

ALIFE 2024



Proceedings of The Artificial Life Conference 2024



Edited by

Andrés Faña, Sebastian Risi, Eric Medvet, Kasper Stoy, Bert Chan,
Karine Miras, Payam Zahadat, Djordje Grbic, Giorgia Nadizar

Preface

This volume presents the proceedings of ALIFE 2024, the 2024 Conference on Artificial Life, held in Copenhagen, Denmark, July, 22th-26th 2024 (<http://2024.alife.org>).

Artificial Life, with its focus on understanding and creating lifelike systems, has always been at the forefront of interdisciplinary research. It strives to uncover the fundamental principles of living systems, investigate the processes that underpin life, and develop new forms that exhibit lifelike behaviors. This conference serves as a platform to highlight these pursuits, showcasing the latest research and developments that bridge biology and technology.

In recent years, the AI community has made tremendous strides in machine learning, neural networks, and autonomous systems. These advancements have not only enhanced our ability to create intelligent systems but also opened up new avenues for integrating ALife concepts. As we look ahead, ALife's role in the future of AI becomes increasingly significant. The principles of evolution, self-organization, adaptation, reproduction, and emergence—core tenets of ALife—are proving invaluable in addressing complex challenges in AI.

We are excited to present a diverse program that includes keynote speeches, technical sessions, workshops, and poster presentations. Our speakers and presenters come from a wide range of disciplines, reflecting the interdisciplinary nature of ALife. We encourage you to engage with the presentations, participate in discussions, and explore the potential of life-as-it-might-be.

The ALIFE 2024 Theme

The theme of this year's conference is 'Exploring new frontiers'. Building on the ancient traditions of the Vikings to travel to new and strange lands and discover new cultures and traditions, we want to dedicate ALIFE 2024 to the weird and wacky of ALife. In particular, in a time where academic research is under pressure to be rational with direct paths to impact, we will instead dedicate ALIFE 2024 to celebrate the fundamental science of artificial life.

Computational artificial intelligence is making victory rounds around the globe, with enormous potential to develop a better world and increase living standards provided we can handle the ethical challenges. The success has also impacted artificial life research, as is evident in this proceeding. However, we must not forget that this success stemmed from many years of fundamental research, initially driven by scientists' interest in understanding these artificial systems and their potential explanation power of biological intelligence. Similarly, artificial life is motivated by the desire to understand the fundamental building blocks of life, primarily through a synthetic approach—developing systems with life-like properties. The aim is to understand these systems and explore what this understanding may imply for other living systems, biological or otherwise.

While computational intelligence represents a significant step forward, it does not solve all our problems. Current technology is taking a toll on our environment. Artificial life research may lead to technologies with unprecedented power-efficiency, recyclability, scalability, robustness, and adaptability—key ingredients for combating climate change. This potential is uncertain, but without the efforts of the ALIFE community, we will never know. Therefore, we dedicate this year's ALIFE conference to 'Exploring new frontiers', maintaining a strong commitment to fundamental science in artificial life, and discovering new insights that inspire individual researchers. In the long term, these discoveries may hold the potential to transform humanity.

The ALIFE 2024 Program

We received a total of 210 full paper and abstract submissions. Our Program Committee reviewed all submissions in a double-blind process. Senior Program Committee members then performed a topic-wide meta-review to recommend acceptance/rejection decisions. As a result, 103 submissions were accepted for oral presentations and 24 for poster presentations, all of which are included in these proceedings.

The conference also hosted two special sessions to focus on specific topics that could expand the ALife landscape. In addition, fourteen workshops and eight tutorials were held. We also invited the authors of papers recently published in the *Artificial Life* journal to present their work in a dedicated session. Moreover, ALIFE 2024 hosted the 10th edition of the Virtual Creatures Competition and a related workshop. Finally, as part of ALIFE 2024, we introduced an art exhibition with international and local artists and an art session. Last, we confirmed Mind-match, an event of scientific matchmaking introduced within ALIFE 2023 to create new opportunities for scientists who have similar approaches to their work or study similar problems to meet.

More specifically, the conference program of this year included the following:

- Five keynote presentations by internationally renowned speakers:
 - Katie Bentley (The Francis Crick Institute, UK – Kings College London, UK – Boston University, USA)
 - Kate Adamala (University of Minnesota, USA)
 - Juan Perez-Mercader (Harvard University, USA)
 - Kohei Nakajima (University of Tokyo, Japan)
 - The winner of the 2024's International Society for Artificial Life (ISAL) award for lifetime achievement
- Two dedicated poster sessions.
- Two special sessions:
 - The Distributed Viking: Cellular Automata, Distributed Dynamical Systems, and Their Applications to Intelligence, organized by Stefano Nichele, Hiroki Sayama, Chrystopher Nehaniv, Eric Medvet and Mario Pavone
 - ALife And Society VIII, organized by Imran Khan and Peter Lewis
- Fourteen workshops:
 - 6th International Workshop on Agent-Based Modelling of Human Behaviour (ABMHuB'24), organized by Soo Ling Lim and Peter J. Bentley
 - The evolution of things, organized by A. E. Eiben and Karine Miras
 - Evolution, Criticality and Creativity in Collective Intelligence (ECCCI), organized by Kazuya Horibe and Michael Crosscombe
 - SB-AI 9. From imitative to biologically plausible synthetic models of natural cognition, organized by Pasquale Stano, Luisa Damiano and Yutetsu Kuruma
 - Molecular Communication Approaches for Wetware Artificial Life, organized by Pasquale Stano, Michael Barros, Malcom Egan, Murat Kuscu, Yutetsu Kuruma and Takashi Nakano
 - Hackathon: ALife, Music & the Visual Arts, organized by Elias Najarro and Claire Glanois
 - Emerging Researchers in Artificial Life, organized by Piotr Walas, Federico Pigozzi, Fernando Rodriguez Vergara, Imy Khan, Ane Kristine Espeseth, Lio Hong and Gabriel J. Severino
 - ALife in Organizations, organized by Alberto Montebelli, Imran Khan and Gary Linnéusson
 - Making Time: Temporality in natural and artificial systems, organized by Adam Rostowski and Fernando Rodriguez
 - ALife Encyclopedia Hackathon, organized by Emily Dolson
 - Goal-Directed behavior in life and non-life, organized by Martin Biehl, Richard Löffler, Miguel Aguilera, Omer Markovitch, Artemy Kolchinsky and Manuel Baltieri
 - CHEMALIFORMS IV: The Fourth Workshop on Chemistry and Artificial Life Forms, organized by Jitka Čejková, Richard Löffler and Steen Rasmussen
 - Nature Inspired Simulation and Visualization of Complex Networks, organized by Oskar Elek, Angus Forbes and Tarin Ziyae
 - Virtual Creatures Competition, organized by Kam Bielawski, Piper Welch, Caitlin Grasso and Karine Miras
- Eight tutorials:
 - Engineering the open-ended evolution of synthetic biology, organized by Michiel Stock, Thomas Goroehowski and Simeon Castle
 - Automatic Design of Robot Bodies and Brains with Evolutionary Algorithms, organized by Kyrre Glette, Kai Olav Ellefsen, Emma Stensby Norstein and Ege de Bruin

- JAX for Scaling Up Artificial Life, organized by Ettore Randazzo and Bert Chan
- SimER: Simulation in Evolutionary Robotics, organized by Anthony J. Clark and Jared M. Moore
- Neuroevolution, organized by Sebastian Risi, Risto Miikkulainen, David Ha and Yujin Tang
- Phylogenies: how and why to track them in artificial life, organized by Emily Dolson, Alexander Lalejini, Matthew Moreno and Jack Garbus
- SwissGL/GPU: tiny libraries for tiny and beautiful programs on the web, organized by Alexander Mordvintsev
- Bio-inspired Data and Network Science with PolyPhy, organized by Oskar Elek
- Artificial Life journal session:
 - Recently published Artificial Life journal papers.
- Virtual creatures competition:
 - 10th edition of the Virtual Creatures Competition and a related workshop, organized by Karine Miras, Caitlin Grasso, Piper Welch and Kam Bielawski.
- Boundary Iconology art exhibition by international and local artists:
 - Brandon Tay
 - Christina König
 - Cody Lukas & Jonas Jørgensen
 - Karlie Zhang & Sara Zebulon Riise
 - Mochu
 - Philip Ullman
- Art session:
 - Performance and Panel with Mochu (artist), an invited scientist, Mandus Ridefelt (curator) and Claire Glanois (Art Chair).
- Mind matching event:
 - Connecting scientists to develop new collaborations and projects based on shared interests.

About the Editors

Andrés Faíña (General Chair)
IT University of Copenhagen

Andrés Faíña is an Associate Professor at the IT University of Copenhagen where he is part of the Robotics, Evolution and Art Lab (REAL). Faíña completed his PhD in robotics from the University of A Coruña, Spain. His interests include modular robots, artificial neural networks, evolutionary robotics, and lab automation and his works combine hardware development with artificial intelligence techniques. He has implemented several robotic prototypes solutions for the industry (e.g., a climbing robot for grit blasting, underwater robot for hull cleaning, and hardware solutions for lab automation) and academia where he tries to evolve "virtual creatures" in hardware by using modular robots. He is a co-founder of Flow Robotics, a company that manufactures affordable and easy to use liquid handling robots.

Sebastian Risi (Vice-general Chair)
IT University of Copenhagen

Sebastian Risi is a Full Professor at the IT University of Copenhagen where he directs the Creative AI Lab, and co-directs the Robotics, Evolution and Art Lab (REAL). Before joining ITU, he did a postdoc at Cornell University and before that, he obtained a Ph.D. from the University of Central Florida. His research currently focuses on the emerging field of collective intelligence for deep learning, investigating how we can make current AI approaches more robust and adaptive. As part of his ERC Consolidator Grant GROW-AI, he and his group are taking insights from artificial life and machine learning to develop novel approaches for growing neural networks that can adapt quickly during their lifetime.

Eric Medvet (Programme Chair)
University of Trieste

Eric Medvet is an Associate Professor of computer engineering at the Department of Engineering and Architecture (DIA), University of Trieste, Italy. He is the head of the Evolutionary Robotics and Artificial Life Lab and the co-head of the Machine Learning Lab. His research focuses on Evolutionary Computation, Machine Learning, and their applications to computer- and engineering-related topics. In recent years, he has concentrated on Grammatical Evolution, a form of Genetic Programming, and on Evolutionary Robotics, within the contexts of Artificial Life and multi-agent systems.

Kasper Stoy (Financial Chair)
IT University of Copenhagen

Kasper Stoy is a Full Professor at the Computer Science Department, IT University of Copenhagen, Denmark where he co-directs the Robotics, Evolution, and Art Lab (REAL) and serves as the Head of PhD School. He is an active player in the international robot research community and reviews for all major journals and conferences in robotics. He has stayed for extended periods at University of Southern California, Harvard University, University of Tarapacá (Chile), and Seam Reap (Cambodia). He holds a Ph.D. degree in computer system engineering from the University of Southern Denmark (2003) where he also worked as assistant professor (2003-2006) and associate professor (2006-2013). He also co-founded Universal Robots, a company that focuses on user-friendly robot arms for industrial applications.

Bert Chan (Special Sessions Chair)
Google DeepMind

Bert Chan is a Researcher Engineer at Google DeepMind based in Tokyo. He was also an external collaborator at Inria, and a visiting researcher at the Institute for Advanced Study. Bert's research interests include artificial life, complex systems and evolutionary computation. He discovered a continuous cellular automata called Lenia, which received the Outstanding Publication of 2019 award by the International Society for Artificial Life.

Karine Miras (Workshop and Tutorial Chair)
Vrije Universiteit Amsterdam

Dr. Karine Miras is an assistant professor in the Computational Intelligence group at the Vrije Universiteit Amsterdam, Netherlands. Her research lies within Evolutionary Robotics and Artificial Life with the objectives of designing better autonomous systems and contributing to biological discussions. Specifically, her research interests involve the influence of the environment on creatures' traits and the incorporation of developmental mechanisms into the DNA structures of evolvable artificial creatures.

Payam Zahadat (Keynote Chair)
IT University of Copenhagen

Payam Zahadat is an Associate Professor at the Computer Science Department, IT University of Copenhagen, Denmark. Before ITU, she was a postdoctoral researcher at the Artificial Life lab, University of Graz, Austria. Her research includes both theoretical studies and applications in the field of swarm robotics, modular robotics, collective intelligence, self-organizing systems, evolutionary robotics, and bio-inspired computation. She studies complex collective systems, and designs computational methods for distributed artificial systems by taking inspiration from natural systems.

Djordje Grbic (Local Chair)
IT University of Copenhagen

Dr. Djordje Grbic is a computer scientist with over 15 years of experience in AI applied to operations research, statistical methods in biology, and reinforcement learning. He received his Ph.D. in Applied Machine Learning to Evolutionary Biology at the University of Geneva in Switzerland. Djordje is currently an Associate Professor at ITU Copenhagen. He published numerous peer-reviewed articles and is teaching at master's level. His current research focuses on meta-reinforcement learning and reinforcement learning applied to operations research.

Giorgia Nadizar (Proceedings Chair)
University of Trieste

Giorgia Nadizar is a third-year Ph.D. student in Applied Data Science and Artificial Intelligence at the University of Trieste. Her research interests lie at the intersection of embodied AI and explainable/interpretable AI, focusing on developing robotic controllers that are both effective and directly human intelligible. Throughout her academic journey, she has gained experience in multiple research environments, like the Centrum Wiskunde & Informatica in Amsterdam, the ISAE-Supaero in Toulouse, and the Massachusetts Institute of Technology (MIT) in the USA.

Acknowledgments

ALIFE 2024 would not have been possible without the help of many dedicated individuals. First and foremost, I (Andrés Faña) would like to extend my heartfelt thanks to all the co-organizers of ALIFE 2024: Sebastian Risi (Vice-general Chair), Eric Medvet (Program Chair), Kasper Stoy (Funding Chair), Morten Roed Frederiksen (Funding Chair), Bert Chan (Special Session Chair), Karine Miras (Workshop & Tutorial Chair), Payam Zahadat (Keynote Speaker Chair), Djordje Grbic (Local Chair), Giorgia Nadizar (Proceedings Chair), Elias Najarro (Community and Communication Chair), Sidney Pontes-Filho (Community and Communication Chair), Jonas Haaugaard Jensen (Social Event Chair), Claire Glanois (Art Chair), Mandus Ridefelt (Art Chair), Nicolas Bessone (Hybrid Chair), and Erwan Plantec (Hybrid Chair). They contributed to various aspects such as the proceedings, conference program, review process, keynote speaker selections, online platforms, social activities, art exhibitions, and countless other tasks to enhance the conference. Their sustained support has been invaluable. In addition, I would also like to thank all the members of the Robots, Evolution, and Art Laboratory (REAL) at the IT University of Copenhagen, especially Rodrigo Moreno (minute taker), and the student volunteers for their work and help.

The organizing committee extends its gratitude to all reviewers and meta-reviewers for their crucial contributions to the review process. We deeply appreciate the reviewers who generously offered their help at different stages, ensuring the highest quality feedback. We also congratulate and thank all the authors for their exceptional submissions.

We are deeply grateful to all the artists who accepted the invitation to exhibit their art pieces at the conference: Brandon Tay, Christina König, Cody Lukas & Jonas Jørgensen, Karlie Zhang & Sara Zebulon Riise, Philip Ullman and Mochu. Your art will inspire and enrich the artificial life community. Thank you for sharing your incredible talents with us.

We would also like to extend our gratitude to the Emerging Researchers in ALife (ERA) and the International Society for Artificial Life (ISAL) for their efforts in organizing numerous events aimed at integrating young scientists into our community.

We would also thank Olaf Witkowski (ISAL President), Harold Fellermann (ISAL Conference Liaison), Jitka Čejková (General Chair of ALIFE 2021), Silvia Holler (General Chair of ALIFE 2022), Hiroyuki Iizuka (General Chair of ALIFE 2023), Manuel Baltieri, Alyssa Adams and Richard Löffler for their suggestions, help and invaluable advice to organize ALIFE 2024.

Finally, we are very grateful to the following sponsors for their financial and in-kind support for the conference, which was essential in making ALIFE 2024 a successful event:

- Carlsberg Foundation
- Google
- IT University of Copenhagen

Senior Program Committee

Aishwaryaprajna
Bailey, Richard
Beer, Randall
Fellermann, Harold
Khan, Imran
Roli, Andrea
Stepney, Susan

Amos, Martyn
Baltieri, Manuel
Cartwright, Julyan
Froese, Tom
Nehaniv, Chrystopher
Sinapayen, Lana
Stoy, Kasper

Aranha, Claus
Banzhaf, Wolfgang
Chan, Bert
Ikegami, Takashi
Nichele, Stefano
Stano, Pasquale
Taylor, Tim

Program Committee

Adami, Christoph
Antonioni, Alberto
Bacardit, Jaume
Capobianco, Silvio
Cheney, Nick
Cully, Antoine
Dolson, Emily
Faulkner Rainford, Penn
Gershenson, Carlos
Goñi-Moreno, Ángel
Hashimoto, Yasuhiro
Holler, Silvia
Ieropoulos, Ioannis A.
Kojima, Hiroki
Lizier, Joseph
Manzoni, Luca
Masumori, Atsushi
Moreno, Rodrigo
Naya, Martin
Ofria, Charles
Philippides, Andrew
Polani, Daniel
Randazzo, Ettore
Sayama, Hiroki
Stanton, Adam
Trefzer, Martin
Villani, Marco
Zarco, Mario

Adams, Alyssa
Arita, Takaya
Bartlett, Stuart
Carletti, Timoteo
Costello, Ben
Cussat-Blanc, Sylvain
Ellefsen, Kai Olav
Fortuna, Miguel A.
Glanois, Claire
Grabowski, Laura
Hauser, Helmut
Hudcová, Barbora
Iizuka, Hiroyuki
Kozdon, Katarzyna
Loeffler, Richard
Markovitch, Omer
Medvet, Eric
Nadizar, Giorgia
Nitschke, Geoff
Oka, Mizuki
Pietropolli, Gloria
Pontes-Filho, Sidney
Risi, Sebastian
Serra, Roberto
Suzuki, Keisuke
Trianni, Vito
Vostinar, Anya

Aguilera, Miguel
Aubert-Kato, Nathanael
Boumaza, Amine
Channon, Alastair
Crosscombe, Michael
Dittrich, Peter
Faina, Andres
Freire-Obregón, David
Glette, Kyrre
Han, The Anh
Hintze, Arend
Iacca, Giovanni
Kalvala, Sara
Letelier, Juan-Carlos
Manoonpong, Poramate
Marriott, Chris
Moreno, Matthew Andres
Najarro, Elias
Nolfi, Stefano
Parrilla Gutierrez, Juan Manuel
Pigozzi, Federico
Powers, Simon
Rossi, Federico
Soros, Lisa
Suzuki, Reiji
Veloz, Tomas
Witkowski, Olaf

The ALIFE 2024 Proceedings Editors:

Andrés Faíña
Sebastian Risi
Eric Medvet
Kasper Stoy
Bert Chan
Karine Miras
Payam Zahadat
Djordje Grbic
Giorgia Nadizar

Conference Program

i Preface

Short Abstracts of Keynote Presentations

1

- 1 Active Perception in cellular adaptive behaviour
Katie Bentley
- 1 Life but not alive: synthetic biological cells
Kate Adamala
- 2 Making bio-chemistry free artificial life in a test tube
Juan Perez Mercader
- 2 Physicalizing Deep Learning
Kohei Nakajima

General Conference

3

Accepted oral presentations

3

- 4 A Computational Model of Developmental Exaptations
Alexa Renner, Émile Marois, Julian Fiorito, Jacob Ashworth and Jason Yoder
- 13 A Macrophage-Inspired Approach To Fault Reaction In The Collective Perception Problem
Patrick Shortall, Chanelle Lee and Jonathan Lawry
- 22 A Saucerful of Secrets: Open-Ended Organizational Closure in the Game of Life
Fernando Rodriguez Vergara and Phil Husbands
- 32 A Swarm Control by Distributed Attractant and Repellent, and its Application for a Stage Effect in Contemporary Ballet
Tatsuo Unemi and Daniel Bisig
- 41 Aevol 4b: Bridging the gap between artificial life and bioinformatics
Hugo Daudey, David P Parsons, Eric Tannier, Vincent Daubin, Bastien Boussau, Vincent Liard, Romain Gallé, Jonathan Rouzaud-Cornabas and Guillaume Beslon
- 50 An Agent-based, Interactive Simulation Model of Root Growth
Marc Mußmann, Daniel Nicolas Hofstadler and Sebastian von Mammen
- 59 Analyzing the Sensibility of Visual Language Models Using an Evolving Image Generation System: Focusing on Color Impressions and Sound Symbolism
Ryoma Shinto and Hiroyuki Iizuka
- 68 Artificial Minimal Self on Free Energy Principle for Autonomous Cooperative Behavior
Kanako Esaki, Tadayuki Matsumura, Takeshi Kato, Shunsuke Minusa, Yang Shao and Hiroyuki Mizuno
- 77 Assessing the ability of the MODES toolbox to detect hallmarks of openendness
Clifford Bohm, Ivy Zhang and Emily Dolson
- 87 Best Practices for Differentiable Soft Robot Modeling and Optimization with the Material Point Method
Kameron Bielawski, Jakub Rozlivek, Matej Hoffman and Josh Bongard
- 95 Breaching the Bottleneck: Evolutionary Transition from Reward-Driven Learning to Reward-Agnostic Domain-Adapted Learning in Neuromodulated Neural Nets
Solvi Arnold, Reiji Suzuki, Takaya Arita and Kimitoshi Yamazaki
- 104 Capital as Artificial Intelligence
Cesare Carissimo and Marcin Korecki
- 113 Classifying the fractal parameter space of the Lenia Orbium
Ivan Yevenko

- 116 Coevolutionary Heuristics for Maximization of Expected Utility
Thomas Willkens and Jordan Pollack
- 126 Collective Innovation in Groups of Large Language Models
Eleni Nisioti, Sebastian Risi, Ida Momennejad, Pierre-Yves Oudeyer and Clément Moulin-Frier
- 137 Commingling Artificial Life and Interactive Machine Learning in Diffractive Artistic Research Practice
Jack Armitage, Victor Shepardson, Miguel Crozzoli and Thor Magnusson
- 146 Comparative Analysis of Heteroclinic Network Construction Algorithms for Simulating *C. elegans* Behavioral and Neural Data
Denizhan Pak and Randall Beer
- 155 Creating Network Motifs with Developmental Graph Cellular Automata
Riversdale Waldegrave, Susan Stepney and Martin Trefzer
- 164 Crossover Destructiveness in Cartesian versus Linear Genetic Programming
Mark Kocherovsky and Wolfgang Banzhaf
- 173 Curious POET: Intrinsic Motivation Improves Exploration Efficiency
Robert Mash, Gregory Castanon and Jared Culbertson
- 183 Deriving Community Models with Evolutionary Robotics: A Case Study of Sensory Pollution
Eden Forbes and Randall Beer
- 192 Deriving The Intrinsic Viability Constraint of an Emergent Individual from First Principles
Randall Beer, Connor McShaffrey and Tom Gaul
- 201 Dissecting Viability in Multi-Agent Systems
Connor McShaffrey and Randall Beer
- 210 Dynamics of human exploration of novel sensorimotor contingencies
Natasha Skuljan, Christopher Erb and Matthew Egbert
- 219 Effective Design and Interpretation in Voxel-Based Soft Robotics: A Part Assembly Approach with Bayesian Optimization
Takumi Saito and Mizuki Oka
- 229 Emergence of autopoietic vesicles able to grow, repair and reproduce in a minimalist particle system
Thomas Cabaret
- 237 Emergent Multiscale Structures and Generative Potential of Isotropic Neural Cellular Automata
Alexander Mordvintsev and Eyvind Niklasson
- 246 Emergent Orchestras: A Modular Framework for Musical Robot Swarms
Lluc Bono Rosselló, Muhanad H.M. Alkilabi, Elio Tuci, Hugues Bersini and Andrea Giovanni Reina
- 255 Energy-Based Models for Virtual Creatures
Junior Rojas
- 258 Environmental connectivity influences the origination of adaptive processes
John Shea, Sydney Leither, Max Foreback, Emily Dolson and Alexander Lalejini
- 268 Evolution and Analysis of Embodied Agents on a Numerical Cognition Task: Sequential Counting
Eduardo J. Izquierdo
- 276 Evolution and Analysis of Respiratory Odor Navigation in Embodied Agents
Gabriel Severino and Ann-Sophie Barwich
- 285 Evolution of Cooperation via Joint Commitments and Reputation
Marcus Krellner and The Anh Han
- 288 Evolution of Rewards for Food and Motor Action by Simulating Birth and Death
Yuji Kanagawa and Kenji Doya
- 297 Evolving Collective Behavior in Self-Organizing Particle Systems
Devendra Parkar, Kirtus Leyba, Raylene Faerber and Joshua Daymude

- 307 Evolving Self-Assembling Neural Networks: From Spontaneous Activity to Experience-Dependent Learning
Erwan Plantec, Joachim Winther Pedersen, Milton Llera Montero, Eleni Nisioti and Sebastian Risi
- 317 Exploring Affordances in Open-Ended Online Collective Robotics
Amine Boumaza
- 326 From Text to Life: On the Reciprocal Relationship between Artificial Life and Large Language Models
Eleni Nisioti, Claire Glanois, Elias Najarro, Andrew Dai, Elliot Meyerson, Joachim Pedersen, Laetitia Teodorescu, Conor Francis Hayes, Shyam Sudhakaran and Sebastian Risi
- 344 Fundamental Vision-Based Navigation Algorithms: Local Sequential, Global Diffusive, and Direct Pathing
Patrick Govoni and Pawel Romanczuk
- 348 Hybrid Modelling of Stimuli-Based and Density-Induced Aggregation in *C. elegans*
Aymeric Vellingier, Nemanja Antonic and Elio Tuci
- 357 I Like to Move It: Using Information Theory to Measure Sensorimotor Coupling in Cells
Sebastian Dohnány, Harriet A. Smith, Josh Bourne, Kelvin Van-Vuuren, Gina Gilpin, Irene M. Aspalter, Pedro A.M. Mediano and Katie Bentley
- 367 Indirect reciprocity with stochastic and dual reputation updates
Yohsuke Murase and Christian Hilbe
- 369 Institutional Incentives for promoting Evolution of Commitment Compliance and Cooperation in Spatial Public Goods Games
Lucas S. Flores and The Anh Han
- 373 Interactive Experiences with a Web-based Drummer Bot for Finger-Tapping
Cagri Erdem and Carsten Griwodz
- 376 ISicell: A Participatory Methodology and Platform for Collaborative Agent-Based Modeling in Cell Biology
Florian Cogoni, David Bernard, Roxana Khazen, Salvatore Valitutti, Valérie Lobjois and Sylvain Cussat-Blanc
- 384 JAXLife: An Open-Ended Agentic Simulator
Chris Lu, Michael Beukman, Michael Matthews and Jakob Foerster
- 394 Large Language Models as In-context AI Generators for Quality-Diversity
Bryan Lim, Manon Flageat and Antoine Cully
- 404 Lattice reservoirs: symmetry breaking and information flow in physical reservoir computing
Thomas Aven, Johannes H. Jensen and Gunnar Tufte
- 414 Meta-Learning an Evolvable Developmental Encoding
Milton Montero, Eleni Nisioti, Erwan Plantec, Joachim Pedersen and Sebastian Risi
- 424 Methods to Estimate Cryptic Sequence Complexity
Matthew Andres Moreno
- 427 Mimicry and the Emergence of Cooperative Communication
Dylan Cope and Peter Mcburney
- 437 Minimal Self in Humanoid Robot "Alter3" Driven by Large Language Model
Takahide Yoshida, Suzune Baba, Atsushi Masumori and Takashi Ikegami
- 444 Modeling language contact with the Iterated Learning Model
Seth Bullock and Conor Houghton
- 452 Modular Growth of Hierarchical Networks: Efficient, General, and Robust Curriculum Learning
Mani Hamidi, Sina Khajehabdollahi, Emmanouil Giannakakis, Tim J. Schäfer, Anna Levina and Charley Wu
- 463 Neural Lindenmayer Systems
Christoph Walter Senn, Carmen Mei-Ling Frischknecht-Gruber, Mathias Weyland and Rudolf Marcel Fuchsli
- 466 NoiseNCA: Noisy Seed Improves Spatio-Temporal Continuity of Neural Cellular Automata
Ehsan Pajouheshgar, Yitao Xu and Sabine Süsstrunk
- 475 Non-Spatial Hash Chemistry
Hiroki Sayama

- 478 On All Fours: A 3D Framework to Study Closed-loop Control of Quadrupedal Mouse Locomotion
Shravan Tata Ramalingasetty, Sergey N. Markin, Andrew B. Lockhart, Jonathan Arreguit, Natalia A. Shevtsova, Auke J. Ijspeert, Ilya A. Rybak and Simon M. Danner
- 489 On Creativity and Open-Endedness
Lisa Soros, Alyssa Adams, Stefano Kalonaris, Olaf Witkowski and Christian Guckelsberger
- 500 On the de novo Evolution of Mutualism in Symbulation
Connor Croft, Kiara Johnson, Anya Vostinar and Emily Dolson
- 509 On the Emergence of Enzymes in an Artificial Chemistry
Giampaolo Folea and Germán Kruszewski
- 512 Ouroborotic Swarm Control: Collective Dynamics as a Reservoir
Koh Uehara, Katsuma Inoue, Yasuo Kuniyoshi and Kohei Nakajima
- 521 Parameterising a Computational model of Supercoiling in Plasmid DNA with Biophysics modelling
Penn Faulkner Rainford, Victor Velasco-Berrelleza, Aalap Mogre, Chares J Dorman, Sarah Harris, Carsten Kroeger and Susan Stepney
- 530 Phenotypic Species Definitions for Genetic Improvement of Source Code
Zsolt Nemeth, Penn Faulkner Rainford and Barry Porter
- 540 Phylogeny-Informed Interaction Estimation Accelerates Co-Evolutionary Learning
Jack Garbus, Thomas Willkens, Alexander Lalejini and Jordan Pollack
- 550 Planetary regulation on the test tube: a Synthetic Daisyworld
Victor Maull, Jordi Pla, Nuria Conde and Ricard Solé
- 553 Planetary Scale Replication as an Agnostic Biosignature
Harrison Smith and Lana Sinapayen
- 556 Predicting the Unpredictable: Using replay experiments to disentangle how evolutionary outcomes are altered by adaptive momentum
Austin Ferguson, Charles Ofria and Clifford Bohm
- 565 Problem-Solving in Language Model Networks
Ciaran Regan, Alexandre Gournail and Mizuki Oka
- 574 Reality of Affordances: A Category-Theoretic Approach
Ryuzo Hirota, Hayato Saigo and Shigeru Taguchi
- 584 Refractive Computation: parallelizing logic gates across driving frequencies in a mechanical polycomputer.
Shawn Beaulieu, Piper Welch, Atoosa Parsa, Corey O'Hern, Rebecca Kramer and Josh Bongard
- 593 Selection Driven Scaling Regimes in Assembly Theory
Gage Siebert, Daniel Czeagal, Leroy Cronin and Sara Walker
- 596 Self-Adaptive Robustness of Applied Neural-Network-Soups
Maximilian Zorn, Philipp Altmann, Gerhard Stenzel, Michael Kölle, Claudia Linnhof-Popien and Thomas Gabor
- 605 Self-propulsion of 3-Phenylpropionaldehyde droplets
Ayase Kawamura, Shinpei Tanaka and Muneyuki Matsuo
- 608 Spatial Community Structure Promotes Extremism in an Adaptive Spatial Social Network Model
Seth Bullock
- 615 Spatial sensitivity of the evolutionary swarm chemistry model
Juste Raimbault
- 618 Specialization or Generalization: investigating NeuroEvolutionary choices via Virtual fMRI
Kevin Godin-Dubois, Sylvain Cussat-Blanc and Yves Duthen
- 627 Strong leaders don't cheat: an evolutionary appraisal of population heterogeneity and leadership
Carlo Longhi, Luis A. Martinez-Vaquero, Fernando Wario Vázquez and Vito Trianni
- 637 Surface directed protocell formation
Silvia Holler and Martin M. Hanczyc

- 640 Takorobo: Towards Closed-Loop Body-Driven Locomotion Processing
Takehiro Hasegawa, Max Austin, Hidenobu Sumioka, Yasuo Kuniyoshi and Kohei Nakajima
- 649 The Digital Gallows: Threat of Institutional Punishment Fosters the Emergence of Cooperation
Theodor Cimpanu and The Anh Han
- 660 The Effect of Noise on the Density Classification Task for Various Cellular Automata Rules
Annajirao Challa, Duxiao Hao, Jordan Rozum and Luis Rocha
- 665 The Simulator: An interactive evolved virtual creature testbed
Michael Finn and Thomas Bräunl
- 673 Toward Artificial Open-Ended Evolution within Lenia using Quality-Diversity
Maxence Faldor and Antoine Cully
- 683 Towards Stepping Stones in Goal-Directed Search: Improving Image Perceptual Similarity by combining Evolution, Backpropagation, and Fourier Transforms
Jackson Dean and Nick Cheney
- 693 Trackable Agent-based Evolution Models at Wafer Scale
Matthew Andres Moreno, Connor Yang, Emily Dolson and Luis Zaman
- 705 Untangling the Effects of Down-Sampling and Selection in Genetic Programming
Ryan Boldi, Ashley Bao, Martin Briesch, Thomas Helmuth, Dominik Sobania, Lee Spector and Alexander Lalejini
- 717 Zealot Detection in Probabilistic Social Learning Using Bounded Confidence
Jonathan Lawry

Accepted poster presentations

726

- 726 A Deep Learning Approach to Understanding the Integration of Tools into Body Schema
Haruna Kasashima, Wataru Noguchi, Yasumasa Tamura, Masahito Yamamoto and Hiroyuki Iizuka
- 729 A Developmentally Morpho-Motor Coupled System Using Neural Cellular Automata
Hiroki Sato and Atsushi Masumori
- 732 A Simulation Environment for the Neuroevolution of Ant Colony Dynamics
Michael Crosscombe, Ilya Horiguchi, Norihiro Maruyama, Shigeto Dobata and Takashi Ikegami
- 735 Adapting to time: why nature chose to evolve a diverse set of neurons
Karim Habashy, Benjamin Evans, Dan Goodman and Jeffrey Bowers
- 738 Asymmetrical division in protocells
Marco Villani and Roberto Serra
- 741 Decentralized cellular timekeeping based on entropy as a regulatory mechanism for growth and aging
Jahan Penny-Dimri
- 744 Emergent Dynamics in Neural Cellular Automata
Yitao Xu, Ehsan Pajouheshgar and Sabine Süsstrunk
- 747 Floatiles: Self-Assembly Based on Cheerios Effect and Aperiodic Monotiles
Georgii Karelin
- 751 Flying Bit in a Model of Excitable Media on Face-centered Cubic Lattice
Shigeru Ninagawa
- 754 Gene-Regulated Neural Cellular Automata
Arend Hintze, Mustafa Al-Hammadi and Eric Libby
- 764 Latent Neural Cellular Automata for Resource-Efficient Image Restoration
Andrea Menta, Alberto Archetti and Matteo Matteucci
- 773 Machine-learning-based prediction of DNA structure volume for Quality-Diversity exploration
Maya Hyakuzuka and Nathanael Aubert-Kato
- 776 Movement of Micrometer Sized Oil Droplets in Aqueous Solutions Follows Ion Movement
Wayne Briner

- 778 Myelin Figures Formation in Surfactants: Insights for Wet Artificial Life
Nina Novakova, Jan Machacek and Jitka Cejkova
- 785 Only ad-hoc solutions can solve the shutdown problem
Daniel Platt
- 788 Plasticity in Swarming Behavior: Introducing Social Networks to Boids Model
Atsushi Masumori, Itsuki Doi and Takashi Ikegami
- 791 Quantifying Autonomy in Ant Colonies Using Non-Trivial Information Closure
Ilya Horiguchi, Norihiro Maruyama, Dobata Shigeto, Michael Crosscombe and Takashi Ikegami
- 794 Reconstruction and analysis of roaming termites' dynamical networks
Louis Devers, Perrine Bonavita and Christian Jost
- 797 Reverse-engineering artificial life under the free-energy principle
Johan Medrano
- 800 Scalable and Stable Patterns Generated by Heterogeneous Boid
Mari Nakamura
- 803 Self-Replicating Prompts for Large Language Models: Towards Artificial Culture
Gerhard Stenzel, Maximilian Zorn, Philipp Altmann, Maximilian Balthasar Mansky, Michael Kölle and Thomas Gabor
- 812 The role of n in the n-mates evaluation method: a quantitative analysis
Paolo Pagliuca and Alessandra Vitanza
- 815 Towards understanding the eco-evo-devo of artificial creatures in 3D physical simulation
Siti Aisyah Binti Jaafar, Reiji Suzuki and Takaya Arita

Special Session: The Distributed Viking: Cellular Automata, Distributed Dynamical Systems, and Their Applications to Intelligence **818**

Accepted oral presentations **818**

- 818 Autoipoiesis in RealLife Euclidean Automata
Thomas Gaul
- 827 Coralai: Intrinsic Evolution of Embodied Neural Cellular Automata Ecosystems
Aidan Barbieux and Rodrigo Canaan
- 830 Enabling Evolvability with Neutral Networks in Multiscale Neural Cellular Automata
Kameron Bielawski and Josh Bongard
- 839 Hash Chemistry on a Cellular Grid: An Open-Ended Artificial Chemistry System with Computational Efficiency and Nontrivial Spatio-Temporal Dynamics
Hiroki Sayama
- 842 Networks of Binary Necklaces Induced by Elementary Cellular Automata Rules
Lapo Frati, Csenge Petak and Nick Cheney
- 851 Non-Platonic Autoipoiesis of a Cellular Automaton Glider in Asymptotic Lenia
Q. Tyrell Davis
- 854 Programming Weakly Connected Chemical Oscillators to Simulate Logic Gates and Cellular Automata
Juan Manuel Parrilla Gutierrez, Abhishek Sharma, Soichiro Tsuda and Leroy Cronin
- 861 Simulating an Artificial Biome of Plants with Biomaker CA
Ettore Randazzo and Alexander Mordvintsev
- 871 Thermina: A minimal model of autonomous agency from the lens of stochastic thermodynamics
Miguel Aguilera and Xabier E. Barandiaran
- 880 Unlocking Nature's Design through Neural Cellular Automata
Frantisek Koutensky, Petr Simanek and Alexander Kovalenko

Special session: ALife And Society VIII **883**

Accepted oral presentations **883**

- 883 (Un)sustainable structural coupling – a testimony for snail intelligence
Floor Schukking, Tomas Veloz and Stefan Leijnen
- 886 Breaking Echo Chambers in Opinion Dynamics with Spatial Information
Mohsen Raoufi, Heiko Hamann and Pawel Romanczuk
- 889 Incorporating Social Expectations into the Expectation Event Calculus
Nathan Lloyd and Peter R. Lewis
- 898 Power-Sensitive Design using Higher-Order Cybernetics Patterns
Asimina Mertzani and Jeremy Pitt
- 908 Speculating on Blockchain as an Unstoppable ‘Nature’ Towards the Emergence of Artificial Life
Botao Amber Hu and Fang Ting

Accepted poster presentations **913**

- 913 AI, Mortality and Existential Risk
Inman Harvey

Special Session: Artificial Life Journal Session **916**

- 916 An Afterword to Rise of the Self-Replicators: Placing John A. Etzler, Frigyes Karinthy, Fred Stahl, and Others in the Early History of Thought About Self-Reproducing Machines
Tim Taylor
- 916 A Spatial Artificial Chemistry Implementation of a Gene Regulatory Network Aimed at Generating Protein Concentration Dynamics
Iliya Miralavy and Wolfgang Banzhaf
- 916 A Survey of Recent Practice of Artificial Life in Visual Art
Zi-Wei Wu, Huamin Qu, and Kang Zhang
- 916 Emergent Resource Exchange and Tolerated Theft Behavior Using Multiagent Reinforcement Learning
Jack Garbus and Jordan Pollack
- 916 Information, Coding, and Biological Function: The Dynamics of Life
Julyan H. E. Cartwright, Jitka Čejková, Elena Fimmel, Simone Giannerini, Diego Luis Gonzalez, Greta Goracci, Clara Grácio, Jeanine Houwing-Duistermaat, Dragan Matić, Nataša Mišić, Frans A. A. Mulder, and Oreste Piro
- 916 Motivations for Artificial Intelligence, for Deep Learning, for ALife: Mortality and Existential Risk
Inman Harvey

Strong leaders don't cheat: an evolutionary appraisal of population heterogeneity and leadership

Carlo Longhi^{1,2}, Luis A. Martinez-Vaquero^{3,4,5}, Fernando Wario Vázquez² and Vito Trianni²

¹Department of Computer, Control and Management Engineering Antonio Ruberti, Sapienza University of Rome, 00185 Rome, Italy

²Institute of Cognitive Sciences and Technologies (ISTC), National Research Council (CNR), 00196 Rome, Italy

³Grupo de Sistemas Complejos, Universidad Politécnica de Madrid, 28040 Madrid, Spain

⁴DEFE, Escuela Técnica Superior de Arquitectura, Universidad Politécnica de Madrid, 28040 Madrid, Spain

⁵Grupo Interdisciplinar de Sistemas Complejos, Universidad Carlos III de Madrid, 28911 Leganés, Madrid, Spain
longhi@diag.uniroma1.it

Abstract

Coordination and cooperation are crucial features of many natural and artificial systems. Among the many mechanisms that have been proposed to support their emergence, leadership can play an important role. In human and other animal groups, inter-individual differences can lead to the emergence of successful leaders, who assume their role thanks to their physical or cognitive capabilities that grant them some influence over the behavior of their peers. Hence, heterogeneity in a population appears as a key element for successful leaders. Here, we present an evolutionary game theoretic model to study the effect of leadership and heterogeneity on cooperative behavior and examine the relationships between the two. We show that the presence of a leader can promote the evolution of cooperation. Moreover, we find that, when there is the possibility for a leader to emerge in the group, heterogeneity benefits cooperation. In our model, players cooperate when they are more likely to become leaders, and defect otherwise. In other words, strong leaders do not defect, but act as exemplar of prosocial behavior that, when followed, lead to full cooperation.

Introduction

A wide range of organisms engage in collective actions to survive, from bacteria to eusocial animals and humans (Miller and Bassler, 2001; Cornforth et al., 2012; Nadell et al., 2016; Queller, 2000; Smith, 2010). In any collective action problem, when individuals perform costly actions that are beneficial to the entire group, a dilemma is established in which each individual has to choose between contributing to the collective benefit or exploiting the common resource. Reciprocity (Axelrod and Hamilton, 1981), signalling (Martinez-Vaquero et al., 2020) and leadership (Johnstone and Manica, 2011; King et al., 2009) are some of the mechanisms that have been proposed as potential solutions to this social dilemma.

Leadership, in particular, is a recurrent aspect in social interactions, and naturally occurs in very different settings. Indeed, the collective behavior of animal groups arises from the combination of individual-level processes that shape and

are shaped by their social structure, movement dynamics, and collective performance (Jolles et al., 2020). Empirical evidence suggests that individuals in these groups are not all equal and that heterogeneity plays a fundamental role in shaping the behavior of the group. Individual differences drive the collective behavior of bird flocks (Aplin et al., 2013; Pettit et al., 2015), fish schools (Jolles et al., 2017), bee colonies (Paleolog et al., 2009), and cattle herds (Šárová et al., 2010). The heterogeneous distribution of phenotypic variations in a group can lead to the emergence of a leadership trait (King and Sueur, 2011; Brent et al., 2015) that may be modulated by physiological characteristics, cognitive capabilities, or social relations. Leaders may display highly directed movements assuming frontal positions in a moving group (Jolles et al., 2017; Pettit et al., 2015), may present greater likelihood of initiating motion (Montiglio et al., 2013), or may influence the behavior of group-mates owing to their greater experience (Sueur and Petit, 2008; King and Sueur, 2011). In all of these cases, leadership entails a trade-off between benefits (e.g., privileged access to food sources, see McClure et al., 2011) and costs (e.g., due to increased exposure to predators, see Ioannou et al., 2019). Humans too evolved as group-living animals that rely on cooperation to survive. The need to coordinate and solve conflicts of interest might have triggered the evolution of leadership (Van Vugt et al., 2008). Leadership is a universal feature of human societies, where leaders initiate, motivate, plan, organise, direct, monitor and punish to achieve group coordination. The social complexity of human groups produced the need for more powerful leaders to manage complex intra- and inter-group relations (King et al., 2009).

A variety of modeling techniques have been used to study individual differences and quantify their effect on the collective outcome (Conradt and Roper, 2009; Couzin et al., 2002; Conradt et al., 2009; Sueur et al., 2010; Ioannou et al., 2015; del Mar Delgado et al., 2018). It has been demonstrated, for instance, that a fraction of informed individuals can modulate the behavior of a group by balancing goal-oriented and

socially-oriented behavior (Couzin et al., 2002, 2005; Ioannou et al., 2015). Despite the results obtained, much work is still needed to understand the effect of leadership and heterogeneity on collective behavior. Moreover, most of these models use many adjustable parameters that make them difficult to study and do not take in consideration the evolution of the individuals' behavior. Evolutionary game theory, on the other hand, provides a simple framework capable of capturing the evolutionary dynamics of a group of interacting individuals, at the cost of abstracting from specific collective behaviors. To account for population heterogeneity and individual personality, evolutionary game theory models have employed diversity in the payoffs (Guo et al., 2024; Qin et al., 2017; Amaral et al., 2016; Amaral and Javarone, 2020), in the ability to spread their strategy (Zhu et al., 2014; Droz et al., 2009; Szolnoki and Szabó, 2007), and in the network structure (Santos et al., 2006, 2012). Few studies explicitly consider leadership in evolutionary game models. Some address the emergence of leadership in repeated coordination games (Johnstone and Manica, 2011) and in different types of networks (Zhang et al., 2014) and provide insights on how leaders might emerge in a population. Others consider the effect of leadership on the cooperation rate of a population of players. Among these, leadership by enforcement or by example has received the most attention. In the former, leaders punish or reward other players if they defect or cooperate and receive a share of the common benefit in exchange (Hooper et al., 2010). In the latter, leaders are unconditional cooperators and can therefore influence both the game and the adoption of their strategy by other players (Zhuang et al., 2012). Similarly, Wang et al. (2017) model leaders as high reputation individuals so that other players are more likely to adopt their strategy.

Similar to the leadership-by-example studies mentioned above, this paper aims to study the effect of leadership and heterogeneity on the evolution of cooperation. Using Evolutionary Game Theory methods, we analyse the effect that having a leader produces on the cooperative behavior of a group of individuals that differ in their tendency to lead and to follow the leader. A model is presented in which leaders, either defectors or cooperators, shape the action choice of the other players in their group. Motivated by the presence of the leader that, by virtue of its characteristics (physical, cognitive, social, etc.), exerts some type of influence on them, players can imitate the available example without any additional mechanism of punishment or reward. Moreover, individual differences are introduced to reproduce the inherent heterogeneity in the leading and following capabilities among natural individuals. This is modelled as a sort of strength, which represents a characteristic of each individual dictated by physiological or psychological conditions, possibly driven by external factors. In the proposed model, strong players are less receptive to the decisions of their peers and have a higher predisposition to lead, whereas weak ones are

less prone to take the lead while being more likely to be influenced by the decision of a leader. By varying the individual behavior, the degree of heterogeneity, and the composition of the population, we quantify the cooperation level and analyse which strategies and conditions promote it.

This work tries to fill a gap in the literature by providing a theoretical model that explains the role of leaders and heterogeneity for cooperation in natural groups. It also attempts to provide a new perspective on the problem of designing artificial systems by looking for the optimal heterogeneity conditions for leadership to promote effective coordination and collaboration. The paper is organised as follows. First, we describe the game played by the individuals and the evolutionary dynamics by which their strategies evolve. Then, we present the outcomes obtained by testing the model on a wide range of parameters. Finally, we analyse the results and provide an interpretation of our findings.

Model

Game definition

We consider the case of a population of Z players interacting in a repeated N -person public goods game (Pacheco et al., 2009). Individuals can either cooperate and invest an amount c into a common pool or defect and not incur any cost. At the end of each round, the total investment is multiplied by a factor r and equally divided among all the players, therefore all N individuals in the group receive a benefit proportional to the fraction of cooperators. The payoff obtained each round by cooperators and defectors, respectively, is the following:

$$\Pi_C = rc \frac{N_C}{N} - c \quad (1)$$

$$\Pi_D = rc \frac{N_C}{N} \quad (2)$$

where N_C is the number of cooperators in the group. If everyone cooperates, each player will obtain a payoff of $(r - 1)c$ but if everyone defects no one receives any benefit. In each round, one of the players is selected as the leader and plays according to its own strategy either cooperating or defecting. The remaining players can play either according to their own strategy or following the leader by copying its action.

We consider an heterogeneous population composed by two types of players: *strong* and *weak*. Each player will behave as strong with a probability p_s and weak with a probability $1 - p_s$. This parameter represents external factors that influence the tendency of individuals to act in one or the other way. Strong and weak players have different strength values, defined by a Fermi function as:

$$S = \begin{cases} \frac{1}{1+e^{-\Delta t}} & \text{if strong} \\ \frac{1}{1+e^{\Delta t}} & \text{if weak} \end{cases} \quad (3)$$

where, the bigger the value of the parameter Δ_l , the bigger the difference in strength between strong and weak players. At each round, each player i within a group of size N has a probability of being selected as the leader, which is proportional to its strength with respect to the rest of the group:

$$p_{l,i} = \frac{S_i}{\sum_{j=0}^N S_j} \quad (4)$$

where S_i is the strength of player i . Therefore, when Δ_l is large, strong players are more likely to be selected than weak ones.

The probability of the other players to follow the leader is determined by their type, strong or weak, and the type of the leader, which we assume can be recognised by the other players. Since we have only two types of players, we distinguish four cases, one for each combination of leader and potential follower. The probabilities of following are expressed as Fermi functions:

$$\begin{aligned} p_{w,w} &= \frac{1}{1 + e^{-\beta_f f}} \\ p_{w,s} &= \frac{1}{1 + e^{-\beta_f (f + \Delta_f)}} \\ p_{s,w} &= \frac{1}{1 + e^{-\beta_f (f - \Delta_f)}} \\ p_{s,s} &= \frac{1}{1 + e^{-\beta_f f}} \end{aligned} \quad (5)$$

where β_f determines the slope of the Fermi function, f represents the tendency of the entire population to follow the leader, and Δ_f is a parameter regulating the difference in the following probability of strong and weak individuals: if $\Delta_f = 0$, the probability of following the leader is independent of the type of the leader and the type of the follower. Increasing the value of Δ_f , weak players become more likely to follow a strong leader and strong players become less likely to follow a weak leader. Figure 1 illustrates the following probabilities determined by the combination of f and Δ_f .

In this model, individuals adopt strategies that are only conditional on their strengths. The strategy set is defined using two bits that represent the action to take in case the player is weak or strong: (a_w, a_s) . When $a_j = 1, j \in \{w, s\}$, the player cooperates, otherwise when $a_j = 0$ the player defects. Hence, there are four possible strategies:

ALLD (0, 0): always defect

WCSD (1, 0): cooperate when weak, defect when strong

WDSC (0, 1): defect when weak, cooperate when strong

ALLC (1, 1): always cooperate

Finally, we introduce execution errors, both for actions and for the possibility of following the leader, as a small probability ϵ that the player performs the opposite of what she intended to.




Figure 1: Example of probabilities to follow for $\beta_f = 1$, $f = 1$ and $\Delta_f = 2$. The value of f determines the central probability of following between leader and follower of the same type. Δ_f determines the probability of following between individuals of different type.

Evolutionary dynamics

Evolution is modeled using a stochastic birth-death process in a well-mixed finite population of Z individuals (Nowak et al., 2004). The members of the population randomly form groups of N players and repeatedly play the game previously described. As a consequence, each individual obtains an average payoff taking into account all possible groups she can be part of. These payoffs are used as proxies for fitness. During reproduction, the fittest individuals have a higher chance of spreading their strategy to the new population, modeled by a Fermi function regulated by an intensity of selection β (Traulsen et al., 2006). This means that the information passed on to the offspring is how to behave in the rounds in which they are weak and in those in which they are strong. The probability for each player to be strong or weak is a parameter of the population that does not change through the evolutionary process.

In order to simplify the evolutionary dynamics analysis, we adopt the small mutation approximation. Using this approximation, a single mutant invading a uniform population can either fixate or disappear. In this way, only monomorphic states are considered and we can describe the prevalence of each strategy using a reduced Markov chain (Fudenberg and Imhof, 2006; Vasconcelos et al., 2017). This technique allows us to compute the invasion diagram among all pairs of strategies, their stationary distributions, and an average cooperation level across the different parameters of the model. The summary of the main parameters of the model can be found in Table 1.

Table 1: Main parameters of the model

| Parameter | Description | Value |
|------------|-----------------------------------------------------------------------------|---------|
| Z | Size of the population | 100 |
| N | Size of the group | 9 |
| c | Cost of cooperation | 1 |
| r | Benefit multiplication factor | [1, 10] |
| p_s | Probability of players to be strong | [0, 1] |
| Δ_l | Difference in strength between strong and weak players | [0, 8] |
| f | Tendency to following the leader of the whole population | [-8, 1] |
| β_f | Following intensity | 1 |
| Δ_f | Difference in tendency to follow the leader between strong and weak players | [0, 8] |
| ϵ | Error probability of each player | 0.01 |
| β | Intensity of selection | 1 |

Results

Evolution of cooperation

We first study the effect of leadership on the evolution of cooperation in the special case in which the population heterogeneity has the same effect on both the probability of acting as a leader and the probability of following, hence setting $\Delta_l = \Delta_f$. Figure 2 shows the expected fraction of players that cooperate with and without a leader and at varying levels of heterogeneity. In the absence of a leader, agents face a standard public goods game in which defection prevails in all conditions for $r < 10$, as expected (Pacheco et al., 2009). The presence of a leader ($\Delta_l = \Delta_f \geq 0$) fosters the emergence of cooperation for lower values of r , that is, when players receive lower returns for cooperating. This is true for any value of p_s , even for the extreme cases in which players always act as weak ($p_s = 0$) or strong ($p_s = 1$). In these extreme conditions, however, the population is homogeneous whatever the value of $\Delta_f = \Delta_l$, and low returns ($r \leq 5$) lead to diffuse defection. When p_s is not extreme, the effects of behavior heterogeneity induced by Δ_l and Δ_f are visible, leading to higher cooperation rates with any value of $r > 1$. What is evident from Figure 2 is that leadership is most effective in promoting cooperation when players do not always act as strong and there is a considerable difference between strong and weak players. At very low benefits ($r \in [2, 4]$), full cooperation is achieved only when the strong players are a minority in the group. When the benefit is higher, groups with several strong players attain full cooperation too. Overall, thanks to the presence of (few) strong leaders, cooperation prevails in the population even when the expected returns are small.

We then quantify the effect of the tendency to follow the leader on cooperation, varying the parameter $f \in [-8, 1]$

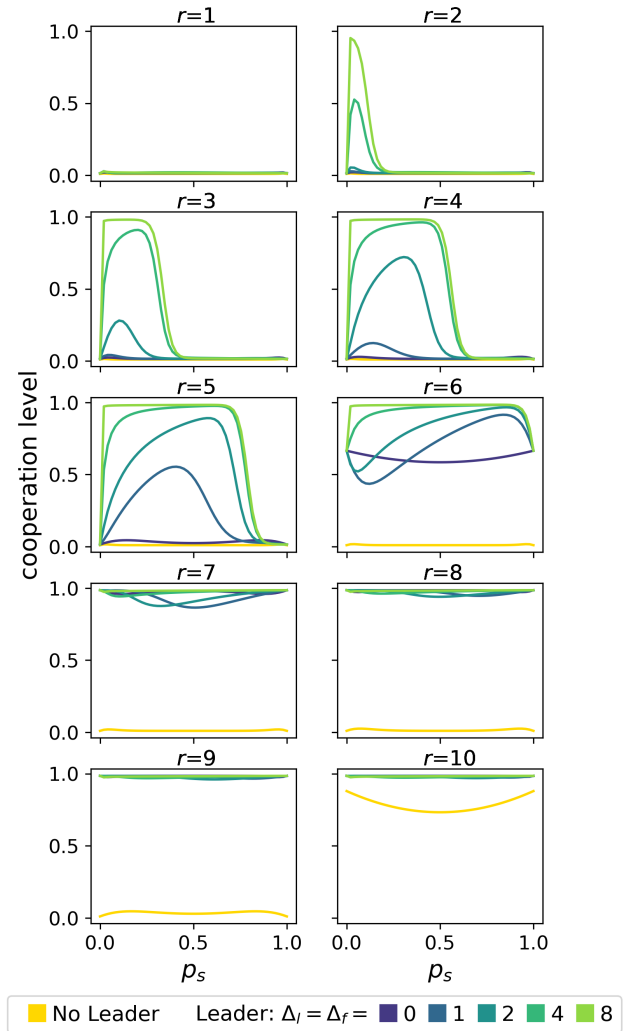


Figure 2: Cooperation level obtained with and without a leader and for different heterogeneity levels. We consider different values of r , Δ_f , Δ_l and p_s and assume $Z = 100$, $N = 9$, $\beta = 1$, $c = 1$, $f = 0$, $\beta_f = 1$ and $\epsilon = 0.01$

(see Figure 3). When f assumes low values, all players—both weak and strong—are less likely to follow the leader and defection is more likely to prevail. In this case, cooperation is achieved only when heterogeneity is high ($\Delta_l = \Delta_f \geq 4$). When f takes higher values, the tendency to follow generally increases driving a growing trend in cooperation rates, even for a less heterogeneous population.

Finally, we test the model in the two cases in which heterogeneity only affects the probability of acting as a leader or the probability of following ($\Delta_l \neq \Delta_f$). Figure 4 shows the expected cooperation level if the probability of following or the probability of acting as a leader are fixed. In the first case ($\Delta_f = 1$), when players receive low returns for cooperating ($r < 6$), cooperation is fostered if Δ_l assumes high




Figure 3: Cooperation level obtained with a leader at different levels of the following behavior. We assume $Z = 100$, $N = 9$, $\beta = 1$, $c = 1$, $r = 5$, $\beta_f = 1$ and $\epsilon = 0.01$

values, meaning that strong players are more likely to be leaders. However, for greater returns ($r \geq 7$), the opposite is true and homogeneous populations achieve more cooperation than heterogeneous ones. Instead, if the probability of becoming the leader is fixed ($\Delta_l = 1$), higher levels of heterogeneity are always more beneficial to cooperation. Interestingly, when cooperation returns are low ($r < 6$), groups with few strong players (low p_s) do not achieve high levels of cooperation anymore. This analysis reveals the complex interplay between the two parameters governing the population heterogeneity. Overall, it is possible to conclude that a mismatch between Δ_l and Δ_f leads to reduced cooperation levels: when the likelihood of behaving as leader is not accompanied by a similar likelihood of being followed, there is more room for exploitation by defectors.

Emergence of behavioral strategies

So far, we have focused on cooperation levels but we have not analysed which strategies are responsible for it. To this end, we compute the fixation probabilities for the four strategies across different conditions, varying the level of heterogeneity ($\Delta_l = \Delta_f$) and the probability of the presence of strong individuals p_s (see Figure 5). Additionally, Figure 6 expands the case of $r = 5$ displaying the full stationary distribution for the single strategies. In the same settings, Figure 7 presents the invasion graphs for the model calculated at particular conditions.




Figure 4: Cooperation level obtained with a leader with different heterogeneity values. We assume $Z = 100$, $N = 9$, $\beta = 1$, $c = 1$, $f = 0$, $\beta_f = 1$ and $\epsilon = 0.01$

The most evident aspect of the plots is the prevalence of the **WDCS** strategy across a wide range of parameters, and conversely the absence of **WCSD**. When the cooperation benefit is low ($r = 5$), pure defection dominates if the population is homogeneous (low Δ_l and Δ_f) or when most of the individuals are strong (high p_s). However, if some heterogeneity is introduced and the population is mostly weak, the strategy **WDCS** becomes dominant. This corresponds to the range of parameters in which we notice the rise in cooperation level in Figure 2. Here, the strategy **WDCS** allows for strong cooperators to prevent weak defectors from exploiting them if they are sufficiently strong. Moreover, the more weak players (hence low values of p_s), the less heterogeneity is necessary to prevent full defection. When **WDCS** prevails, an increasing population heterogeneity leads to strong players cooperating and weak players uniformly following the cooperating leader, producing high levels of cooperation. When this happens, **WDCS** becomes equivalent to **ALLC** and the two strategies are equally successful (see Figure 6). Note that for low probabilities of having strong players, more heterogeneity is needed to reach full cooperation. Indeed, if very few individuals are strong, they have to be much different from the weak ones in order to foster cooperation.

The switch from full defection to partial cooperation is also evident in the invasion graphs in Figure 7. When transitioning from an homogeneous population to an heterogeneous one, **WDCS** becomes gradually more dominant and pure defection less so. **WDCS** gets weakly exploited by

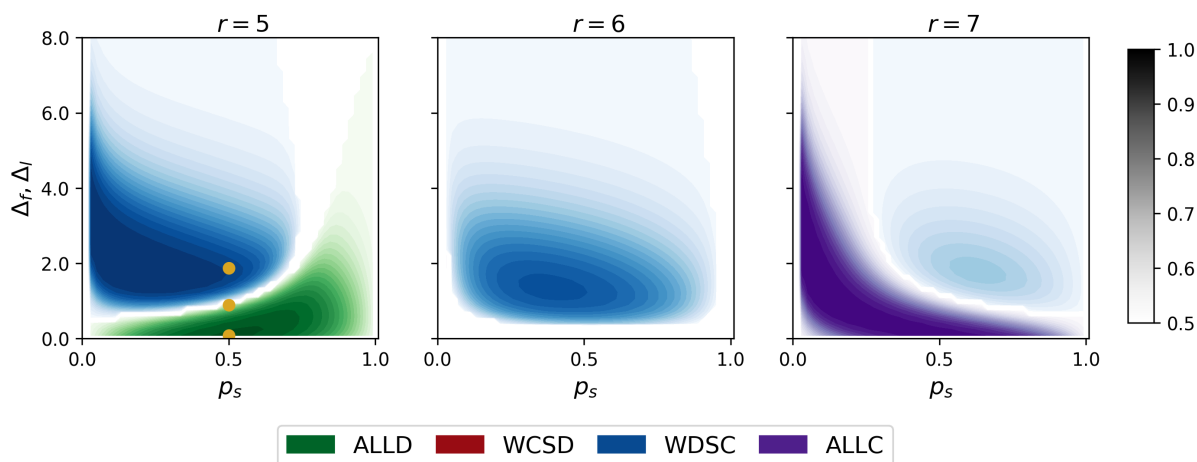


Figure 5: Stationary distribution of the strategies in the parameter space defined by p_s , Δ_f and Δ_l . Colors indicate the types of strategies that are prevalent in certain conditions and color darkness is the intensity of the probability in the stationary distribution. To make the visualization easier in overlapping areas, only probabilities higher than 0.5 are displayed. The yellow markers correspond to the coordinates of the invasion graphs of Figure 7. Parameters: $Z = 100$, $N = 9$, $\beta = 1$, $c = 1$, $f = 0$, $\beta_f = 1$ and $\epsilon = 0.01$

ALLD only in a homogeneous population. However, this exploitation is already avoided at low heterogeneity values ($\Delta_f = \Delta_l = 1$). Moreover, **ALLC** and **WCSD** are never dominant in these conditions.

When the cooperation benefit increases ($r = 6$ or $r = 7$ in Figure 5), the advantage of adopting **WDSC** is reduced and **ALLC** becomes progressively more dominant, especially when the population heterogeneity is low. The higher benefit promotes cooperation for all values of p_s , as also shown in Figure 2. This means that either strong leaders cooperate, benefiting also from following peers (hence the strategy **WDSC** prevails), or that every individual cooperates following the **ALLC** strategy (especially when heterogeneity is low). Overall, the relevance of the strategy **WDSC** in promoting cooperation is confirmed also when the conditions for cooperation are less demanding.

Finally, we consider the case in which $\Delta_l \neq \Delta_f$ (see Figure 8). When $\Delta_f = 1$, the **WDSC** strategy prevails in conditions where previously dominance was shared with other strategies. In particular, when Δ_l is large, **WDSC** completely replaces **ALLC**. When $\Delta_f = 1$, the probabilities of following are not extremely different among all possible pairs, meaning that weak individuals similarly follow strong and weak leaders. In such conditions, **ALLC** gets invaded by **WDSC** and **ALLD**, as it was happening for $\Delta_f = \Delta_l = 1$ (see Figure 7b), leading to **WDSC** prevailing. For $r = 7$, this implies that a lower cooperation rate is expected when the heterogeneity is high, as shown in Figure 4, due to the prevalence of **WDSC**. Conversely, when $\Delta_l = 1$, weak individuals act as leaders with non-negligible probability, especially when p_s is small. The bottom-left




Figure 6: Stationary distribution of the four strategies in the parameter space defined by p_s , Δ_f and Δ_l . Color darkness indicates the intensity of the probability in the stationary distribution. The yellow markers correspond to the coordinates of the invasion graphs of Figure 7. Parameters: $Z = 100$, $N = 9$, $\beta = 1$, $c = 1$, $\beta_f = 1$, $f = 0$, $r = 5$ and $\epsilon = 0.01$

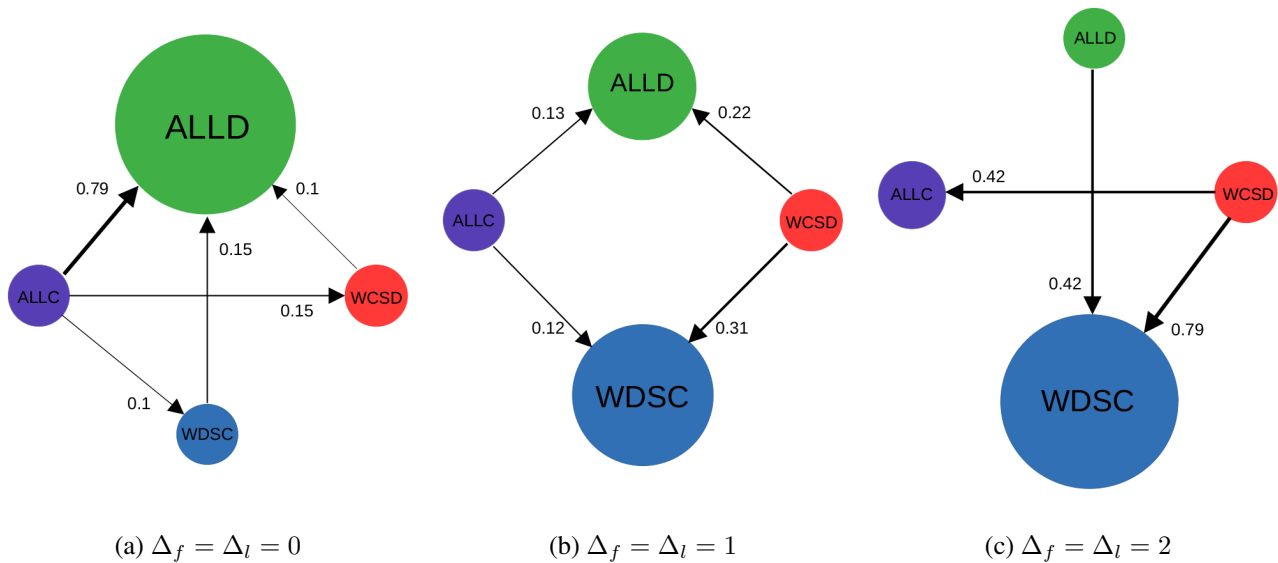


Figure 7: Graphs of invasion among the four strategies at different heterogeneous levels. The node size is proportional to the probability in the stationary distribution. The values on the arrows and their thickness represent the logarithm of the normalised transition probability among the two strategies. Only transition probabilities greater than neutral drift are displayed. Parameters used: $Z = 100$, $N = 9$, $\beta = 1$, $c = 1$, $f = 0$, $\beta_f = 1$, $\epsilon = 0.01$, $p_s = 0.5$ and $r = 5$

panel in Figure 8 shows that, when Δ_f is high and $r = 5$, weak individuals have no advantage in cooperating because they would not be followed, hence granting an advantage to **WDSC** over **ALLC**. In the case of $r = 7$, instead, the balance between **WDSC** and **ALLC** is shifted because of the higher expected returns from cooperating when leading.

Discussion and conclusions

This study investigates the role of leadership and heterogeneity in promoting cooperation in a collective action problem. Our results suggest that the presence of a leader is beneficial especially when the returns from cooperating are low. The presented model suggests that in groups where individuals can become leaders and obtain more influence over others, cooperation is favored over defection.

In a homogeneous population, however, when returns for cooperating are particularly low, the presence of a leader fails to promote full cooperation and defectors end up dominating. In this case, the results obtained here hint at a beneficial effect of heterogeneity. Indeed, when the population contains both strong and weak individuals, we find that a high level of cooperation is achieved when few strong leaders are sufficiently influential to promote cooperation in the whole group. Stronger individuals have particular characteristics that make them more likely to act as leaders and have more influence over their peers. When very strong players are only a minority in a population composed of much weaker ones, the former can establish their dominance by usually adopting the leading role and having the rest of the players follow them. Our analysis shows that such groups

can achieve full cooperation, even in adverse conditions.

The analysis of the probability of following (see Figure 3) reveals that, in a generally individualistic population ($f = -8$), only very strong heterogeneity can promote cooperation. If the members of the population, instead, have a higher probability of following the leader ($f = 1$), less heterogeneous populations achieve cooperation too. Our results also suggest that homogeneous populations cooperate more than heterogeneous ones only when returns for cooperation are high and strong players have more chances of being leaders but not of being followed. To promote cooperation in such conditions, strong individuals need to have the ability to become leaders but should also be able to inspire others to follow them.

Our analysis reveals that the strategy responsible for the increase in cooperation is **WDSC**. Individuals adopting this strategy cooperate only when they are strong and therefore have more chances of becoming leaders. In this situation, strong individuals avoid exploitation by defectors since, unless the population is homogeneous, they have a cohort of individuals that, by following them, cooperate too. When **WDSC** individuals are weak, however, we can distinguish between the two cases in which the leader is a defector or a cooperator. When the leader defects, **WDSC** players always defect, avoiding being exploited by the defecting leader and their acolytes. When the leader is a cooperator, instead, they either follow her and cooperate or defect, becoming exploiters. Adopting these behaviors allows **WDSC** individuals to resist the invasion of pure defectors and, at the same time, grants them an advantage over pure

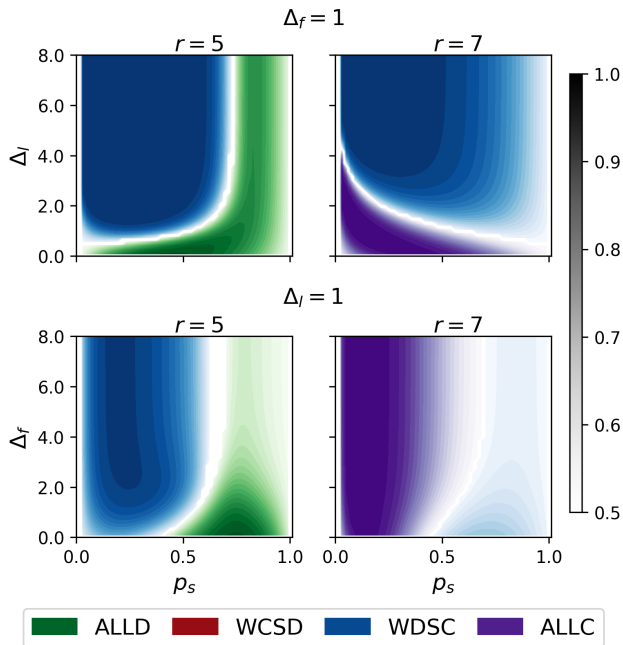


Figure 8: Stationary distribution of the four strategies when either $\Delta_f = 1$ (top panels) or $\Delta_l = 1$ (bottom panels). See also Figure 4 for the corresponding cooperation level. Parameters: $Z = 100$, $N = 9$, $\beta = 1$, $c = 1$, $\beta_f = 1$, $f = 0$, $r = 5$ and $\epsilon = 0.01$

cooperators. A similar effect has been reported in the context of cooperative agreements, where individuals willing to participate and cooperate in commitments but not collaborating to their establishment are a necessary evil to sustain cooperation (Martinez-Vaquero et al., 2017). Such strategy turns out to be prevailing also when there is a misalignment between the tendency to become a leader and the following behavior ($\Delta_l \neq \Delta_f$): in these conditions, weak individuals have more chances to exploit strong individuals, either because they do not follow more often strong than weak leaders (i.e., when Δ_f is small) or because they lead and defect (i.e., when Δ_l is small). Overall, we show that cooperation can be hampered when individuals that are often selected as leaders are not followed sufficiently often.

The model tries to capture the fundamental features of leadership and heterogeneity in an interacting group. Further work will be done to verify the assumptions made here and expand the study of the evolution of leadership. First of all, this model assumes the presence of only one leader, which is not always the case in natural groups. We will consider the possibility of having multiple leaders, addressing possible conflicts in the proposed group actions. On the other hand, a future in-depth study will add the possibility that leaders obtain benefits and/or incur costs derived from their leading role, conditions that are widely reported in ex-

perimental studies. Finally, the spatial distribution and the social configuration of the population will be also taken into account.

The present study highlights the key role of leadership and heterogeneity on the evolution of cooperation and the collective behavior. This is also relevant for the design of artificial systems in which multiple entities coexist and possibly cooperate. For instance, collective behavior is increasingly implemented in robot swarms, where members coordinate in a decentralized way to solve a shared task (Dorigo et al., 2021). While traditional studies often assume robots to be identical (e.g., to contain manufacturing costs), increasing evidence points to the importance of heterogeneity (Prorok et al., 2017) and demonstrates how individual differences can be leveraged to boost the collective performance (Sion et al., 2022; Feola et al., 2023; Raoufi et al., 2023). The present study provides useful insights on the level of heterogeneity that should be introduced in a population to promote cooperative interactions among robots, e.g., when the cooperative behavior should be balanced with individual needs such as the necessity to reduce energy expenditure (Notomista et al., 2022; Cai et al., 2023).

Acknowledgments

This research work has received funding from the Horizon Europe, PathFinder Open European Innovation Council Work Programme for the project BABOTS under grant agreement No 101098722.

References

- Amaral, M. A. and Javarone, M. A. (2020). Heterogeneity in evolutionary games: an analysis of the risk perception. *Proceedings of the Royal Society A*, 476(2237):20200116.
- Amaral, M. A., Wardil, L., Perc, M., and da Silva, J. K. (2016). Evolutionary mixed games in structured populations: Cooperation and the benefits of heterogeneity. *Physical Review E*, 93(4):042304.
- Aplin, L. M., Farine, D. R., Morand-Ferron, J., Cole, E. F., Cockburn, A., and Sheldon, B. C. (2013). Individual personalities predict social behaviour in wild networks of great tits (*Parus major*). *Ecology letters*, 16(11):1365–1372.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211(4489):1390–1396.
- Brent, L. J., Franks, D. W., Foster, E. A., Balcomb, K. C., Cant, M. A., and Croft, D. P. (2015). Ecological knowledge, leadership, and the evolution of menopause in killer whales. *Current biology*, 25(6):746–750.
- Cai, X., Schlotfeldt, B., Khosoussi, K., Atanasov, N., Pappas, G. J., and How, J. P. (2023). Energy-aware, collision-free information gathering for heterogeneous robot teams. *IEEE Transactions on Robotics*, 39(4):2585–2602.
- Conradt, L., Krause, J., Couzin, I. D., and Roper, T. J. (2009). “leading according to need” in self-organizing groups. *The American Naturalist*, 173(3):304–312.

- Conradt, L. and Roper, T. J. (2009). Conflicts of interest and the evolution of decision sharing. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1518):807–819.
- Cornforth, D. M., Sumpter, D. J., Brown, S. P., and Brännström, Å. (2012). Synergy and group size in microbial cooperation. *The American Naturalist*, 180(3):296–305.
- Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516.
- 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.
- del Mar Delgado, M., Miranda, M., Alvarez, S. J., Gurarie, E., Fagan, W. F., Penteriani, V., di Virgilio, A., and Morales, J. M. (2018). The importance of individual variation in the dynamics of animal collective movements. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1746):20170008.
- Dorigo, M., Theraulaz, G., and Trianni, V. (2021). Swarm robotics: Past, present, and future [point of view]. *Proceedings of the IEEE*, 109(7):1152–1165.
- Droz, M., Szwabiński, J., and Szabó, G. (2009). Motion of influential players can support cooperation in prisoner's dilemma. *The European Physical Journal B*, 71:579–585.
- Feola, L., Sion, A., Trianni, V., Reina, A., and Tuci, E. (2023). Aggregation through adaptive random walks in a minimalist robot swarm. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '23*, page 21–29, New York, NY, USA. Association for Computing Machinery.
- Fudenberg, D. and Imhof, L. A. (2006). Imitation processes with small mutations. *Journal of Economic Theory*, 131(1):251–262.
- Guo, H., Shen, C., Zou, R., Tao, P., Shi, Y., Wang, Z., and Xing, J. (2024). Complex pathways to cooperation emergent from asymmetry in heterogeneous populations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 34(2).
- Hooper, P. L., Kaplan, H. S., and Boone, J. L. (2010). A theory of leadership in human cooperative groups. *Journal of Theoretical Biology*, 265(4):633–646.
- Ioannou, C. C., Rocque, F., Herbert-Read, J. E., Duffield, C., and Firth, J. A. (2019). Predators attacking virtual prey reveal the costs and benefits of leadership. *Proceedings of the National Academy of Sciences*, 116(18):8925–8930.
- Ioannou, C. C., Singh, M., and Couzin, I. D. (2015). Potential leaders trade off goal-oriented and socially oriented behavior in mobile animal groups. *The American Naturalist*, 186(2):284–293.
- Johnstone, R. A. and Manica, A. (2011). Evolution of personality differences in leadership. *Proceedings of the National Academy of Sciences*, 108(20):8373–8378.
- Jolles, J. W., Boogert, N. J., Sridhar, V. H., Couzin, I. D., and Manica, A. (2017). Consistent individual differences drive collective behavior and group functioning of schooling fish. *Current Biology*, 27(18):2862–2868.
- Jolles, J. W., King, A. J., and Killen, S. S. (2020). The role of individual heterogeneity in collective animal behaviour. *Trends in ecology & evolution*, 35(3):278–291.
- King, A. J., Johnson, D. D., and Van Vugt, M. (2009). The origins and evolution of leadership. *Current biology*, 19(19):R911–R916.
- King, A. J. and Sueur, C. (2011). Where next? Group coordination and collective decision making by primates. *International journal of primatology*, 32:1245–1267.
- Martinez-Vaquero, L. A., Han, T. A., Pereira, L. M., and Lenaerts, T. (2017). When agreement-accepting free-riders are a necessary evil for the evolution of cooperation. *Scientific reports*, 7(1):2478.
- Martinez-Vaquero, L. A., Santos, F. C., and Trianni, V. (2020). Signalling boosts the evolution of cooperation in repeated group interactions. *Journal of the Royal Society Interface*, 17(172):20200635.
- McClure, M., Ralph, M., and Despland, E. (2011). Group leadership depends on energetic state in a nomadic collective foraging caterpillar. *Behavioral Ecology and Sociobiology*, 65:1573–1579.
- Miller, M. B. and Bassler, B. L. (2001). Quorum sensing in bacteria. *Annual Reviews in Microbiology*, 55(1):165–199.
- Montiglio, P.-O., Ferrari, C., and Reale, D. (2013). Social niche specialization under constraints: personality, social interactions and environmental heterogeneity. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1618):20120343.
- Nadell, C. D., Drescher, K., and Foster, K. R. (2016). Spatial structure, cooperation and competition in biofilms. *Nature Reviews Microbiology*, 14(9):589–600.
- Notomista, G., Mayya, S., Emam, Y., Kroninger, C., Bohannon, A., Hutchinson, S., and Egerstedt, M. (2022). A resilient and energy-aware task allocation framework for heterogeneous multirobot systems. *IEEE Transactions on Robotics*, 38(1):159–179.
- Nowak, M. A., Sasaki, A., Taylor, C., and Fudenberg, D. (2004). Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428(6983):646–650.
- Pacheco, J. M., Santos, F. C., Souza, M. O., and Skyrms, B. (2009). Evolutionary dynamics of collective action in n-person stag hunt dilemmas. *Proceedings of the Royal Society B: Biological Sciences*, 276(1655):315–321.
- Paleolog, J. et al. (2009). Behavioural characteristics of honey bee (*Apis mellifera*) colonies containing mix of workers of divergent behavioural traits. *Animal Science Papers and Reports*, 27(3):237–248.
- Pettit, B., Akos, Z., Vicsek, T., and Biro, D. (2015). Speed determines leadership and leadership determines learning during pigeon flocking. *Current Biology*, 25(23):3132–3137.
- Prorok, A., Hsieh, M. A., and Kumar, V. (2017). The impact of diversity on optimal control policies for heterogeneous robot swarms. *IEEE Transactions on Robotics*, 33(2):346–358.

- Qin, J., Chen, Y., Kang, Y., and Perc, M. (2017). Social diversity promotes cooperation in spatial multigames. *Europhysics Letters*, 118(1):18002.
- Queller, D. C. (2000). Relatedness and the fraternal major transitions. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 355(1403):1647–1655.
- Raoufi, M., Romanczuk, P., and Hamann, H. (2023). Individuality in swarm robots with the case study of kilobots: Noise, bug, or feature? In *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*, page 35. MIT Press.
- Santos, F. C., Pacheco, J. M., and Lenaerts, T. (2006). Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proceedings of the National Academy of Sciences*, 103(9):3490–3494.
- Santos, F. C., Pinheiro, F. L., Lenaerts, T., and Pacheco, J. M. (2012). The role of diversity in the evolution of cooperation. *Journal of theoretical biology*, 299:88–96.
- Šárová, R., Špinka, M., Panamá, J. L. A., and Šimeček, P. (2010). Graded leadership by dominant animals in a herd of female beef cattle on pasture. *Animal Behaviour*, 79(5):1037–1045.
- Sion, A., Reina, A., Birattari, M., and Tuci, E. (2022). Controlling robot swarm aggregation through a minority of informed robots. In *International Conference on Swarm Intelligence*, pages 91–103. Springer.
- Smith, E. A. (2010). Communication and collective action: language and the evolution of human cooperation. *Evolution and human behavior*, 31(4):231–245.
- Sueur, C., Deneubourg, J.-L., Petit, O., and Couzin, I. D. (2010). Differences in nutrient requirements imply a non-linear emergence of leaders in animal groups. *PLoS Computational Biology*, 6(9):e1000917.
- Sueur, C. and Petit, O. (2008). Organization of group members at departure is driven by social structure in macaca. *International Journal of Primatology*, 29(4):1085–1098.
- Szolnoki, A. and Szabó, G. (2007). Cooperation enhanced by inhomogeneous activity of teaching for evolutionary prisoner’s dilemma games. *Europhysics Letters*, 77(3):30004.
- Traulsen, A., Nowak, M. A., and Pacheco, J. M. (2006). Stochastic dynamics of invasion and fixation. *Physical Review E*, 74(1):011909.
- Van Vugt, M., Hogan, R., and Kaiser, R. B. (2008). Leadership, followership, and evolution: some lessons from the past. *American psychologist*, 63(3):182.
- Vasconcelos, V. V., Santos, F. P., Santos, F. C., and Pacheco, J. M. (2017). Stochastic dynamics through hierarchically embedded markov chains. *Physical Review Letters*, 118(5):058301.
- Wang, Z., Chen, T., and Wang, Y. (2017). Leadership by example promotes the emergence of cooperation in public goods game. *Chaos, Solitons & Fractals*, 101:100–105.
- Zhang, J., Zhang, C., Cao, M., and Chu, T. (2014). Cooperation with potential leaders in evolutionary game study of networking agents. In *2014 IEEE Congress on Evolutionary Computation (CEC)*, pages 918–923. IEEE.
- Zhu, C.-j., Sun, S.-w., Wang, L., Ding, S., Wang, J., and Xia, C.-y. (2014). Promotion of cooperation due to diversity of players in the spatial public goods game with increasing neighborhood size. *Physica A: Statistical Mechanics and its Applications*, 406:145–154.
- Zhuang, Q., Wang, D., Fan, Y., and Di, Z. (2012). Evolution of cooperation in a heterogeneous population with influential individuals. *Physica A: Statistical Mechanics and its Applications*, 391(4):1735–1741.