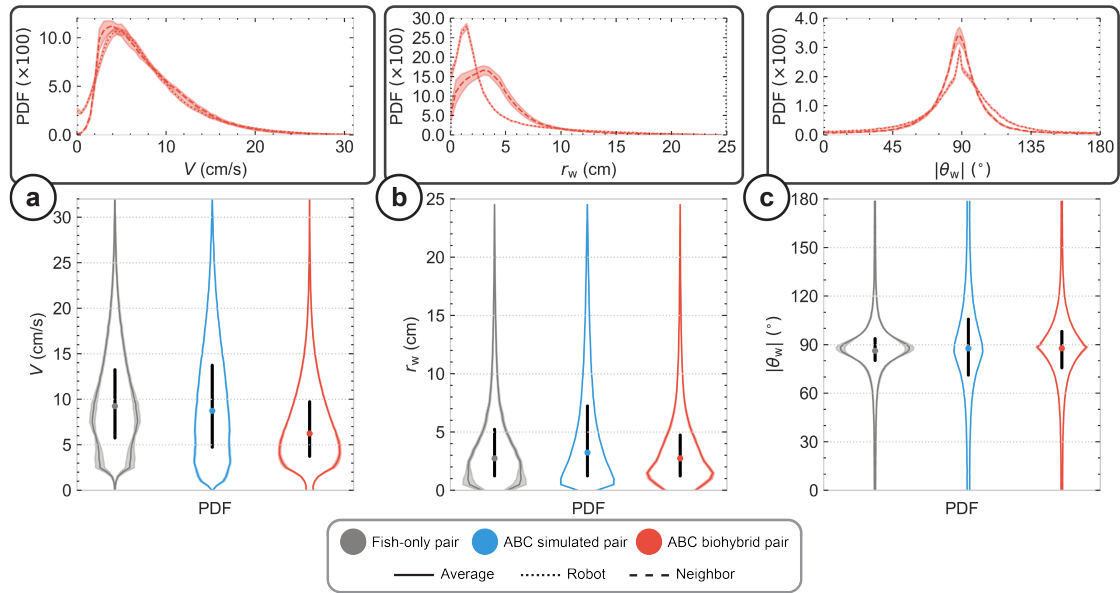


Figure 7.4: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Absolute angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the ABC simulated pairs, and the ABC biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

Fig. 7.4a shows the speed PDF for fish-only, ABC simulated pairs (ABC-SP), and ABC biohybrid pairs (ABC-BP). Fish pairs swim at a mean speed of 10.5 cm/s associated to a standard deviation (SD) of 5.73 cm/s. ABC-SP produce fairly similar speed PDF (Hellinger distance $H = 0.18$), although slightly wider with a SD of 6.39 cm/s, and a very similar mean of 10.38 cm/s. Conversely, ABC-BP speed PDF differ than those of fish-only ($H = 0.26$) and ABC-SP ($H = 0.27$), with a much lower mean of 7.93 cm/s and narrower distribution (SD of 4.89 cm/s). In fact, from the inset plot in Fig. 7.4a, which shows the speed PDF of the robot (dotted line) and that of the neighboring fish (dashed line), we observe that both agents produce this slower and narrower PDF (see also Table E.3).

In Fig. 7.4b, we plot the PDF of the distance to the wall, r_w , for the three cases. Fish pairs, as discussed in Chapter 6, swim very close to the wall and produce a mean distance of 4.39 cm and SD of 3.86 cm, a — already discussed — consequence of the burst-and-coast swimming. ABC-SP reproduce this PDF fairly well ($H = 0.15$) with mean distance of 5.49 cm and SD of 4.91 cm, that is, the pairs swim slightly farther from the wall and with a greater SD than fish. On the contrary, ABC-BP perform very similarly to fish experiments ($H = 0.06$) than ABC-SP ($H = 0.18$). Biohybrid pairs swim with a mean distance of 4.10 cm from the wall and a SD of 3.70 cm, both values almost identical to fish-only experiments. The inset plot in Fig. 7.4b, however, reveals that the neighboring fish swims on average much farther from the wall with a mean distance of 4.46 cm, compared to the robot that produces a sharp peak at 3.74 cm with narrower tail (SD of 0.07 cm; see Table E.3). That is, despite the agreement when the agent distributions are superimposed (averaged), there exist some dissimilarities at the individual level.

Finally, Fig. 7.4c shows the absolute angle of incidence to the wall, $|\theta_w|$. Similarly to the result of single individuals, fish pairs moving close to the wall naturally produce (absolute) angles of incidence slightly below 90° . Here, fish pairs typically swim with an absolute angle of 87.42° and SD of 21.91° . ABC-SP produce fairly a similar distribution ($H = 0.19$), characterized, albeit with a slightly greater mean of 89.95° and wider tail (SD of 33.63°). Once again, ABC-BP show an excellent better agreement to fish-only experiments ($H = 0.10$), but also with respect to ABC-SP ($H = 11$). Biohybrid pairs produced distributions with a mean of 88.12° and SD of 25.93° , both marginally different from fish-only experiments. The inset plot of Fig. 7.4c, confirms that both robot and fish in the biohybrid pair produce similar distributions, albeit the robot exhibits a marginally wider distribution (see Table E.3).

All in all, ABC-SP and ABC-BP are in fair agreement (see Table E.4), indicating that the ABC model was transposed back to real life with fair accuracy. In fact, ABC-BP are in better agreement to fish-only experiments than ABC-SP, indicating that there is a bigger gap in the simulated dynamics than in real experiments with biohybrid pairs. However, the lower velocities exhibited by both fish and robot in biohybrid experiments, indicate a general deviation from spontaneous interactions of fish, which directly stems from the interplay of the artificial lure and the living individual(s).

Instantaneous collective observables

In Fig. 7.5a, we plot the interindividual distance of the agent pair, d_{ij} . Fish pairs swim in close formation, with a mean interindividual distance of 8.05 cm and SD of 5.71 cm. Also, note that the median interindividual distance in fish pairs is less than 2 body lengths (see Fig. 7.5a). ABC simulated pairs produce markedly different dynamics ($H = 0.38$), with a mean interindividual distance of 5.49 cm and SD of 4.15 cm. Conversely, ABC biohybrid pairs are in fair agreement with fish-only experiments ($H = 0.23$), and produce dynamics where the interindividual distance is 8.10 cm and the SD 7.41 cm, that is, the resulting distribution exhibits a very similar peak with fish pairs, but has a wider tail (see Fig. 7.5a), meaning that the

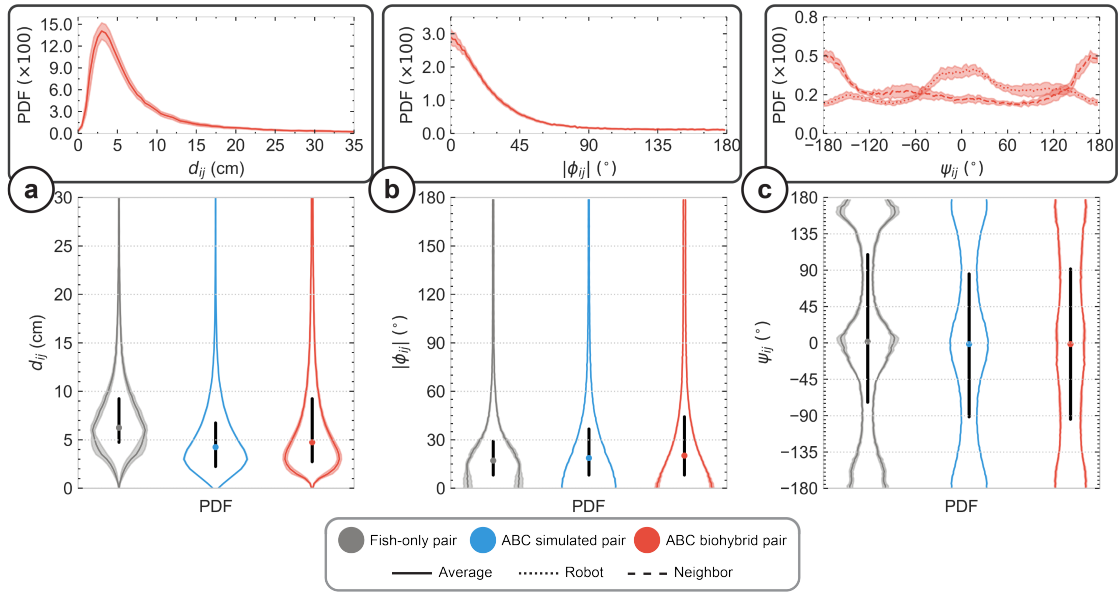


Figure 7.5: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles $|\phi_{ij}|$ probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, ABC simulated pairs and ABC biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

pairs swim more often at a great distances, *e.g.*, at $d_{ij} > 15$ cm.

Fig. 7.5b shows the absolute value of the heading angle difference of agent pairs, $|\phi_{ij}|$, that is, we quantify the pair alignment. Fish pairs tend to form strong interactions that consist of highly aligned movement. Indeed, their absolute heading angle difference is on average 26.72° with a SD of 29.81° . ABC-SP produce a significantly similar PDF ($H = 0.06$) with a mean value of 29.35° and a SD of 30.85° . Notably, ABC biohybrid pairs are also very similar to fish pairs ($H = 0.10$) and in excellent agreement with their simulated counterpart ($H = 0.07$). The mean absolute heading difference in biohybrid pairs greater than that of fish, with a value of 36.50° and SD of 40.60 , thus they produce a wider PDF than both the fish pairs and ABC-SP. That is a consequence of the higher probability for the pair to swim in large distances, at which instances the robot and fish attempt to close the distance by swimming anti-aligned, ergo there is a noticeable weight at $|\phi_{ij}| = 180^\circ$.

We conclude the collective observables by plotting the angle of perception (or viewing angle) in Fig. 7.5c, ψ_{ij} . Fish pair PDF present clear peaks near $\psi_{ij} = 0^\circ$ and $|\psi_{ij}| = 180^\circ$. This suggests that fish typically follow each other in an aligned manner, instead of swimming side by side. ABC simulated pairs present excellent agreement in their dynamics compared to fish

($H = 0.05$), albeit with less pronounced peaks. This difference is mainly derived from the fact that ABC-SP swim farther from the wall. ABC-BP also present similar peaks, but even less pronounced than in ABC-SP. Regardless, the ABC-BP are in excellent agreement with fish pairs ($H = 0.09$) and ABC-SP ($H = 0.04$). In the inset of Fig. 7.5c, we observe that the peak at $\psi_{ij} = 0^\circ$ is dominated by the contribution of the fish (almost a flat line; see dashed curve), indicating that the fish spends more time following the robot than the inverse.

In summary, both ABC-SP and ABC-BP failed to capture the interindividual distance profile produced by spontaneously interacting *H. rhodostomus* pairs. However, the remaining collective observables are in excellent agreement with fish pairs. Furthermore, ABC-SP and ABC-BP are in very good agreement with respect to all collective observables, supporting the claim that the robot is well reproducing the commands of the ABC model, hence closing the biomimicry gap, albeit with marginal differences.

Temporal correlation functions

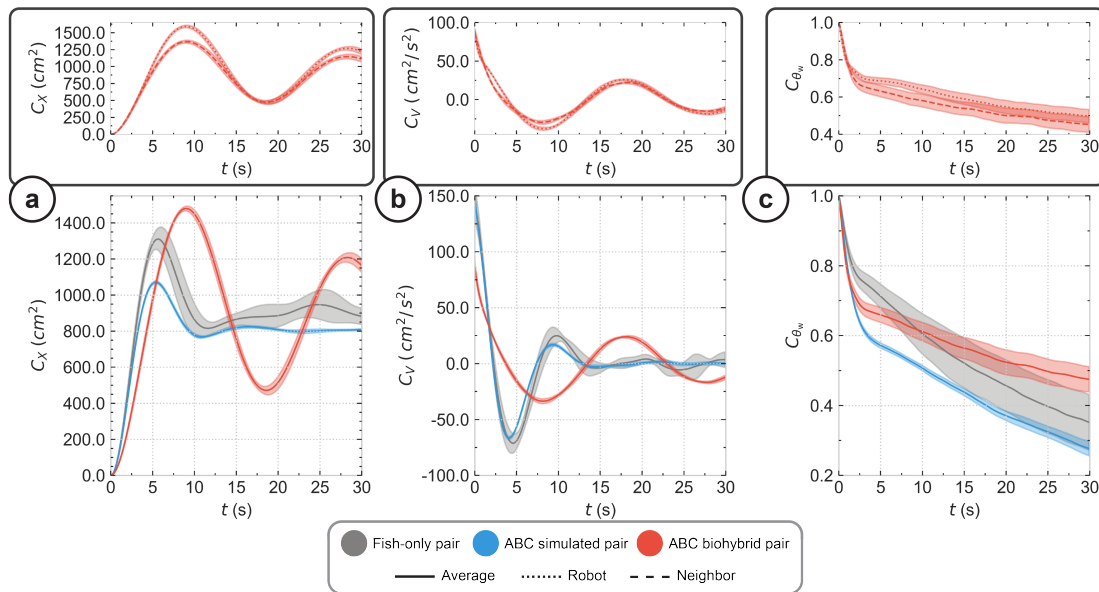


Figure 7.6: Temporal correlation quantities. (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, ABC simulated pairs and ABC biohybrid pairs, respectively. Dotted, dashed and solid lines indicate the robot, neighbor and average agent distributions, respectively. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

Finally, in Fig. 7.6, we plot the three temporal correlation observables, which probe the long-term emergent dynamics of the system.

Fig. 7.6a shows the mean square displacement, $C_X(t)$, of the agent pairs in each case. $C_X(t)$ presents a rapid growth, which eventually results to a peak, followed by decay to a mean level. ABC-BP swim on average closer to the wall, thus creating a more correlated displacement in the long-term timescale. On the contrary, both fish pairs and ABC-SP produce curves that are damped faster, and systematically present smaller values than ABC-BP. This corresponds to the very close movement of ABC-BP to the wall. However, ABC-BP have on average smaller speed than fish and ABC-SP, which is noticeable in the delayed peaks $C_X(t)$ peaks of the ABC-BP. We note that, the position of the peaks roughly corresponds to the time for an agent to travel half the of the tank perimeter. The inset plot of Fig. 7.6a, shows that the robot and fish motion is strongly correlated.

For the velocity autocorrelation, $C_V(t)$, shown in Fig. 7.6b, the peaks of the first two oscillations roughly correspond to the inflection points before and after the main peak in $C_X(t)$. Fish pairs and ABC-SP are in excellent agreement, presenting only a marginal difference. Conversely, ABC-BP, similarly to the mean square displacement correlation, present a markedly different curve, that does not decay as fast and is not damped to the mean throughout the 30 s window. This suggests that the biohybrid pair generally moves in a more coordinated and correlated manner.

In Fig. 7.6c, we plot the final and most subtle of temporal correlations, the correlations function of the heading of an agent relative to the wall, $C_{\theta_w}(t)$. Although for very large times t , $C_{\theta_w}(t)$ decays, we observe that in all cases the pairs still maintain relatively high values for $C_{\theta_w}(t = 30\text{ s})$. More specifically, the fish pair presents a value $C_{\theta_w}(t = 30\text{ s}) \approx 0.35$, ABC-SP presents a value of $C_{\theta_w}(t = 30\text{ s}) \approx 0.28$, and ABC-BP presents a value of $C_{\theta_w}(t = 30\text{ s}) \approx 0.44$. ABC-SP and ABC-BP present a slightly faster decay rate, but all three correlation functions maintain the fairly strong correlations for the duration of 30 s. Furthermore, in the case of ABC-BP, the inset plot in Fig. 7.6c, shows that both robot and fish contribute equally to the average $C_{\theta_w}(t)$, also observed for the first two correlation observables.

In summary, ABC-SP remains fairly faithful to the temporal dynamics observed in fish pairs. On the contrary, ABC-BP fails to reproduce any of the temporal dynamics (ABC-SP or fish pairs), and presents very strong correlations with the fish that swims in the tank. Naturally, the differences are also partly attributed to the biohybrid pair's significantly lower average speed (see Fig. 7.2a), and tendency to move closer to the wall (see Fig. 7.2b). However, this difference is not apparent in ABC-SP, which is based on the same ABC model formulation that commands the robot.

7.3.3 Experiments with groups of five agents (1 agent in perception field)

In the following paragraphs, we present the results of simulated and biohybrid experiments using the ABC model, which is explicitly set to generate trajectories by taking into account only the first most influential neighbor. We note that, as discussed in Chapter 4 and shown in Lei et al. (2020), fish groups seem to take into account multiple individuals in the decision-making

process, with $k = 2$ individuals producing good agreement compared to experimental data of fish. This section presents the results produced when we used ABC with $k = 1$, and in the following section the results when $k = 2$.

Instantaneous individual observables

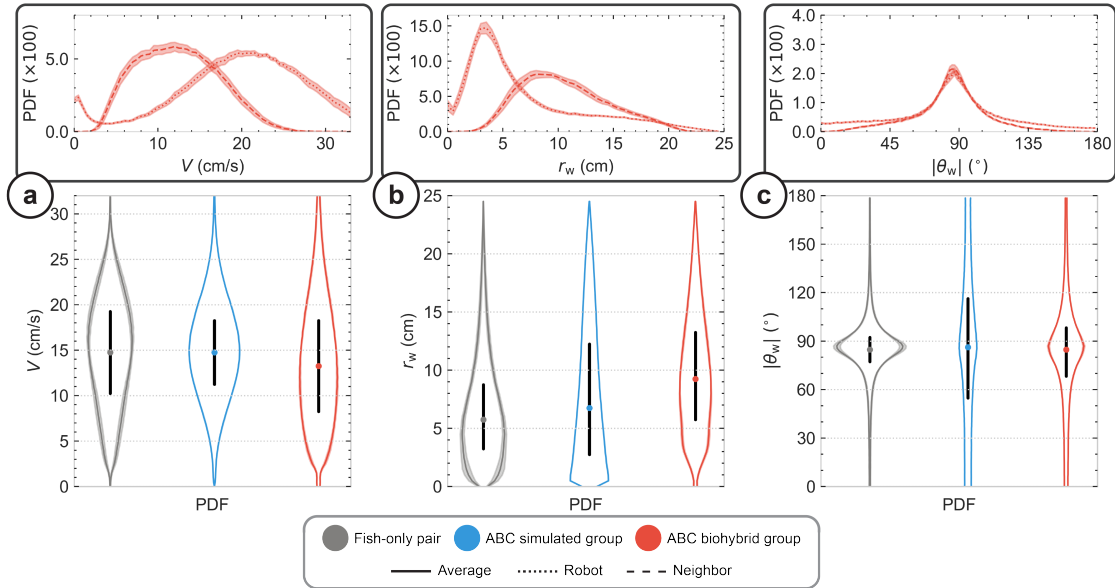


Figure 7.7: Instantaneous individual quantities. (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Absolute angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the ABC simulated groups, and the ABC biohybrid groups, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

In Fig. 7.7, we plot the three instantaneous observables for all cases of groups with five agents, where the simulated and real-time ABC model is taking into account the $k = 1$ most influential individuals.

Fig. 7.7a shows the speed PDF for fish, ABC simulated groups (ABC-SG), and ABC biohybrid groups (ABC-BG). Fish swim on average faster than in smaller groups, with a mean speed of 15.23 cm/s and an SD of 5.94 cm/s. ABC simulated groups produce very similar speed PDF ($H = 0.08$), characterized by a mean of 13.65 cm/s and an SD of 5.47 cm/s. ABC biohybrid groups also produce fairly similar speed PDF, albeit with a lower mean of 14.15 cm/s and a wider PDF tail (SD of 6.92), resulting in a Hellinger distance $H = 0.13$. The inset plot in Fig. 7.7a, shows the speed PDF of the four fish (dashed line) compared to the speed PDF of the robot

(dotted line), from which we observe that the *H. rhodostomus* agents contribute to the lower end of the speed weight, while the robot tends to move much faster (see also Table E.5).

In Fig. 7.7b, we plot the PDF of the distance to the wall, r_w , for which we observe that fish generally swim farther from the wall than smaller groups due to their, on average, higher speed. In fact, they swim with a mean distance of 7.02 cm and a SD of 4.53 cm. ABC-SG once again produces a fairly similar distance PDF, albeit with a higher mean value of 8.38 cm and wider tail (SD of 5.99 cm), eventually leading to a Hellinger distance of $H = 15$. ABC-BG produce greater distances of 10.06 cm with SD of 5.13 cm, that are in relatively good agreement with fish experiments ($H = 0.23$). However, the distance to the wall PDF of the robot in the inset plot in Fig. 7.7b, highlights that the robot is much more likely to swim close to the wall than fish, that, in fact, swim almost 20% farther from the wall than in fish-only experiments (see Table E.5).

Finally, in Fig. 7.7c, we plot the absolute heading angle relative to the normal to the wall, $|\theta_w|$. Similarly to smaller groups, the (close) proximity at which the agents are swimming to the wall, causes a mean value $|\theta_w|$ that is slightly below 90° , that is, agents tend to swim more often towards the wall than away from it. The fish group PDF is measured to have the narrowest distribution out of the three cases, with a SD of 20.02° , compared to a SD of 43.97° for ABC-SG, and 32.18° for ABC-BG ($H = 0.37$ and $H = 0.20$, respectively). Notably, the robot tends to swim towards the wall almost as often as the neighboring fish group, but has higher probability to swim parallel to it (see the inset plot in Fig. 7.7c), a potential effect of its less precise angular control.

Similarly to previous sections, the ABC-SG and ABC-BG do not reproduce the dynamics at their full extent, but generally maintain a good agreement with the spontaneous social dynamics exhibited in groups of *H. rhodostomus*. However, the biohybrid experiments are in good agreement with the simulated results, both commanded by the same social interaction model. This verifies that the closed-loop system presented in Chapter 5 is capable of bridging the biomimicry gap to an adequate degree to reproduce realistic dynamics at the individual level.

Instantaneous collective observables

Fig. 7.8a shows the cohesion PDF, $P(C_i)$, of the group (see Eq. 7.1). Fish produce a cohesion PDF with a mean of 7.37 cm and SD of 2.53 cm, which demonstrates that fish typically swim in close formation with respect to all their neighbors, forming a tightly packed group. ABC-SG produce fairly similar cohesion PDF ($H = 0.23$), albeit with a greater mean value of 8.93 and wider tail (SD of 4.46). ABC-BG produce even less similar PDF compared to fish groups ($H = 0.25$), with a mean of 9.65 cm and SD of 5.54. The inset plot in Fig. 7.8a, confirms that although fish agents produce a similar PDF to the fish-only experiments, the robot contributes to the ultimately wider PDF, which is caused by its tendency to, on average, swim farther than the fish group.

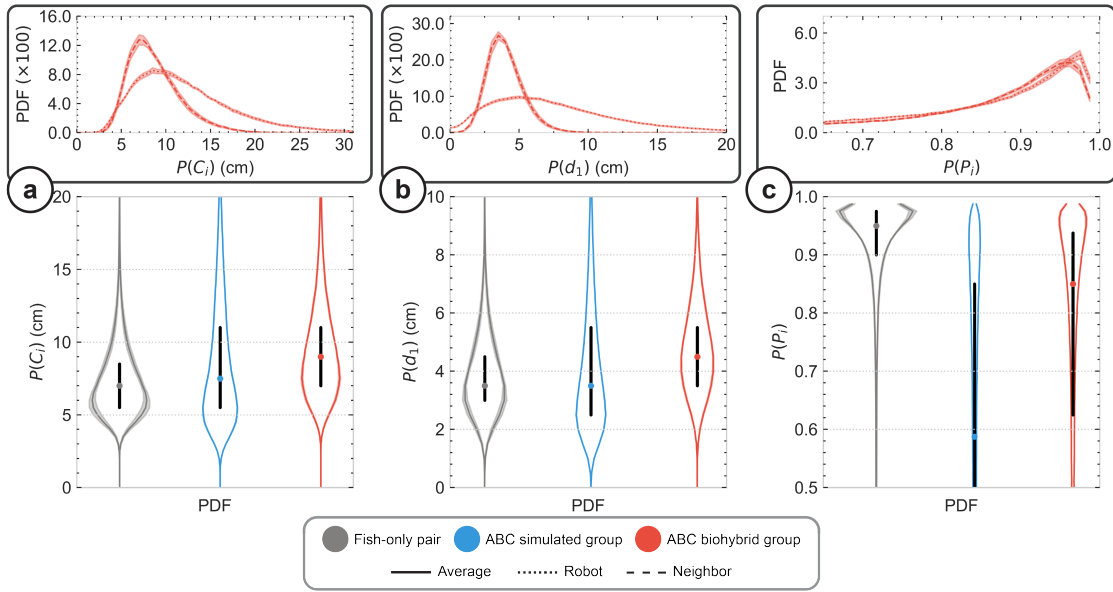


Figure 7.8: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles $|\phi_{ij}|$ probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, ABC simulated groups and ABC biohybrid groups, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

Focusing on Fig. 7.8b, which shows the average distance to the nearest neighbor, $P(d_i)$, we observe that fish groups maintain an average distance of 4.37 cm with SD of 2.27 cm from their nearest individual. For reference, this value is in the scale of one to two body lengths (average body length of the *H. rhodostomus* used is ≈ 3.5 cm). ABC-SG produce relatively similar PDF ($H = 21$) with a mean distance from the nearest neighbor equal to 4.37 cm and SD of 2.27 cm. ABC-BG slightly more dissimilar PDF ($H = 0.22$) with a greater mean value compared to fish, equal to 5.03 cm and a SD of 1.59 cm. Indeed, from the inset plot in Fig. 7.8b, we observe that the robot produces a much wider PDF, swimming on average much farther from its nearest neighbor than fish do (mean 8.15 cm; SD 5.10 cm; see Table E.5).

In Fig. 7.8c, we plot the PDF of the polarization measure defined in Eq. 7.2. As it happens for smaller groups of fish of previous sections, groups of five fish typically swim in an aligned manner. Fish groups produce an average alignment PDF with mean value 0.91 and SD of 0.15, reflecting a highly aligned movement of the group. Conversely, both ABC-SG and ABC-BG move in a significantly less aligned manner ($H = 0.53$ and $H = 0.32$, respectively). Here, we note that the PDF of simulated and biohybrid pairs are in fair agreement $H = 0.25$, strongly indicating that the social interaction model parametrization is not calibrated for aligned

movement, but this is also a direct effect of opting to take into account the social influence of the most influential individual alone. Notably, the fish in the biohybrid group seem to produce a generally less aligned movement even without the contribution of the robot (see inset plot in Fig. 7.8c), suggesting that the robot’s motion affects the group’s typical movement at a collective level.

In summary of this section, ABC-SG and ABC-BG are shown to produce dynamics that are overall dissimilar compared to fish dynamics. On one hand, this may be attributed to considering only the most influential individual in the decision-making process. On the other hand, the greater average distances reported are may well be correlated with the robot’s slightly delayed response and less accurate angular control. However, we note that the PDF produced by ABC-SG are generally in good agreement with the PDF of ABC-BG (see Table E.6), once again showing that the robot is fairly good at transposing the model dynamics back to real life, but ABC requires further calibration to bridge the biomimicry gap.

Temporal correlation functions

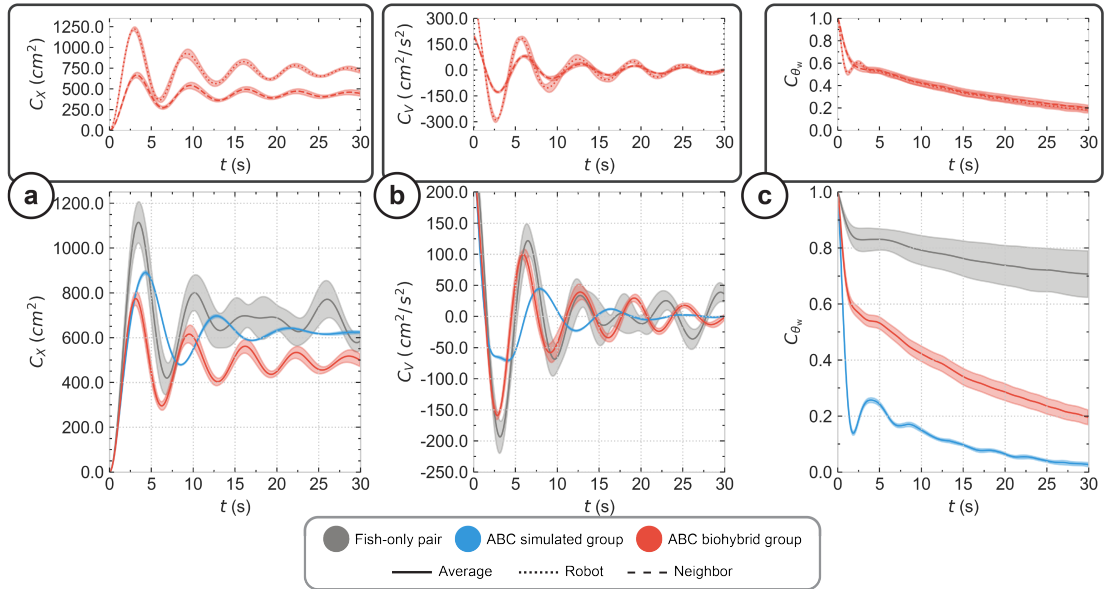


Figure 7.9: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, ABC simulated groups and ABC biohybrid groups, respectively. Dotted, dashed and solid lines indicate the robot, neighbor and average agent distributions, respectively. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

Fig. 7.9 shows plots for the three temporal correlation functions, namely, the mean square displacement $C_X(t)$, velocity autocorrelation $C_V(t)$, and the temporal correlation of the angle of incidence to the wall $C_{\theta_w}(t)$.

First, in Fig. 7.9a, we plot the mean square displacement, $C_X(t)$. For fish groups, the function is characterized by a rapid growth, which eventually leads to a first peak at time ≈ 4 s followed by a decaying oscillation around a mean level, equal to twice the mean square distance to the center of the tank. For large time t , the positions eventually become uncorrelated, and the curve is damped as the agents are guided by the tank walls until the correlation vanishes. The resulting ABC-SG correlation curve has the same qualitative form as that of fish, while the ABC-BG curve presents more oscillations along the 30 s window depicted in Fig. 7.9. Notably, all three curves present differences in the location and height of their peaks, as well as their asymptotic values. The latter correlated to the swimming distance from the wall in such a way that the closer agents swim to the wall, the higher the means square displacement $C_X(t)$ is observed. This mirrors the measurements in Table E.5 for the distance to the wall, r_w , which show that ABC-BG swim farther from the wall than the other two cases. A breakdown of the ABC-BG case, shown in the inset plot in Fig. 7.9a, justifies the robot's asymptotic behavior, which linked to its shorter distance from the wall compared to its neighboring fish, but the average curve results in a lower mean $C_X(t)$ value. Furthermore, the slowly damped periodic oscillations of the ABC biohybrid groups demonstrates that their movement remains correlated for a longer amount of time, compared to the other two cases. Finally, the location of the peaks roughly corresponds to the time it takes for the agents to travel half the tank's perimeter, and is, therefore, directly correlated with their speed and distance to the wall. Here, all three cases present their (first) peak with a small-time difference, primarily owed to their different distances to the wall.

In Fig. 7.9b, we plot the velocity autocorrelation, $C_V(t)$. The peaks of the first two oscillations correspond to the two inflection points right before and after the main peak in $C_X(t)$. $C_V(t = 0)$ corresponds to the mean square velocity, which is in agreement with the results presented in previous paragraphs. We note that, quantitatively, the ABC-BG is in better agreement with the fish-only group velocity autocorrelation curve. The inset plot in Fig. 7.9b confirms that the robot initially contributes to a typically higher velocity than that of its neighboring fish, but for large enough t the curves are almost identical.

In Fig. 7.9c, we quantitatively observe that all three $C_{\theta_w}(t)$ curves markedly differ. However, this is arguably the subtlest among the correlation functions and the most difficult to reproduce. Despite the fact that for very large time t , $C_{\theta_w}(t)$ will decay, we observe that fish groups we still have a relatively high value $C_{\theta_w}(t = 30 \text{ s}) = 0.7$ (almost twice what was observed in fish pairs). For ABC simulated groups, $C_{\theta_w}(t)$ vanishes very rapidly, and has a value $C_{\theta_w}(t = 30 \text{ s}) < 0.05$. ABC-BG also show a rapid decay, and maintain a value $C_{\theta_w}(t = 30 \text{ s}) = 0.25$, fairly lower than that of fish groups. For reference, as discussed in previous paragraphs, the decay rate is correlated to the sharpness of the peak near θ_w , which indeed justifies the very fast decay of ABC-SG that produced a very wide PDF θ_w (see Fig. 7.7).

Both ABC-SG and ABC-BG are fairly capable of reproducing the mean square displacement and the velocity autocorrelation, but fail to recover the angle of incidence to the wall correlation. Similarly to previous observables, the biohybrid group produce better agreement for the correlations compared to the simulated groups, suggesting that the simulation, at its current parametrization, is not adequately reproducing the short- or long-term social dynamics. On the other hand, the robot, perhaps partly masked by the fish group's behavior, is overall good at approximating the spontaneous social dynamics of fish-only *H. rhodostomus* groups.

7.3.4 Experiments with groups of five agents (2 agents in perception field)

In this section, we repeat the experiments with groups of five fish, only this time we consider the $k = 2$ most influential individuals.

Instantaneous individual observables

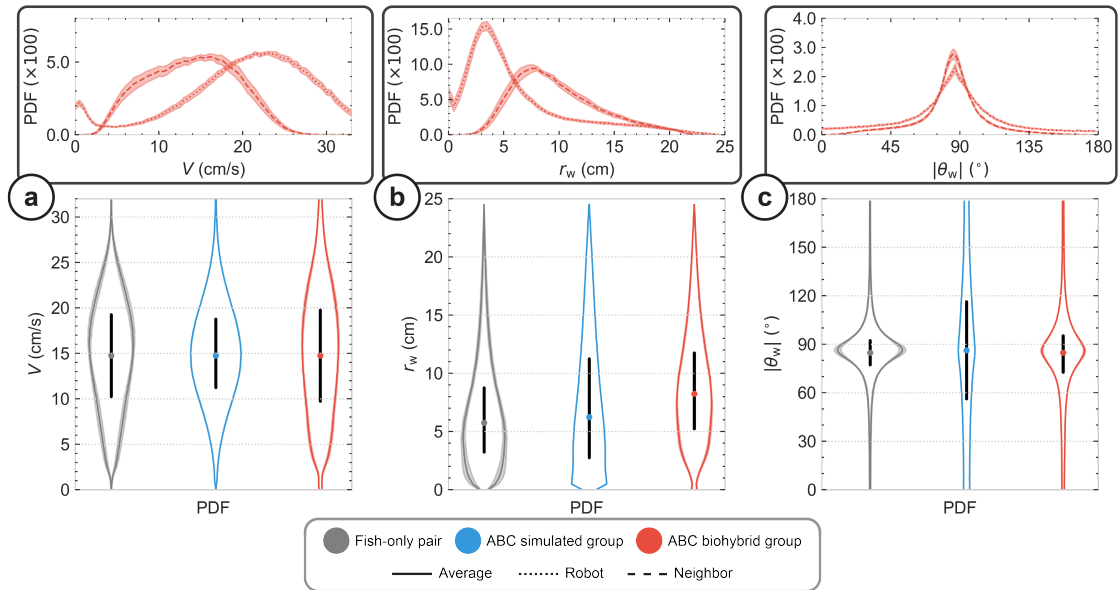


Figure 7.10: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Absolute angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the ABC simulated groups, and the ABC biohybrid groups, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

In Fig 7.10 we plot the individual quantity observables for all three cases, where the ABC model considers the top two most influential neighbors. First, Fig 7.10a shows the speed

PDF. For fish groups, we use the same dataset, thus the mean and SD remain identical (mean 15.23 cm/s; SD 5.94 cm/s). The ABC simulated groups produce a significantly similar PDF ($H = 0.08$) with a mean value of 15.35 cm/s and SD of 5.15 cm/s. The ABC-BG with $k = 2$ produce a PDF that shows an improved, with respect to $k = 1$, agreement compared to fish groups ($H = 0.10$). However, similarly to the $k = 1$ case, the inset plot in Fig 7.10a, reveals that the fish contribution masks the higher speed of the robot and its wider speed PDF.

When considering the top two most influential individuals, ABC-SG produces a distance to the wall PDF that is slightly improved compared to the $k = 1$ case, with a mean value of 8.10 cm and SD of 5.65 cm, resulting in a Hellinger distance $H = 0.12$ (see Fig 7.10b). ABC biohybrid groups also present a sizeable improvement over the $k = 1$ case, with a mean and SD values of 9.32 cm and 4.85 cm, respectively, and an improved Hellinger distance $H = 0.19$. The inset plot Fig 7.10b, reveals that the fish group (excluding the robot's contribution), maintains a greater distance from the wall compared to fish-only groups, but when ABC is considering the top two most influential neighbors, the group swims ≈ 0.8 cm closer to the wall than when considering only the most influential. This shift, although small and fully explained from this observable alone, already suggests the ABC-BG engage in more life-like interactions when $k = 2$.

Finally, in Fig 7.10c we plot the absolute heading angle difference to the normal to the wall, $|\theta_w|$. The ABC-SG with $k = 2$ produces dissimilar angle PDF ($H = 0.38$), equally dissimilar to the $k = 1$ case. However, the biohybrid groups, show a 30% similarity improvement in the Hellinger distance ($H = 0.14$), compared to the $k = 1$ case. More specifically, the ABC-BG produce a PDF with a sharper peak and narrower tails (SD of 28.40°), albeit they are still wider than those produced by fish-only groups (SD of 20.02°).

Allowing the model to consider the top two most influential neighbors to produce motion predictions, showed a small increase in similarity at the level of the individual quantity observables, both in the case of simulated and biohybrid pairs. Finally, the ABC-SG and ABC-BG also showed an improvement of the order of 10% in their Hellinger distance (see Table E.8), demonstrating how high-fidelity models with proper calibration can effectively bridge the biomimicry gap.

Instantaneous collective observables

In Fig. 7.11a, we show the cohesion PDF for the two ABC cases that consider the top two most influential neighbors. ABC simulated groups exhibit a 20% deterioration in the Hellinger distance compared to the $k = 1$ case. Conversely, ABC biohybrid groups, exhibit a 40% improvement ($H = 0.15$) over the $k = 1$ case. The improvement is caused by the on average smaller intergroup distance both when considering the robot alone and the neighboring fish alone (*i.e.*, by excluding the robot). However, similarly to the $k = 1$ case, the fish contribute to the high probability of maintaining a small distance to their neighbors, rather than the robot (see inset plot in Fig. 7.11a).

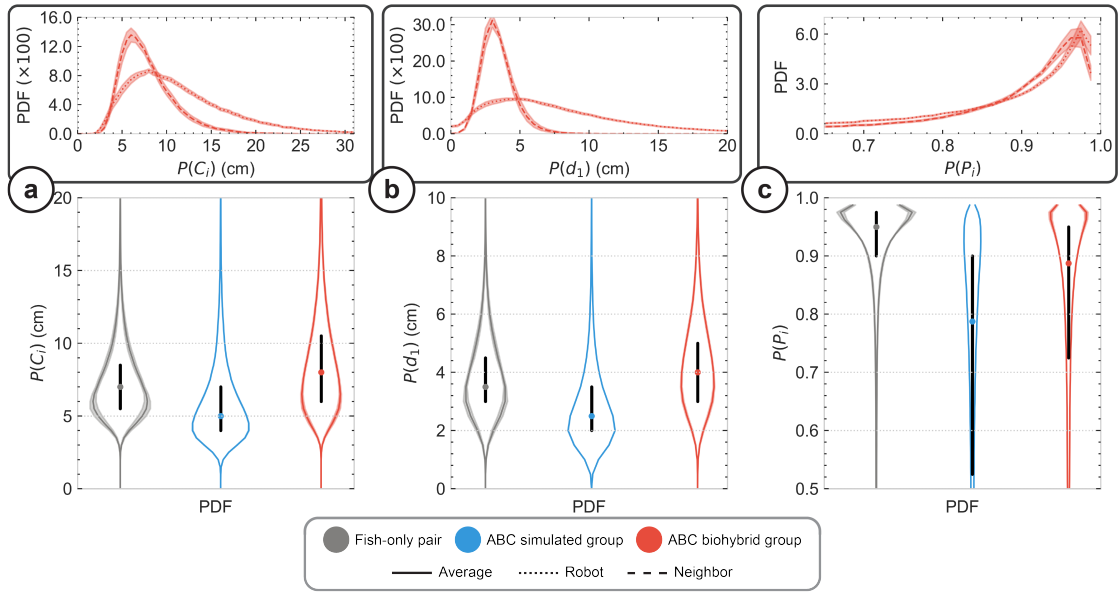


Figure 7.11: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles $|\phi_{ij}|$ probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, ABC simulated groups and ABC biohybrid groups, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

Fig. 7.11b shows the average distance to the nearest neighbor, for which ABC-SG with $k = 2$, once again, show a small deterioration in their corresponding Hellinger distance ($H = 0.31$), while the ABC-BG improve by approximately 60% ($H = 0.11$), showing very good agreement with fish groups. However, the inset plot in Fig. 7.11b highlights that this improvement is owed to the higher probability of fish to swim closer to their neighbor, while the robot, similarly to the $k = 1$ case, struggles to stay very close to its nearest neighbor.

Finally, in Fig. 7.11c we plot the polarization PDF of the group. Contrary to the results of the last two observables, here, ABC-SG show an improved polarization agreement with fish groups, albeit the PDF remains significantly dissimilar ($H = 0.43$). Conversely, ABC-BG produce a polarization PDF that showcases an approximately 30% improvement over the $k = 1$ case, with a value of $H = 0.20$, and achieve a relatively good agreement with fish group polarization PDF. Notably, the inset plot in Fig. 7.11c shows that both the robot and the four fish contribute equally to the average polarization PDF. However, the robot has higher probability to be unaligned with the rest of the group, and more notably, the fish group produces higher probability of being unaligned as well (see the Table E.7).

With respect to the collective quantity observables, we observed a markedly large improvement when the ABC model considers the top two most influential neighbors. However, we note that, the similarity between ABC-SG and ABC-BG with $k = 2$ improves only for the individual and deteriorates for the collective quantities, with respect to how the $k = 1$ groups compared to fish-only groups.

Temporal correlation functions

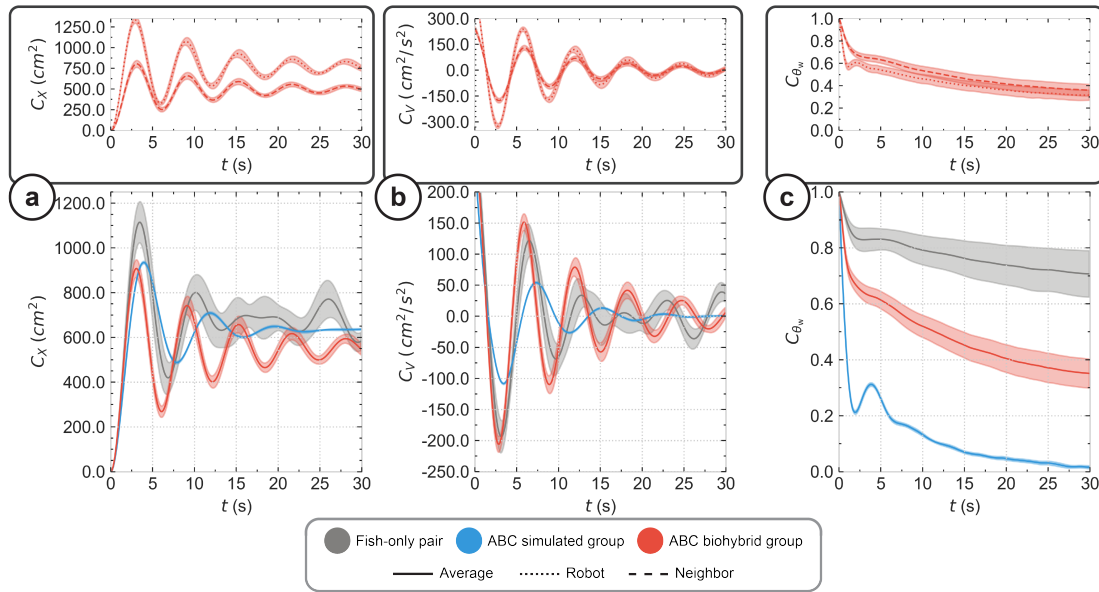


Figure 7.12: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, ABC simulated groups and ABC biohybrid groups, respectively. Dotted, dashed and solid lines indicate the robot, neighbor and average agent distributions, respectively. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the standard deviation.

In Fig. 7.12, we plot the three temporal correlation functions: mean square displacement $C_X(t)$, velocity autocorrelation $C_V(t)$, and the temporal correlation of the angle of incidence to the wall $C_{\theta_w}(t)$. We note that, the ABC variants are set to consider the top two most highly influential neighbors to generate future motion primitives.

Fig. 7.12a shows the mean square displacement, $C_X(t)$. The correlation curve for fish groups remains identical to the one of the previous section, with a rapid growth phase, a first peak at time ≈ 4 s, followed by a decaying oscillation. The resulting ABC-SG correlation curve is very similar to the one observed for ABC-SG with $k = 1$. However, the ABC-BG curve presents a measurable improvement, with a first peak height more similar to fish groups than in the case of $k = 1$, as well as a more similar asymptotic rate compared to fish. The improvement

is primarily owed to the improved (smaller) distance to the wall produced by ABC-BG (see Table E.7 for further details), while ABC-SG do not produce improved results.

Then, in Fig. 7.12b, we plot the velocity autocorrelation, $C_V(t)$. Here, considering the top two most influential neighbors drives a significant improvement in the agreement of the velocity autocorrelation functions. In the case of ABC-SG, the location and height of the first peak are markedly improved from its $k = 1$ counterpart, albeit the inability of the simulated group to produce similar speeds to fish groups, leads to delayed peaks. On the contrary, ABC-BG present a small but noticeable improvement compared to ABC-BG with $k = 1$, exhibiting very similar peak heights and location for the first 15 s, after which point, the curves diverge, but maintain roughly the same asymptotic value.

Finally, Fig. 7.12c remains a difficult quantity to reproduce, despite the multi-level improvements presented in previous paragraphs. Indeed, the ABC simulated pairs produce only a slightly improved decay rate, but nevertheless still almost entirely vanish within the observation time-window with a value $C_{\theta_w}(t = 30 \text{ s}) < 0.05$. For the ABC-BG, the angle of incidence to the wall, θ_w , is improved with respect to ABC-BG with $k = 1$ and compared to fish groups (see Fig. 7.10c; Table E.7) leads to more slowly decaying $C_{\theta_w}(t)$, closing the gap between ABC-BG and fish groups, but still decay fairly fast to the value $C_{\theta_w}(t = 30 \text{ s}) = 0.37$.


In summary, considering $k = 2$ most influential neighbors in the ABC compute cycle, produced more similar correlation curves than when the ABC considers only $k = 1$. That is, as discussed in detail in Lei et al. (2020), the artificial agent is much more able to follow life-like social dynamics, and as presented in this section, is better integrated in the fish group.

7.4 Conclusion


Chapter 6 delves into the transferability of social interaction models from simulated to real-world environments, while providing insight on the biomimicry gap that surfaces in this transition. While it sets a precedent for the formal investigation of this gap, it does not investigate the scalability of such interaction models with respect to the number of individuals, nor does it investigate the biological notion of the number of individuals that are taken into account during an individual's decision-making process, and as a result reduces the biomimicry gap. This chapter tackles these limitations by applying the analytical burst-and-coast model as outlined in Calovi et al. (2018) and Chapter 4 to command the LureBot (as described in Chapter 5) in real-time, allowing it to interact with varying group sizes of fish. Our goal, in this chapter, was to examine how such models scales, and how the biomimicry gap either widens or narrows when social dynamics evolve from the interactions of one to five individuals, as compared to a pair of individuals (one robot and one fish) as presented in Chapter 6. We also investigated the significance of taking into account multiple individuals when taking decisions (for groups of $N = 5$ individuals), which while it has been explored in simulation and robot swarms Lei et al. (2020), has never been validated in biohybrid groups. Despite the unprecedented, yet partial, success on bridging the biomimicry gap, we also

observed discrepancies that did not allow for fully bridging the biomimicry gap.


(A) $N = 1$



(B) $N = 2$



(C) $N = 5, k = 1$



(D) $N = 5, k = 2$




Figure 7.13: **Quantification of biomimicry gap (root mean squared Hellinger distance across all observables) values between different experiments.** (A) values for single agent experiments ($N = 1$), (B) values for pairs of agent experiments ($N = 2$), (C) values for group of five agent experiments ($N = 5$) where the robot considers the most influential neighbor ($k = 1$), and (D) values for single agent experiments ($N = 5$) where the robot considers the two most influential neighbor ($k = 2$).

At a first level, the social interaction gap stems from the imperfect, and sometimes insufficient, performance of the ABC model in simulations, where self-propelled agents respond instantly without any physical constraints, thereby revealing missing dynamics in the model. We found that, in many instances, biohybrid groups performed better compared to the simulated model which displayed a larger discrepancy relative to real experiments (see root mean squared Hellinger distances for single, pairs, and groups of five agents in Fig. 7.13A, Fig. 7.13B, and

Fig. 7.13C-D, respectively). This naturally leads to the expectation that better social interaction models would substantially improve biohybrid pairs' performance. These limitations become more evident as we increase the group size (see Tables E.6, E.8; and Fig. 7.13C-D), potentially caused by the superposition of discrepancies between each individual and the robot. Notably, in the 5-agent groups, we validated for the first time in a biohybrid system that these discrepancies are significantly reduced, by approximately 30%, when the robot considers the top two most influential neighbors (see Table E.8; Fig. 7.13C-D). This demonstrates how missing elements in social interaction models at the individual level can scale to discrepancies at the collective level. In all group size cases, much like in Chapter 6, the biohybrid groups displayed poor performance in alignment, leading to subpar performance in the corresponding correlation observable for the angle of incidence. The inability of our system, comprising the social interaction model and biohybrid interaction framework, to reproduce alignment has inadvertently propagated in the remaining observables. However, despite the discrepancies, the majority of distributions we studied for the biohybrid case showed reasonable agreement with real experiments, falling within an acceptable range of biomimetic performance, with means squared Hellinger distance values of $H = 0.15$ for $N = 1$, $H = 0.18$ for $N = 2$, $H = 0.23$ for $N = 5$, $k = 1$, and $H = 0.16$ for $N = 5$, $k = 2$. That is, except for biohybrid experiments where we consider only the most influential neighbor, our system (model and robotic framework) achieved Hellinger values of $H < 0.2$, which, as previously, discussed, indicates very good agreement with real groups of fish.

At the second level, the observed discrepancies in the physics and communication cue gaps can be traced back to unaccounted physical factors, such as simulation mechanics, robot motion, and computational methods that limit robot response time, as well as subtle patterns that the social interaction model may be missing. These discrepancies can be attributed to issues such as the loss of magnetic coupling between the LureBot and lure, lossy and delayed wireless communication with the robot, imperfect fish tracking, and robot inertia. Regarding the physical gap, the experiments in this chapter and Chapter 6 now allow us to identify limitations in our current system. These limitations could be addressed in future work through software or hardware developments. However, as discussed in previous chapters, the communication gap remains particularly challenging to understand, since ideally, the task of bridging it would require the elimination of the other two sources of discrepancies. Nonetheless, it can be intuitively assumed that the lure's lack of vertical axis movement may make the lure itself less noticeable to neighbors in some occasions, thereby decreasing its impact on their decision-making and causing short- and long-term discrepancies. Furthermore, despite our best efforts, replicating the colors and shape of the lure to resemble the *H. rhodostomus* remains only an approximation of the living animal.

Despite these setbacks, similarly to Chapter 6, we establish that the fish engage into meaningful social interactions with the robot throughout the experiments, as the root mean squared Hellinger distances demonstrate (depicted in detail in Fig. 7.13). For instance, even when the robot moves faster than the fish group, significant changes in the behavior of the remaining (live) group members are not observed. Yet, there is potential for further improvements. For

example, having structured this multi-level issue, attention could be shifted towards optimization algorithms that calibrate the social interaction model in real time or simulations to align with the target dynamics. This approach mirrors those found in Cazenille et al. (2015, 2016, 2018b, 2019), but with validation in both simulation and biohybrid experiments, and utilizing high-fidelity models. Naturally, this introduces a new trade-off between computational time and system responsiveness.

Together with the work presented in Chapter 6, this chapter, to the best of our knowledge, represents the first (along with that of Chapter 6) end-to-end investigation concentrating on the biomimicry gap, that utilize high-fidelity models that control robots in real time. In particular, we executed an experiment involving varying fish group sizes, leveraging the same end-to-end approach established in Chapters 4. To achieve this, we make use of the specialized equipment presented in Chapter 5 and scale up the initial biomimicry gap experiment outlined in Chapter 6 to include different sizes of fish groups. We believe that the culmination of these elements of work, leading to this chapter's experiment, could potentially set a benchmark for similar biohybrid experiments within this field. With this research, we hope to establish groundwork for future end-to-end methodologies that verify behavioral models both in simulations and real life, and extend the work of Chapter 6 by assessing the scalability of such experiments at all levels (simulation and biohybrid) when additional individuals contribute to the collective dynamics. In this way, we are making strides towards genuinely understanding and bridging the biomimicry gap as defined in Chapter 1.

Conclusion Part V

8 Contributions and research outlook

The investigation of collective behavior phenomena spans multiple disciplines, robotics being among them. While the domain of robot-animal interaction is not a new field, as discussed throughout this thesis, it nonetheless harbors numerous yet-to-be-explored research inquiries. The advancements in electronics manufacturing, the development of modern algorithms, and the widespread access to computational power have collectively transformed this field. These progressions have furnished researchers with increasingly efficient tools for experimentation, thereby enabling more sophisticated and comprehensive studies. The ultimate aim is to provide insight into the underlying interaction rules governing collective systems.

In Chapter 1, we presented the primary limitations inherent to the various disciplines involved in studying collective behavior and creating biohybrid groups. We also discussed the research questions derived from these limitations (Chapter 2). In Chapter 3, we provided evidence supporting the notion that fish groups are more likely to integrate a robot engaging in active bidirectional interactions (RQ1). This finding not only establishes the baseline behavior we strive for in robot-animal experiments, but also underscores the need for additional work on modeling social interactions. In Chapter 4, we tackled this modeling challenge by implementing and rigorously evaluating two social interaction models (one analytical and one machine learning) that exhibit high fidelity in mimicking spontaneous fish movements (RQ2). We further stressed the necessity of benchmarking such models over short and long timescales to examine the emergence of temporal correlations, rather than merely fitting short-term spatial metrics. In Chapter 5, we introduced a novel open-source robotic system equipped to replicate the demanding dynamics of the models presented in Chapter 4. Lastly, in Chapters 6 and 7, we explored how the high-fidelity models outlined in this thesis can be transposed back into real-world scenarios via the Behavioral Observation & Biohybrid Interaction (BOBI) framework (refer to Chapter 5; RQ3). The successful transposition of these models is subject to potential discrepancies arising from both physical realities and gaps in social interaction modeling. We coin this discrepancy the “biomimicry gap” (RQ4), a concept analogous to the reality gap Jakobi et al. (1995); Jakobi (1997); Mouret and Chatzilygeroudis (2017), but differentiated by its emphasis on the social interaction component.

This thesis' work was purposely structured to investigate the three components of the biomimicry gap (*social interaction, physics, and communication cue gap*) were investigated incrementally and in accordance to the general mixed society construction methodology introduced in Chapter 6 and summarized in Fig. 8.1. In the following sections, we summarize the key outcomes of the thesis that align and provide solutions to our research questions, covering the majority of mixed society study areas.

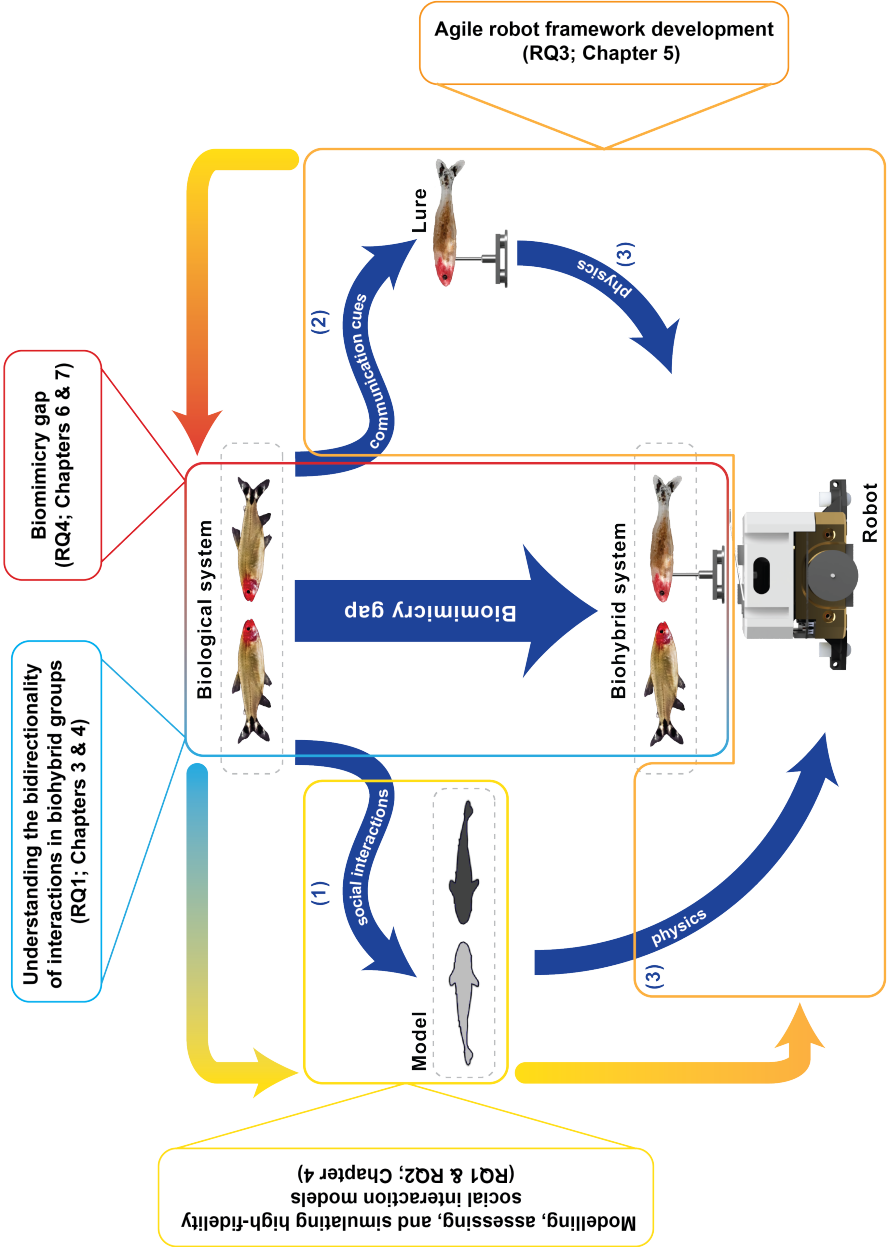


Figure 8.1: **Biomimicry gap study plan.** The light blue, yellow, orange, and red frames depict the areas of study covered in this thesis, as well as, their corresponding relationship to our research questions and chapters.

8.1 How can social interactions be modeled in an informative and expressive manner, with the purpose of constructing cohesive mixed (*i.e.*, biohybrid) groups of animals and robots?

8.1.1 (RQ1) How can we explore the bidirectionality of fish-robot interactions to measure how they facilitate the integration of a robot into fish groups?

While it is recognized that fish groups engage in bidirectional communication, the extent to which this bidirectional interaction is affected when artificial agents are present remains largely unknown. A host of researchers have delved into the concept of eliciting responses via artificial lures (*e.g.*, Faria et al. (2010); Swain et al. (2011); Landgraf et al. (2013); Cazenille et al. (2018c); Porfiri (2018); Chemtob et al. (2020); Romano and Stefanini (2021, 2022b)), yet none have undertaken a comparative study. In Chapter 3, we ventured into this domain, experimenting with models that range from: 1) open-loop models, which remain unresponsive to stimuli from neighbors, to 2) closed-loop models that can be further subdivided into i) passive models, which indiscriminately follow group-level movements, and ii) active models, which alternate between following and leading.

Our results indicated that the active approach facilitated the robot's integration into the biohybrid group, and we even observed that the fish were more inclined to follow the robot's lead. However, the low complexity of the model used limited us from gaining a deeper understanding of how this model made the robotic agent more appealing to the fish. This served as motivation to transition towards more intricate and accurate models of interaction, which could then be compared against spontaneous fish interactions (RQ2).

8.1.2 (RQ2) How can we design and assess bidirectional models of high fidelity (*w.r.t.*, to observations from the spontaneous animal interactions)?

To gain a deeper understanding of how fish adapt their behavior in the presence of an artificial lure, we shifted our focus to the design of more detailed models of social interaction (refer to Chapter 4). We aimed to ascertain whether machine learning methodologies could yield such models, in contrast to using traditional analytical approaches, and whether substantial improvements could be observed when employing either of these methods. However, in this endeavor, we uncovered a more substantial challenge. While it is well established that (fish) collective behavior operates over multiple temporal scales Ioannou and Laskowski (2023), most models of social interaction predominantly depend on short-term observations and predictions of spatial quantities. During the development phase, it quickly became apparent that a model performing well over short intervals does not necessarily translate into equivalent success over long periods.

We tackled these limitations in a comprehensive research work in Chapter 4, where:

1. we proposed a robust spatial and temporal benchmarking metric set;
2. we developed a novel deep learning model of social interaction and contrasted it against the state-of-the-art analytical model and a similarly purposed trajectory forecasting algorithm.

Through this process, we identified key design choices that were instrumental in enabling the machine learning algorithm to learn the dynamics. Namely:

- Models that perform well in short timescales do not necessarily generate realistic long timescale social dynamics;
- Neural networks with memory significantly outperform their non-memory counterparts;
- modelling the data as a sequence-to-sequence problem aids the model convergence (*i.e.*, to a model that can generate realistic social dynamics);
 - Input sequence lengths that are biologically informed are trained faster and perform better (*i.e.*, their length is selected based on biological knowledge, here 0.5 s is the average kick length).
- there is an imperative need for standardized benchmarking metric sets that validate a social interaction model's generative abilities.

We trust that this work will establish a more methodical design and evaluation standard, and stimulate the research community to validate the generative capabilities of social interaction models across both short and long timescales to ensure their accurate generative capabilities of the targeted social dynamics.

8.1.3 (RQ3) Which key elements make for robotic devices that can closely mimic fish motion and dynamics to subsequently transfer such models “back to reality”?

The development of high-fidelity social interaction models in Chapter 4 further highlighted three limitations in the realms of software and hardware:

1. In combination with the rapid real-time control of multiple devices, including a mobile robot, these models necessitate real-time identification of all individuals. This functionality was neither available in our previous experimental framework (CATS) Bonnet et al. (2017) nor is it typically featured in existing fish-robot frameworks;

2. The models demand robots capable of matching or exceeding the agility of fish, a criterion that the majority of current robotic systems fails to meet (refer to Appendix C.1);
3. The models call for faithful (fish) replicas to ensure minimal visual bias.

These limitations served as the motivation for the design and development of a novel framework, encompassing an experimental setup, a mobile robot, and the ancillary software, as presented in Chapter 5. We exploited the BOBI framework to showcase how the amalgamation of these advancements succeeded in achieving unprecedented levels of biomimicry across varying sizes of fish groups.

8.1.4 (RQ4) How can we measure, and what is the gap between real experiments and simulations with such models?

Despite numerous animal-robot interaction studies, including the subset of fish-robot studies which are the focal point of this thesis, to our knowledge, no study performs a comprehensive, end-to-end analysis of how social interaction models transition from simulation to reality. Analogous to the reality gap [Jakobi et al. \(1995\)](#); [Jakobi \(1997\)](#); [Mouret and Chatzilygeroudis \(2017\)](#), we introduced the term “biomimicry gap”, which represents a similar discrepancy resulting from unaccounted physics when controlling the robot, and subtle complexities missing from the social interaction models commanding the robots. Echoing the challenges of the first research question (RQ1), this gap not only constrains our understanding of collective behavior (within the context of animal-robot interaction), but also inherently complicates identifying the root causes of such a gap, whether they lie in physical or modeling elements. In this thesis, we strive to enhance our understanding and bridge this gap by integrating the work of (RQ1-3; Chapters 3-5). The culmination of these findings lays the groundwork for Chapters 6 and 7.

In Chapter 6, we introduced the concept of the biomimicry gap to the research community by transferring a deep learning model of social interactions from simulation to reality. We demonstrated the feasibility of this approach, observed significant similarity between simulations and biohybrid experiments, although some deviation from the spontaneous behavior of fish groups was noted. To our knowledge, this represents the first validation of a neural network model for social interactions in both simulation and reality.

Similarly, Chapter 7 involves a larger-scale experiment. Using an analytical model of social interactions with varying fish group sizes, we accurately measured the discrepancies between fish-only, simulated, and biohybrid fish groups. More specifically, we demonstrated that our robotic system is capable of bridging the physics and communication gaps to unprecedented levels, with our models causing only minor discrepancies contributing to the social interaction gap. Notably, we observed that these discrepancies widen as the group size increases. Both our models and the biohybrid system exhibited a larger discrepancy for larger groups, which could be attributed to the superposition of the complex dynamics at the individual level

generating increasingly diverse dynamics at the collective level as the number of individuals involved increases. Our models and robot are currently constrained within a certain amount of variance, indicating potential directions for future work. We posit that this experimental approach, to the best of our knowledge, is pioneering in its kind, bridging the gap between simulated and real-life studies of robot-assisted social interactions.

8.1.5 Technical limitations and future work

This section's content is adapted from the following article:

- Papaspyros, V., Burnier, D., Cherfan, R., Theraulaz, G., Sire, C. and Mondada, F., 2023. A biohybrid interaction framework for the integration of robots in animal societies. IEEE Access. <https://doi.org/10.1109/ACCESS.2023.3290960> (reuse authorised under the CC 4.0 licence) – Papaspyros et al. (2023a)
 - Vaïos Papaspyros's contribution: Conception, data curation, formal analysis, investigation, methodology, software, hardware, visualization, writing – original draft, writing – review & editing.

While significant progress has been made in comparison to existing robotic platforms, certain aspects of our robotic platform still warrant further exploration and development. Narrowing the gap between the spontaneous movements of fish and those generated by biohybrid groups depends on the continuous development of motion models and the hardware supporting them.

The inherent limitations of mechanical devices, such as the LureBot, which lack the plasticity of living organisms, pose a challenge. For instance, fish, including *H. rhodostomus* and *D. rerio*, can execute extremely rapid directional changes that the LureBot can only approximate. In particular, the LureBot cannot fully express the burst-and-coast dynamics (sudden acceleration followed by a gliding period) observed in several fish species. Moreover, due to the inherent design of the BOBI and LureBot, the lure effectively moves in two dimensions. While the majority of fish experiments in the context of the study of collective phenomena consider fish swimming in shallow water (an effective two-dimensional setup), let us mention that some fish robotic platforms have the ability to move in 3D Ruberto et al. (2016); Bartolini et al. (2016). In many cases, such as in this study, the biomimetic lure, while visually very similar to the considered fish, is passive: it does not actively beat its tail or bend its body, nor can one control its posture. Whereas lures with an actively beating tail have been proposed Bonnet et al. (2016), such lures are bound to be larger than the considered species, to accommodate the necessary electronics and mechanics. Therefore, additional research on improved motors and LureBot configurations that could potentially allow it to perform a true burst-and-coast motion is a valid future research direction.

8.1.6 Modelling limitations and future work

Modeling social interactions remains a considerable challenge within the scientific community. Various methodologies exist for creating models, but within the context of this thesis, we categorize them into two primary types: 1) analytical models, and 2) machine learning models. Each type comes with its own limitations, as we detailed in Chapters 1 and 4. Generally, these limitations manifest as a failure of analytical models to detect subtle patterns Marques et al. (2018), and a lack of explainability in machine learning models. In turn, those discrepancies between models and real-life dynamics, lead to a wider social interaction gap (see Fig. 1.1). One potential solution might lie in the utilization of transformer architectures Vaswani et al. (2017), which can provide high-level insights into the learned decision-making process, while automatically learning the dynamics from data. In addition to these limitations, current social interactions are typically validated with a relatively small number of individuals to allow for easier interpretation of the results and also to discern the effects of individuals on other individuals. However, further growing our understanding of collective behavior, inherently requires scaling up behavioral experiments to more species and large fish groups. In Chapter 4, we showed that the Deep Learning Interaction (DLI) model can be extended to two fish species that swim in a burst-and-coast manner, but future work should include 1) extending the models to more burst-and-coast species, like the guppies, and 2) implementing, adapting, extending, or validating general purpose social interaction models that can accelerate scalability studies across markedly different species.

8.1.7 Behavioral experiment, scalability limitations, and future work

These scalability limitations are also prominent in biohybrid experiments, where they typically arise from the complexity of real-time information fusion and action. Tracking multiple individuals, especially when maintaining the identities of study animals for extended experiments, is inherently difficult, but the difficulty varies based on the defined real-time requirements of the task. While studies on slower-moving animals or long-term pattern emergence might only require lower frame rates (*e.g.*, 5-10 frames per second), fish-robot experiments typically demand at least 25 frames per second for accurate animal tracking. Moreover, to mitigate physical discrepancies due to the robot's slow response times, owed to its mechanical components, that contribute to the biomimicry gap (see Fig. 1.1), processing times should be significantly reduced (*i.e.*, requiring higher frame rates), allowing the robot to adapt almost instantaneously to individual and collective dynamics.

Another prevalent limitation that is observed, is the relatively brief duration of experiments (refer to Table 1.1), which hinders advancements in understanding the evolution within long-term timescale experiments, as discussed throughout this thesis. Before the introduction of the robotic system in Chapter 5, and as evidenced in Table 1.1, many robotic systems either were not engineered for prolonged operation or were utilized in studies specifically exploring short-term timescale phenomena. Particularly in contexts of social interaction,

that influence the biomimetic efficacy of robotic systems, future research should embrace experiments of considerably longer durations. In this thesis, we managed to double the duration of the average biohybrid experiment, but posit that extending experiments even further could yield additional behavioral insights. Furthermore, the present technological state and understanding in this field could arguably accommodate studies on even longer timescales, such as exploring the disparities in social interactions between juvenile and adult fish in laboratory conditions, with and/or without robotic systems, similarly to recent behavioral research on fish in the wild Monk et al. (2023).

On a similar tangent, Table 1.1 illustrates that, for over a decade, biohybrid interaction experiments have typically involved small groups, ranging from 5 to 10 individuals. While we have identified in previous chapters that even within such modest group sizes, substantial discrepancies occasionally generate a broad biomimicry gap, it is imperative that larger groups gain a more focal role in biohybrid experiments moving forward. Consequently, it is pivotal to quantify how effectively high-fidelity social interaction models can scale to accommodate larger groups. Although this concept is occasionally probed in simulation studies Calovi et al. (2018); Wang et al. (2022); Xue et al. (2023), yielding significant biomimetic outcomes, its exploration in biohybrid experiments has been scarce.

Finally, the large majority of experiments, are conducted in notably similar, and typically symmetrical experimental setups, frequently designed to explore the binary choices of fish Bonnet et al. (2018); Chemtob et al. (2020); Bonnet et al. (2019), and less often to explore the fidelity of biohybrid systems with respect to biological ones Cazenille et al. (2017); Landgraf et al. (2016). Regardless of the setup design, its effect on the social dynamics is typically not discussed, despite its measurable effect. For instance, in square arenas Cazenille et al. (2017), individuals often demonstrate a pronounced spatial preference for corners, which proves challenging to interpret or model, thereby motivating the selection of symmetrically circular arenas, as seen in Chapter 4. In fact, fish in open circular setups exhibit much more diverse social dynamics, primarily due to the more available space compared, for example, to the circular ring setup in Chapter 3. However, even with circular arenas, we observe the setup's curvature influencing the emergent social dynamics. Additionally, we pinpoint another limitation of analytical models (explored in Chapter 4), which is the lack of environmental bound encoding. In simpler terms, in the absence of the rejection procedure (see Chapter 4), analytical burst-and-coast simulated agents (or robots in Chapter 7) would escape the confines of the wall since it is not perceptually present for individuals. Conversely, DLI simulated agents in Chapter 4 appear to have encoded sufficient implicit information during training so that, even without the rejection procedure, they consistently remain within the wall confines for a minimum of 1 h (refer to Fig. B.1) on average. Similar observations were made in experiments by Cazenille et al. (2017), and suggest that the deployed models have also learned pivotal aspects of environment interaction, enabling them to mirror spatial distributions akin to those of fish, though explicit information regarding the presence of a rejection procedure remains undisclosed.

8.2 Final words

Establishing cohesive biohybrid groups that emulate realistic social dynamics necessitates a multi-faceted approach aimed at bridging the biomimicry gap. At the fundamental level, this calls for the development of high-fidelity bidirectional models of social interaction, supplemented with benchmarking tools to confirm their generative capabilities. Transferring such models to physical systems necessitates agile robotic counterparts, in this case, a mobile robot capable of closely emulating the motion profile (*e.g.*, velocity, acceleration) of the target species. We have presented substantial strides toward bridging the biomimicry gap and creating genuinely biohybrid groups, and we have open-sourced our complete set of software and hardware tools. This allows the scientific community to reproduce, enhance, and scale up animal-robot interaction studies. While a comprehensive understanding of the governing rules of collective behavior remains elusive, the robotic tools, including those presented in this thesis, prove to be a highly promising, if not indispensable, resource for future ethological studies revolving around and beyond social interaction modeling.

A Appendix for Chapter 3 “Establishing a baseline of fish-robot dynamics”

Text 1. The experiments were conducted between the 30th of April and 5th of June 2018 at the premises of École Polytechnique Fédérale de Lausanne. The earliest recorded experiment was conducted at 9:30 and the latest at 18:30. During the experimentation time, the zebrafish were housed with a 14 to 10 day to night ratio.

Table A.1: Post hoc analysis for the average angular distance distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion.

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	-36.1312	-22.7000	-9.2688	0.0001
fish-only	Despotic	-33.2312	-19.8000	-6.3688	0.0009
fish-only	Feedback-Initiative	-26.1312	-12.7000	0.7312	0.0717
Follower	Despotic	-10.5312	2.9000	16.3312	0.9453
Follower	Feedback-Initiative	-3.4312	10.0000	23.4312	0.2226
Feedback-Initiative	Despotic	-6.3312	7.1000	20.5312	0.5258

Table A.2: Post hoc analysis for the average angular distance distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion.

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	0.2707	13.7000	27.1293	0.0435
fish-only	Despotic	2.2707	15.7000	29.1293	0.0142
fish-only	Feedback-Initiative	-14.0293	-0.6000	12.8293	0.9995
Follower	Despotic	11.4293	2.0000	15.4293	0.9810
Follower	Feedback-Initiative	-27.7293	-14.3000	-0.8707	0.0316
Feedback-Initiative	Despotic	-29.7293	-16.3000	-2.8707	0.0098

S1 Video. Recordings (segments) of the models in action. A short video depicting each model’s resulting dynamics for a duration of 60 seconds. The full videos are available upon request due to the large file size. <https://doi.org/10.1371/journal.pone.0220559.s020>

Table A.3: Post hoc analysis for the U-turn initiation success rate distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion.

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	-5.0272	4.2000	13.4272	0.5347
fish-only	Feedback-Initiative	-15.8272	-6.6000	2.6272	0.2143
Follower	Feedback-Initiative	-20.0272	-10.8000	-1.5728	0.0168

Table A.4: Post hoc analysis for the highest U-turn initiation success rate by any one individual distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion.

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	-4.3272	4.9000	14.1272	0.4269
fish-only	Feedback-Initiative	-12.9272	-3.7000	5.5272	0.6151
Follower	Feedback-Initiative	-17.8272	-8.6000	0.6272	0.0738

Table A.5: Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average outgoing TE including all agents)

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	3.9688	17.4000	30.8312	0.0048
fish-only	Despotic	11.3688	24.8000	38.2312	0.0000
fish-only	Feedback-Initiative	-6.8312	6.6000	20.0312	0.5870
Follower	Despotic	-6.0312	7.4000	20.8312	0.4896
Follower	Feedback-Initiative	-24.2312	-10.8000	2.6312	0.1644
Feedback-Initiative	Despotic	-31.6312	-18.2000	-4.7688	0.0028

Table A.6: Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average outgoing TE only for the robot’s contribution)

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	-1.5312	11.9000	25.3312	0.1036
fish-only	Despotic	12.8688	26.3000	39.7312	0.0000
fish-only	Feedback-Initiative	-6.4312	7.0000	20.4312	0.5380
Follower	Despotic	0.9688	14.4000	27.8312	0.0300
Follower	Feedback-Initiative	-18.3312	-4.9000	8.5312	0.7848
Feedback-Initiative	Despotic	-32.7312	-19.3000	-5.8688	0.0013

Table A.7: Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average outgoing TE only for the fish contribution)

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	5.0688	18.5000	31.9312	0.0023
fish-only	Despotic	-0.3312	13.1000	26.5312	0.0590
fish-only	Feedback-Initiative	-10.2312	3.2000	16.6312	0.9283
Follower	Despotic	-18.8312	-5.4000	8.0312	0.7301
Follower	Feedback-Initiative	-28.7312	-15.3000	-1.8688	0.0180
Feedback-Initiative	Despotic	-23.3312	-9.9000	3.5312	0.2307

Table A.8: Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average incoming TE including all agents – identical to outgoing case for all agents)

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	3.9688	17.4000	30.8312	0.0048
fish-only	Despotic	11.3688	24.8000	38.2312	0.0000
fish-only	Feedback-Initiative	-6.8312	6.6000	20.0312	0.5870
Follower	Despotic	-6.0312	7.4000	20.8312	0.4896
Follower	Feedback-Initiative	-24.2312	-10.8000	2.6312	0.1644
Feedback-Initiative	Despotic	-31.6312	-18.2000	-4.7688	0.0028

Table A.9: Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average incoming TE only for the robot’s contribution)

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	-0.7312	12.7000	26.1312	0.0717
fish-only	Despotic	16.0688	29.5000	42.9312	0.0000
fish-only	Feedback-Initiative	2.3688	15.8000	29.2312	0.0134
Follower	Despotic	3.3688	16.8000	30.2312	0.0072
Follower	Feedback-Initiative	-10.3312	3.1000	16.5312	0.9342
Feedback-Initiative	Despotic	-27.1312	-13.7000	-0.2688	0.0436

Table A.10: Post hoc analysis for the average outgoing transfer entropy distributions preceded by a Kruskal–Wallis test and using Tukey’s honest significant difference criterion. (Average outgoing TE only for the fish contribution)

Model	Model	Lower CI	Estimate	Upper CI	p-value
fish-only	Follower	4.6688	18.1000	31.5312	0.0030
fish-only	Despotic	6.6688	20.1000	33.5312	0.0007
fish-only	Feedback-Initiative	-9.6312	3.8000	17.2312	0.8864
Follower	Despotic	-11.4312	2.0000	15.4312	0.9810
Follower	Feedback-Initiative	-27.7312	-14.3000	-0.8688	0.0317
Feedback-Initiative	Despotic	-29.7312	-16.3000	-2.8688	0.0099




Figure A.1: **Duration of consecutive movement towards a direction (across all 10 replicates)** (fish-only case).




Figure A.2: **Duration of consecutive movement towards a direction (across all 10 replicates)**. For the fish-only case, a random individual is excluded.




Figure A.3: **Duration of consecutive movement towards a direction (across all 10 replicates).** Follower model.




Figure A.4: **Duration of consecutive movement towards a direction (across all 10 replicates).** Follower model with the robot excluded from the analysis.




Figure A.5: **Duration of consecutive movement towards a direction (across all 10 replicates).** Despotic model.




Figure A.6: **Duration of consecutive movement towards a direction (across all 10 replicates).** Despotic model with the robot excluded from the analysis.

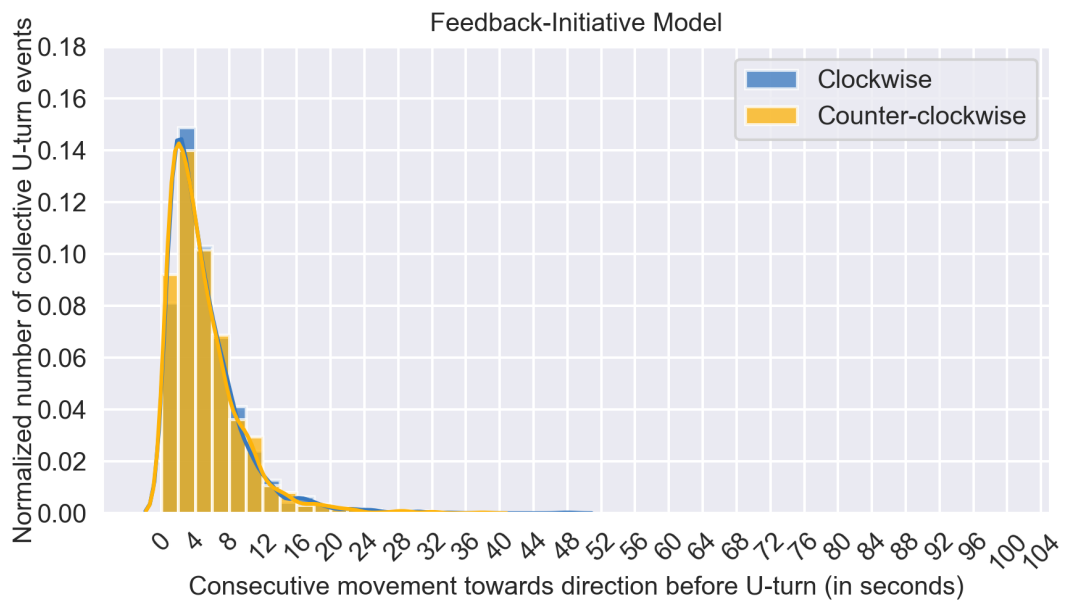


Figure A.7: **Duration of consecutive movement towards a direction (across all 10 replicates).** Feedback-Initiative model.

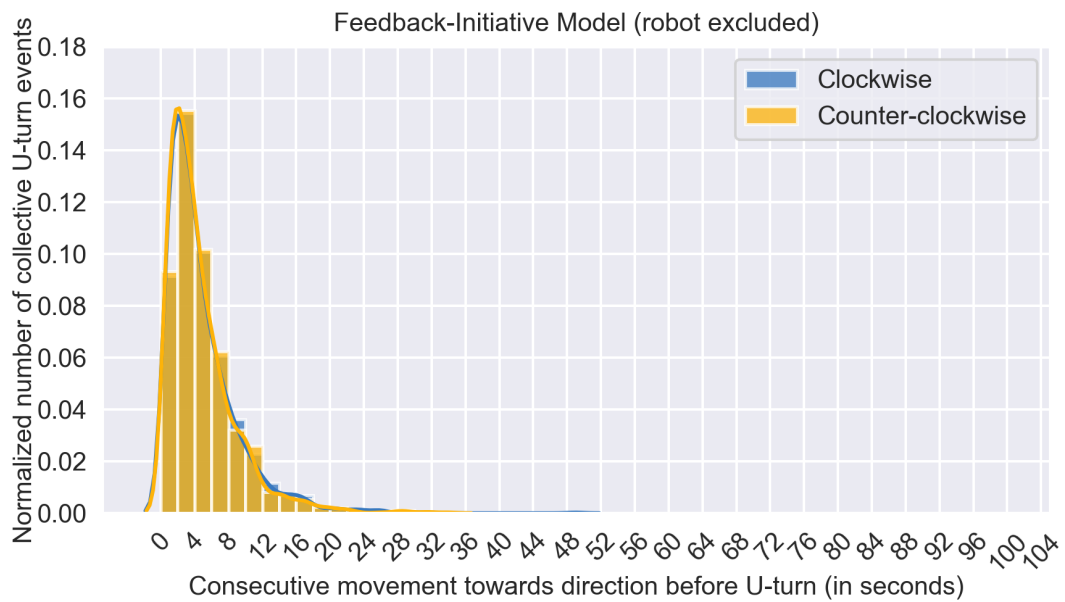


Figure A.8: **Duration of consecutive movement towards a direction (across all 10 replicates).** Feedback-Initiative model with the robot excluded from the analysis.

B Appendix for Chapter 4 “Designing and benchmarking social interaction models: from analytical to machine learning models, and from short to long timescale dynamics generation”

Table B.1: Mean and standard deviation values for the experiment.

Quantity	Pair	Leader	Follower
V (cm/s)	11.04 ± 6.69	11.54 ± 7.05	10.53 ± 6.27
r (cm)	4.58 ± 3.53	3.93 ± 3.56	5.23 ± 3.38
θ_w (°)	12.44 ± 91.40	12.07 ± 90.25	12.81 ± 92.54
d (cm)	8.16 ± 5.76	—	—
ϕ_{ij} (°)	0 ± 53.47	—	—
ψ (°)	-8.54 ± 107.47	-15.66 ± 143.82	-1.42 ± 48.12

Table B.2: Mean and standard deviation values for the ABC model.

Quantity	Pair	Leader	Follower
V (cm/s)	11.42 ± 6.13	11.62 ± 6.18	11.21 ± 6.08
r (cm)	4.93 ± 4.20	4.46 ± 4.18	5.40 ± 4.17
θ_w (°)	19.92 ± 90.12	19.48 ± 88.02	20.37 ± 92.18
d (cm)	7.92 ± 5.44	—	—
ϕ_{ij} (°)	0 ± 40.72	—	—
ψ (°)	-5.45 ± 107.41	-9.87 ± 144.03	-1.02 ± 47.85

Video 1. Examples of trajectories obtained in experiments with *H. rhodostomus* (left), for the Analytical burst-and-coast (ABC) model (center), and for the Deep Learning Interaction (DLI) model (right). This video illustrates the qualitative agreement between trajectories generated by the ABC and DLI models and experimental trajectories, while the quantitative agreement between the models and experiments is studied in detail in the Result section. The video can be downloaded at https://github.com/epfl-mobots/preddl_2023/tree/v1.0.4/Videos/S1_Video.mp4.

Table B.3: Mean and standard deviation values for the DLI model.

Quantity	Pair	Leader	Follower
V (cm/s)	11.07 ± 7.04	11.56 ± 7.44	10.58 ± 6.58
r (cm)	5.65 ± 4.42	5.24 ± 4.42	6.07 ± 4.26
θ_w (°)	6.70 ± 94.37	6.19 ± 92.15	7.21 ± 96.54
d (cm)	7.42 ± 4.38	—	—
ϕ_{ij} (°)	0 ± 54.23	—	—
ψ (°)	-4.11 ± 107.13	-7.20 ± 142.12	-1.02 ± 52.30

Table B.4: Mean and standard deviation values for D-LSTM model.

Quantity	Pair	Leader	Follower
V (cm/s)	9.17 ± 4.71	9.41 ± 4.88	8.93 ± 4.51
r (cm)	6.96 ± 5.84	6.71 ± 5.88	7.20 ± 5.79
θ_w (°)	12.40 ± 91.41	12.03 ± 90.28	12.77 ± 92.53
d (cm)	5.58 ± 3.22	—	—
ϕ_{ij} (°)	0 ± 62.59	—	—
ψ (°)	-1.20 ± 107.40	-2.44 ± 140.59	0.04 ± 57.46

Table B.5: Mean and standard deviation values for the Euclidean distance between prediction and real trajectory (focal individual).

	Focal individual		
Future time-point	0.12 s	0.24 s	0.36 s
DLI	0.12 ± 0.11	0.29 ± 0.26	0.46 ± 0.46
D-LSTM	0.11 ± 0.13	0.25 ± 0.22	0.40 ± 0.32

Table B.6: Mean and standard deviation values for the Euclidean distance between prediction and real trajectory (neighboring individual).

	Neighboring individual		
Future time-point	0.12 s	0.24 s	0.36 s
DLI	0.11 ± 0.15	0.26 ± 0.27	0.41 ± 0.39
D-LSTM	0.10 ± 0.12	0.22 ± 0.20	0.35 ± 0.39

Video 2. Example of a generated trajectory simulation for the D-LSTM model. Already at the qualitative level, the D-LSTM model fails at reproducing realistic trajectories (compare with Video 1). The video can be downloaded at https://github.com/epfl-mobots/preddl_2023/tree/v1.0.4/Videos/S2_Video.mp4.

Video 3. Example of collective behavior in a group of 5 DLI agents, *without any retraining*. For a given focal agent, we compute the predicted acceleration and noise which would be produced by each of the 4 other agents. Following Lei et al. (2020), we define the two most influential neighbors as the neighbors leading to the two highest predicted acceler-

Table B.7: Implementation details of the DLI per layer. The 5 columns correspond to the increasing layer count, the type of layer, activation function, number of inputs, and number of outputs, respectively.

N ^o Layer	Type	Activation function	N ^o Inputs	N ^o Outputs
0	LSTM	ReLU	256	128
1	Fully connected	ReLU	128	64
2	Fully connected	tanh	64	64
3	LSTM	ReLU	256	128
4	Fully connected	ReLU	128	64
5	Fully connected	tanh	64	64
6	Fully connected	None	64	4

Table B.8: Mean and standard deviation values for *D. rerio* experiment.

Quantity	Pair	Leader	Follower
V (cm/s)	11.20 ± 7.37	11.22 ± 7.47	11.19 ± 7.26
r (cm)	2.49 ± 3.42	2.45 ± 3.43	2.52 ± 3.40
θ_w (°)	-20.26 ± 90.42	-20.51 ± 88.60	-20.00 ± 92.20
d (cm)	14.49 ± 13.30	—	—
ϕ_{ij} (°)	0 ± 85.65	—	—
ψ (°)	16.45 ± 111.59	26.19 ± 143.86	6.70 ± 63.40

Table B.9: Mean and standard deviation values for DLI model (*D. rerio*).

Quantity	Pair	Leader	Follower
V (cm/s)	9.87 ± 7.37	10.00 ± 7.27	9.74 ± 6.94
r (cm)	3.56 ± 4.24	3.56 ± 4.30	3.57 ± 4.17
θ_w (°)	-15.24 ± 93.67	-15.23 ± 89.62	-15.24 ± 97.56
d (cm)	17.77 ± 13.96	—	—
ϕ_{ij} (°)	0 ± 93.87	—	—
ψ (°)	11.69 ± 107.60	15.53 ± 135.31	7.86 ± 69.40

ations. Ultimately, the focal fish speed and position are updated according to equations (9-11), using the sum of these two highest accelerations and the average predicted noise. This video illustrates the fact that, although the DLI was only trained to mimic the social interactions between pairs of fish, it produces cohesive and aligned trajectories for 5 agents, in good qualitative agreement with corresponding trajectories for 5 rummy-nose tetra Lei et al. (2020). In the future, we will address the quantitative comparison between long-term trajectories for groups of DLI agents and real fish, in particular, in connection to our robotic platform Papaspyros et al. (2019); Bonnet et al. (2016, 2018). The video can be downloaded at https://github.com/epfl-mobots/preddl_2023/tree/v1.0.4/Videos/S3_Video.mp4.

Video 4. Example of a generated trajectory simulation for the Multi-layered Perceptron Interaction (MLI) model. Already at the qualitative level, the MLI model fails at reproducing realistic trajectories (compare with Video 1). The video can be downloaded at https://github.com/epfl-robots/preddl_2023/tree/v1.0.4/Videos/S4_Video.mp4.

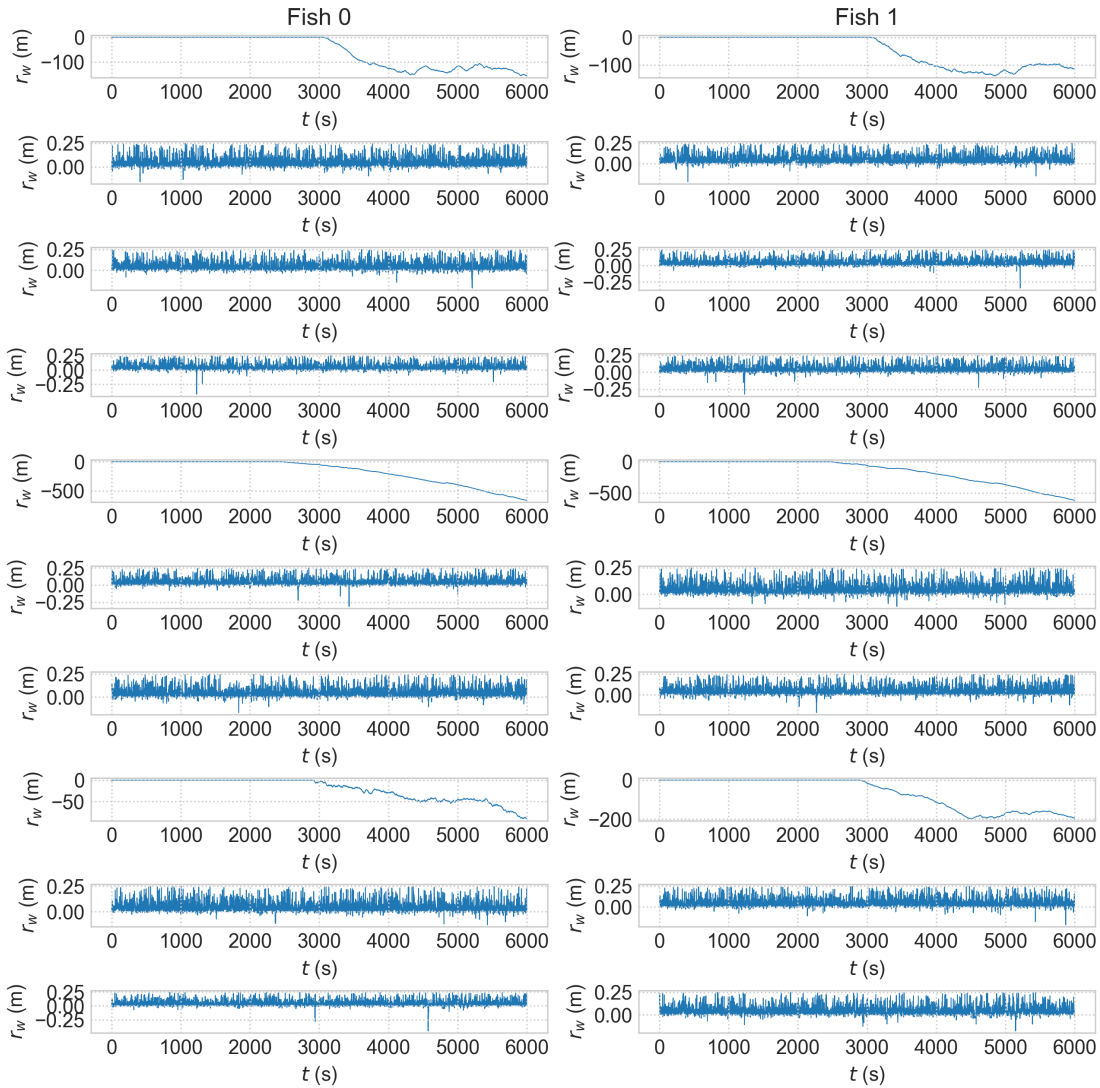


Figure B.1: DLI simulations without a rejection procedure in place. The rows depict different simulations of DLI pairs, and the columns depict the behavior of each individual during those simulations. More specifically, each sub-figure depicts the distance to the wall of each individual as a function of time. Overall, the figure depicts the underlying dynamics learned by the DLI during training, that allowed it to implicitly learn the location of the wall and stay within its confines for the majority of the experiments. Alternatively, to guarantee that the pair stays within the wall confines, a rejection procedure is put in place, resampling the predicted distributions until they yield a valid motion inside the tank walls.

Table B.10: Neural network grid. We depict the 82 neural networks considered for the social interaction modelling problem. The first column indicates the type of network (with or without memory and sequence input), the second corresponds to the number of inputs, the third indicates the use of collective information (along with individual used in all networks), the fourth depicts the sequence of layers (where L, F, D correspond to LSTM, fully connected, and dropout layers, respectively), and the fifth depicts the activation function sequence (where T, S, G, and N correspond to tanh, sigmoid, Gaussian, and no activation, respectively).

N	Memory	N ^o inputs	Neurons	Layers	Activations
1	✓	11	256, -, 128, 64, 256, -, 96, 20, 4	LDFFLDFFF	G[0.1]GGG[0.1]GGN
2	✓	11	256, 128, 64, -, 256, 96, 20, -, 4	LFFDLFFDF	GGG[0.1]GGG[0.1]N
3	✓	11	256, 128, 64, -, 128, 64, 20, -, 4	LFFDLFFDF	GGG[0.1]GGG[0.1]N
4	✓	11	256, 192, 128, 96, 32, 24, 12, -, 4	LFFFFFFDF	GGGGGGG[0.1]N
5	✓	11	256, -, 128, 64, 256, -, 96, 20, 4	LDFFLDFFF	G[0.4]GGG[0.4]GGN
6	✓	11	256, 128, 64, -, 256, 96, 20, -, 4	LFFDLFFDF	GGG[0.4]GGG[0.4]N
7	✓	11	256, 128, 64, -, 128, 68, 20, -, 4	LFFDLFFDF	GGG[0.4]GGG[0.4]N
8	✓	11	256, 192, 128, 96, 32, 24, 12, -, 4	LFFFFFFDF	GGGGGGG[0.4]N
9	✓	11	256, -, 128, 64, 256, -, 96, 20, 4	LDFFLDFFF	R[0.4]RTR[0.4]RTN
10	✓	11	256, 192, 128, 96, 32, 24, 12, -, 4	LFFFFFFDF	GGGGGGG[0.4]N
11	✓	11	256, 128, 64, 256, 96, 20, 4	LFFLFFF	RRTRRTN
12	✓	11	128, 80, 50, 128, 80, 20, 4	LFFLFFF	RRTRRTN
13	✓	11	128, 80, 50, 96, 64, 20, 4	LFFLFFF	RRTRRTN
14	✓	11	64, 32, 16, 64, 32, 16, 4	LFFLFFF	RRTRRTN
15	✓	11	256, 128, 64, 256, 96, 20, 4	LFFLFFF	SSTSSTN
16	✓	11	128, 80, 50, 128, 80, 20, 4	LFFLFFF	SSTSSTN
17	✓	11	128, 80, 50, 96, 64, 20, 4	LFFLFFF	SSTSSTN
18	✓	11	64, 32, 16, 64, 32, 16, 4	LFFLFFF	SSTSSTN
19	✓	11	256, -, 128, 64, 256, -, 96, 20, 4	LDFFLDFFF	R[0.1]RTR[0.1]RTN
20	✓	11	128, -, 80, 50, 128, -, 80, 20, 4	LDFFLDFFF	R[0.1]RTR[0.1]RTN
21	✓	11	128, -, 80, 50, 96, -, 64, 20, 4	LDFFLDFFF	R[0.1]RTR[0.1]RTN
22	✓	11	64, -, 32, 16, 64, -, 32, 16, 4	LDFFLDFFF	R[0.1]RTR[0.1]RTN
23	✓	11	256, 128, 64, 256, 96, 20, 4	FFFFFFF	SSTSSTN
24	✓	11	128, 80, 50, 128, 80, 20, 4	FFFFFFF	SSTSSTN
25	✓	11	128, 80, 50, 96, 64, 20, 4	FFFFFFF	SSTSSTN
26	✓	11	64, 32, 16, 64, 32, 16, 4	FFFFFFF	SSTSSTN
27	✓	11	256, 128, 64, 256, 96, 20, 4	LFFLFFF	GGGGGGN
28	✓	11	128, 80, 50, 128, 80, 20, 4	LFFLFFF	GGGGGGN
29	✓	11	128, 80, 50, 96, 64, 20, 4	LFFLFFF	GGGGGGN
30	✓	11	64, 32, 16, 64, 32, 16, 4	LFFLFFF	GGGGGGN
31	✓	11	256, -, 128, 64, 256, -, 96, 20, 4	LDFFLDFFF	G[0.1]GGG[0.1]GGN
32	✓	11	128, -, 80, 50, 128, -, 80, 20, 4	LDFFLDFFF	G[0.1]GGG[0.1]GGN
33	✓	11	128, -, 80, 50, 96, -, 64, 20, 4	LDFFLDFFF	G[0.1]GGG[0.1]GGN
34	✓	11	64, -, 32, 16, 64, -, 32, 16, 4	LDFFLDFFF	G[0.1]GGG[0.1]GGN
35	✓	11	512, 256, 128, 512, 256, 64, 4	LFFLFFF	RRTRRTN
36	✓	11	512, -, 256, 128, 512, -, 256, 64, 4	LDFFLDFFF	R[0.1]RTR[0.1]RTN
37	✓	11	512, 256, 128, 512, 256, 64, 4	LFFLFFF	GGTGTN
38	✓	11	512, -, 256, 128, 512, -, 256, 64, 4	LDFFLDFFF	G[0.1]GGG[0.1]GGN
39	✓	11	512, 256, 128, 64, 256, 128, 64, 32, 4	LFFFFFFF	RRTRRTN
40	✓	11	512, -, 256, 128, 64, 256, -, 128, 64, 32, 4	LDFFLDFFFF	R[0.1]RTR[0.1]RTN
41	✓	11	512, 256, 128, 64, 256, 128, 64, 32, 4	LFFFFFFF	GGGTGTN
42	✓	11	512, -, 256, 128, 64, 256, -, 128, 64, 32, 4	LDFFLDFFFF	G[0.1]GGT[0.1]GGTN
43	✓	11	512, -, 256, 128, 64, 256, -, 128, 64, 32, 12, 4	LDFFLDFFFLF	R[0.1]RTR[0.1]RTRN
44	✓	11	512, -, 256, 128, 64, 256, -, 128, 64, 32, 12, 4	LDFFLDFFFLF	G[0.1]GGT[0.1]GGTN

N	Memory	N° inputs	Neurons	Layers	Activations
45	✓	11	256, 64, 4	FFF	SSN
46	✓	11	128, 48, 4	FFF	SSN
47	✓	11	96, 32, 4	FFF	SSN
48	✓	11	64, 16, 4	FFF	SSN
49	✓	11	32, 8, 4	FFF	SSN
50	✓	11	11, 6, 4	FFF	SSN
51	✓	11	11, 256, 128, 4	LFFF	SSSN
52	✓	11	11, 128, 64, 4	LFFF	SSSN
53	✓	11	11, 96, 48, 4	LFFF	SSSN
54	✓	11	11, 64, 32, 4	LFFF	SSSN
55	✓	11	11, 32, 8, 4	LFFF	SSSN
56	✓	11	11, 256, -, 128, -, 4	LFDFDF	SS[0.1]S[0.1]N
57	✓	11	11, 128, -, 64, -, 4	LFDFDF	SS[0.1]S[0.1]N
58	✓	11	11, 96, -, 48, -, 4	LFDFDF	SS[0.1]S[0.1]N
59	✓	11	11, 64, -, 32, -, 4	LFDFDF	SS[0.1]S[0.1]N
60	✓	11	11, 32, -, 16, -, 4	LFDFDF	SS[0.1]S[0.1]N
61	✓	11	256, 96, -, 32, -, 4	FFDFDF	SS[0.1]S[0.1]N
62	✓	11	128, 64, -, 16, -, 4	FFDFDF	SS[0.1]S[0.1]N
63	✓	11	96, 48, -, 16, -, 4	FFDFDF	SS[0.1]S[0.1]N
64	✓	11	64, 32, -, 16, -, 4	FFDFDF	SS[0.1]S[0.1]N
65	✓	11	32, 16, -, 8, -, 4	FFDFDF	SS[0.1]S[0.1]N
66	✓	11	11, 8, -, 6, -, 4	FFDFDF	SS[0.1]S[0.1]N
67	✓	11	256, 96, -, 32, -, 4	LFDFDF	SS[0.1]S[0.1]N
68	✓	11	128, 64, -, 16, -, 4	LFDFDF	SS[0.1]S[0.1]N
69	✓	11	96, 48, -, 16, -, 4	LFDFDF	SS[0.1]S[0.1]N
70	✓	11	64, 32, -, 16, -, 4	LFDFDF	SS[0.1]S[0.1]N
71	✓	11	32, 16, -, 8, -, 4	LFDFDF	SS[0.1]S[0.1]N
72	✓	11	11, 8, -, 6, -, 4	LFDFDF	SS[0.1]S[0.1]N
73	✓	11	256, 96, 32, 4	FFFF	SSSN
74	✓	11	128, 64, 16, 4	FFFF	SSSN
75	✓	11	96, 48, 16, 4	FFFF	SSSN
76	✓	11	64, 32, 16, 4	FFFF	SSSN
77	✓	11	32, 16, 8, 4	FFFF	SSSN
78	✓	11	11, 8, 6, 4	FFFF	SSSN
79	✓	11	256, 96, 32, 4	LFFF	SSSN
80	✓	11	128, 64, 16, 4	LFFF	SSSN
81	✓	11	96, 48, 16, 4	LFFF	SSSN
82	✓	11	64, 32, 16, 4	LFFF	SSSN
83	✓	11	32, 16, 8, 4	LFFF	SSSN
84	✓	11	11, 8, 6, 4	LFFF	SSSN
85	✓	11	256, 128, 64, 32, 4	FFFFD	SSSSN
86	✓	11	128, 64, 32, 16, 4	FFFFF	SSSSN
87	✓	11	96, 48, 24, 12, 4	FFFFF	SSSSN
88	✓	11	64, 32, 16, 8, 4	FFFFF	SSSSN
89	✓	11	32, 16, 8, 4, 4	FFFFF	SSSSN
90	✓	11	11, 8, 6, 4, 4	FFFFF	SSSSN
91	✓	11	256, 128, 64, 32, 4	FFFFD	SSSSN
92	✓	11	128, 64, 32, 16, 4	LFFFF	SSSSN
93	✓	11	96, 48, 24, 12, 4	LFFFF	SSSSN
94	✓	11	64, 32, 16, 8, 4	LFFFF	SSSSN
95	✓	11	32, 16, 8, 4, 4	LFFFF	SSSSN
96	✓	11	11, 8, 6, 4, 4	LFFFF	SSSSN

C Appendix for Chapter 5 “Introducing a novel biohybrid interaction framework”

To quantify the (dis)similarity of the PDF presented in the results section, we consider the Hellinger distance $D(F|G)$ between two PDF $F(x)$ and $G(x)$ for the same observable x Beran (1977); Basu et al. (1997):

$$\begin{aligned} H(F|G) &= \frac{1}{2} \int \left(\sqrt{F(x)} - \sqrt{G(x)} \right)^2 dx \\ &= 1 - \int \sqrt{F(x)} \sqrt{G(x)} dx, \end{aligned} \quad (\text{C.1})$$

where we have used the normalization of the PDF, $\int F(x) dx = \int G(x) dx = 1$, to obtain the last equality. The first definition of $H(F|G)$ makes clear that it measures the overall difference between $F(x)$ and $G(x)$, while the second equivalent definition has a nice interpretation in terms of the *overlap* of both PDF. Indeed, the second definition measures the distance from unity of the scalar product of $\sqrt{F(x)}$ and $\sqrt{G(x)}$ seen as vectors of unit Euclidean norm (a consequence of the normalization, $\int \sqrt{F(x)}^2 dx = 1$).

The Hellinger distance is zero if and only if $F(x) = G(x)$, and it always satisfies $H(F|G) \leq 1$. The upper bound $H(F|G) = 1$ is reached whenever the supports of the two PDF are not intersecting, so that $F(x) \times G(x) = 0$, for all values of x . In practice, a value of $H(F|G) \geq 0.2$ points to the two PDF being markedly dissimilar.

In Tables C.1 to C.4 below, we have computed the Hellinger distance for the PDF for the two open-loop experiments (circular and eightfold rose trajectory for the DS and BM lures), for a fish or a BM lure alone in the tank, and for the closed-loop experiments for pairs of individuals (2 fish; fish 0 and DS lure; fish 0 and BM lure).

Video 1. Video segments of experiments and overview of BOBI. <https://doi.org/10.1109/ACCESS.2023.3290960/mm1>

Table C.1: Hellinger distance for pairs of individuals (the LureBot follows a passive circular trajectory).

Agent	Quantity	Hellinger distance
Fish <i>vs</i> DS lure	V	0.591
	α	0.296
Fish <i>vs</i> BM lure	V	0.193
	α	0.202

Table C.2: Hellinger distance for pairs of individuals (the LureBot follows a passive eightfold rose trajectory).

Agents	Quantity	Hellinger distance
Fish <i>vs</i> DS lure	V	0.432
	α	0.172
Fish <i>vs</i> BM lure	V	0.352
	α	0.134

Table C.3: Hellinger distance for single individuals (the LureBot is commanded by the behavioral model).

Agent	Quantity	Hellinger distance
Fish <i>vs</i> BM lure	V	0.172
	α	0.086

Table C.4: Hellinger distance for pairs of individuals (the LureBot is commanded by the behavioral model).

Agents	Quantity	Hellinger distance
2 fish	V	0.06
	α	0.076
Fish <i>vs</i> DS lure	V	0.087
	α	0.135
Fish <i>vs</i> BM lure	V	0.138
	α	0.202
Pair of Fish <i>vs</i> Pair of DS and fish	d	0.455
Pair of Fish <i>vs</i> Pair of BM and fish	d	0.411

C.1 Scaling up from the FishBot v4.4 to the LureBot

While attempting to transfer the models presented in Chapter 4 back to reality, we initially relied on the FishBot v4.4 Bonnet et al. (2012, 2014) and CATS Bonnet et al. (2017) (see Sec. 3.2.4). However, as extensively discussed in Chapter 5, the FishBot v4.4 lacked the agility to reproduce complex social interaction in real time. In fact, we implemented the same model as in Sec. 5.6.3, to find that the robot lacks the combined speed and acceleration to follow the model.

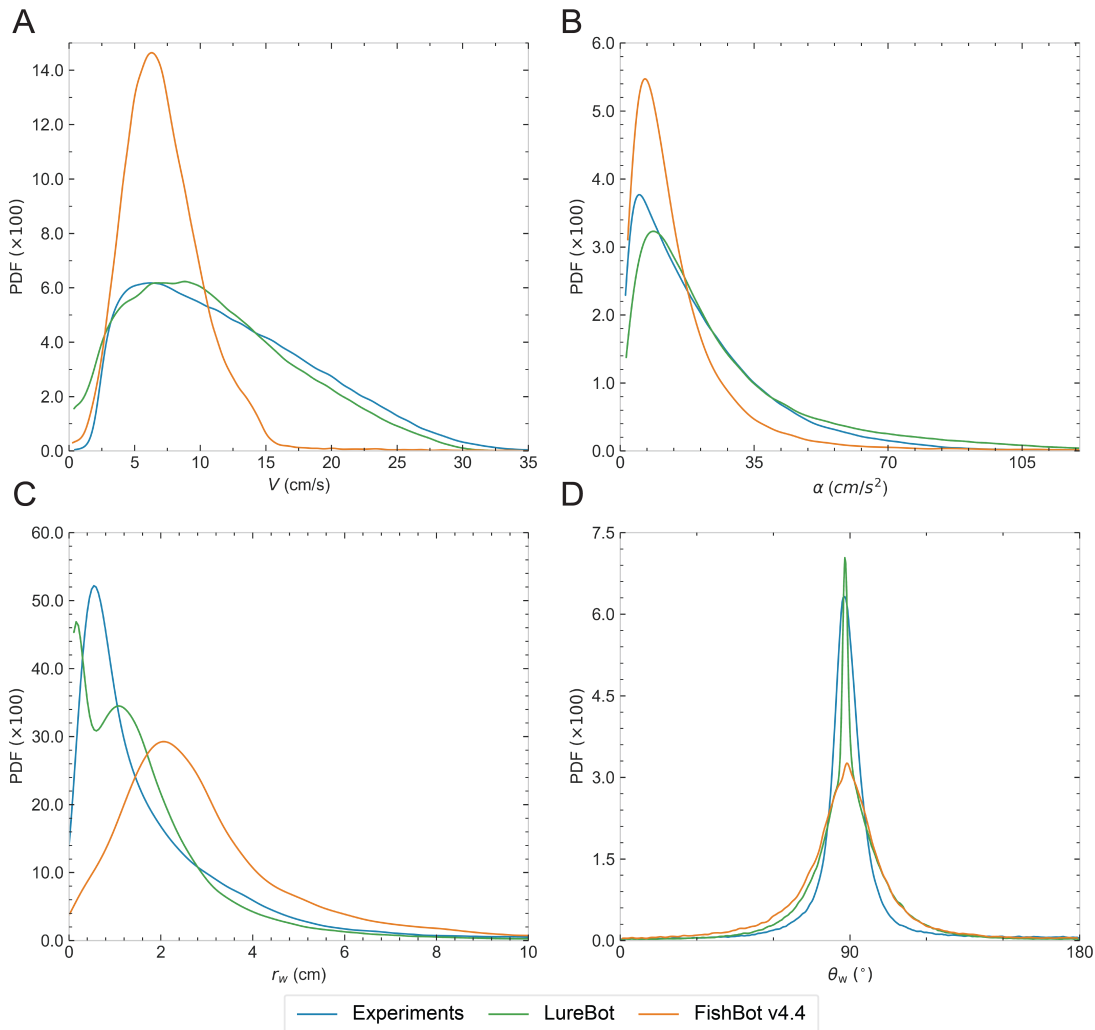


Figure C.1: **LureBot versus FishBot v4.4 performance when interacting biomimetically with the tank wall.** Probability density functions (PDFs) for: **A.** velocity V (cm/s), **B.** acceleration α (cm/s²), **C.** distance to wall d_w (cm), and **D.** angle of incidence to the wall θ_w (°).

In Fig. C.1 we compare the performance of the FishBot compared to that of the LureBot, with respect to the dynamics exhibited by a single *H. rhodostomus*. The primary weakness of the FishBot is already apparent in Fig. C.1A, which shows the PDF of velocity for each agent. Whereas the two robots typically move with a speed of 6.15 cm/s, the FishBot is hard-limited by

its top velocity (approximately 20 cm/s). We note that although the FishBot can in theory move with up to 50 cm/s, the small size of the setup and the robot's narrow design causes severely unstable movement with velocities greater than 25 cm/s when the robot needs to perform rapid direction changes. The difference between the two robots is smaller with respect to their typical acceleration (see Fig. C.1B), albeit still markedly different to that of the LureBot. Both robots typically swim with an acceleration of 4 cm/s^2 when being commanded by the same biomimetic model. However, the FishBot is once again limited by its ability to accelerate fast (see the narrower tail of the PDF in Fig. C.1). The FishBot's inability to move close to the wall (see Fig. C.1C), similarly to the LureBot and the fish in the experiment, is owed to the weaker magnetic coupling between robot and lure. Due to the FishBot's lighter frame (which weights approximately three times less than the LureBot), it would often detach when in contact or collision with the setup wall, even if such a contact was weak. Therefore, we often have to settle with commanding the robot to farther from the wall. Conversely, the LureBot's coupling to the lure can withstand a significantly increased amount of collision strength before we observe a detachment of the lure. Finally, both robots perform marginally different when it comes to their ability to align with the wall during their movement, as show in Fig. C.1. The LureBot can typically move more consistently (see the sharpness of the peak in Fig. C.1) compared to the FishBot, but both perform marginally different (with more angular variability – thicker PDF tails in Fig. C.1) from the fish in the experiment.

D Appendix for Chapter 6 “Evaluating the biomimicry gap in biohybrid pairs driven by a deep learning model”

Video 1. Video segments of the experiment, DLI-SP, and DLI-BP (<https://doi.org/10.5281/zenodo.8253256>).

D.1 DLiv2 Supplementary figures (standard deviation)

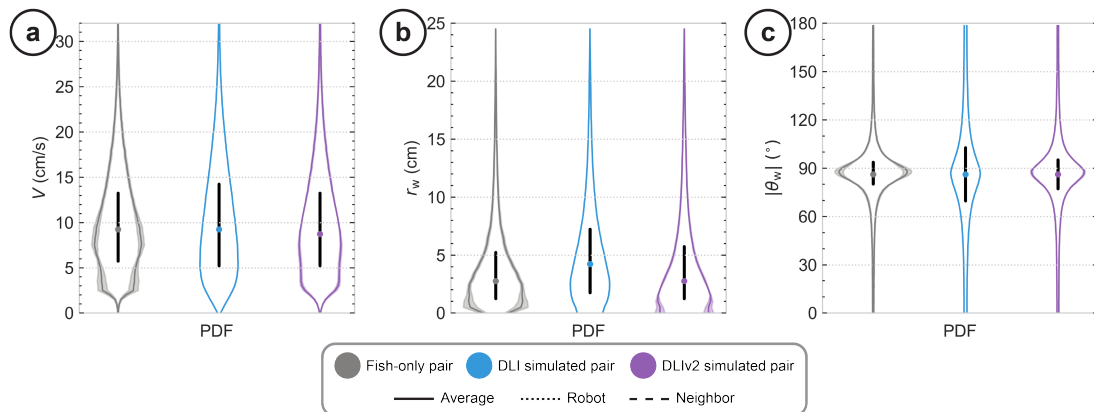


Figure D.1: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Angle of incidence to the wall θ_w probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the DLI simulated pairs, and the DLiv2 simulated pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The shaded areas correspond to the standard deviation.

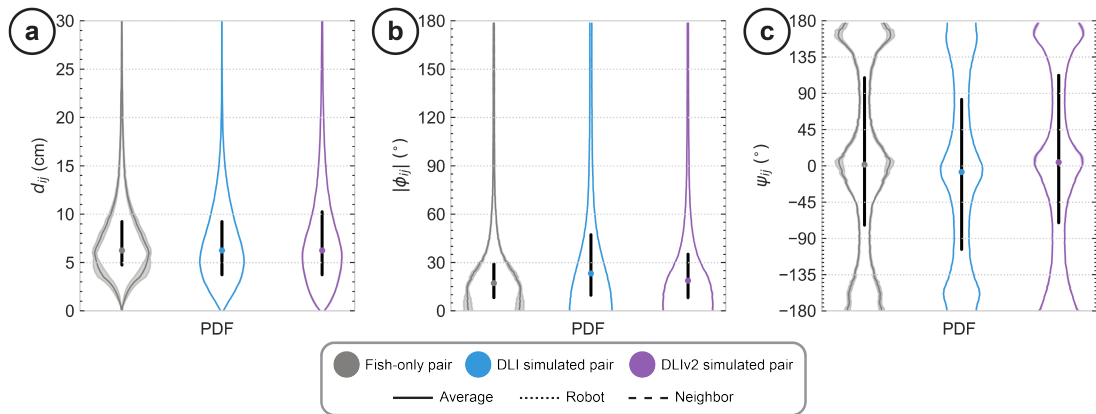


Figure D.2: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles $|\phi_{ij}|$ probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, DLI simulated pairs and DLI biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The shaded areas correspond to the standard deviation.

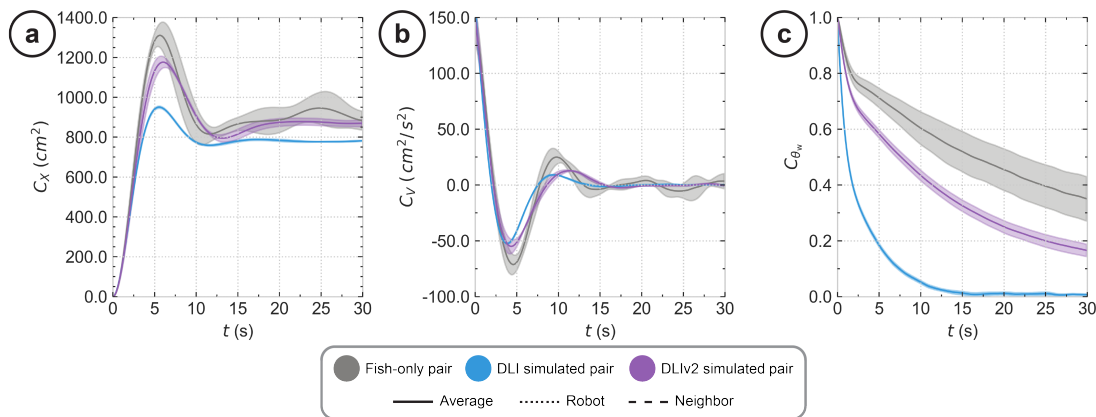


Figure D.3: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, DLI simulated pairs and DLI biohybrid pairs, respectively. The shaded areas correspond to the standard deviation.

D.2 DLiv2 Supplementary figures (confidence interval)

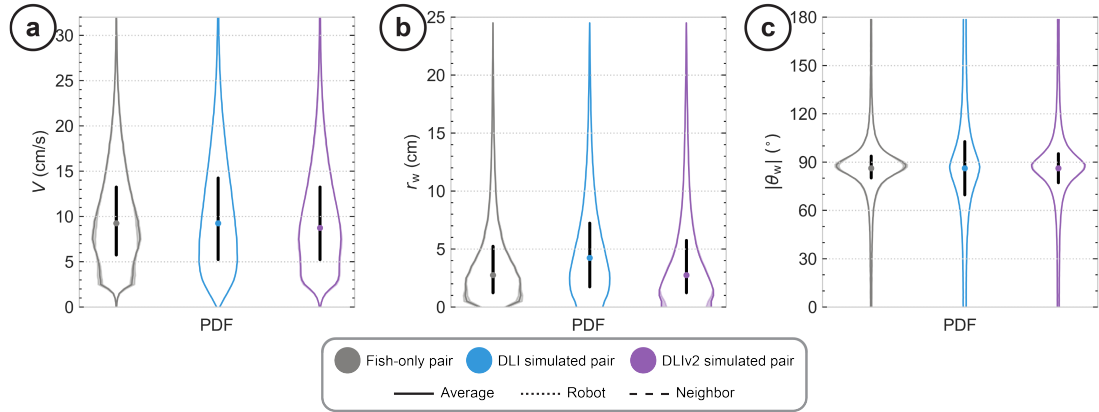


Figure D.4: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Angle of incidence to the wall θ_w probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the DLI simulated pairs, and the DLI biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The shaded areas correspond to the 68% confidence interval.

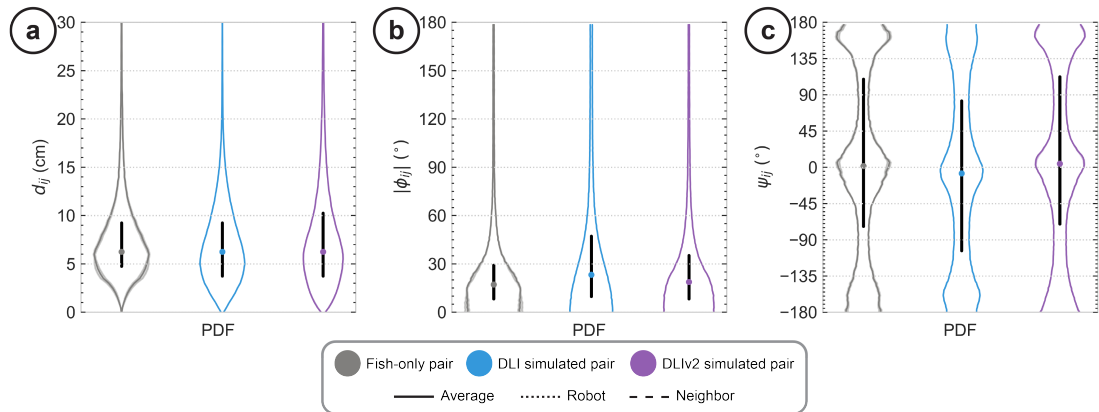


Figure D.5: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles $|\phi_{ij}|$ probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, DLI simulated pairs and DLI biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The shaded areas correspond to the 68% confidence interval.

D.3 DLI biohybrid pair supplementary figures (confidence interval)

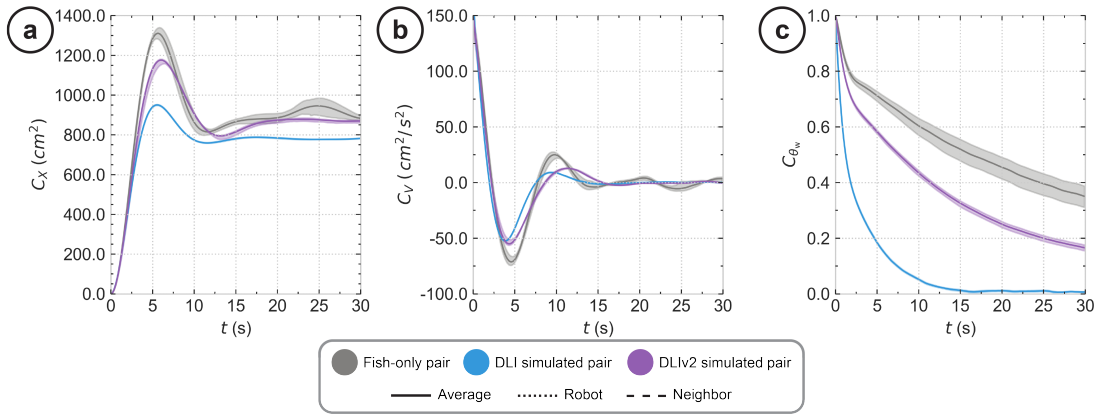


Figure D.6: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, DLI simulated pairs and DLI biohybrid pairs, respectively. The shaded areas correspond to the 68% confidence interval.

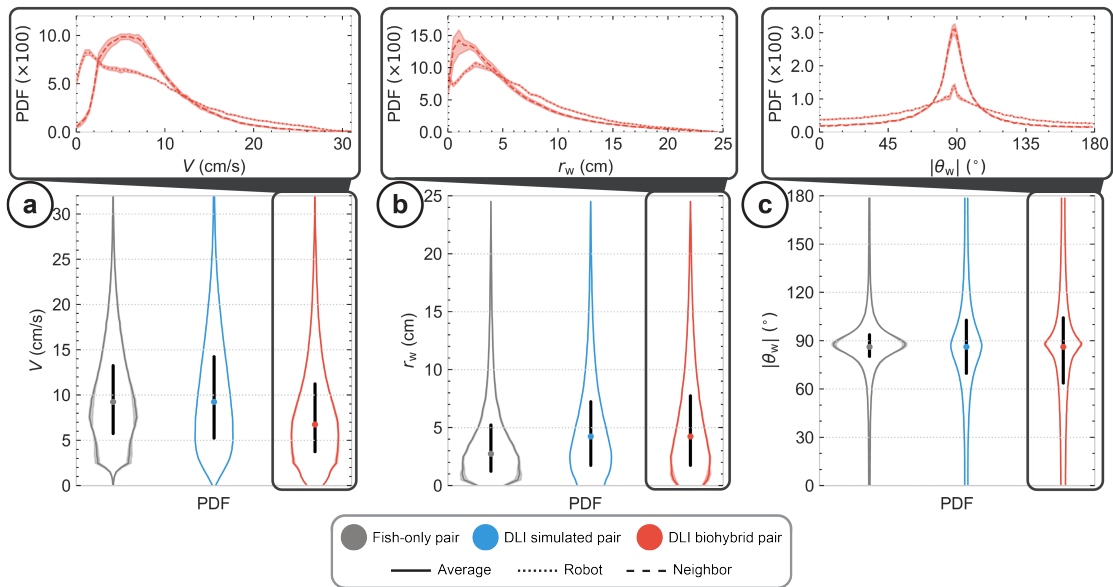


Figure D.7: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Angle of incidence to the wall θ_w probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the DLI simulated pairs, and the DLI biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the DLI biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

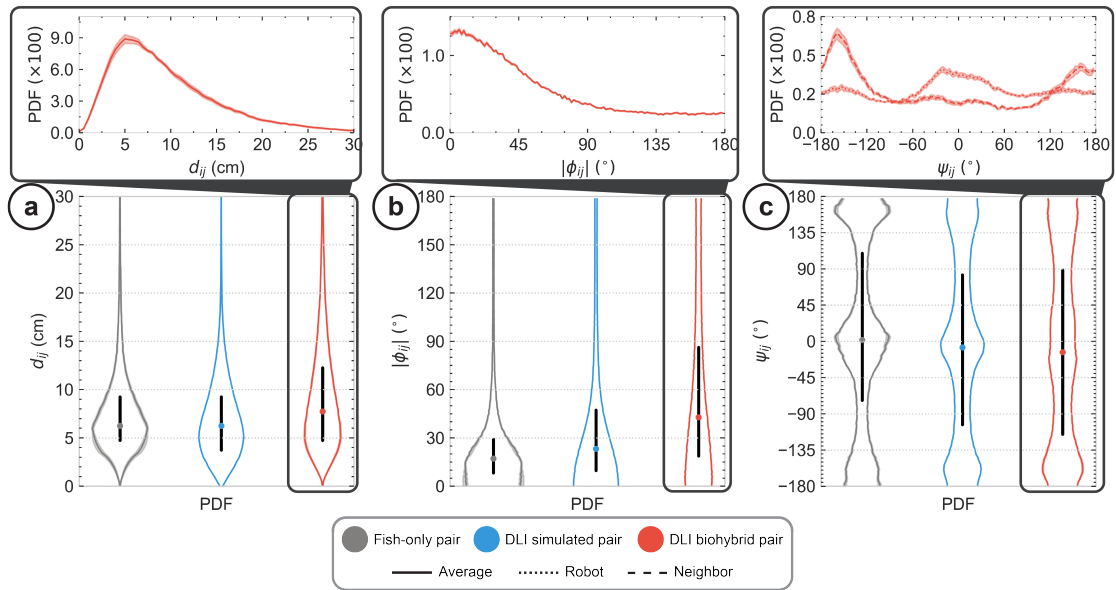


Figure D.8: **Instantaneous collective quantities.** **(a)** Interindividual distance d_{ij} probability density function. **(b)** Difference in heading angles ϕ_{ij} probability density function. **(c)** Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, DLI simulated pairs and DLI biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the DLI biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

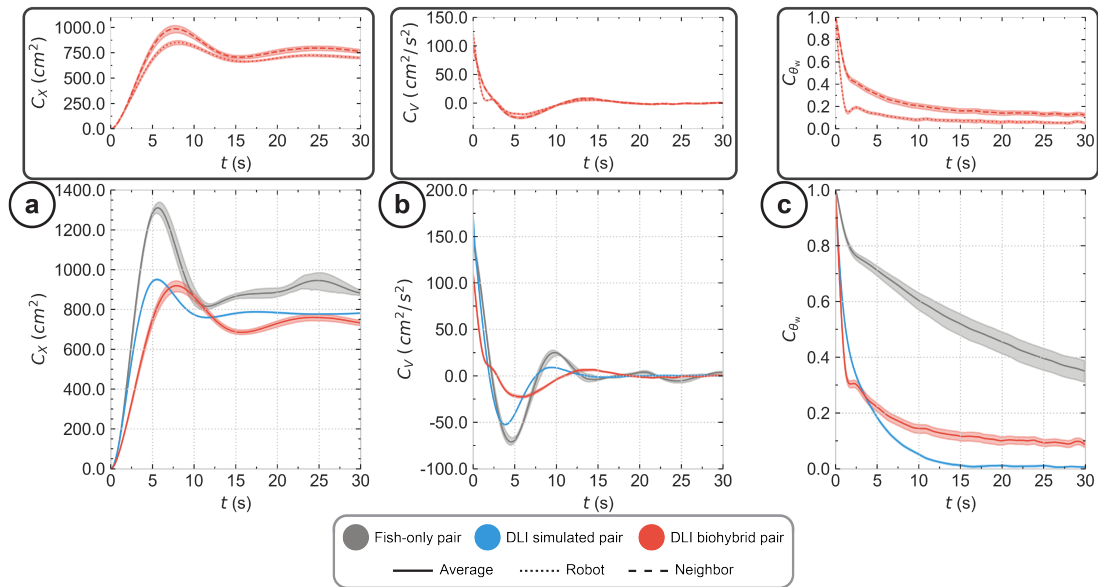


Figure D.9: **Temporal correlation quantities.** **(a)** Mean squared displacement $C_X(t)$. **(b)** Velocity autocorrelation $C_V(t)$. **(c)** Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, DLI simulated pairs and DLI biohybrid pairs, respectively. Dotted, dashed and solid lines indicate the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

D.4 Supplementary Tables

Table D.1: **Means and standard deviations.** For the case of fish-only experiments, DLI simulated pairs (DLI-SP), and biohybrid pairs (DLI-SP), we report the mean and the standard deviation (SD) of the 6 observables introduced in Section 6.2.3, along with their respective standard error. The speed V is given in cm/s, the distances r_w and d_{ij} are given in cm, and the angles $|\theta_w|$, $|\phi_{ij}|$, and ψ_{ij} are in degrees. Note the small standard error in the case of the (DLI-SP) resulting from extensive simulations (16.6h long, almost twice the amount of data collected for other cases) and the fact that the 2 agents are statistically identical. For the biohybrid experiments, we report the mean and SD for V , r_w , and $|\theta_w|$, averaged over the fish and the LureBot, as well as for each of them. Finally, we present the corresponding results for a DLI model retrained on the present fish experiments (DLiv2-SP).

Pair	Quantity	Mean	Standard deviation
Fish-only	V	10.50 ± 0.60	5.73 ± 0.36
	r_w	4.39 ± 0.43	3.86 ± 0.22
	$ \theta_w $	87.42 ± 0.39	21.91 ± 1.46
	d_{ij}	8.05 ± 0.71	5.11 ± 0.43
	$ \phi_{ij} $	26.72 ± 1.91	29.81 ± 1.24
	ψ_{ij}	7.96 ± 4.73	108.98 ± 1.19
DLI-SP	V	11.06 ± 0.04	7.04 ± 0.02
	r_w	5.66 ± 0.03	4.42 ± 0.03
	$ \theta_w $	88.07 ± 0.06	34.55 ± 0.16
	d_{ij}	7.43 ± 0.03	4.38 ± 0.04
	$ \phi_{ij} $	38.06 ± 0.19	38.63 ± 0.17
	ψ_{ij}	-4.11 ± 0.33	107.13 ± 0.06
DLI-BP	V	8.60 ± 0.22	5.93 ± 0.12
	r_w	6.05 ± 0.25	4.76 ± 0.06
	$ \theta_w $	86.44 ± 0.17	38.07 ± 0.73
	d_{ij}	9.96 ± 0.48	6.27 ± 0.33
	$ \phi_{ij} $	58.60 ± 0.91	48.38 ± 0.24
	ψ_{ij}	-7.42 ± 4.16	110.41 ± 0.51
DLI-BP (fish)	V	8.44 ± 0.26	5.13 ± 0.21
	r_w	5.54 ± 0.35	4.54 ± 0.09
	$ \theta_w $	87.46 ± 0.19	32.76 ± 1.25
DLI-BP (robot)	V	8.74 ± 0.16	6.62 ± 0.12
	r_w	6.59 ± 0.15	4.91 ± 0.05
	$ \theta_w $	85.42 ± 0.24	42.78 ± 0.79
DLiv2-SP	V	10.53 ± 0.48	6.18 ± 0.28
	r_w	4.64 ± 0.23	4.37 ± 0.05
	$ \theta_w $	87.56 ± 0.11	26.47 ± 0.47
	d_{ij}	8.39 ± 0.07	6.15 ± 0.11
	$ \phi_{ij} $	30.54 ± 0.30	33.11 ± 0.29
	ψ_{ij}	11.72 ± 0.87	109.08 ± 0.19

Table D.2: **Hellinger distances.** We exploit the Hellinger distance between two PDF to compare the PDF of the 6 observables introduced in Section 6.2.3, for fish-only experiments, DLI simulated pairs (DLI-SP), and biohybrid pairs (DLI-SP).

Pair	Quantity	Hellinger distance
Fish-only <i>vs</i> DLI-SP	V	0.09
	r_w	0.13
	$ \theta_w $	0.23
	d_{ij}	0.12
	$ \phi_{ij} $	0.14
	ψ_{ij}	0.09
Fish-only <i>vs</i> DLI-BP	V	0.18
	r_w	0.15
	$ \theta_w $	0.25
	d_{ij}	0.16
	$ \phi_{ij} $	0.30
	ψ_{ij}	0.15
DLI-SP <i>vs</i> DLI-BP	V	0.14
	r_w	0.04
	$ \theta_w $	0.04
	d_{ij}	0.18
	$ \phi_{ij} $	0.17
	ψ_{ij}	0.07
Fish-only <i>vs</i> DLIv2-SP	V	0.05
	r_w	0.08
	$ \theta_w $	0.08
	d_{ij}	0.14
	$ \phi_{ij} $	0.06
	ψ_{ij}	0.04

Table D.3: Statistical significance tests for the experiment, DLI-SP, and DLI-BP. We ran a Kruskal-Wallis test to measure whether the distributions differ significantly and a Tukey's honestly significant difference (T-HSD) test to test which distributions differ significantly.

Quantity	Kruskal-Wallis	Pair	T-HSD
V	$p = 0.80, \chi^2 = 0.45$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLI-BP DLI-SP <i>vs</i> DLI-BP	$p = 0.83$ $p = 1.00$ $p = 0.83$
r_w	$p = 0.41, \chi^2 = 1.77$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLI-BP DLI-SP <i>vs</i> DLI-BP	$p = 0.62$ $p = 0.40$ $p = 0.94$
$ \theta_w $	$p = 0.03, \chi^2 = 6.73$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLI-BP DLI-SP <i>vs</i> DLI-BP	$p = 0.15$ $p = 0.81$ $p = 0.03$
d_{ij}	$p = 0.03, \chi^2 = 6.73$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLI-BP DLI-SP <i>vs</i> DLI-BP	$p = 0.15$ $p = 0.81$ $p = 0.03$
$ \phi_{ij} $	$p = 0.00, \chi^2 = 53.38$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLI-BP DLI-SP <i>vs</i> DLI-BP	$p = 0.00$ $p = 0.00$ $p = 0.00$
ψ_{ij}	$p = 0.06, \chi^2 = 5.56$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLI-BP DLI-SP <i>vs</i> DLI-BP	$p = 0.35$ $p = 0.05$ $p = 0.60$

Table D.4: Statistical significance tests for the experiment, DLI-SP, and DLIv2-SP. We ran a Kruskal-Wallis test to measure whether the distributions differ significantly and a Tukey's honestly significant difference (T-HSD) test to test which distributions differ significantly.

Quantity	Kruskal-Wallis	Pair	T-HSD
V	$p = 0.89, \chi^2 = 0.24$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLIv2-SP DLI-SP <i>vs</i> DLIv2-SP	$p = 0.88$ $p = 0.94$ $p = 0.99$
r_w	$p = 0.67, \chi^2 = 0.80$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLIv2-SP DLI-SP <i>vs</i> DLIv2-SP	$p = 0.64$ $p = 0.87$ $p = 0.92$
$ \theta_w $	$p = 0.00, \chi^2 = 44.58$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLIv2-SP DLI-SP <i>vs</i> DLIv2-SP	$p = 0.00$ $p = 0.03$ $p = 0.00$
d_{ij}	$p = 0.01, \chi^2 = 8.47$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLIv2-SP DLI-SP <i>vs</i> DLIv2-SP	$p = 0.11$ $p = 0.68$ $p = 0.01$
$ \phi_{ij} $	$p = 0.00, \chi^2 = 28.15$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLIv2-SP DLI-SP <i>vs</i> DLIv2-SP	$p = 0.00$ $p = 0.03$ $p = 0.01$
ψ_{ij}	$p = 0.09, \chi^2 = 4.87$	Fish-only <i>vs</i> DLI-SP Fish-only <i>vs</i> DLIv2-SP DLI-SP <i>vs</i> DLIv2-SP	$p = 0.08$ $p = 0.81$ $p = 0.28$

E Appendix for Chapter 7 “Investigating the biomimicry gap with varying sizes of biohybrid groups driven by a high-fidelity model”

E.1 Experiments with a single agent

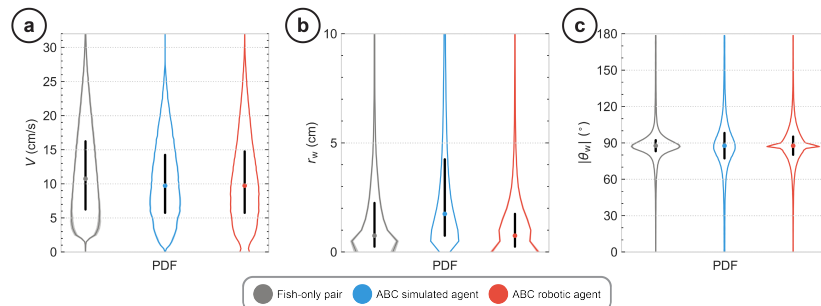


Figure E.1: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Absolute angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the ABC simulated agents, and the ABC robotic agents, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The shaded areas correspond to the 68% confidence interval.

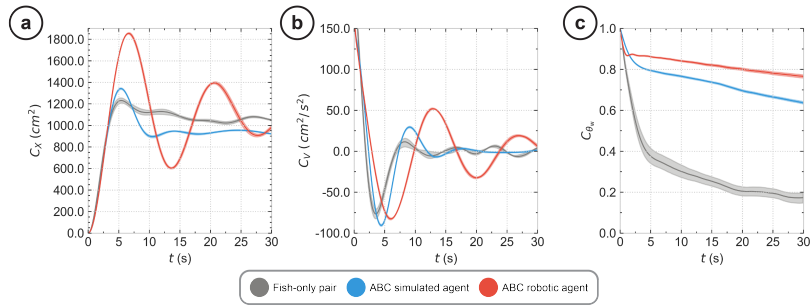


Figure E.2: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, ABC simulated agents and ABC robotic agents, respectively. Solid lines indicate the average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

E.2 Experiments with agent pairs

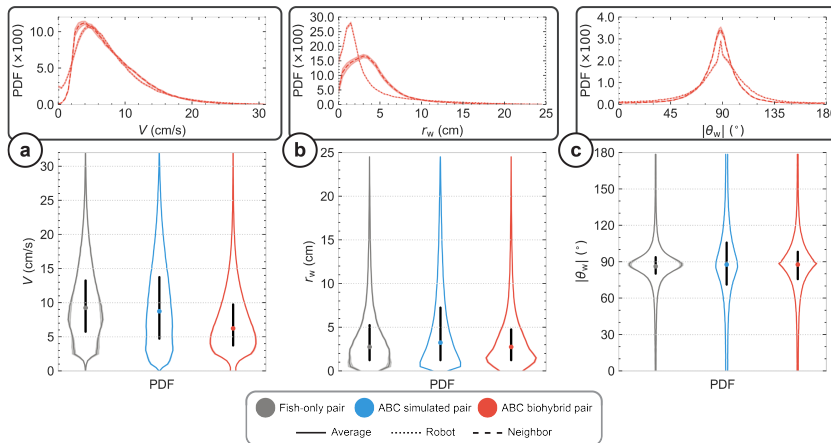


Figure E.3: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Absolute angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the ABC simulated pairs, and the ABC biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

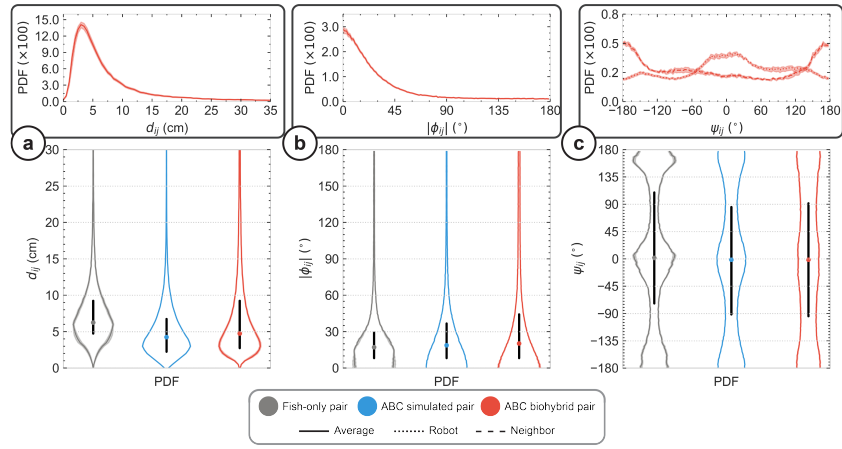


Figure E.4: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles $|\phi_{ij}|$ probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, ABC simulated pairs and ABC biohybrid pairs, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

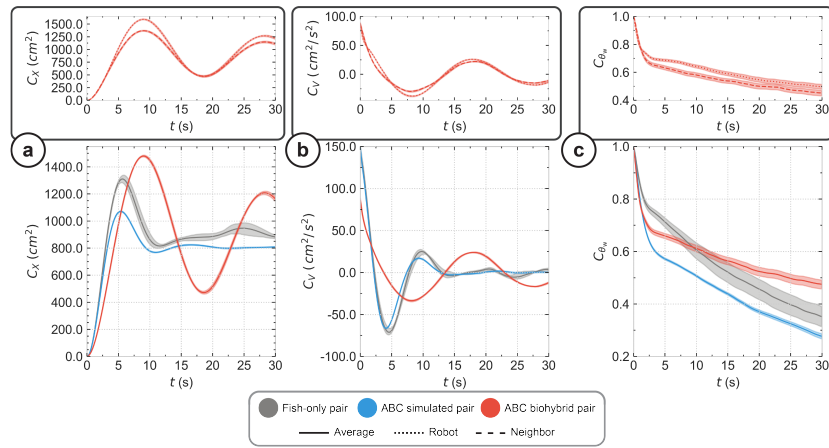


Figure E.5: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, ABC simulated pairs and ABC biohybrid pairs, respectively. Dotted, dashed and solid lines indicate the robot, neighbor and average agent distributions, respectively. The inset plots depict the distributions of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

E.3 Experiments with groups of five agents

E.3.1 Experiments with groups of five agents (1 agent in perception field)

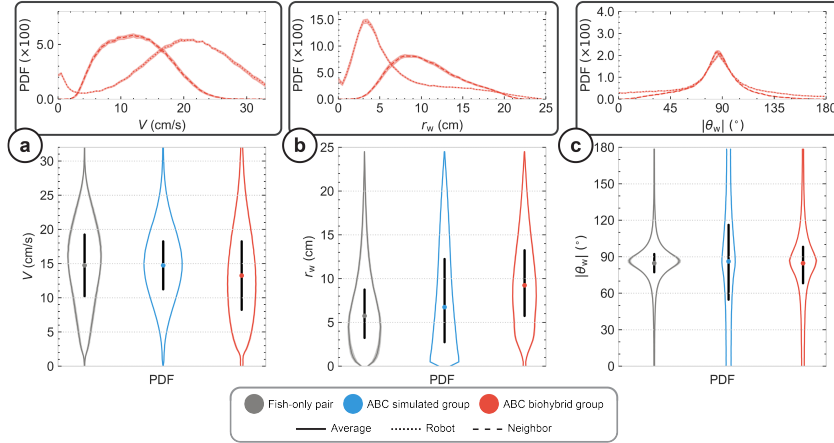


Figure E.6: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Absolute angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the ABC simulated groups, and the ABC biohybrid groups, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

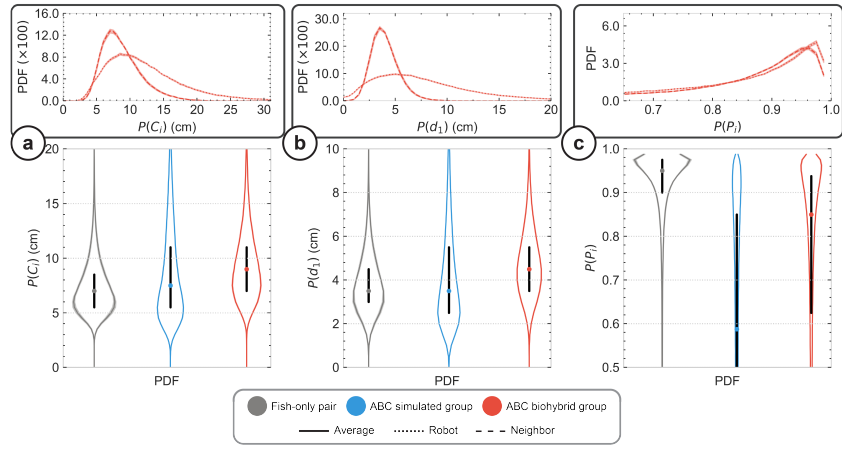


Figure E.7: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles $|\phi_{ij}|$ probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, ABC simulated groups and ABC biohybrid groups, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

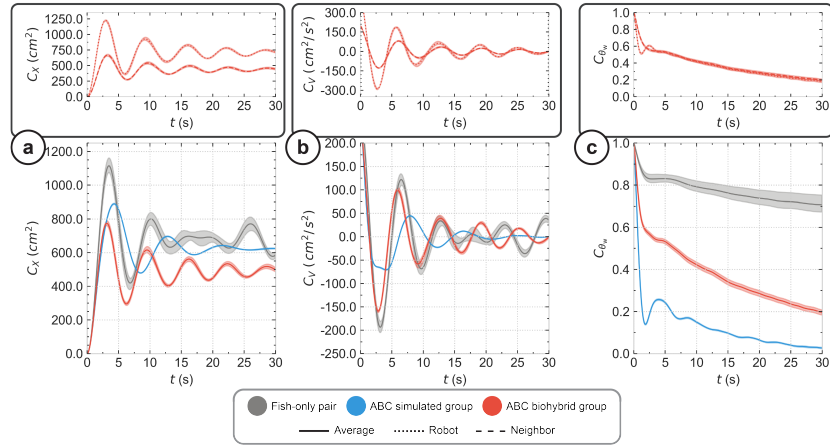


Figure E.8: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, ABC simulated groups and ABC biohybrid groups, respectively. Dotted, dashed and solid lines indicate the robot, neighbor and average agent distributions, respectively. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

E.3.2 Experiments with groups of five agents (2 agents in perception field)

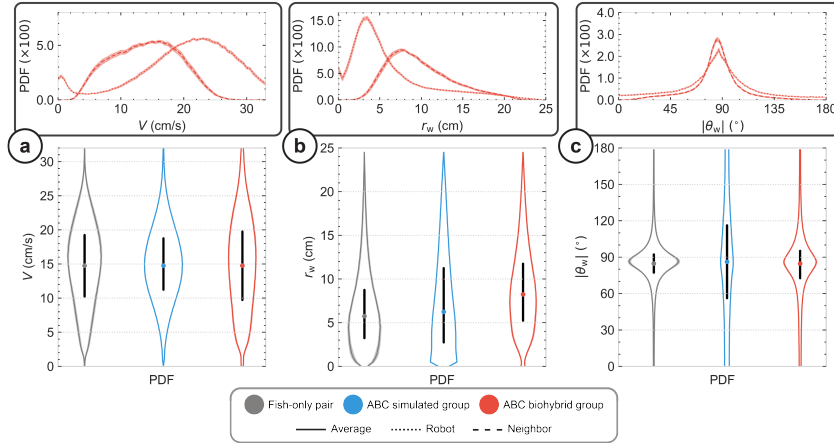


Figure E.9: **Instantaneous individual quantities.** (a) Speed V probability density function. (b) Distance to the wall r_w probability density function. (c) Absolute angle of incidence to the wall $|\theta_w|$ probability density function. Dark gray, blue, and red colors correspond to the distributions of the fish-only experiment, the ABC simulated groups, and the ABC biohybrid groups, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

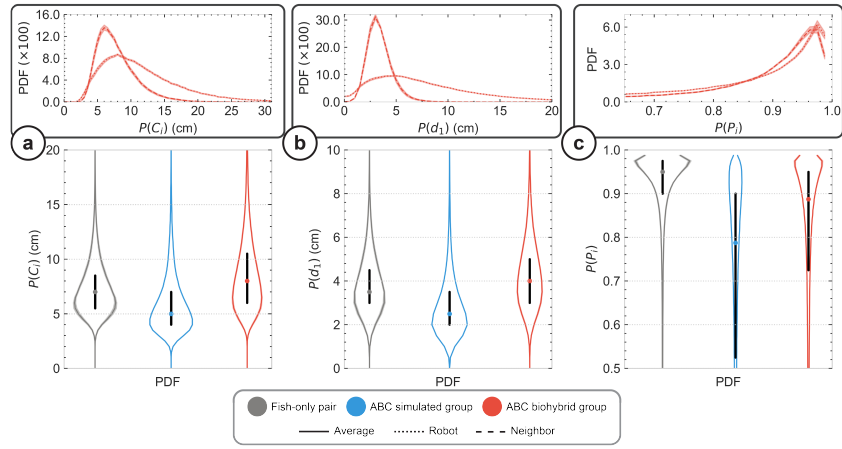


Figure E.10: **Instantaneous collective quantities.** (a) Interindividual distance d_{ij} probability density function. (b) Difference in heading angles $|\phi_{ij}|$ probability density function. (c) Viewing angle ψ_{ij} probability density function. Dark gray, blue, and red colors correspond to the distributions of the experiment, ABC simulated groups and ABC biohybrid groups, respectively. In all PDFs, the colored dot corresponds to the median, and the thick horizontal black line corresponds to the limits of the first and third quartile. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

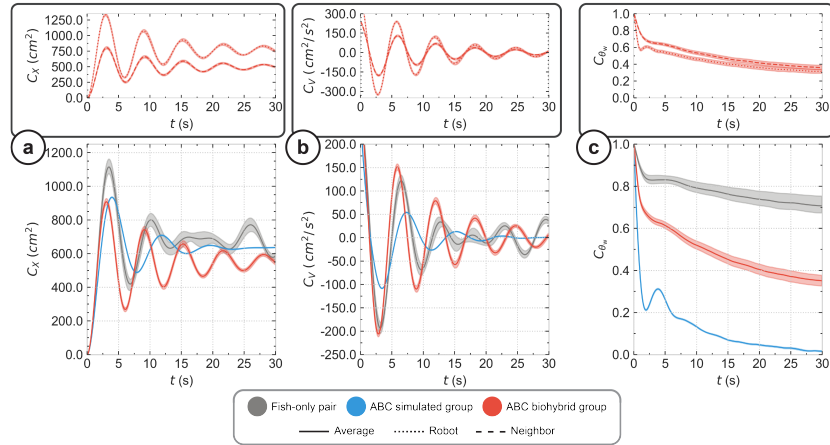


Figure E.11: **Temporal correlation quantities.** (a) Mean squared displacement $C_X(t)$. (b) Velocity autocorrelation $C_V(t)$. (c) Temporal correlations of the angle of incidence to the wall $C_{\theta_w}(t)$. Dark gray, blue and red colors correspond to the distributions of the experiment, ABC simulated groups and ABC biohybrid groups, respectively. Dotted, dashed and solid lines indicate the robot, neighbor and average agent distributions, respectively. The inset plots depict the PDFs of the ABC biohybrid pair experiments where the dotted, dashed, and solid lines correspond to the robot, neighbor and average agent distributions, respectively. The shaded areas correspond to the 68% confidence interval.

E.4 Tables for Chapter 7

Table E.1: Means and standard deviations of quantities for single agents

Pair	Quantity	Mean	Standard deviation
Fish-only	V	12.16 ± 0.92	6.68 ± 0.29
	r_w	2.03 ± 0.24	2.49 ± 0.19
	$ \theta_w $	89.46 ± 0.32	18.16 ± 1.10
ABC (simulated agent)	V	10.96 ± 0.13	5.77 ± 0.04
	r_w	3.72 ± 0.59	3.99 ± 0.04
	$ \theta_w $	88.85 ± 0.13	25.12 ± 0.46
ABC (robot)	V	10.99 ± 0.17	6.31 ± 0.08
	r_w	1.73 ± 0.09	2.04 ± 0.08
	$ \theta_w $	89.33 ± 0.04	18.39 ± 0.29

Table E.2: Hellinger distances for single agents

Pair	Quantity	Hellinger distance
Fish-only <i>vs</i> ABC-SP	V	0.20
	r_w	0.31
	$ \theta_w $	0.18
Fish-only <i>vs</i> ABC-BP	V	0.18
	r_w	0.15
	$ \theta_w $	0.11
ABC-SP <i>vs</i> ABC-BP	V	0.13
	r_w	0.39
	$ \theta_w $	0.11

Table E.3: Means and standard deviations of quantities for agent pairs.

Agent	Quantity	Mean	Standard deviation
Fish	V	10.50 ± 0.60	5.73 ± 0.36
	r_w	4.39 ± 0.43	3.86 ± 0.22
	$ \theta_w $	87.42 ± 0.39	21.91 ± 1.46
	d_{ij}	8.05 ± 0.71	5.11 ± 0.43
	$ \phi_{ij} $	26.72 ± 1.91	29.81 ± 1.24
	ψ_{ij}	7.96 ± 4.73	108.98 ± 1.19
ABC-SP	V	10.38 ± 0.07	6.39 ± 0.02
	r_w	5.49 ± 0.02	4.91 ± 0.03
	$ \theta_w $	89.95 ± 0.10	33.63 ± 0.22
	d_{ij}	5.63 ± 0.04	4.15 ± 0.08
	$ \phi_{ij} $	29.35 ± 0.31	30.85 ± 0.47
	ψ_{ij}	0.54 ± 0.61	106.77 ± 0.06
ABC-BP	V	7.93 ± 0.18	4.89 ± 0.15
	r_w	4.10 ± 0.11	3.70 ± 0.06
	$ \theta_w $	88.12 ± 0.16	25.93 ± 0.60
	d_{ij}	8.10 ± 0.46	7.41 ± 0.73
	$ \phi_{ij} $	36.50 ± 2.05	40.60 ± 1.27
	ψ_{ij}	0.94 ± 2.36	106.41 ± 0.63
ABC-BP (fish)	V	8.00 ± 0.23	4.67 ± 0.21
	r_w	4.46 ± 0.26	3.32 ± 0.10
	$ \theta_w $	87.60 ± 0.26	23.21 ± 0.88
ABC-BP (robot)	V	7.87 ± 0.15	5.09 ± 0.12
	r_w	3.74 ± 0.07	4.00 ± 0.08
	$ \theta_w $	88.64 ± 0.11	28.38 ± 0.57

Table E.4: Hellinger distances for agent pairs

Agent	Quantity	Hellinger distance
Fish <i>vs</i> ABC-SP	V	0.18
	r_w	0.15
	$ \theta_w $	0.19
	d_{ij}	0.38
	$ \phi_{ij} $	0.06
	ψ_{ij}	0.05
	Fish <i>vs</i> ABC-BP	V
r_w		0.06
$ \theta_w $		0.10
d_{ij}		0.32
$ \phi_{ij} $		0.10
ψ_{ij}		0.09
ABC-SP <i>vs</i> ABC-BP		V
	r_w	0.18
	$ \theta_w $	0.11
	d_{ij}	0.23
	$ \phi_{ij} $	0.07
	ψ_{ij}	0.04

Table E.5: Means and standard deviations of quantities for groups of five agents (ABC-BG, ABC-SG with $k = 1$)

Pair	Quantity	Mean	Standard deviation
Fish-only	V	15.23 ± 0.83	5.94 ± 0.20
	r_w	7.02 ± 0.72	4.53 ± 0.38
	$ \theta_w $	86.33 ± 0.41	20.02 ± 1.87
	$P(C_i)$	7.37 ± 0.39	2.53 ± 0.22
	$P(d_i)$	4.10 ± 0.22	1.37 ± 0.10
	$P(P_i)$	0.91 ± 0.01	0.15 ± 0.02
ABC-SG	V	15.35 ± 0.06	5.47 ± 0.02
	r_w	8.38 ± 0.05	5.99 ± 0.02
	$ \theta_w $	87.72 ± 0.06	43.97 ± 1.19
	$P(C_i)$	8.93 ± 0.06	4.46 ± 0.03
	$P(d_i)$	4.37 ± 0.03	2.27 ± 0.02
	$P(P_i)$	$0.50 \pm 1e-3$	$0.38 \pm 1e-3$
ABC-BG	V	14.15 ± 0.37	6.92 ± 0.15
	r_w	10.06 ± 0.27	5.10 ± 0.04
	$ \theta_w $	84.39 ± 0.32	32.18 ± 0.65
	$P(C_i)$	9.65 ± 0.22	3.31 ± 0.09
	$P(d_i)$	5.03 ± 0.10	1.59 ± 0.04
	$P(P_i)$	0.73 ± 0.01	0.29 ± 0.01
ABC-BG (fish)	V	12.68 ± 0.37	4.78 ± 0.15
	r_w	10.83 ± 0.31	3.89 ± 0.08
	$ \theta_w $	84.44 ± 0.39	23.85 ± 0.70
	$P(C_i)$	8.98 ± 0.20	2.88 ± 0.09
	$P(d_i)$	4.25 ± 0.11	1.30 ± 0.04
	$P(P_i)$	0.76 ± 0.01	0.27 ± 0.01
ABC-BG (robot)	V	20.02 ± 0.43	7.77 ± 0.18
	r_w	7.00 ± 0.13	5.21 ± 0.06
	$ \theta_w $	84.19 ± 0.16	36.11 ± 0.57
	$P(C_i)$	12.36 ± 0.19	5.54 ± 0.14
	$P(d_i)$	8.15 ± 0.13	5.10 ± 0.14
	$P(P_i)$	0.61 ± 0.02	0.46 ± 0.01

Table E.6: Hellinger distances for groups of five agents (ABC-BG, ABC-SG with $k = 1$)

Pair	Quantity	Hellinger distance
Fish-only <i>vs</i> ABC-SG	V	0.08
	r_w	0.15
	$ \theta_w $	0.37
	$P(C_i)$	0.23
	$P(d_i)$	0.21
	$P(P_i)$	0.53
	Fish-only <i>vs</i> ABC-BG	V
r_w		0.23
$ \theta_w $		0.20
$P(C_i)$		0.25
$P(d_i)$		0.22
$P(P_i)$		0.32
ABC-SG <i>vs</i> ABC-BG		V
	r_w	0.20
	$ \theta_w $	0.19
	$P(C_i)$	0.24
	$P(d_i)$	0.28
	$P(P_i)$	0.25

Table E.7: Means and standard deviations of quantities for groups of five agents (ABC-BG, ABC-SG with $k = 2$)

Pair	Quantity	Mean	Standard deviation
Fish-only	V	15.23 ± 0.83	5.94 ± 0.20
	r_w	7.02 ± 0.72	4.53 ± 0.38
	$ \theta_w $	86.33 ± 0.41	20.02 ± 1.87
	$P(C_i)$	7.37 ± 0.39	2.53 ± 0.22
	$P(d_i)$	4.10 ± 0.22	1.37 ± 0.10
	$P(P_i)$	0.91 ± 0.01	0.15 ± 0.02
	ABC-SG	V	15.34 ± 0.05
r_w		8.10 ± 0.03	5.65 ± 0.02
$ \theta_w $		88.06 ± 0.05	43.73 ± 1.44
$P(C_i)$		6.19 ± 0.03	2.81 ± 0.03
$P(d_i)$		3.08 ± 0.01	1.53 ± 0.02
$P(P_i)$		$0.67 \pm 3e-3$	$0.31 \pm 1e-3$
ABC-BG		V	15.53 ± 0.42
	r_w	9.32 ± 0.29	4.85 ± 0.08
	$ \theta_w $	84.80 ± 0.18	28.40 ± 0.98
	$P(C_i)$	8.81 ± 0.31	3.40 ± 0.19
	$P(d_i)$	4.54 ± 0.13	1.58 ± 0.07
	$P(P_i)$	0.79 ± 0.02	0.26 ± 0.01
	ABC-BG (fish)	V	14.21 ± 0.47
r_w		10.08 ± 0.33	3.86 ± 0.05
$ \theta_w $		84.79 ± 0.21	21.43 ± 0.71
$P(C_i)$		8.09 ± 0.31	2.91 ± 0.21
$P(d_i)$		3.65 ± 0.15	1.18 ± 0.10
$P(P_i)$		0.82 ± 0.02	0.23 ± 0.02
ABC-BG (robot)		V	20.80 ± 0.38
	r_w	6.27 ± 0.23	4.80 ± 0.13
	$ \theta_w $	84.84 ± 0.16	34.29 ± 0.97
	$P(C_i)$	11.70 ± 0.32	5.76 ± 0.21
	$P(d_i)$	8.12 ± 0.22	5.41 ± 0.17
	$P(P_i)$	0.65 ± 0.02	0.45 ± 0.01

Table E.8: Hellinger distances for groups of five agents (ABC-BG, ABC-SG with $k = 2$)

Pair	Quantity	Hellinger distance
Fish-only <i>vs</i> ABC-SG	V	0.08
	r_w	0.12
	$ \theta_w $	0.38
	$P(C_i)$	0.29
	$P(d_i)$	0.31
	$P(P_i)$	0.43
Fish-only <i>vs</i> ABC-BG	V	0.10
	r_w	0.19
	$ \theta_w $	0.14
	$P(C_i)$	0.15
	$P(d_i)$	0.11
	$P(P_i)$	0.23
ABC-SG <i>vs</i> ABC-BG	V	0.13
	r_w	0.18
	$ \theta_w $	0.25
	$P(C_i)$	0.37
	$P(d_i)$	0.34
	$P(P_i)$	0.20

Bibliography

- Abaid, N., Bartolini, T., Macrì, S., and Porfiri, M. (2012a). Zebrafish responds differentially to a robotic fish of varying aspect ratio, tail beat frequency, noise, and color. *Behavioural brain research*, 233(2):545–553.
- Abaid, N., Marras, S., Fitzgibbons, C., and Porfiri, M. (2013). Modulation of risk-taking behaviour in golden shiners (*notemigonus crysoleucas*) using robotic fish. *Behavioural processes*, (100):9–12.
- Abaid, N. and Porfiri, M. (2010). Fish in a ring: spatio-temporal pattern formation in one-dimensional animal groups. *Journal of The Royal Society Interface*, page rsif20100175.
- Abaid, N., Spinello, C., Laut, J., and Porfiri, M. (2012b). Zebrafish (*danio rerio*) responds to images animated by mathematical models of animal grouping. *Behavioural brain research*, 232(2):406–410.
- Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., and Savarese, S. (2016). Social lstm: Human trajectory prediction in crowded spaces. In O’Conner, L., editor, *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 961–971.
- Anderson, D. J. and Perona, P. (2014). Toward a science of computational ethology. *Neuron*, 84(1):18–31.
- Aoki, I. (1982). A simulation study on the schooling mechanism in fish. *Bulletin of the Japanese Society of Scientific Fisheries (Japan)*.
- Ball, P. (2011). *Flow: Nature’s patterns: A tapestry in three parts*, volume 2. Oxford University Press.
- Barmak, R., Stefanec, M., Hofstadler, D. N., Piotet, L., Schönwetter-Fuchs-Schistek, S., Mondada, F., Schmickl, T., and Mills, R. (2023). A robotic honeycomb for interaction with a honeybee colony. *Science Robotics*, 8(76):eadd7385.
- Bartolini, T., Mwaffo, V., Showler, A., Macrì, S., Butail, S., and Porfiri, M. (2016). Zebrafish response to 3d printed shoals of conspecifics: the effect of body size. *Bioinspiration & biomimetics*, 11(2):026003.

- Basu, A., Harris, I. R., and Basu, S. (1997). 2 minimum distance estimation: The approach using density-based distances. *Handbook of Statistics*, 15:21–48.
- Beran, R. (1977). Minimum hellinger distance estimates for parametric models. *The annals of Statistics*, pages 445–463.
- Bertin, E., Droz, M., and Grégoire, G. (2006). Boltzmann and hydrodynamic description for self-propelled particles. *Physical Review E*, 74(2):022101.
- Bode, N. W., Franks, D. W., and Wood, A. J. (2010). Limited interactions in flocks: relating model simulations to empirical data. *Journal of The Royal Society Interface*, 8(55):301–304.
- Bonnet, F., Binder, S., de Oliveria, M. E., Halloy, J., and Mondada, F. (2014). 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.
- Bonnet, F., Cazenille, L., Gribovskiy, A., Halloy, J., and Mondada, F. (2017). Multi-robot control and tracking framework for bio-hybrid systems with closed-loop interaction. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4449–4456. IEEE.
- Bonnet, F., Gribovskiy, A., Halloy, J., and Mondada, F. (2018). Closed-loop interactions between a shoal of zebrafish and a group of robotic fish in a circular corridor. *Swarm Intelligence*, 12(3):227–244.
- Bonnet, F., Kato, Y., Halloy, J., and Mondada, F. (2016). Infiltrating the zebrafish swarm: design, implementation and experimental tests of a miniature robotic fish lure for fish-robot interaction studies. *Artificial Life and Robotics*, 21(3):239–246.
- Bonnet, F., Mills, R., Szopek, M., Schönwetter-Fuchs, S., Halloy, J., Bogdan, S., Correia, L., Mondada, F., and Schmickl, T. (2019). Robots mediating interactions between animals for interspecies collective behaviors. *Science Robotics*, 4(28):eaau7897.
- Bonnet, F., Rétoznaz, P., Halloy, J., Gribovskiy, A., and Mondada, F. (2012). Development of a mobile robot to study the collective behavior of zebrafish. In *2012 4th IEEE RAS & EMBS international conference on Biomedical Robotics and Biomechatronics (BioRob)*, pages 437–442. Ieee.
- Boujard, T. and Leatherland, J. F. (1992). Circadian rhythms and feeding time in fishes. *Environmental Biology of Fishes*, 35:109–131.
- Bradski, G. (2000). The opencv library. *Dr. Dobb's Journal: Software Tools for the Professional Programmer*, 25(11):120–123.
- Branson, K., Robie, A. A., Bender, J., Perona, P., and Dickinson, M. H. (2009). High-throughput ethomics in large groups of drosophila. *Nature methods*, 6(6):451–457.
- Buhl, J., Sumpter, D. J., Couzin, I. D., Hale, J. J., Despland, E., Miller, E. R., and Simpson, S. J. (2006). From disorder to order in marching locusts. *Science*, 312(5778):1402–1406.

- Bumann, D. and Krause, J. (1993). Front individuals lead in shoals of three-spined sticklebacks (*Gasterosteus aculeatus*) and juvenile roach (*Rutilus rutilus*). *Behaviour*, 125(3):189–198.
- Butail, S., Bartolini, T., and Porfiri, M. (2013). Collective response of zebrafish shoals to a free-swimming robotic fish. *PLoS One*, 8(10):e76123.
- Butail, S., Ladu, F., Spinello, D., and Porfiri, M. (2014a). Information flow in animal-robot interactions. *Entropy*, 16(3):1315–1330.
- Butail, S., Mwaffo, V., and Porfiri, M. (2016). Model-free information-theoretic approach to infer leadership in pairs of zebrafish. *Physical Review E*, 93(4):042411.
- Butail, S., Polverino, G., Phamduy, P., Del Sette, F., and Porfiri, M. (2014b). Influence of robotic shoal size, configuration, and activity on zebrafish behavior in a free-swimming environment. *Behavioural brain research*, 275:269–280.
- Calovi, D. S., Litchinko, A., Lecheval, V., Lopez, U., Escudero, A. P., Chaté, H., Sire, C., and Theraulaz, G. (2018). Disentangling and modeling interactions in fish with burst-and-coast swimming reveal distinct alignment and attraction behaviors. *PLoS computational biology*, 14(1):e1005933.
- Calovi, D. S., Lopez, U., Ngo, S., Sire, C., Chaté, H., and Theraulaz, G. (2014). Swarming, schooling, milling: phase diagram of a data-driven fish school model. *New journal of Physics*, 16(1):015026.
- Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Bonabeau, E., and Theraulaz, G. (2003). *Self-organization in biological systems*. Princeton university press.
- Cavagna, A., Giardina, I., and Grigera, T. S. (2018). The physics of flocking: Correlation as a compass from experiments to theory. *Physics Reports*, 728:1–62.
- Cazenille, L., Bredeche, N., and Halloy, J. (2015). Multi-objective optimization of multi-level models for controlling animal collective behavior with robots. pages 379–390.
- Cazenille, L., Bredeche, N., and Halloy, J. (2016). Automated optimisation of multi-level models of collective behaviour in a mixed society of animals and robots. *arXiv preprint arXiv:1602.05830*.
- Cazenille, L., Bredeche, N., and Halloy, J. (2018a). Evolutionary optimisation of neural network models for fish collective behaviours in mixed groups of robots and zebrafish. In *Biomimetic and Biohybrid Systems: 7th International Conference, Living Machines 2018, Paris, France, July 17–20, 2018, Proceedings 7*, pages 85–96. Springer.
- Cazenille, L., Bredeche, N., and Halloy, J. (2019). *Automatic Calibration of Artificial Neural Networks for Zebrafish Collective Behaviours using a Quality Diversity Algorithm*.

- Cazenille, L., Chemtob, Y., Bonnet, F., Gribovskiy, A., Mondada, F., Bredeche, N., and Halloy, J. (2017). Automated calibration of a biomimetic space-dependent model for zebrafish and robot collective behaviour in a structured environment. pages 107–118.
- Cazenille, L., Chemtob, Y., Bonnet, F., Gribovskiy, A., Mondada, F., Bredeche, N., and Halloy, J. (2018b). *How to blend a robot within a group of zebrafish: Achieving social acceptance through real-time calibration of a multi-level behavioural model.*
- Cazenille, L., Collignon, B., Chemtob, Y., Bonnet, F., Gribovskiy, A., Mondada, F., Bredeche, N., and Halloy, J. (2018c). How mimetic should a robotic fish be to socially integrate into zebrafish groups? *Bioinspiration & biomimetics*, 13(2):025001.
- Chaté, H., Ginelli, F., Grégoire, G., Peruani, F., and Raynaud, F. (2008). Modeling collective motion: variations on the vicsek model. *The European Physical Journal B*, 64(3-4):451–456.
- Chemtob, Y., Cazenille, L., Bonnet, F., Gribovskiy, A., Mondada, F., and Halloy, J. (2020). Strategies to modulate zebrafish collective dynamics with a closed-loop biomimetic robotic system. *Bioinspiration & Biomimetics*, 15(4):046004.
- Chen, Y., Angulo, M. T., and Liu, Y.-Y. (2019). Revealing complex ecological dynamics via symbolic regression. *BioEssays*, 41(12):1900069.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. (2018). Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, pages 4754–4765.
- Collignon, B., Séguret, A., Chemtob, Y., Cazenille, L., and Halloy, J. (2017). Collective departures in zebrafish: profiling the initiators. *arXiv preprint arXiv:1701.03611*.
- Collignon, B., Séguret, A., and Halloy, J. (2016). A stochastic vision-based model inspired by zebrafish collective behaviour in heterogeneous environments. *Royal Society open science*, 3(1):150473.
- Correll, N., Schwager, M., and Rus, D. (2008). Social control of herd animals by integration of artificially controlled congeners. In *International Conference on Simulation of Adaptive Behavior*, pages 437–446. Springer.
- Costa, T., Laan, A., Heras, F. J., and de Polavieja, G. G. (2020). Automated discovery of local rules for desired collective-level behavior through reinforcement learning. *Front. Phys.* 8: 200. doi: 10.3389/fphys.
- Couzin, I. D., Krause, J., et al. (2003). Self-organization and collective behavior in vertebrates. *Advances in the Study of Behavior*, 32(1):10–1016.
- Couzin, I. D., Krause, J., James, R., Ruxton, G. D., and Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11.

- Crosato, E., Jiang, L., Lecheval, V., Lizier, J. T., Wang, X. R., Tichit, P., Theraulaz, G., and Prokopenko, M. (2018). Informative and misinformative interactions in a school of fish. *Swarm Intelligence*, 12(4):283–305.
- Czeglédi, I., Sály, P., Takács, P., Dolezsai, A., Vitál, Z., Nagy, S. A., and Erős, T. (2016). Do diel variations in stream fish assemblages depend on spatial positioning of the sampling sites and seasons? *Acta Zoologica Academiae Scientiarum Hungaricae*, 62(2):175–190.
- Czeglédi, I., Specziár, A., and Erős, T. (2022). Temporal dynamics of freshwater fish assemblages, their background and methods of quantifications—a synthesis. *Fish and Fisheries*, 23(1):93–108.
- de Lourdes Ruiz-Gomez, M., Kittilsen, S., Höglund, E., Huntingford, F. A., Sørensen, C., Pottinger, T. G., Bakken, M., Winberg, S., Korzan, W. J., and Øverli, Ø. (2008). Behavioral plasticity in rainbow trout (*oncorhynchus mykiss*) with divergent coping styles: when doves become hawks. *Hormones and behavior*, 54(4):534–538.
- Dell, A. I., Bender, J. A., Branson, K., Couzin, I. D., de Polavieja, G. G., Noldus, L. P., Pérez-Escudero, A., Perona, P., Straw, A. D., Wikelski, M., et al. (2014). Automated image-based tracking and its application in ecology. *Trends in ecology & evolution*, 29(7):417–428.
- Deutsch, A., Friedl, P., Preziosi, L., and Theraulaz, G. (2020). Multi-scale analysis and modelling of collective migration in biological systems. *Philosophical Transactions of the Royal Society B*, 375(1807):20190377.
- Dijkgraaf, S. (1963). The functioning and significance of the lateral-line organs. *Biological Reviews*, 38(1):51–105.
- Donati, E., Worm, M., Mintchev, S., Van Der Wiel, M., Benelli, G., Von Der Emde, G., and Stefanini, C. (2016). Investigation of collective behaviour and electrocommunication in the weakly electric fish, *mormyrus rume*, through a biomimetic robotic dummy fish. *Bioinspiration & biomimetics*, 11(6):066009.
- Escobedo, R., Lecheval, V., Papaspyros, V., Bonnet, E., Mondada, F., Sire, C., and Theraulaz, G. (2020). A data-driven method for reconstructing and modelling social interactions in moving animal groups. *Philosophical Transactions of the Royal Society B*, 375(1807):20190380.
- Faria, J. J., Dyer, J. R., Clément, R. O., Couzin, I. D., Holt, N., Ward, A. J., Waters, D., and Krause, J. (2010). A novel method for investigating the collective behaviour of fish: introducing ‘robofish’. *Behavioral Ecology and Sociobiology*, 64(8):1211–1218.
- Folkertsma, G. A., Straatman, W., Nijenhuis, N., Venner, C. H., and Stramigioli, S. (2017). Robird: a robotic bird of prey. *IEEE robotics & automation magazine*, 24(3):22–29.
- Gallois, B. and Candelier, R. (2021). Fasttrack: an open-source software for tracking varying numbers of deformable objects. *PLoS computational biology*, 17(2):e1008697.

- Gautrais, J., Ginelli, F., Fournier, R., Blanco, S., Soria, M., Chaté, H., and Theraulaz, G. (2012). Deciphering interactions in moving animal groups. *Plos computational biology*, 8(9):e1002678.
- Gilpin, W., Huang, Y., and Forger, D. B. (2020). Learning dynamics from large biological datasets: machine learning meets systems biology. *Current Opinion in Systems Biology*.
- Gribovskiy, A., Halloy, J., Deneubourg, J.-L., Bleuler, H., and Mondada, F. (2010). Towards mixed societies of chickens and robots. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4722–4728. IEEE.
- Gribovskiy, A., Halloy, J., Deneubourg, J.-L., and Mondada, F. (2018). Designing a socially integrated mobile robot for ethological research. *Robotics and Autonomous Systems*, 103:42–55.
- Griparić, K., Haus, T., Miklić, D., Polić, M., and Bogdan, S. (2017). A robotic system for researching social integration in honeybees. *PloS one*, 12(8):e0181977.
- Halloy, J., Sempo, G., Caprari, G., Rivault, C., Asadpour, M., Tâche, F., Saïd, I., Durier, V., Canonge, S., Amé, J. M., et al. (2007). Social integration of robots into groups of cockroaches to control self-organized choices. *Science*, 318(5853):1155–1158.
- Hemelrijk, C. K. and Hildenbrandt, H. (2008). Self-organized shape and frontal density of fish schools. *Ethology*, 114(3):245–254.
- Heras, F. J. H., Romero-Ferrero, F., Hinz, R. C., and de Polavieja, G. G. (2019). Deep attention networks reveal the rules of collective motion in zebrafish. *PLOS Computational Biology*, 15(9):1–23.
- Herbert-Read, J. E. (2016). Understanding how animal groups achieve coordinated movement. *Journal of Experimental Biology*, 219(19):2971–2983.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ilgün, A., Angelov, K., Stefanec, M., Schönwetter-Fuchs, S., Stokanic, V., Vollmann, J., Hofstadler, D. N., Kärcher, M. H., Mellmann, H., Taliaronak, V., et al. (2021). Bio-hybrid systems for ecosystem level effects. In *ALIFE 2021: The 2021 Conference on Artificial Life*, number CONF. MIT Press.
- Ioannou, C. C. and Laskowski, K. L. (2023). A multi-scale review of the dynamics of collective behaviour: from rapid responses to ontogeny and evolution.
- Jakobi, N. (1997). Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive behavior*, 6(2):325–368.
- Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: The use of simulation in evolutionary robotics. In *European Conference on Artificial Life*, pages 704–720. Springer.

- Jayles, B., Escobedo, R., Pasqua, R., Zanon, C., Blanchet, A., Roy, M., Trédan, G., Theraulaz, G., and Sire, C. (2020). Collective information processing in human phase separation. *Philosophical Transactions of the Royal Society B*, 375(1807):20190801.
- Jiang, L., Giuggioli, L., Perna, A., Escobedo, R., Lecheval, V., Sire, C., Han, Z., and Theraulaz, G. (2017). Identifying influential neighbors in animal flocking. *PLoS computational biology*, 13(11):e1005822.
- Jolly, L., Pittet, F., Caudal, J.-P., Mouret, J.-B., Houdelier, C., Lumineau, S., and de Margerie, E. (2016). Animal-to-robot social attachment: initial requisites in a gallinaceous bird. *Bioinspiration & biomimetics*, 11(1):016007.
- KaewTraKulPong, P. and Bowden, R. (2002). An improved adaptive background mixture model for real-time tracking with shadow detection. In *Video-based surveillance systems*, pages 135–144. Springer.
- Kim, C., Ruberto, T., Phamduy, P., and Porfiri, M. (2018). Closed-loop control of zebrafish behaviour in three dimensions using a robotic stimulus. *Scientific reports*, 8(1):657.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kopman, V., Laut, J., Polverino, G., and Porfiri, M. (2013). Closed-loop control of zebrafish response using a bioinspired robotic-fish in a preference test. *Journal of the Royal Society Interface*, 10(78):20120540.
- Kopman, V. and Porfiri, M. (2013). Design, modeling, and characterization of a miniature robotic fish for research and education in biomimetics and bioinspiration. *IEEE/ASME Transactions on mechatronics*, 18(2):471–483.
- Kothari, P., Kreiss, S., and Alahi, A. (2021). Human trajectory forecasting in crowds: A deep learning perspective. *IEEE Transactions on Intelligent Transportation Systems*.
- Krause, J. and Ruxton, G. D. (2002). *Living in groups*. Oxford University Press.
- Kruusmaa, M., Rieucau, G., Montoya, J. C. C., Markna, R., and Handegard, N. O. (2016). 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.
- Kuhn, H. W. (1955). The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97.
- Ladu, E., Bartolini, T., Panitz, S. G., Chiarotti, F., Butail, S., Macrì, S., and Porfiri, M. (2015a). Live predators, robots, and computer-animated images elicit differential avoidance responses in zebrafish. *Zebrafish*, 12(3):205–214.
- Ladu, E., Mwaffo, V., Li, J., Macrì, S., and Porfiri, M. (2015b). Acute caffeine administration affects zebrafish response to a robotic stimulus. *Behavioural brain research*, 289:48–54.

- Landgraf, T., Bierbach, D., Nguyen, H., Muggelberg, N., Romanczuk, P., and Krause, J. (2016). Robofish: increased acceptance of interactive robotic fish with realistic eyes and natural motion patterns by live trinidadian guppies. *Bioinspiration & biomimetics*, 11(1):015001.
- Landgraf, T., Nguyen, H., Forgo, S., Schneider, J., Schröer, J., Krüger, C., Matzke, H., Clément, R. O., Krause, J., and Rojas, R. (2013). Interactive robotic fish for the analysis of swarm behavior. In *Advances in Swarm Intelligence: 4th International Conference, ICSI 2013, Harbin, China, June 12-15, 2013, Proceedings, Part I 4*, pages 1–10. Springer.
- Landgraf, T., Nguyen, H., Schröer, J., Szengel, A., Clément, R. J., Bierbach, D., and Krause, J. (2014). Blending in with the shoal: robotic fish swarms for investigating strategies of group formation in guppies. In *Biomimetic and Biohybrid Systems: Third International Conference, Living Machines 2014, Milan, Italy, July 30–August 1, 2014. Proceedings 3*, pages 178–189. Springer.
- Landgraf, T., Oertel, M., Rhiel, D., and Rojas, R. (2010). A biomimetic honeybee robot for the analysis of the honeybee dance communication system. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3097–3102. IEEE.
- Lecheval, V., Jiang, L., Tichit, P., Sire, C., Hemelrijk, C. K., and Theraulaz, G. (2018). Social conformity and propagation of information in collective u-turns of fish schools. *Proceedings of the Royal Society B: Biological Sciences*, 285(1877):20180251.
- Lei, L., Escobedo, R., Sire, C., and Theraulaz, G. (2020). Computational and robotic modeling reveal parsimonious combinations of interactions between individuals in schooling fish. *PLoS computational biology*, 16(3):e1007194.
- Li, G., Ashraf, I., François, B., Kolomenskiy, D., Lechenault, F., Godoy-Diana, R., and Thiria, B. (2021). Burst-and-coast swimmers optimize gait by adapting unique intrinsic cycle. *Communications biology*, 4(1):40.
- Lizier, J. T. (2014). Jidt: An information-theoretic toolkit for studying the dynamics of complex systems. *Frontiers in Robotics and AI*, 1:11.
- Lopez, U., Gautrais, J., Couzin, I. D., and Theraulaz, G. (2012). From behavioural analyses to models of collective motion in fish schools. *Interface focus*, 2(6):693–707.
- Lukeman, R., Li, Y.-X., and Edelstein-Keshet, L. (2010). Inferring individual rules from collective behavior. *Proceedings of the National Academy of Sciences*, 107(28):12576–12580.
- Mammadli, R., Wolf, F., and Jannesari, A. (2019). The art of getting deep neural networks in shape. *ACM Transactions on Architecture and Code Optimization (TACO)*, 15(4):1–21.
- Marchand, E., Uchiyama, H., and Spindler, F. (2015). Pose estimation for augmented reality: a hands-on survey. *IEEE transactions on visualization and computer graphics*, 22(12):2633–2651.

- Marques, J. C., Lackner, S., Félix, R., and Orger, M. B. (2018). Structure of the zebrafish locomotor repertoire revealed with unsupervised behavioral clustering. *Current Biology*, 28(2):181–195.
- Marras, S. and Porfiri, M. (2012). Fish and robots swimming together: attraction towards the robot demands biomimetic locomotion. *Journal of The Royal Society Interface*, page rsif20120084.
- Martín, J. A. H., de Lope, J., and Santos, M. (2007). Evolution of neuro-controllers for multi-link robots. In Emilio Corchado, Juan M. Corchado, A. A., editor, *Innovations in Hybrid Intelligent Systems*, pages 175–182. Springer.
- Martinoli, A., Ijspeert, A. J., and Mondada, F. (1999). Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots. *Robotics and Autonomous Systems*, 29(1):51–63.
- Maxeiner, M., Hocke, M., Moenck, H., Gebhardt, G., Weimar, N., Musiolek, L., Krause, J., Bierbach, D., and Landgraf, T. (2023). Social competence improves the performance of biomimetic robots leading live fish. *Bioinspiration & Biomimetics*.
- Mondada, F., Bonani, M., Riedo, F., Briod, M., Pereyre, L., Rétornaz, P., and Magnenat, S. (2017). Bringing robotics to formal education: The thymio open-source hardware robot. *IEEE Robotics & Automation Magazine*, 24(1):77–85.
- Mondada, F., Martinoli, A., Correll, N., Gribovskiy, A., Halloy, J. I., Siegwart, R., and Deneubourg, J.-L. (2013). A general methodology for the control of mixed natural-artificial societies. *Handbook of collective robotics*, pages 547–586.
- Monk, C. T., Aslak, U., Brockmann, D., and Arlinghaus, R. (2023). Rhythm of relationships in a social fish over the course of a full year in the wild. *Movement Ecology*, 11(1):1–16.
- Mouret, J.-B. and Chatzilygeroudis, K. (2017). 20 years of reality gap: a few thoughts about simulators in evolutionary robotics. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1121–1124.
- Mwaffo, V., Anderson, R. P., Butail, S., and Porfiri, M. (2015). A jump persistent turning walker to model zebrafish locomotion. *Journal of The Royal Society Interface*, 12(102):20140884.
- Mwaffo, V., Butail, S., and Porfiri, M. (2017). Analysis of pairwise interactions in a maximum likelihood sense to identify leaders in a group. *Frontiers in Robotics and AI*, 4:35.
- Mwaura, J. and Keedwell, E. (2015). Evolving robotic neuro-controllers using gene expression programming. In Kellenberger, P., editor, *2015 IEEE symposium series on computational intelligence*, pages 1063–1072. IEEE.
- Nagai, K. H., Sumino, Y., Montagne, R., Aranson, I. S., and Chaté, H. (2015). Collective motion of self-propelled particles with memory. *Physical review letters*, 114(16):168001.

- Nasevicius, A. and Ekker, S. C. (2000). Effective targeted gene 'knockdown' in zebrafish. *Nature genetics*, 26(2):216.
- Niwa, H.-S. (1996). Newtonian dynamical approach to fish schooling. *Journal of Theoretical Biology*, 1(181):47–63.
- Olin, M. and Malinen, T. (2003). Comparison of gillnet and trawl in diurnal fish community sampling. *Hydrobiologia*, 506:443–449.
- Papaspyros, V., Bonnet, F., Collignon, B., and Mondada, F. (2019). Bidirectional interactions facilitate the integration of a robot into a shoal of zebrafish *danio rerio*. *PloS one*, 14(8):e0220559.
- Papaspyros, V., Burnier, D., Cherfan, R., Theraulaz, G., Sire, C., and Mondada, F. (2023a). A biohybrid interaction framework for the integration of robots in animal societies. *IEEE Access*, 11:67640–67659.
- Papaspyros, V., Escobedo, R., Alahi, A., Theraulaz, G., Sire, C., and Mondada, F. (2023b). Predicting long-term collective animal behavior with deep learning.
- Papaspyros, V., Theraulaz, G., Sire, C., and Mondada, F. (2023c). Quantifying the biomimicry gap in biohybrid systems.
- Pérez-Escudero, A., Vicente-Page, J., Hinz, R. C., Arganda, S., and De Polavieja, G. G. (2014). idtracker: tracking individuals in a group by automatic identification of unmarked animals. *Nature methods*, 11(7):743–748.
- Phamduy, P., Polverino, G., Fuller, R., and Porfiri, M. (2014). Fish and robot dancing together: bluefin killifish females respond differently to the courtship of a robot with varying color morphs. *Bioinspiration & biomimetics*, 9(3):036021.
- Pillot, M. H., Gautrais, J., Gouello, J., Michelena, P., Bon, R., et al. (2010). Moving together: Incidental leaders and naïve followers. *Behavioural Processes*, 83(3):235–241.
- Polverino, G., Abaid, N., Kopman, V., Macrì, S., and Porfiri, M. (2012). Zebrafish response to robotic fish: preference experiments on isolated individuals and small shoals. *Bioinspiration & biomimetics*, 7(3):036019.
- Polverino, G., Phamduy, P., and Porfiri, M. (2013). Fish and robots swimming together in a water tunnel: robot color and tail-beat frequency influence fish behavior. *PloS one*, 8(10):e77589.
- Polverino, G. and Porfiri, M. (2013a). Mosquitofish (*gambusia affinis*) responds differentially to a robotic fish of varying swimming depth and aspect ratio. *Behavioural brain research*, 250:133–138.
- Polverino, G. and Porfiri, M. (2013b). Zebrafish (*danio rerio*) behavioural response to bioinspired robotic fish and mosquitofish (*gambusia affinis*). *Bioinspiration & biomimetics*, 8(4):044001.

- Porfiri, M. (2018). Inferring causal relationships in zebrafish-robot interactions through transfer entropy: a small lure to catch a big fish. *Animal Behavior and Cognition*, 5(4):341–367.
- Porfiri, M. and Marín, M. R. (2017). Symbolic dynamics of animal interaction. *Journal of theoretical biology*, 435:145–156.
- Porfiri, M., Spinello, C., Yang, Y., and Macrì, S. (2019). Zebrafish adjust their behavior in response to an interactive robotic predator. *Frontiers in Robotics and AI*, 6:38.
- Quade, M., Abel, M., Shafi, K., Niven, R. K., and Noack, B. R. (2016). Prediction of dynamical systems by symbolic regression. *Physical Review E*, 94(1):012214.
- Reynolds, C. W. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34.
- Roberts, A. C., Chornak, J., Alzagatiti, J. B., Ly, D. T., Bill, B. R., Trinkeller, J., Pearce, K. C., Choe, R. C., Campbell, C., Wong, D., et al. (2019). Rapid habituation of a touch-induced escape response in zebrafish (*danio rerio*) larvae. *PloS one*, 14(4):e0214374.
- Romano, D., Bloemberg, J., Tannous, M., and Stefanini, C. (2020). Impact of aging and cognitive mechanisms on high-speed motor activation patterns: evidence from an orthoptera-robot interaction. *IEEE Transactions on Medical Robotics and Bionics*, 2(2):292–296.
- Romano, D., Donati, E., Benelli, G., and Stefanini, C. (2019). A review on animal–robot interaction: from bio-hybrid organisms to mixed societies. *Biological cybernetics*, 113(3):201–225.
- Romano, D. and Stefanini, C. (2021). Unveiling social distancing mechanisms via a fish-robot hybrid interaction. *Biological Cybernetics*, 115(6):565–573.
- Romano, D. and Stefanini, C. (2022a). Any colour you like: fish interacting with bioinspired robots unravel mechanisms promoting mixed phenotype aggregations. *Bioinspiration & Biomimetics*, 17(4):045004.
- Romano, D. and Stefanini, C. (2022b). Robot-fish interaction helps to trigger social buffering in neon tetras: The potential role of social robotics in treating anxiety. *International Journal of Social Robotics*, 14(4):963–972.
- Romero-Ferrero, F., Bergomi, M. G., Hinz, R. C., Heras, F. J., and De Polavieja, G. G. (2019). Idtracker. ai: tracking all individuals in small or large collectives of unmarked animals. *Nature methods*, 16(2):179–182.
- Ruberto, T., Clément, R. J., Spinello, C., Neri, D., Macri, S., and Porfiri, M. (2018). The tagging procedure of visible implant elastomers influences zebrafish individual and social behavior. *Zebrafish*, 15(5):433–444.

- Ruberto, T., Mwaffo, V., Singh, S., Neri, D., and Porfiri, M. (2016). Zebrafish response to a robotic replica in three dimensions. *Royal Society open science*, 3(10):160505.
- Schmickl, T., Szopek, M., Mondada, F., Mills, R., Stefanec, M., Hofstadler, D. N., Lazic, D., Barmak, R., Bonnet, F., and Zahadat, P. (2021). Social integrating robots suggest mitigation strategies for ecosystem decay. *Frontiers in Bioengineering and Biotechnology*, 9:612605.
- Schreiber, T. (2000). Measuring information transfer. *Physical review letters*, 85(2):461.
- Sekaj, I., Ciferský, L., and Hvozdk, M. (2019). Neuro-evolution of mobile robot controller. *MENDEL*, 25(1):39–42.
- Shannon, C. E. and Weaver, W. (1949). The mathematical theory of information (urbana, il.
- Shi, J. et al. (1994). Good features to track. In *1994 Proceedings of IEEE conference on computer vision and pattern recognition*, pages 593–600. IEEE.
- Shi, Q., Ishii, H., Sugahara, Y., Takanishi, A., Huang, Q., and Fukuda, T. (2014). Design and control of a biomimetic robotic rat for interaction with laboratory rats. *IEEE/ASME Transactions on Mechatronics*, 20(4):1832–1842.
- Simon, R., Varkevisser, J., Mendoza, E., Hochradel, K., Elsinga, R., Wiersma, P. G., Middelburg, E., Zoeter, E., Scharff, C., Riebel, K., et al. (2023). Robofinch: A versatile audio-visual synchronised robotic bird model for laboratory and field research on songbirds. *Methods in Ecology and Evolution*.
- Spinello, C., Macrì, S., and Porfiri, M. (2013). Acute ethanol administration affects zebrafish preference for a biologically inspired robot. *Alcohol*, 47(5):391–398.
- Stanford Artificial Intelligence Laboratory et al. (2018). Robotic operating system.
- Stefanec, M., Szopek, M., Schmickl, T., and Mills, R. (2017). Governing the swarm: Controlling a bio-hybrid society of bees & robots with computational feedback loops. In *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8. IEEE.
- Strömbom, D., Mann, R. P., Wilson, A. M., Hailes, S., Morton, A. J., Sumpter, D. J., and King, A. J. (2014). Solving the shepherding problem: heuristics for herding autonomous, interacting agents. *Journal of the royal society interface*, 11(100):20140719.
- Sumpter, D. J. (2010). *Collective animal behavior*. Princeton University Press.
- Suriyampola, P. S., Zúñiga-Vega, J. J., Jayasundara, N., Flores, J., Lopez, M., Bhat, A., and Martins, E. P. (2023). River zebrafish combine behavioral plasticity and generalized morphology with specialized sensory and metabolic physiology to survive in a challenging environment. *Scientific Reports*, 13(1):16398.
- Swain, D. T., Couzin, I. D., and Leonard, N. E. (2011). Real-time feedback-controlled robotic fish for behavioral experiments with fish schools. *Proceedings of the IEEE*, 100(1):150–163.

- Takens, F. (1981). Detecting strange attractors in turbulence. In *Dynamical systems and turbulence, Warwick 1980*, pages 366–381. Springer.
- Team, G. (2016). Gstreamer: open source multimedia framework.
- Theraulaz, G., Gautrais, J., Camazine, S., and Deneubourg, J.-L. (2003). The formation of spatial patterns in social insects: from simple behaviours to complex structures. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 361(1807):1263–1282.
- Valletta, J. J., Torney, C., Kings, M., Thornton, A., and Madden, J. (2017). Applications of machine learning in animal behaviour studies. *Animal Behaviour*, 124:203–220.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Vaughan, R., Sumpter, N., Henderson, J., Frost, A., and Cameron, S. (2000). Experiments in automatic flock control. *Robotics and autonomous systems*, 31(1-2):109–117.
- Vicente, R., Wibral, M., Lindner, M., and Pipa, G. (2011). Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *Journal of computational neuroscience*, 30(1):45–67.
- Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., and Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Physical review letters*, 75(6):1226.
- Vicsek, T. and Zafeiris, A. (2012). Collective motion. *Physics reports*, 517(3-4):71–140.
- Walter, T. and Couzin, I. D. (2021). Trex, a fast multi-animal tracking system with markerless identification, and 2d estimation of posture and visual fields. *Elife*, 10:e64000.
- Wang, W., Escobedo, R., Sanchez, S., Sire, C., Han, Z., and Theraulaz, G. (2022). The impact of individual perceptual and cognitive factors on collective states in a data-driven fish school model. *PLoS computational biology*, 18(3):e1009437.
- Weiss, K., Khoshgoftaar, T. M., and Wang, D. (2016). A survey of transfer learning. *Journal of Big data*, 3(1):1–40.
- Wibral, M., Pampu, N., Priesemann, V., Siebenhühner, F., Seiwert, H., Lindner, M., Lizier, J. T., and Vicente, R. (2013). Measuring information-transfer delays. *PloS one*, 8(2):e55809.
- Worm, M., Landgraf, T., Prume, J., Nguyen, H., Kirschbaum, F., and von der Emde, G. (2018). Evidence for mutual allocation of social attention through interactive signaling in a mormyrid weakly electric fish. *Proceedings of the National Academy of Sciences*, page 201801283.
- Xue, T., Li, X., Lin, G., Escobedo, R., Han, Z., Chen, X., Sire, C., and Theraulaz, G. (2023). Tuning social interactions’ strength drives collective response to light intensity in schooling fish. *bioRxiv*, pages 2023–04.

Zhuang, F, Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., and He, Q. (2020). A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76.

Zienkiewicz, A., Barton, D. A., Porfiri, M., and Di Bernardo, M. (2015). Data-driven stochastic modelling of zebrafish locomotion. *Journal of mathematical biology*, 71(5):1081–1105.

Vaios Papaspyros

Robotics & AI Software Engineer

General Information

Mail:

vaios.papaspyros@gmail.com






Professional Disciplines

Computer Science & Engineering;
Robotics;
Machine/Deep Learning;
Collective Behaviour;
Swarm Intelligence

Technical skills

C/C++(20), Python,
Robot Operating System (ROS),
OpenCV, shell, Qt,
L^AT_EX, Software design,
Tensorflow/Keras,
PyTorch (C++/Python API)

Web Presence

Personal Website 
Twitter 
LinkedIn 
Github 
Bitbucket 

OS Experience

Linux ★★★★★
MacOS ★★★★★
Windows ★★★★★

Languages

Greek (native) ★★★★★
English (C2) ★★★★★
French (B2) ★★★★★

Experience

06/18 - 11/23 **Doctoral Assistant** [EPFL, Lausanne, Switzerland](#)

Description: I conducted and published research on collective behavior models and their transferability to mixed groups of animals and robots. I designed and implemented high-fidelity social interaction models based on analytical and machine/deep learning approaches. I also implemented the computer vision, robot planning and control, networking interface, and the low level robot software that allows for transferring the models into closed-loop real-life systems. Finally, I assisted in and taught courses and supervised various student projects and exams.

03/18 - 05/18 **Research Intern** [EPFL, Lausanne, Switzerland](#)

Description: I conducted research on social interaction models that allowed a robot to engage in meaningful bidirectional interactions with groups of fish.

06/17 - 11/17 **Research Engineer** [MEAD, Univ. of Patras, Patras, Greece](#)

Description: I led the research team in organizing the project timeline, designing, implementing, and testing the communication protocols and software tools (e.g., with ROS) that supported the operation of a heterogeneous swarm of drones.

05/16 - 10/16 **Research Intern** [Inria Nancy Grand-Est, Nancy, France](#)

Description: I conducted and published research on intelligent algorithms that allow robots to adapt from damage. More specifically, I developed a safety-aware trial-and-error algorithm (based on constrained Bayesian optimization and Gaussian processes) that allowed a humanoid to adapt to its damages and crawl again (in simulation).

Education

06/18 - 11/23 **Doctor of Philosophy - Ph.D.** [EPFL, Lausanne, Switzerland](#)

Robotics, Control, and Intelligent Systems
Supervisor: Francesco Mondada.

09/12 - 11/17 **M.Eng in Computer Engineering & Science** [Univ. of Patras, Patras, Greece](#)

GPA: 7.35 / 10
Thesis title: Safety-Aware Intelligent Trial-and-Error for Robot Damage Recovery.
Supervisors: Ioannis Hatzilygeroudis, Jean-Baptiste Mouret.

09/10 - 06/12 **High School** [Costeas-Geitonas School, Athens, Greece](#)

GPA: 19.2 / 20

Open-source project contributions

- C/C++ **Author to Behavioural Observation & Biohybrid Interaction (BOBI) framework** (<https://github.com/epfl-mobots/bobi>)
BOBI is ROS-based code that supports a robot-animal experimentation setup.
- C/C++ **Author to Lurebot low-level control code** (https://github.com/epfl-mobots/lurebot_low_control)
The repository contains the low-level code that allows the LureBot to communicate with high-level systems (e.g., BOBI)
- Python **Author to Fish Interaction moDeling framework (find)** (<https://github.com/epfl-mobots/find>)
“find” contains analysis and modelling tools (primarily) aimed at fish behaviour
- C/C++ **Co-author to robot_dart** (https://github.com/resibots/robot_dart)
robot_dart is a flexible and generic C++11 wrapper for DART and is suitable for evolutionary computation.
- C/C++ **Contributor to limbo** (<https://github.com/resibots/limbo>)
limbo is a highly templated C++11 Bayesian optimization framework.

Reviewing Experience

I have been repeatedly invited to review for various top-tier venues including:

- Nature Communications
- IROS
- ICRA
- NeurIPS (BayesOpt)
- RO-MAN
- AAMAS

Teaching Experience

• Winter Semester

2020-2022 **Basics of Mobile Robotics** EPFL
2h / week - 1st year Master of Robotics

Description: Assisted in the design and teaching of introductory exercises on mobile robot systems. The exercises included topics such as sensing, computer vision, robot control, localization, SLAM, neural networks & genetic algorithms, graph algorithms. I also contributed in guiding teams during their semester projects spanning those topics and exam corrections.

• Spring Semester

2020-2022 **Robotics practicals | Robot Operating System (ROS) basics** [EPFL](#)
4h / week - 1st year Master of Robotics

Description: I created, organized, and taught an introductory course on the Robot Operating System (ROS). During the course, I taught students the fundamentals of ROS and general good practices in robotics (software) projects. I also supervised and graded team projects that required the students to control a differential drive mobile robot first in simulation, then in real life, and conduct an assessment of the reality gap.

Publications

• Journals

- Aug 2023 *“Quantifying the biomimicry gap in biohybrid systems”*,
Papaspyros V, Theraulaz G, Sire C, Mondada F.
[Preprint/Under review](#)
- June 2023 *“A biohybrid interaction framework for the integration of robots in animal societies”*,
Papaspyros V, Burnier D, Cherfan R, Theraulaz G, Sire C, Mondada F.
[IEEE Access](#)
- Apr 2023 *“Predicting long-term collective animal behavior with deep learning”*,
Papaspyros V, Escobedo R, Alahi A, Theraulaz G, Sire C, Mondada F.
[Preprint/Under review](#)
- Apr 2022 *“The role of feedback and guidance as intervention methods to foster computational thinking in educational robotics learning activities for primary school”*,
Chevalier M, Giang C, El-Hamamsy L, Bonnet E, **Papaspyros V**, Pellet JP, Audrin C, Romero M, Baumberger B, Mondada F.
[Computers & Education](#)
- Sept 2020 *“A data-driven method for reconstructing and modelling social interactions in moving animal groups”*,
El-Hamamsy L, **Papaspyros V**, Kangur T, Mathex L, Giang C, Skweres M, Bruno B, Mondada F.
[Philosophical Transactions of the Royal Society B](#)
- Aug 2019 *“Bidirectional interactions facilitate the integration of a robot into a shoal of zebrafish Danio rerio”*,
Papaspyros V, Bonnet F, Collignon B, Mondada F.
[PLOS One](#)

• Peer-Reviewed Conferences/Workshops

Dec 2021 “Exploring a handwriting programming language for educational robots”,
El-Hamamsy L, **Papaspyros V**, Kangur T, Mathex L, Giang C, Skweres M,
Bruno B, Mondada F.

[RiE](#)

Dec 2016 “Safety-aware robot damage recovery using constrained bayesian optimiza-
tion and simulated priors”,

Papaspyros V, Chatzilygeroudis K, Vassiliades V, Mouret JB.

[BayesOpt NIPS](#)

Interests

- Machine Learning & AI
- Robotics
- Programming
- Basketball, Photography & Music