

# MOTION PLANNING FOR SOCIALLY COMPETENT ROBOT NAVIGATION

A Dissertation

Presented to the Faculty of the Graduate School  
of Cornell University

in Partial Fulfillment of the Requirements for the Degree of  
Doctor of Philosophy

by

Christoforos Christos Mavrogiannis

May 2019

© 2019 Christoforos Christos Mavrogiannis

# MOTION PLANNING FOR SOCIALLY COMPETENT ROBOT NAVIGATION

Christoforos Christos Mavrogiannis, Ph.D.

Cornell University 2019

Crowded human environments such as pedestrian scenes constitute challenging domains for mobile robot navigation, for a variety of reasons including the heterogeneity of pedestrians' decision-making mechanisms, the lack of channels of explicit communication among them and the lack of universal rules or social conventions regulating traffic. Despite these complications, humans exhibit socially competent navigation through coordination, realized with implicit communication via a variety of modalities such as path shape and body posture. Sophisticated mechanisms of inference and decision-making allow them to understand subtle communication signals and encode them into their own actions. Although the problem of planning socially competent robot navigation has received significant attention over the past three decades, state-of-the-art approaches tend to explicitly focus on reproducing selected social norms or directly imitating observed human behaviors, while often lack of extensive and thorough validation procedures, thus raising questions about their generalization and reproducibility.

This thesis introduces a family of planning algorithms, inspired by studies on human navigation. Our algorithms are designed to produce socially competent robot navigation behaviors by leveraging the existing mechanisms of implicit coordination in humans. We model multi-agent motion coordination through a series of data structures, based on mathematical abstractions from low-dimensional topology and physics, that capture fundamental properties of

multi-agent collision avoidance. These models enable a robot to anticipate the effects of its actions on the inference and decision-making processes of nearby agents and allow for the generation of motion that is compliant with the unfolding evolution of the scene and consistent with the robot's intentions.

The introduced planning algorithms are supported by extensive simulated and experimental validation. Key findings include: (1) evidence extracted from a series of simulated studies, suggesting that the outlined planning architecture indeed results in effective coordination within groups of non-communicating agents in a variety of simulated scenarios; (2) evidence extracted from an online, video-based user study with more than 180 participants, indicating that humans perceive the motion generated by our framework as intent-expressive; (3) evidence extracted from an experimental study, conducted in a controlled lab environment with 105 human participants, suggesting that humans follow low-acceleration paths when navigating next to a robot running our framework.

## **BIOGRAPHICAL SKETCH**

Christoforos Christos Mavrogiannis was born in Athens, Greece in 1989. He grew up in Glyfada, a suburb of Athens, and graduated from the 6th public High School of the area in 2007. He then enrolled in the National Technical University of Athens, from which he graduated with a Diploma in Mechanical Engineering in February 2013. In August of the same year, he entered the Ph.D. program of the Sibley School of Mechanical and Aerospace Engineering at Cornell University, where he conducted research on Robotics, Artificial Intelligence and Human-Robot Interaction.

To my parents, Aspasia and Ioannis.

## ACKNOWLEDGEMENTS

I would like to express my deep gratitude to my advisor, Ross Knepper. I could have never asked for a more inspirational and caring figure to advise me during my PhD studies. Our weekly meetings always touched upon the big questions and triggered sequences of deep dives into a variety of different fields that helped shape all pieces of this work. The interaction with Ross helped me develop a computer scientific way of thinking that complements my mechanical engineering foundation. His enthusiasm, flexibility and open mind encouraged me to take independent and creative directions that helped me develop my research identity. Overall, I learned a lot and matured scientifically and personally but most importantly I had fun during the past 5 years of collaboration. Thank you, Ross for being a great advisor, mentor and friend! I hope we can find ways to continue collaborating.

I have been honored and privileged to have a multidisciplinary team of world experts to serve as members of my Special Committee, all of whom I look up to and admire their work. I would like to thank all of them for their valuable feedback. The insightful comments of Kilian Weinberger and Shimon Edelman when I started working on my thesis proposal were instrumental to the development of my research directions and to the identification of my research goals. The work of Anca Dragan constituted a foundational inspiration for this thesis, and her on-point comments on my conference and exam presentations always shifted me in the right directions and inspired me to do my best. The advice of Guy Hoffman on the design, evaluation and presentation of my user studies has been crucial and I am sure it will prove invaluable over the years, as I continue to work in the field of human-robot interaction.

I have also been fortunate to be part of the Robotic Personal Assistants Lab, a

community of super smart and supportive lab mates, all of whom I consider as true friends. I have especially enjoyed my close collaboration with fellow PhD students Wil Thomason, Valts Blukis, Claire Liang and Julia Proft and my mentorship of undergraduate members Alena Hutchinson, and John Macdonald. I also had fun and learned a lot during my teaching assistantships alongside Julia, Wil and Claire. Besides, being part of the lab since its conception, I have enjoyed sharing Thai food and participating in engaging discussions on our Monday lab meetings with many bright students including Ian Lenz, Minae Kwon, Nathan Adara, Rebecca Bernstein, Niko Grupen, Elizabeth Ricci and Soham Sankaran. In the beginning of my PhD, I had fun sharing the lab space in Gates 323 with Ian and his insightful feedback on my machine learning projects proved invaluable.

Cornell is home to some excellent instructors, and I was privileged to have interacted with a few of them. I learned a lot about dynamics from Andy Ruina, acquired a broad background on all areas of robotics through the class of Hadas Kress-Gazit, solidified my robotic manipulation and motion planning foundations through the class of Ross Knepper, enjoyed the deep machine learning insights of Kilian Weinberger and Thorsten Joachims, and appreciated the AI anecdotes of Bart Selman.

I am thankful to Marcia James Sawyer, the Mechanical Engineering Graduate Field Administrator at Cornell. Marcia's experience is invaluable and her advice on all matters, from enrollment issues to tax and immigration issues has helped me save hundreds of hours over the years. Marcia: BIG thanks – this department is lucky to have you!

I would like to thank my family for the support and understanding that it has continuously provided me. I have been privileged to come from a home

with a great amount of emotional stability and intellectual curiosity that always gave me peace of mind and inspired me to pursue my interests. This thesis is dedicated to my parents: Aspasia Gkioka and Ioannis Mavrogiannis, who have always encouraged me to do what I love, even at the cost of not getting to see each other often. Furthermore, I am thankful to my talented brother, Angelos Mavrogiannis, for being a fun and inspiring figure, consistently my closest friend and music co-adventurer ever since I can remember. Finally, I have been privileged to have an abnormally high number of relatives in both coasts of the United States, all of whom have supported me in every possible way. I am especially thankful to my cousin, Eleni Magoulas, who has allowed me to crash her couch in New York countless times and has introduced me to some of the best food in the City. I am also thankful to my aunt Maria Magoulas, uncle Tasos Magoulas and cousin Gary Andriotis, for always looking after me over the years. I have been extremely lucky to have several people on the other coast too. Special thanks to my aunt Grigoria Geokezas and cousins Sylvia and Nota Lucas, uncle Meletis Geokezas and cousins Maria and Dimitris Geokezas, and David Howard for their great hospitality every time I have visited Seattle.

Over the past 5 years in Ithaca, I have enjoyed a personal balance that has provided me with the foundations for happiness and productivity. First, I have been extremely lucky to meet my girlfriend Patrícia who has been an amazing source of support, my best friend and travel partner. I have also had a lot of fun with the people from my class and especially Bryan Peele, Nozomi Hitomi, Ravi Patel and Robert Carson, learning more about the American culture and history and discussing science, technology, startups and politics over IPAs at the Chapter House, the Westy and the Big Red Barn. I've had great times with Pierre-Luc Pradier and Marcos Guimarães, talking about college hockey, French cheese and

Brazilian (Portuguese) language. I have also been fortunate to be able to speak (more than) a bit of Greek and hang out with my roommates George Valogianis and Thodoris Diamantopoulos, but also with Pantelis Pipergias-Analysis, Konstantinos Spyrou, Vasilis Fountoulakis and Thodoris Lykouris. Finally, I've had good times in Ithaca and New York City with my almost-roommate Jason Koutoudis, checking out art exhibitions and ethnic food.

I have also been privileged to have a very strong circle of friends from Greece, who have kept in touch over all these years despite the distance. The countless hours of chatting with Kostas Braimakis, Christos Athanasiou, Antonis Chatzistefanou and Dimitris Myrisiotis have kept me sane, even at the most difficult and stressful times, whereas the occasional hangouts in Athens over the summer and winter breaks with Kostas Bourchas, Gerasimos Makris, Giannis Konstantinidis, Aris Kalgreadis, Manos Kaloeidas, and Valentino Redaelli always helped me recharge my batteries and go back to work feeling inspired and motivated.

I would like to especially thank a few people who have inspired me to pursue my research path. I would like to express my gratitude to my Diploma thesis advisor, Prof. Kostas J. Kyriakopoulos for conveying to me his passion about robotics, inspiring me to pursue a PhD in the US and for being an invaluable mentor since 2011. In his lab, I was also fortunate to meet my friend and research collaborator Prof. Minas Liarokapis who convinced me that a PhD in a top University in the US is within my reach and worth it despite the distance (I still agree!). Last but not least, I would like to thank Haris Aronis, my high-school Physics teacher for showing me the beauty of applied math and scientific thought in a way that I still recall whenever I write down a set of equations or draw a diagram.

On a technical note, I would like to thank some of the scientists whose contributions inspired this work. This thesis would not have been possible without the foundational work of Artin [10] and Birman [20] on topological braid groups, the inspiring notes of Ghrist [49] on the use of braids for planning problems in robotics, the work of Thiffeault [131] on the use of braids as abstractions of multi-particle systems in fluids and their Braidlab package [132], the definition of topological complexity by Dynnikov and Wiest [42], and the work of Berger [18] on the generation of multi-particle trajectories from topological invariants. Finally, this journey perhaps would not have happened if it had not been for the blog post of Richeson [117], which reassured my vision that the ribbon pattern generated by the dancers of the greek dance called γαϊτανάκι (similar to the maypole dance found in some other European countries) carries some properties that could be relevant to the problem of modeling multi-agent collision avoidance.

On an official note, I would like to gratefully acknowledge the support provided by the National Science Foundation, through the grants IIS-1526035 and IIS-1563705.

On a lighter note, I would like to thank some of the bands and musicians that kept me company and contributed to my inspiration and relaxation over all these years. This was a fun ride thanks in part to some notes, beats, rimes and riffs by OM, Sleep, Lamb of God, Mastodon, Norma Jean, the Dillinger Escape Plan, Alice in Chains, Soundgarden, Mobb Deep, Skalpel, Slayer, Helmet, Tool, Miles Davis, Meshuggah, Bowery Electric, Nine Inch Nails, Radiohead, Lee Burton, Deftones, Black Sabbath, Whitechapel, and more.

## TABLE OF CONTENTS

Biographical Sketch . . . . .	iii
Dedication . . . . .	iv
Acknowledgements . . . . .	v
Table of Contents . . . . .	x
List of Tables . . . . .	xv
List of Figures . . . . .	xvi
<b>1 Introduction</b> . . . . .	<b>1</b>
1.1 Preface . . . . .	1
1.2 Related Work . . . . .	4
1.2.1 Human Navigation . . . . .	5
1.2.2 Simulating Artificial Pedestrian Flows . . . . .	7
1.2.3 Tracking Pedestrian Flows . . . . .	8
1.2.4 Social Robot Navigation . . . . .	8
1.2.4.1 Simulation Studies . . . . .	9
1.2.4.2 Experimental Demonstrations . . . . .	10
1.2.4.3 Experimental Studies . . . . .	11
1.2.5 Topological Representations for Multi-Agent Navigation . . . . .	13
1.2.6 Implicit Communication in Human-Robot Interaction . . . . .	16
1.2.6.1 Implicit Communication in Multi-Agent Navigation . . . . .	17
1.2.7 Exploiting Structure . . . . .	18
1.3 Contributions . . . . .	19
1.3.1 Key Results . . . . .	20
1.3.2 Technical Developments . . . . .	21
1.3.3 A Broader Perspective . . . . .	23
1.4 Published Content . . . . .	26
1.5 Overview . . . . .	27
<b>I Implicit Communication for Human-Robot Interaction</b> . . . . .	<b>30</b>
<b>2 A Mathematical Framework for Implicit Communication in Human-Robot Interaction Applications</b> . . . . .	<b>31</b>
2.1 Implicit Communication in Human Activities . . . . .	34
2.2 Framework . . . . .	36
2.2.1 Definitions . . . . .	36
2.2.2 Foundations . . . . .	37
2.2.3 Implicit Communication Criteria . . . . .	39
2.2.4 Understanding and Generation . . . . .	41
2.3 Achieving Joint Goals . . . . .	42
2.3.1 Trust . . . . .	44

2.3.2	Consensus . . . . .	45
2.3.3	Receptivity . . . . .	46
2.4	Case Studies . . . . .	47
2.4.1	Implicit Communication through Natural Language . . . . .	48
2.4.1.1	Implicature . . . . .	48
2.4.1.2	Inverse Semantics . . . . .	51
2.4.2	Communicative Motion . . . . .	52
2.4.2.1	Legibility . . . . .	53
2.4.2.2	Dynamic Legibility . . . . .	54
2.5	Other Examples . . . . .	56
2.5.1	Tact . . . . .	57
2.6	Practical Implementation . . . . .	58
2.7	Discussion . . . . .	60
2.7.1	A Call to Action . . . . .	60

## **II Planning by Reasoning about Multi-Agent Navigation Strategies 62**

<b>3</b>	<b>Multi-Agent Path Topology in Support of Socially Competent Navigation Planning 63</b>
3.1	Foundations . . . . . 64
3.1.1	Game-Theoretic Setup . . . . . 66
3.1.2	A Topological Model of Joint Strategies . . . . . 67
3.1.2.1	Background on Braids . . . . . 68
3.1.2.2	Abstracting Joint Strategies Using Braids . . . . . 71
3.2	Inference of Collective Behaviors . . . . . 73
3.2.1	Teleological Reasoning in Multi-Agent Navigation . . . . . 74
3.2.2	Inferring Joint Strategies from Context . . . . . 74
3.2.2.1	Inferring the Final Permutation of the System . . . . . 76
3.2.2.2	Inferring the Path Entanglement . . . . . 79
3.3	Decision Making . . . . . 83
3.3.1	Modeling Agents' Cost Functions . . . . . 84
3.3.2	Planning Joint Strategies . . . . . 85
3.3.3	Permutation Graph Search . . . . . 87
3.3.4	Online Algorithm . . . . . 88
3.3.4.1	A Discussion of Complexity . . . . . 90
3.4	Evaluation . . . . . 93
3.4.1	Setup . . . . . 93
3.4.2	Simulation Results . . . . . 95
3.4.2.1	Qualitative Behavior . . . . . 95
3.4.2.2	Performance . . . . . 97
3.5	Discussion . . . . . 100

<b>4</b>	<b>Socially Competent Navigation Planning Using a Learned Model for Predicting the Multi-Agent Path Topology</b>	<b>102</b>
4.1	Socially Competent Navigation . . . . .	103
4.2	Inferring System Path Topologies from Context . . . . .	104
4.3	Learning Collective Navigation Behaviors . . . . .	105
4.3.1	Generating a Dataset of Diverse Collective Behaviors . . . . .	106
4.3.1.1	The Social Force Model . . . . .	106
4.3.1.2	Experimental Setup . . . . .	107
4.3.2	Learning Setup . . . . .	107
4.3.3	Learning Architecture . . . . .	108
4.3.4	Implementation Details and Performance . . . . .	109
4.4	Socially Competent Motion Generation . . . . .	111
4.4.1	Decision Making Policy . . . . .	111
4.4.2	Generating a Set of System Path Topologies . . . . .	113
4.4.3	Generating a Set of Actions . . . . .	114
4.4.4	Online Algorithm . . . . .	114
4.5	Evaluation . . . . .	115
4.5.1	Experimental Setup . . . . .	115
4.5.2	Results . . . . .	117
4.6	Discussion . . . . .	120

### **III Planning by Reasoning about Pairwise Collision Avoidance Intentions** **122**

<b>5</b>	<b>Social Momentum: A Framework for Legible Navigation in Dynamic Multi-Agent Environments</b>	<b>123</b>
5.1	Multi-Agent Trajectory Analysis . . . . .	125
5.1.1	Complexity of Braids . . . . .	127
5.1.1.1	Curve Diagrams . . . . .	128
5.1.1.2	The Complexity Index . . . . .	129
5.1.2	Complexity of Multi-Agent Planning . . . . .	129
5.2	The Social Momentum Planning Framework . . . . .	131
5.2.1	Angular Momentum for Collision Avoidance . . . . .	132
5.2.2	Social Momentum for Legible Collision Avoidance . . . . .	133
5.2.3	Decision Making . . . . .	136
5.3	Evaluation . . . . .	137
5.3.1	Simulations . . . . .	137
5.3.1.1	Experimental Setup . . . . .	138
5.3.1.2	Quality Measures . . . . .	139
5.3.1.3	Performance Comparison . . . . .	139
5.3.2	User Study . . . . .	142
5.4	Discussion . . . . .	146

<b>6</b>	<b>Effects of Distinct Robot Navigation Strategies on Human Behavior in a Crowded Environment</b>	<b>147</b>
6.1	Contributions . . . . .	150
6.2	User Study . . . . .	151
6.2.1	General Experiment Procedure . . . . .	152
6.2.2	Background Scenario and Task Description . . . . .	154
6.2.3	Trial Description . . . . .	155
6.2.4	Conditions . . . . .	155
6.2.5	Hypotheses . . . . .	158
6.3	Analysis . . . . .	159
6.3.1	Effect of Navigation Strategy on Robot Behavior . . . . .	162
6.3.2	Effect of Navigation Strategy on Human Behavior . . . . .	163
6.3.3	Effect of Navigation Strategy on Group Behavior . . . . .	164
6.3.4	Effect of Navigation Strategy on Human Impressions . . . . .	165
6.4	Discussion . . . . .	167
6.4.1	Limitations . . . . .	168
6.4.2	Broader Impact . . . . .	169

## **IV Planning with Topologically Robust Multi-Agent Trajectory Prediction 170**

<b>7</b>	<b>Multi-Agent Trajectory Prediction and Generation with Topological Invariants Enforced by Hamiltonian Dynamics</b>	<b>171</b>
7.1	Foundations . . . . .	173
7.1.1	Hamiltonian Motion for Multi-Particle Trajectory Braiding	174
7.1.2	Topological Invariants of Particle Trajectories . . . . .	175
7.1.2.1	The Winding Number . . . . .	176
7.1.3	Two-Particle Vortex Motion . . . . .	176
7.1.4	Two-Agent Collision Avoidance as Vortex Motion . . . . .	178
7.1.5	Multi-Agent Trajectory Generation from Topological Specifications . . . . .	179
7.1.6	TANP: Topologically Adaptive Navigation Planning . . . . .	181
7.1.6.1	Destination Prediction . . . . .	182
7.1.6.2	Outcome Evaluation . . . . .	184
7.1.6.3	Decision Making . . . . .	185
7.1.6.4	Pseudocode . . . . .	186
7.1.7	Complexity and Practical Considerations . . . . .	188
7.2	Results . . . . .	189
7.2.1	Offline Performance . . . . .	189
7.2.2	Comparison with Trajectory Optimization . . . . .	192
7.2.3	Online Performance . . . . .	194
7.3	Discussion . . . . .	198

<b>V</b>	<b>Discussion</b>	<b>199</b>
<b>8</b>	<b>Conclusion</b>	<b>200</b>
8.1	Summary . . . . .	200
8.2	Assumptions, Evaluation and Limitations . . . . .	202
8.2.1	Reasoning about Multi-Agent Navigation Strategies . . . . .	203
8.2.2	Reasoning about Pairwise Collision Avoidance Intentions . . . . .	205
8.2.3	Topologically Robust Multi-Agent Trajectory Prediction . . . . .	209
8.3	Future Work . . . . .	211
8.3.1	The Role of Environment Geometry . . . . .	211
8.3.2	Scalability . . . . .	212
8.3.3	Learning Models of Multi-Agent Interactions . . . . .	213
8.3.4	Multi-Modal Legibility . . . . .	213
8.3.5	Mechanisms Underlying Human Decision Making . . . . .	214
8.3.6	Topological Representations for Robotics . . . . .	215
8.4	The Need for Benchmarking . . . . .	216

## LIST OF TABLES

4.1	Generated dataset sizes (number of examples) . . . . .	109
4.2	Braid prediction accuracies for the whole topology $\tau$ and next, second and third events $\tau_1, \tau_2, \tau_3$ respectively. . . . .	110
4.3	Per-permutation prediction accuracies for the next braid generator $\tau_1$ , compared against random guessing and guessing with probability proportional to the prior distribution (frequency of the generator). . . . .	111
5.1	Statistics of paired t-tests between SM and SF, ORCA for different agent numbers. We considered $N - 1$ degrees of freedom, where $N = 200$ is the number of scenarios per class. . . . .	142
6.1	Effect of Navigation Strategy on Robot Behavior . . . . .	162
6.2	Effect of Navigation Strategy on Human Behavior . . . . .	163
6.3	Effect of Navigation Strategy on Group Behavior . . . . .	165
6.4	Effect of Nav. Strategy on Human Impressions . . . . .	167
7.1	Success rate of HTTG in generating the desired, topologically distinct executions for each of the 100 scenarios consider per condition. . . . .	192
7.2	Success rates and computation times of HTTG and CHOMP over 500 randomly generated 2-agent scenarios. . . . .	193

## LIST OF FIGURES

2.1	Robots that collaborate with humans, such as in an assembly task [83], must consider the correctness of both the functional and communicative aspects of their actions. . . . .	32
2.2	Robots operating in close proximity with humans, such as in pedestrian environments need to be cognizant of the communicative effects of their actions in order to ensure human comfort and social compliance. . . . .	33
2.3	These plots illustrate the inference mechanism described in Section 2.2.4 and its effect on $P(a M_{pub})$ . Among the set of actions $a \in A_f$ that accomplish a task, each can be assigned a likelihood of being observed in context. Actions with high likelihoods of $P(a M_{pub})$ are deemed <i>predictable</i> , low ones <i>legible</i> . . . . .	43
2.4	The red, navigating agent (human or robot) selects an action $\hat{a}$ . Out of context (top), the red agent (human or robot) is not avoiding an obstacle, and so the probability of expending needless extra energy is low. In the case of an oncoming blue agent ( $\hat{m}$ ), the likelihood of the oblivious action $P(a_0 M_{pub})$ is low due to social norms, despite being low energy. Conversely, the normally-improbable act of spending extra energy becomes probable in this context. An observer who sees only the red agent's motion can infer $\hat{m}$ from observing $a_3$ . . . . .	56
3.1	A human and a robot are navigating towards opposing directions of a hallway. To avoid collision, they need to agree on an avoidance protocol (passing from the right or left hand side of each other). The jerky behavior of the robot so far and the smooth but non-committal –with respect to a passing side– path of the human agent yield a high-entropy belief distribution over an emerging avoidance protocol from the perspective of both agents. The goal of our planner is to generate a sequence of highly-informative actions that will rapidly reduce the entropy and break a potential livelock or deadlock situation. . . . .	65
3.2	A set of 4 particles are initially ( $z = 0$ ) arranged along the $x$ -axis, on the points 1, 2, 3 and 4. Through a sequence of rearrangements, the particles finally ( $z = 1$ ) reach a final arrangement on the points 2, 3, 1 and 4 respectively. The pattern of their trajectories is a <i>geometric braid</i> . . . . .	68
3.3	The generators of the Braid Group $B_n$ . . . . .	69
3.4	Example of the Composition operation $\sigma_1 \cdot \sigma_2^{-1}$ for $\sigma_1, \sigma_2^{-1} \in B_n$ . . . . .	70

3.5	A space-time representation of a system path in a workspace with four agents (left) along with its corresponding braid diagram (right) and braid word (top right), defined with respect to the path's $x$ -projection. The visualization of the braid diagram and the extraction of the braid word was done using BraidLab [132]. . . . .	71
3.6	Schematic representation of the inference mechanism from the perspective of a robot, navigating in a workspace with 3 other agents. From the perspective of the robot, the system state at time $t$ corresponds to a permutation $p_t$ (derived upon projecting on the $x$ -axis of the robot's body frame), represented graphically with the color permutation at the bottom (the robot order in the permutation is denoted with red color). A set of $M$ final permutations, taking the robot to its destination are considered and a set of 3 compatible joint strategies-braids are planned for each final permutation. The robot reasons over the set corresponding to the union of all sets of joint strategies $\mathcal{T} = \bigcup \mathcal{T}^m, m = \{1, \dots, M\}$ .	75
3.7	Demonstration of the Momentum heuristic for predicting a generator superscript: The $x$ -projections of the agents' paths are about to cross, forming a $\sigma_1^{-1}$ generator. The $z$ component of their angular momentum is negative, indicating a tendency for a counterclockwise rotation, which indicates a negative braid exponent; likewise, in case the $z$ component of the angular momentum were positive, the emerging exponent would be positive. . . . .	81
3.8	A multi-agent scene from the perspective of the planning agent (blue color). At time $t$ , the agent arranges all agents in the scene in an order of increasing $x$ -coordinates with respect to the $x$ axis of its body frame $\{B\}$ and derives a corresponding permutation $p_t$ . Based on observation of all agents' past trajectories (solid lines) and given knowledge of existing landmarks in the scene, the blue agent makes a prediction of everyone's destination (colored pointers) and derives a corresponding final permutation $p^d$ . Transitioning from $p_t$ to $p^d$ may be implemented with a joint strategy $\tau \in \mathcal{T}$ . . . . .	87
3.9	A permutatohedron of order 4 for a scene with four agents. Three alternative paths implementing the transition from the permutation 1234 to the permutation 3412 are depicted in different colors. Each path consists of a sequence of transitions, each of which can be implemented topologically with a braid generator or its inverse.	89

3.10	A game with 3 agents in a $6 \times 6$ workspace. Figure 3.10a and Figure 3.10b depict partial executions of the same scenario (same start and end positions for all agents) with SCN and GREEDY respectively. The current system state is denoted with non-transparent system circles, whereas faded configurations correspond to configurations of past time steps. . . . .	96
3.11	Comparative diagrams, generated upon running 200 experiments with 3 agents in a workspace of size $6 \times 6$ . The red curves correspond to agents running SCN and the blue curves to agents running GREEDY. The compromise between <i>Efficiency</i> and <i>Consensus</i> was set to $\lambda = 0.2$ and the number of paths per permutation to 3. Student's t-tests were performed on all rounds to determine the statistical significance of the profiles difference. The symbols *, ** and *** denote rounds on which the difference in the performance between SCN and GREEDY was found to be significant to a degree described by p-values $< 0.05$ , $0.01$ and $0.001$ respectively, according to a paired Student's t-test. Due to space constraints, the significance symbols were place vertically.	98
3.12	Comparative diagrams, generated upon running 200 experiments with 4 agents in a workspace of size $6 \times 6$ . The red curves correspond to agents running SCN and the blue curves to agents running GREEDY. The compromise between <i>Efficiency</i> and <i>Consensus</i> was set to $\lambda = 0.2$ and the number of paths per permutation to 3. Student's t-tests were performed on all rounds to determine the statistical significance of the profiles difference. The symbols *, ** and *** denote rounds on which the difference in the performance between SCN and GREEDY was found to be significant to a degree described by p-values $< 0.05$ , $0.01$ and $0.001$ respectively, according to a paired Student's t-test. Due to space constraints, the significance symbols were place vertically.	99
3.13	Average Time to destination and average time to get free, i.e., reach a configuration at which no agents are ahead, generated after running 200 experiments involving 3 (Figure 3.13a) and 4 (Figure 3.13b) agents, in a workspace of size $6 \times 6$ . Red bars correspond to agents running SCN and blue bars to agents running GREEDY. The error bars indicate 25-75 percentiles. For these experiments, the compromise between efficiency and consensus was set to $\lambda = 0.2$ and the number of paths per permutation to 3. The *** symbol denotes a highly significant timing difference, according to Student's t-test (p-value $< 0.001$ ). . . . .	100
4.1	Context and time flow around a planning step. . . . .	106
4.2	Learning Architecture. . . . .	108

4.3	The robot is reasoning about different actions leading to qualitatively distinct scene evolutions (left), corresponding to distinct system path topologies (right). . . . .	113
4.4	Swept volumes of 4 agents navigating a circular workspace. The red agent runs <i>SCN</i> whereas the rest of the agents run a separate instance of the social force model. . . . .	116
4.5	Average time to destination per experiment per agent (excluding agent #1). The error bars indicate <i>25th</i> and <i>75th</i> percentiles over 50 experiments. $1SCN \vee 3SF$ is shown to terminate faster than both $4SF$ and $3SF$ , (according to a Student's t-test with p-value $< 0.001$ ). . . . .	118
4.6	Entropy profiles averaged across experiments and agents (excluding agent #1). The black circles indicate timesteps where the entropy measured for $1SCN \vee 3SF$ is lower than for $3SF$ with statistical significance (Student's t-test, with p-value $\leq 0.022$ ). The gray area in the plot corresponds to the first time frame $T_p$ , during which the agent running <i>SCN</i> was moving efficiently and observing the context. . . . .	119
5.1	A human and a robot move towards opposing sides of a hallway. The initial configurations of the two agents make it hard for the human to predict the emerging avoidance strategy ("right" or "left"). The robot detects a slight advantage towards the "right" strategy and acts towards amplifying it and thus facilitating the inference of the human regarding the emerging (right) strategy of avoidance. . . . .	124
5.2	Two braids of the same length, $b_1 = \sigma_2^{-1}\sigma_1$ and $b_2 = \sigma_2^{-1}\sigma_1^{-1} \in B_3$ , with different entanglements. Qualitatively, it may be observed that $b_2$ is more complex than $b_1$ . Figure 5.3 formalizes the complexity measurement, using the Topological Complexity Index, defined on curve diagrams. . . . .	128
5.3	Curve Diagrams of braids from $B_3$ on a disk with 3 punctures. From left to right, curve diagrams of: a trivial braid $\beta_a = e$ (Figure 5.3a), $\beta_b = \sigma_1^{-1}$ (Figure 5.3b), $\beta_c = \sigma_2^{-1}\sigma_1^{-1}$ (Figure 5.3c) and $b_d = \sigma_2\sigma_1^{-1}$ (Figure 5.3d), with complexities $c(\beta_a) = 0$ , $c(\beta_b) = 1$ , $c(\beta_c) = 1.5850$ , $c(\beta_d) = 2$ respectively. . . . .	130
5.4	Social Momentum: The planning agent (red color) is moving towards the red target $X$ , while complying with its pairwise momenta with all other agents. . . . .	134

5.5	Average <i>Complexity Index</i> of trajectories generated by executing 200 scenarios with 3, 4, 5 and 6 agents with the <i>Social Momentum</i> (SM), <i>Social Force</i> (SF) and <i>Optimal Reciprocal Collision Avoidance</i> (ORCA) models. A theoretical lower bound baseline is also included for reference. Datapoints marked black correspond to significantly lower average <i>Complexity</i> of SM than both SF and ORCA, whereas the datapoint marked green indicates significantly lower average <i>Complexity</i> of SM than SF, according to paired Student's T-test. Test statistics can be found at table 5.1. . . . .	140
5.6	Average <i>Path Irregularity</i> per agent extracted by executing 200 scenarios with 3, 4, 5 and 6 agents. Datapoints marked black denote different irregularity of SM than both SF and ORCA, whereas the datapoint marked green indicates lower irregularity of SM than SF, according to paired Student's T-tests. Test statistics can be found at table 5.1. . . . .	141
5.7	Study interface: A video of a scenario execution is shown and users predict how the red agent is going to avoid the blue agent by pressing the corresponding button at the bottom. The display of user's score and performance statistics aim to incentivize fast and accurate responses. . . . .	143
5.8	Relation between the <i>Complexity Index</i> and (a) time until two specific agents pass each other (green points/line) and (b) median time until users give a correct prediction of the passing (blue crosses/line). Times are normalized to the total length of the relevant video. . . . .	145
6.1	This study examines the performance of autonomous robot navigation algorithms in crowded pedestrian spaces. . . . .	148
6.2	Study apparatus: (a) Beam Pro robot [1]; (b) an easel, representing a machine with sticky notes, representing completed maintenance tasks; (c) tracking helmets, sticky notes and markers distributed to participants. . . . .	153
6.3	Top view of the workspace along with example human and robot trajectories, corresponding to transitions between easels (blue objects). . . . .	154
6.4	Expected means and confidence intervals for robot trajectory quality criteria. Quantities labeled with distinct letters (A, B, C) come from significantly different distributions (Tukey's HSD test, $p < 0.05$ ). . . . .	161
6.5	Expected means and confidence intervals for human (Figure 6.5a, Figure 6.5b, Figure 6.5c) and group trajectory quality criteria (Figure 6.5d). Quantities labeled with distinct letters (A, B) are significantly different (Tukey's HSD test, $p < 0.05$ ). . . . .	166

7.1	Spacetime plot of the trajectories of two agents, navigating in a circular workspace (left) and projection of their trajectories until time $t_1$ , onto the $xy$ plane, along with the definition of their pairwise winding angle and winding number (right). . . . .	177
7.2	The destination prediction mechanism. The red agent makes destination predictions for all agents, lying within its circular sensing disk and in front of it. . . . .	183
7.3	Illustration of the planning scheme. At every replanning cycle, the planning agent generates a set of diverse (topologically distinct) predictions about the joint future behavior of all agents, evaluates them with respect to a cost function $C$ and executes the action assigned to it from the prediction of lowest cost. . . .	187
7.4	Top view of trajectories generated by executing the same 3-agent scenario with all possible topological specifications. The subcaptions denote the topology tuple that was used as a specification for each execution. . . . .	190
7.5	Top view of trajectories generated by executing the same 4-agent scenario with all possible topological specifications. The subcaptions denote the topology tuple that was used as a specification for each execution. . . . .	191
7.6	Trajectory Quality for all experiment configurations considered. For group size, the same 200 randomly generated scenarios are executed under each of the conditions considered with both planners. For each condition and measure, we perform a paired Student's t-test to compare the populations yielded by TANP and ORCA. Points with black circular boundaries indicate rejection of the null hypothesis with p-value $< 0.001$ whereas points with star boundaries indicate rejection of the null hypothesis with p-value $< 0.05$ . . . . .	195
7.7	Overlaid predictions made by a TANP agent (red color) as it navigates towards the red landmark in environments with 2, 3 and 4 agents. . . . .	196

# CHAPTER 1

## INTRODUCTION

### 1.1 Preface

The traffic flow in human environments, such as crowded hallways, sidewalks, and rooms may often be characterized as unstructured and even unpredictable, for a variety of reasons, including the heterogeneity of pedestrians' decision-making mechanisms, the lack of formal rules to regulate traffic and the lack of channels of explicit communication among pedestrians. Nonetheless, humans are capable of traversing such environments with remarkable efficiency, without hindering one another's motion. Human navigation not only achieves collision avoidance; it does so while respecting several social considerations, such as the passing preference of others and their personal space [58]. This behavior has largely been attributed to *trust*, with pedestrians trusting that others will behave *competently* [148]. This form of trust enables humans to infer the intentions of others, under the assumption of rational action [33] but also effectively communicate their own intentions by leveraging various implicit communication channels that allow them to broadcast and receive information through their path shape, their body posture, their gaze etc. This exchange of information through nonverbal, implicit communication channels enables humans to effectively negotiate and achieve consensus over mutually preferable joint plans for collision avoidance.

Deploying a mobile robot in a dynamic environment with the aforementioned protocols of communication poses a significant challenge. First, encoding the outlined mechanisms of implicit communication that humans employ to

navigate socially in a crowded workspace into a set of concrete rules and policies is not a trivial task and it remains an active area of research. Second, hardware design limitations significantly constrain a robot's perception and action capabilities, enabling only a limited scene understanding and affording only a limited set of possible behaviors to be generated. Besides, non-expert humans have not yet developed mental models, adequate for accurate interpretation of observed robot behaviors, inference of their intentions, anticipation of their future behaviors or broadly, estimation of their capabilities [89].

The described challenges have resulted in the formulation of interesting and significant research questions that have fueled robotics and human-robot interaction (HRI) research over the past three decades. Traditional approaches have decomposed the problem of robot navigation into subproblems including human motion prediction and robot path planning and control [16, 134]. Over the years, it has become empirically observed that these problems are coupled [47, 137, 138], i.e., the robot's actions influence human behaviors and vice versa. Thus, attention has been shifted towards the design of socially aware navigation planners that incorporate models of behavior prediction into the robot's decision making mechanism [28, 34, 57, 78, 81, 86, 114, 122, 123, 140]. This design concept enables a robot to understand the communicative implications of its own actions to any observing humans but also adjust to human behavior in an informed and socially acceptable fashion.

Despite the significant contributions of the past decades in the area of socially aware navigation, we may still observe several gaps in the literature. These mostly concern: (1) the lack of generalization—approaches tend to be context-dependent, employing heuristically derived, hand-tuned cost functions;

(2) the lack of a sufficiently expressive representation of the joint interaction phenomena, taking place in a dynamic multi-agent environment; (3) the lack of a thorough and in-depth validation in the presence of humans, considering also feedback from them.

This thesis aims at contributing towards addressing these gaps. We strive for generalizability by following the insights of studies on human behavior. We leverage the observation that human navigation tends to be cooperative [148] by introducing a novel, symbolic representation of collective multi-agent behavior, inspired by the Braid Groups [10, 11, 20, 73, 110]. Furthermore, we leverage the insight that human inference tends to be goal-driven [33] and the paradigm of legible robot motion planning [40], by designing an inference mechanism that connects action and goal in multi-agent navigation towards producing context-aware, intention-aware and intent-expressive robot behaviors. We use this mechanism for both human motion prediction and robot motion generation, enabling implicit communication of the robot’s intentions through its path shape—the most universal modality that any mobile robot has access to. This architecture represents our perspective of social compliance as intention-compliance; our family of planning algorithms is designed towards facilitating and accelerating a consensus among multiple navigating agents over a mutually preferable joint navigation strategy that contributes collision-avoidance and progress towards agents’ destinations. We conduct and present an extensive validation of our models, methods and algorithms through simulations and user studies with human participants.

## 1.2 Related Work

Contributing to the field of social robot navigation requires an understanding of human pedestrian navigation. Human navigation is a highly complex behavior, incorporating a multitude of sophisticated mechanisms related to human decision making, inference, and respect of social norms. Researchers from highly diverse research areas including sociology, psychology, computer graphics have contributed towards modeling these mechanisms. Over the past few decades, roboticists have focused on incorporating these models into the processes underlying robot perception and planning. This thesis belongs to this field of research, bringing together the insights of studies on human behavior into the design of computational tools for social robot navigation.

In this section, we introduce a set of works from literature focused on understanding the computational processes underlying human navigation, that have inspired our approach. We also review relevant literature from computer graphics, focused on modeling and reproducing empirically observed pedestrian navigation patterns. We then introduce works that represent the current state-of-the-art in social robot navigation, classified with respect to their validation instruments. Subsequently, we highlight a few features that make our approach unique and expand on them by introducing relevant literature. These include the introduction of topological representations for modeling emerging behaviors in multi-agent navigation and leveraging the communicative component of robot behaviors as a tool for natural and effective blending of mobile robots in pedestrian environments. Finally, we take a step back and note that the aforementioned features essentially constitute tools that allow us to exploit our understanding of the problem structure. This is a common pattern in robotics

literature that has resulted in a number of contributions across a series of challenging classes of robotics problems.

### 1.2.1 Human Navigation

Understanding and modeling human navigation accurately has been the focus of researchers from various fields for a long time. Hall [58] introduced the foundational ideas underlying the field of Proxemics, highlighting the importance of personal space, as a foundation for social comfort in modern urban environments, including pedestrian domains. Karp et al. [72], in their definition of the *Mini-Max Hypothesis of Urban Life*, specify that “urbanites seek to minimize involvement and to maximize social order”. This idea is also present in Wolfinger’s definition of the *Pedestrian Bargain* [148], a concise, high-level protocol of foundational social rules that regulate pedestrian navigation: (1) “people must behave like competent pedestrians” and (2) “people must trust co-present others to behave like competent pedestrians”. Trust to the rules of the bargain constitutes the basis of smooth co-navigation in human environments, as it enables pedestrians to plan with the expectation that others will also behave competently and thus cooperate to resolve potential conflicts. Based on our everyday-life experience, this is clearly not always the case: sometimes, some pedestrians might violate the rules of the bargain, by behaving erratically or antisocially in the presence of others. However, as Wolfinger [148] observes –and as we might recall from our own experiences as well– humans tend to make use of remedial mechanisms in such cases, i.e., they tend to acknowledge their fault, apologize and move on. These mechanisms restore and reinforce the established notion of trust to the rules of the bargain, that is necessary to preserve the social order in

pedestrian environments.

The aforementioned mechanisms are made possible through a sophisticated mechanism of perception and action [145], enabled through information exchange mostly via path shape, body posture and gaze [51] that has been widely studied from a number of different fields. For example, Carton et al. [24] studied the trajectory planning horizon of humans in locomotion tasks towards informing the design of models for the prediction of human walking behaviors. Warren [145] proposed a model that may describe organization in human behavior in a number of tasks—including lane formation in navigation—by treating an agent and its environment as a pair of coupled interacting dynamical systems. More broadly, Csibra and Gergely [33] highlighted the teleological nature of human inference, suggesting that humans tend to attribute goals to observed actions.

Human navigation which is heavily based on a multi-modal information exchange. However, there is evidence that even low-degree-of-freedom robots are capable of increased expressiveness [85]. Leveraging the most universally available modality that all mobile robots possess—the path shape—our approach focuses on expressing a robot’s navigation intentions by encoding its communicative signals in its actions. Our planning framework aims at generating motion that maximizes social order through a local, collision-free action selection with a global lookahead. This is made possible with a principled design of a goal-driven inference that connects individual and collective behavior towards enabling artificial agents to understand the effects of their actions on the behaviors of others.

## 1.2.2 Simulating Artificial Pedestrian Flows

The problem of generating smooth, collision-free, multi-agent simulations has been central in a number of applications, ranging from city planning to the study of evacuation scenarios and computer game design. To this end, a class of papers proposes frameworks that model agents as interacting particles, attracted to their destinations and repulsed by others. Within this class, the Social Force model [62] has been one of the first and most influential approaches, whereas several works have employed similar models with additional considerations such as discomfort fields [139], local predictive processes [67, 69, 70] and cognitive heuristics [21, 45, 109]. Some works have employed data-driven techniques to learn the parameters of human navigation in different contexts from simulated [63] or real world demonstrations [71]. Finally, van den Berg et al. [142] and Knepper and Rus [81] have proposed decentralized motion planners that explicitly leverage the expectation of cooperation.

Our proposed planners do not aim at directly replicating observed human behavior. Instead, they focus on prescribing actions corresponding to the perceived flow of a pedestrian scene, by detecting and complying with human intentions. Our planners are also cooperative by design. However, in contrast to most of the aforementioned approaches that either treat agents as moving obstacles or make purely local motion prediction, our inference mechanism enables agents to understand that they are part of a crowd of intelligent agents that cooperate to reach consensus over a collision avoidance protocol.

### 1.2.3 Tracking Pedestrian Flows

The computer vision community has contributed a number of methods focused on tracking the local or global behavior of pedestrians in different environments. For example, Scovanner and Tappen [121], Pellegrini et al. [115], Alahi et al. [4] and Ma et al. [95] present data-driven models for local short-term trajectory predictions for interacting pedestrians, that make use of models of social interactions, whereas Zhou et al. [150] predict large-scale, global pedestrian interactions. These contributions are of particular significance both for offline labeling of pedestrian datasets but also for online tracking of multi-agent behaviors for robotics applications.

Similarly to some of the aforementioned approaches, our proposed inference mechanisms enable agents to perceive the global nature of observed actions. Under the assumption of rationality (agents move efficiently towards intended destinations), our prediction models essentially score the set of possible multi-agent collision avoidance strategies that the agents could follow to reach their destinations in a collision-free fashion.

### 1.2.4 Social Robot Navigation

Social robot navigation constitutes a significant thrust of human-robot interaction research over the past few decades [27, 88, 133]. Particular emphasis has been placed on the design of autonomous socially aware navigation planning algorithms and on the study of the interaction between navigating robots and humans. Researchers have been inspired by the mechanisms underlying human navigation [148] and general human behavior in public spaces [51, 58]. This has

led to the adoption of theory, models and methods from the field of cognitive science [109], psychology [145] and sociology [148] into the design and evaluation of proposed navigation frameworks. However, the complexity and cost of building and testing an autonomous robotic system often prohibits a systematic and thorough experimental validation of navigation frameworks. This section presents a classification of state-of-the-art approaches with respect to their type of validation.

#### 1.2.4.1 Simulation Studies

Recent advances in the fields of graphics [63, 71, 142] and crowd dynamics [62, 67] were based on physics-inspired models of the interactions among multiple navigating agents: socially compliant and humanlike motion is generated as the result of multiple interacting potential fields, representing agents' objectives and intentions. This foundational idea has set the conceptual basis behind the design of a number of approaches in the field of social robot navigation.

Luber et al. [94] learn a set of dynamic navigation prototypes from a human trajectory dataset and use them for trajectory prediction and generation. They demonstrate the performance of their planner with respect to the efficiency and humanlikeness of generated paths on 182 scenarios of the same dataset and show how it outperforms a rule-based, proxemics-theory enforcing baseline. Vasquez et al. [143] learn a cost function to represent the dynamics of social navigation by training on a dataset extracted by teleoperating a robot in different real-world scenarios. They evaluate the ability of their learned model to reproduce trajectories of social compliance, modeled as a composite score of cost functions representing human comfort. Finally, Bera et al. [17] make use of

concepts from Personality Trait theory to classify the behavior of pedestrians towards informing their motion models and a robot’s path planning. Simulation results demonstrate improved trajectory prediction and more socially compliant on a number of human datasets.

#### 1.2.4.2 Experimental Demonstrations

A number of works have presented important experimental demonstrations in human environments to validate their approaches. Althaus et al. [6] focus on the problem of social engagement. They build a robot designed to approach humans and engage in a conversation with them. They present control laws for approaching a person and maintaining a socially acceptable distance. A recorded experiment with three participants demonstrates the efficacy of their approach. Sisbot et al. [123] presents a cost-based planner that considers a set of social heuristics at the planning stage to generate motion that is visible and safe around humans. A series of documented interactions between the robot and a human in a lab environment demonstrate the capabilities of the framework. Lam et al. [90] present a motion planner, designed according to a set of “harmonious” rules, inspired by the theory of proxemics [58]. Their framework is shown to respect human zones of personal space in a series of experiments of close interaction. Park et al. [114] build an automated wheelchair and design a model-predictive control law for smooth motion generation in crowded environments. Their approach treats humans as dynamic obstacles and focuses on avoiding them smoothly. They test their framework in an indoor environment and report a set of successful collision avoidance processes under crowded settings. Kretzschmar et al. [86] employ an inverse reinforcement learning ap-

proach to learn the social components of human navigation by training on an hour-long lab dataset of four navigating humans and on a public dataset on a crowded scene. They deploy the model on a robotic wheelchair that is able to navigate socially next to navigating humans in a narrow hallway. Finally, Chen et al. [28] present a deep reinforcement learning approach to learn social norms (passing from the right-hand side and overtaking on the left) from a synthetic, simulated dataset. They report an experimental demo, run in a large, crowded academic building.

#### 1.2.4.3 Experimental Studies

A significant body of work has employed field studies in crowded environments such as museums, malls or academic hallways. Thrun et al. [134] present a tour-guide robot equipped with a set of probabilistic algorithms for mapping, localization, people-tracking, and planning. The robot interacted successfully with thousands of visitors for two weeks in a busy museum. The authors present a comprehensive report of the robot's log and a classification of observed types of interaction between the robot and visitors. Bennewitz et al. [16] cluster a dataset of observed human trajectories into a set of classes and use it for on-line prediction on a robot. A series of 10 experiments indicates increased time-efficiency resulting from their approach, compared to a linear prediction baseline. Pacchierotti et al. [113] implement and test a proxemics-based control framework on an autonomous robot through a user study, conducted in a corridor. A total of 10 participants were exposed to three different conditions corresponding to the robot passing next to them with a different lateral distance each time. Users' ratings showed that humans felt uncomfortable

when the robot was closer to them. Foka and Trahanias [48] present a probabilistic algorithm that makes predictions about future human paths to plan collision-free motion. They report logs and performance aspects upon running the robot for 70 hours in an indoor academic building. Kirby et al. [80] present a constrained optimization-based algorithm that incorporates a series of social conventions, such as passing-side conventions and respect of humans' personal space into the robot's decision-making. A user study involving 27 human subjects navigating alongside a robot in an academic hallway demonstrated evidence that humans interpret the robot's behavior as socially appropriate [79]. Shiomi et al. [122] present a planner, based on the social force model [62] for generating humanlike collision avoidance navigation behaviors. A 4-hour field study in a shopping mall demonstrated that the proposed approach achieves safer and more comfortable interaction than a baseline. Trautman et al. [138] present a navigation framework that explicitly incorporates the assumption of human cooperation into their learned trajectory prediction model to enable a robot to navigate among dense human crowds. They report the performance of a real robot in terms of safety and efficiency in a large-scale field study (488 runs), conducted in a crowded cafeteria. Kato et al. [74] learn a model of human intent inference to generate social approaching navigation behaviors. They test their approach on a humanlike robot employee in a crowded mall and record interactions with 130 people, suggesting that a compromise between proactive and passive approaching behavior is preferred by humans. Kim and Pineau [78] learn a model of socially compliant robot motion from human demonstrations and robot teleoperations in crowded environments. They test their approach on a robotic wheelchair in a crowded hallway and report humanlike and efficient performance in 10 field runs. Truong and Ngo [140] fuse elements of

the social force model [62] and the Reciprocal Velocity Obstacle model [142] to generate socially aware robot motion in crowded scenes. Examples from experiments, conducted in an office environment demonstrate smooth operation against static or moving obstacles.

### 1.2.5 Topological Representations for Multi-Agent Navigation

In the area of decentralized multi-agent navigation, a significant portion of the literature typically assumes that navigating agents employ policies with shared architecture, with notable examples including the Social Force [62] and Reciprocal Velocity Obstacle [142] frameworks as well as more recent algorithms of similar basis (e.g. Moussaïd et al. [109], Karamouzas et al. [71]). These works do not make explicit predictions of other agents' trajectories but make decisions under strong assumptions on their behaviors. The problem of predicting the trajectories of multiple navigating agents in real time and for an adequate horizon to allow for motion planning is challenging, as suggested by the literature on tracking (e.g. [95]). This has motivated roboticists to look for more practical alternatives for multi-agent motion prediction. In particular, several approaches have leveraged the coupling of agents' decision-making in multi-agent navigation as a way to guide the motion planning process. Some of them have employed learning techniques [28, 78, 86, 138] to develop models for prediction and generation of humanlike trajectories whereas others have employed heuristics to directly exploit the topological structure of the problem [81] as a more tractable alternative to explicit trajectory prediction. A topological representation provides several benefits including robustness of computation, increased expressiveness and a formal enumeration of possible classes of joint behavior.

Our work aims at directly leveraging these benefits; we propose a series of planners that focus on predicting the topology of an emerging multi-agent behavior rather than its geometry. This type of prediction is robust to minor variance in the assumed decision-making mechanisms of other agents, as it ignores local features, in favor of global emerging patterns.

More specifically, it can be observed that multi-agent navigation of rational agents in a bounded environment has an interesting property: agents' decision making is spatiotemporally constrained. This property arises from: (1) the fact that agents cannot occupy the same configuration at the same time and (2) driven by their rational decision-making mechanisms, they have an incentive to move efficiently towards their destinations, while avoiding collisions with the workspace boundary and each other. This property may be traced at the entanglement of agents' trajectories throughout the execution of a multi-agent scene. Agents' trajectories may be thought of as strings that become knitted around each other, according to agents' navigation strategies. This knitting has topological properties which can be studied with tools from the low-dimensional topology (the field of topology that studies topological spaces of four or fewer dimensions). This thesis has been heavily inspired by the described structure and makes use of a series of formalisms, metrics and computational tools to leverage its existence. Imbuing artificial agents with a model of this structure enables them to incorporate a notion of social compliance into their decision making. More broadly, we find that the use of topological representations for modeling, perception and planning problems in robotics could offer significant benefits, including a greater potential for generalizability across diverse domains and inherent model explainability through the affordance of a qualitative reasoning.

More specifically, the foundation of this work is the topological construction of braids from the field of low-dimensional topology [73]. The formalism of braids, first presented by Artin [10, 11] and extensively studied by Birman [20] has been an inspiration for applications in various disciplines. Our approach is specifically inspired by the use of braids as a model that captures the entanglement of particle trajectories in a fluid [131]; in a similar fashion, we employ braids to model the entanglement of the trajectories of navigating agents.

The idea of employing braids in planning problems in robotics is not new. It may be traced at least as back as to Ghrist [49], who drew a parallel between braids and configuration spaces of robotic systems. Later, a few works proposed planning frameworks for explicit coordination of multi-agent systems, using braids as prototypes of multi-robot collision avoidance maneuvers. In particular, Hu et al. [68] presented an optimization-based framework for determining low-energy conflict resolutions among multiple coordinating agents, navigating on a plane, whereas Diaz-Mercado and Egerstedt [36] developed a framework for centralized multi-robot mixing.

Although we are also making use of braids to model multi-robot behaviors, the scope of our approach is inherently different, since our target application concerns navigation in dynamic environments where no explicit communication takes place among agents. In our case the agents do not follow a pre-specified or decided-upon braid pattern, but rather employ a braid-based probabilistic reasoning to reach a topological consensus that best complies with everyone’s intentions or preferences. For our purposes, braids provide an abstraction of the complex spatiotemporal multi-agent dynamics of interaction among agents. This abstraction enables artificial agents to reason about uncertainty

in a principled fashion, as the dual geometric and algebraic representation of braids enables them to enumerate a set of diverse, topologically distinct scene evolutions.

### 1.2.6 Implicit Communication in Human-Robot Interaction

Humans tend to attribute meaning and intentions to observed actions, executed by others. This process has been well studied in the cognitive science and psychology literature Baker et al. [12], Csibra and Gergely [32, 33], Wiese et al. [146], with researchers highlighting the teleological nature of human inference. Humans tend to interpret observed actions as approximately rational and hence to attribute context-specific, goal-driven intentions to them. This mechanism allows humans to leverage channels of implicit communication to communicate more fluently by encoding communicative signals to their goal-driven actions.

Over the past few years, a significant body of work has focused on the development of computational frameworks that leverage the existence of this mechanism in humans towards achieving safe, natural and effective human-robot interaction [40, 66, 112, 126, 128]. At the core of this field of research is the development of mathematical models that connect intention and action. These models may be used directly for human action prediction (e.g. for Unhelkar et al. [141] show how biomechanical turn indicators may improve human trajectory prediction) but they can also be inverted to generate informative robot actions that implicitly communicate the robot's objectives. Engineering systems capable of the latter requires a characterization of actions in terms of their communicative properties. To this end, Dragan et al. [41] introduced the the notion of legibility

as a property of action that allows an observer to make a correct inference of an actor’s goal and presented a computational framework for generating legible actions [40]. Legible actions have been shown to result in reduced planning effort for humans [23, 25] but also increased efficiency in human-robot interaction applications such as human-robot handovers [126], legible robot pointing [66], legible reaching [40], effective robot recovery from failure through verbal communication [128], effective collaborative manipulation [112]. Such applications have also resulted in a growing interest in incorporating intention-aware and intent-expressive systems in autonomous cars. For example, Bandyopadhyay et al. [13] and Ferguson et al. [46] presented navigation frameworks for planning autonomous car navigation, based on models for predicting pedestrian intentions, whereas Sadigh et al. [120] proposed a planning framework that reasons about human drivers’ mental models to plan actions that influence their decision making towards desired outcomes.

#### 1.2.6.1 Implicit Communication in Multi-Agent Navigation

The family of planning algorithms presented in this thesis is designed towards leveraging the sophisticated implicit communication channels that humans have developed [33], towards enabling mobile robots to broadcast clearly their intentions through their actions, so that they are intuitively understood by humans. In particular, it aims at manipulating an agent’s path shape to convey its intentions to others. Our inference mechanism allows our agents to monitor the uncertainty over future multi-agent behaviors and take information-rich or legible actions that aim at reducing this uncertainty.

Existing works on legible motion generation tend to associate the notion of

a goal or intention with a point in a configuration space (e.g. [40, 87]). In a static and structured environment, where the dynamics of interaction among agents is predictable or known a priori, this is a well-motivated modeling decision, as the observers' belief could be assumed to be an isolated relationship between an observed motion and a potential destination. However, in dynamic and unstructured environments, such as typical pedestrian navigation domains, where the dynamics of interaction among multiple agents are complex, sole knowledge of an agent's destination may be insufficient to inform others of the agent's immediate behaviors. This highlights the need for a new consideration of legibility that captures the interactions with neighboring agents. We contribute towards addressing this gap through the introduction of a series of novel topological representations that captures the collective behavior of multiple navigating agents. This allows an artificial agent to understand and anticipate the complex emerging couplings among agents' navigation strategies and inform its decision making towards exhibiting socially compliant behaviors.

### **1.2.7 Exploiting Structure**

From a broader point of view, this thesis makes a series of contributions through leveraging domain knowledge about the problem. From a high-level perspective, our planners seek to achieve a succinct communication of the robot's intent (navigation strategy) to any observing (human) agents, by broadcasting a corresponding signal through its action selection. This machinery directly leverages the existence of information that is known to be publicly available to all agents (commonly referred to as common ground, or context) and also the existence of goal-driven inference mechanisms of the assumed-rational agents navigating

next to the robot (for more details, please see Part I). Furthermore, the introduction of topological braids into the study of problems of multi-agent collision avoidance could also be thought of as a way to exploit their underlying structure (see Part II). This concept is not new in engineering and especially in robotics; in fact, the same foundation has resulted in a number of revolutionary ideas in different research areas. In the field of navigation, both Bennewitz et al. [16] and Ziebart et al. [151] exploit the structure of an indoor environment to extract the patterns of human motion. The geometry of static obstacles in a pedestrian workspace naturally gives rise to a set of clusters of collision-free human trajectories. Both of these works leverage this structure to make human trajectory prediction. In manipulation, Erdmann and Mason [43] exploit contact mechanics to reorient planar objects without the use of sensors. In particular, they show how a sequence of impacts with the environment could lead to a desired reorientation of a planar part. This idea has been further developed into the research area of non-prehensile manipulation, in which a robotic manipulator exploits the surrounding –sometimes cluttered [39]– environment to achieve a desired manipulation task [35]. Finally, Choudhury [29] presents a family of sampling-based planning algorithms that are designed to output high-quality solutions by exploiting the domain knowledge about the structure of the planning problem in consideration.

### 1.3 Contributions

This thesis presents a family of motion planning algorithms for socially competent robot navigation in crowded, dynamic human environments with no explicit communication. The components of our algorithms are guided by prin-

principles extracted from studies on human behavior, whereas our extensive simulated and experimental validation presents a thorough, in-depth comparative look at the benefits and weaknesses of our approach. In the following sections we enlist and expand on our main contributions, grouping them into three main classes: (a) an explicit summary of the key results presented in this thesis; (b) a summary of some technical developments that arose in the process of accomplishing these results; (c) a broader perspective, arising from the consideration of this thesis as a whole.

### 1.3.1 Key Results

A foundational idea underlying this thesis is the incorporation of models of multi-agent collision avoidance into the motion planning process of a robot. Our goal is to approach the notion of socially competent navigation, as presented by Wolfinger [148], by (1) understanding and respecting the navigation intentions and preferences of nearby humans and (2) by implicitly communicating its own intentions to them through its actions. The following results could be interpreted as direct consequences of this architecture:

- We show that groups of artificial navigating agents, capable of understanding the spatiotemporal, topological structure of multi-agent collision avoidance are able to coordinate more efficiently with each other to avoid collisions than groups of purely efficient (greedy) agents in 3-agent and 4-agent scenarios with no communication on discretized domains [101, 104]. This allows them to achieve higher time/energy efficiency.
- We show that the introduction of a single navigating agent employing the

aforementioned inference mechanism in a scene with a group of (3 and 4) other agents navigating with a purely reactive baseline (Social Force [62]) in scenarios with no communication on a continuous domain results in accelerated uncertainty decrease and higher efficiency for the latter ones [105]. This result implies that the decision making of agents running our planner simplifies the inference and planning processes for other agents.

- We show that multi-agent navigation behaviors, generated by our planning framework appear to be more legible, from the perspective of human observers [107] than other baselines in the area of multi-agent collision avoidance ([62, 142]). Humans are able to predict the evolution of multi-agent scenes more confidently, accurately and faster.
- We show that human subjects walking around a robot navigating with our architecture in a controlled lab environment appear to follow paths of lower acceleration than another autonomous [142] and a teleoperated baseline [108].

### 1.3.2 Technical Developments

In the process of producing the results above, we also arrive at the following findings and technical developments:

- We introduce a formal framework for symbolically enumerating and classifying multi-agent navigation behaviors with respect to their topological properties [97], through the incorporation of topological braids [20].
- We present a learned model, using the LSTM architecture [64] that outputs a probability distribution over future topological classes of multi-

agent collision avoidance, given observations of agents' past trajectories in a multi-agent navigation scenarios [105].

- We introduce the index of Topological Complexity [42] as a measure for multi-agent trajectory analysis [107]. Topological Complexity quantifies the intensity of the entanglement of a multi-agent trajectory over time.
- We show that low Topological Complexity correlates with high Legibility in multi-agent navigation scenarios [107]. In particular, according to the findings of an online, video-based user study, the lower the topological complexity of a multi-agent trajectory, the more accurate and faster the predictions of human subjects about the evolution of a multi-agent navigation scene are.
- We show that the introduction of a teleoperated robot into a workspace with a group of (3) humans results in lower overall mixing among them (lower Topological Complexity), compared to two identical conditions under which the robot is running autonomously [108]. More experiments are required to extract a generalizable conclusion but this finding appears to highlight fundamental differences in human decision making for navigation (see Chapter 8).
- We develop the first –to the best of our knowledge– method for automatically synthesizing multi-agent trajectories from symbolic topological specifications [101]. Our method takes as input a sequence of symbols, specifying the ways with which a set of agents should be avoiding each other and outputs a set of trajectories, the entanglement of which satisfies the provided specification. We show how our method is superior to a trajectory-optimization baseline [153].

### 1.3.3 A Broader Perspective

In this section, we provide a deeper look on some of the developments described in the previous section.

**Topological Structure of multi-agent collision avoidance:** We observe that in a shared workspace, where multiple agents are navigating towards their destinations, their navigation planning is *spatiotemporally coupled*, under the geometric constraints imposed by the workspace bounds and the assumption of rationality in their decision making. This coupling has topological properties; in the simple case of two agents walking towards opposing sides of the same hallway, two topologically distinct outcomes arise: collision avoidance from the right or left hand-side. This thesis proposes a novel representation, based on the formalism of braid groups [20] that captures the described decision making coupling. This representation naturally provides a formal model for enumerating the set of classes of topologically distinct, multi-agent navigation behaviors that could emerge under any scenario, involving any number of agents in any type of environment. Each behavior is mapped into a unique symbol, which allows for the construction of mechanisms for inference and symbolic planning.

**Social Compliance:** We leverage the described structure to introduce a novel consideration of social compliance for multi-agent navigation. We view social compliance as a state of consensus over a cooperative, joint, collision-free navigation strategy among a set of navigating agents. Pedestrians negotiate this strategy nonverbally, through signals broadcast by a set of modalities, such as path shape, body posture, eye gaze, gestures etc. This perspective is motivated and supported by studies highlighting the cooperative nature of human navigation (e.g. [148]). The family of planning algorithms presented in this thesis

aim at enabling a mobile robot to leverage an understanding of this mechanism towards accelerating consensus with co-navigating humans.

**Nonverbal Communication via Path Shape:** A central consideration in the design of our algorithms is the notion of intention-awareness. We have proposed a family of inference mechanisms, that attribute intentions to navigating agents, according to their observed past behaviors. These mechanisms are used both for predicting the future motion of others but also for generating compliant and consistently intent-expressive ego motion for a robot. This design enables non-verbal communication of the robot’s intentions via careful selection of its path shape and allows artificial agents to be cognizant of how their actions might be interpreted by humans around. Our framework is inspired and supported by studies on the mechanisms of goal attribution in humans (e.g. [33]), highlighting the natural tendency of humans to assign goals to observed actions, executed by other agents.

**Legibility in Multi-Agent Navigation:** We cast the problem of achieving consensus over a navigation strategy as the problem of minimizing their corresponding uncertainty. We model uncertainty as the Information Entropy of the probability distribution over joint navigation strategies. Minimizing Entropy corresponds to maximizing agents’ Information Gain regarding the unfolding joint navigation strategy. Our algorithms make use of the Entropy cost as a heuristic towards the selection of intent-expressive actions. This mechanism for generating informative actions is in parallel to the definition of *Legibility* cost by Dragan and Srinivasa [40] but represents a broader, model-free consideration of Legibility as the property of intent-expressiveness.

**Potential for Generalization:** The use of topological representations provides

the potential for generalization to any environment geometry with any number of agents. Furthermore, the incorporation of principles from studies on human navigation into the design of inference mechanisms provides the potential for generalization to any type of context. However, in the context of this thesis, we have not conducted relevant studies to validate these hypotheses in practice.

**Evidence that Topological Planning is Sufficient:** The majority of works in multi-agent navigation employ short-term geometric trajectory predictions to account for collision avoidance. We argue that the cooperative nature of human navigation renders detailed trajectory prediction unnecessary, as long as consensus among agents is reached over a joint navigation strategy. Although we have not directly proven this statement, the performance of our algorithms –all of which are based on the prediction of trajectory topology– attest to this. Intent expressiveness, combined with inference and trust allow for a qualitative type of prediction, similar to the type of prediction that humans employ.

**Extensive Validation:** To the best of our knowledge, our work is unique in terms of depth and thoroughness of validation within the area of social robot navigation and social robotics more broadly. In particular, we present evidence extracted from: (1) a series of simulation studies, demonstrating the performance of our framework with respect to a set of trajectory quality measures of social compliance; (2) an online, video-based user study with more than 180 participants, indicating that humans perceive the motion generated by our framework as more intent-expressive; (3) an experimental study, featuring a principled and original study design, conducted in a controlled lab environment with 105 human participants, suggesting that human acceleration tends to be lower when navigating next to a robot running our framework.

## 1.4 Published Content

The work presented in this thesis is based on material that has been published in peer-reviewed workshop papers and proceedings of international conferences and journals.

The foundation of our research direction was our general mathematical framework for implicit communication via Entropy minimization [84]. This work motivated the idea of constructing a probabilistic mechanism for inference of multi-agent navigation strategies to be used as a tool for accelerating consensus in multi-agent navigation scenarios with no explicit communication. To do so, we introduced a representation for enumerating joint navigation strategies, based on the formalism of topological braids [97] (preliminary versions [98–100]) and incorporated it into the Socially Competent Navigation (SCN) planner. We presented evidence from extensive simulated evaluations in a discretized workspace with homogeneous agents, suggesting that SCN results in increased efficiency [104]. We then extended our framework to account for more realistic, continuous domains, through the incorporation of a learned mechanism for inferring navigation strategies. We conducted extensive simulated evaluations on a circular workspace and showed that our approach enables heterogeneous agents (not running our planner) to achieve increased efficiency [105] (preliminary version [106]).

The high computational costs of SCN motivated us to develop Social Momentum (SM), a cost-based planner that reasons about pairwise collision avoidance maneuvers instead of joint strategies. We presented evidence, extracted from an online, video-based user study with more than 180 participants, sug-

gesting that humans find the motion generated by SM as more legible, compared to widely used baselines for multi-agent simulation [107]. Furthermore, we conducted an extensive experimental validation of SM in an experimental study involving a Beam Pro telepresence robot, navigating around groups of human participants in a controlled lab environment. We found evidence suggesting that humans navigating around a robot running SM exhibit smoother motion than two other considered baselines (a teleoperated one and another, widely employed autonomous one) [108].

Finally, we extended the SM planner to perform online multi-agent trajectory prediction for robust navigation in environments with multiple agents with changing intentions or agents without collision avoidance capabilities/intentions. This framework –named Topologically Adaptive Navigation Planning (TANP)– is based on the simultaneous generation of a set of topologically distinct multi-agent trajectory predictions, generated by a physics-inspired mechanism, based on a system of virtual vortices. Extensive, comparative simulations demonstrated the efficacy of our approach [101] (preliminary versions [102, 103]).

## 1.5 Overview

This thesis is organized into five parts. We explain this organization and describe how parts relate to each other.

Part I provides a mathematical framework for implicit communication in Human-Robot Interaction (HRI) and motivates the technical development underlying this thesis. Then, the three following parts present three generations of motion planning frameworks, building on the theoretical basis of Part I and on

the observed topological foundations of multi-agent collision avoidance.

Part II introduces the concept of planning by reasoning about joint navigation strategies. We model joint strategies as topological braids and design a probability distribution that predicts a joint strategy from observation of agents' past behaviors. We then present a cost-based planner that generates intent-expressive robot motion by minimizing the entropy of the distribution over joint strategies. We conduct an extensive simulated validation of the proposed approach, considering a series of challenging multi-agent navigation scenarios.

Part III builds on part II, introducing our Social Momentum planning framework. The basis of the Social Momentum framework is the decomposition of a joint navigation strategy into a set of pairwise collision avoidance maneuvers between the robot and all other agents. Our framework makes use of the physical quantity of angular momentum as a heuristic for inference of pairwise collision avoidance intentions between the robot and other agents. A cost-based planner generates motion of maximal compliance between the intentions of the robot and the intentions of surrounding agents. Extensive simulated trials and two user studies (one online, video-based study and one experimental, lab study) demonstrate the efficacy of our approach.

Part IV combines the central ideas from part Part II and Part III and presents a framework for online, multi-agent trajectory prediction. This framework is based on the generation of multiple, topologically distinct multi-agent navigation strategies at every replanning cycle towards enabling a robot to react robustly to changes of a dynamic environment. Simulation results demonstrate the ability of the robot to maintain a smooth behavior under challenging conditions involving the emergence of agents with changing intentions or agents

with heterogeneous policies.

Finally, Part [V](#) summarizes the findings of this thesis and provides conclusions and directions for future work.

## **Part I**

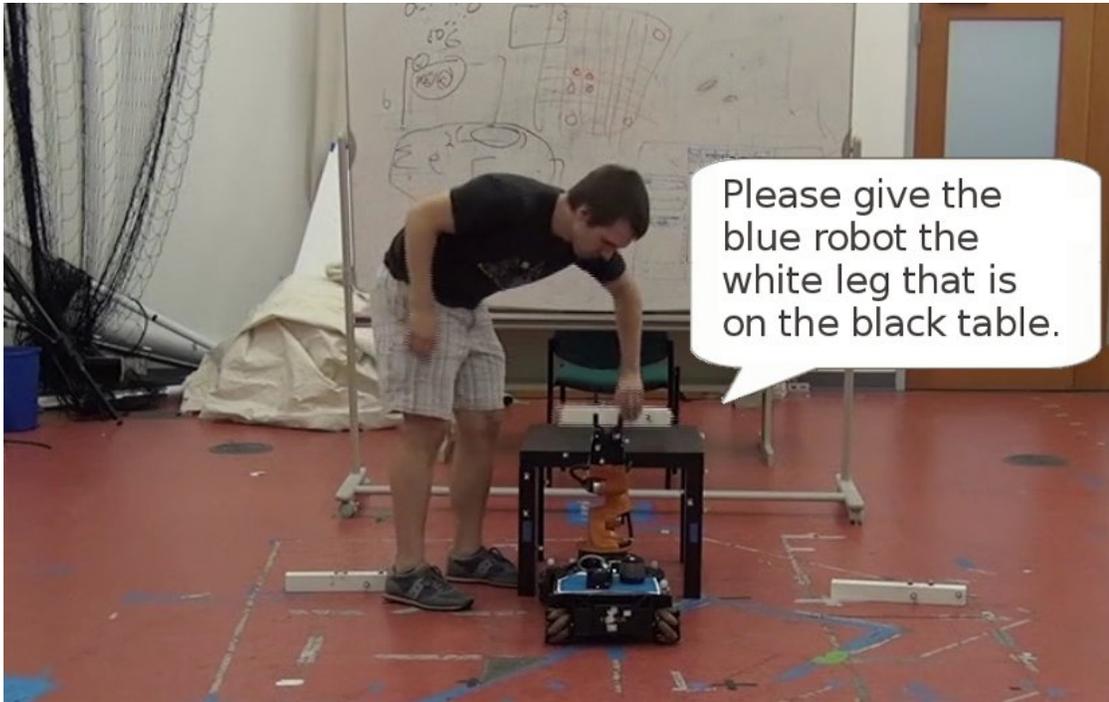
# **Implicit Communication for Human-Robot Interaction**

## CHAPTER 2

### A MATHEMATICAL FRAMEWORK FOR IMPLICIT COMMUNICATION IN HUMAN-ROBOT INTERACTION APPLICATIONS

Actions performed by an agent (actor) in the presence of another agent (observer) comprise two aspects: functional and communicative. The functional component achieves the goal of the actor, whereas its communicative component expresses additional information to the observer. In the absence of explicit communication channels between the agents (e.g. verbal communication), the interpretation of the communicative component is restricted to be done by leveraging information that is public to both, known as common ground. Much of human communication is performed through this implicit mechanism, and humans cannot help but infer some meaning – whether or not it was intended by the actor – from most actions. In fact, humans instinctively perform this inference, thus reading additional meaning about the intent of an action [33], and it is not even uncommon to treat information gleaned implicitly through inference as though it had been stated outright.

We argue that it is crucial for robots collaborating with humans (e.g. collaborative furniture assembly Figure 2.1) or operating in close proximity with them (e.g. navigating pedestrian environments Figure 2.2) to model humans' inference mechanisms. Humans will instinctively make use of them for both inferring the meaning of observed actions but also for generating their own actions and they will also expect others to do so. We further argue that if a robot fails to attend to a human's interpretation of its own actions through the implicit communication mechanisms, then people will perceive the robot's purely functional actions as sending random implicit signals, sowing confusion. Therefore,



**Figure 2.1:** Robots that collaborate with humans, such as in an assembly task [83], must consider the correctness of both the functional and communicative aspects of their actions.

social compliance, human comfort and robot task performance in a variety of settings and contexts rely heavily on robots' understanding of humans' implicit communication mechanisms. As a result, the area of implicit communication in human-robot interaction applications has received a significant amount of attention over the recent years, complementing the more established area of explicit human-robot collaboration [5, 38, 53, 59, 60, 96, 147]. Implicit communication is identified by various terms in differing contexts. For example, in robot motion, including reaching [41] and social navigation [107], it has been termed *legibility*, whereas in natural language generation for HRI, it has been called *inverse semantics* [83]. In both of these cases, the meaning is extracted by leveraging common ground. Common ground constitutes also the core of *conversational implicature* [56], a well-studied area in the field Linguistics. Observing the simi-



**Figure 2.2: Robots operating in close proximity with humans, such as in pedestrian environments need to be cognizant of the communicative effects of their actions in order to ensure human comfort and social compliance.**

larities among these works, the goal of this chapter is to unify them by explicating a common mathematical framework that underlies all of them. In particular, we introduce: (a) a unifying mathematical framework describing how and why people implicitly communicate information on top of functional behaviors; (b) formal expressions for encoding and decoding communicative actions; (c) collected example applications to illustrate the theory.

## 2.1 Implicit Communication in Human Activities

Humans are able to express a multitude of ideas “in code”, by means other than explicit natural language statements. Motivations for implicit communication include efficiency, tact, group cohesion, and social bonding. In this section, we give examples of several categories of implicit communication. Message categories include expressing intent, coordinating plans, and conveying information. Broadly, these categories all fulfill the role of setting expectations, and we consider each separately.

Social navigation is among the most superficial forms of interaction, yet it is rife with implicit communication. In social navigation, the objective is to avoid collision with co-inhabitants of the space and reach one’s destination. Combined, these objectives comprise the navigator’s intent. Collision avoidance without intent expression is only the barest definition of correct navigation – it alone would not be judged as competent behavior by fellow pedestrians [123]. Competence demands that we convey our intended trajectory to nearby observers. We trust in return that they will convey their intent to us. Such intent-expressive actions minimize the global uncertainty about future motions of the agents (humans or robots) in the scene, leading to smooth and stable motion.

We borrow from Barbalet [14] the definition of trust as “the confidence that another’s actions will correspond with one’s expectations.” In the absence of social trust, people begin to behave defensively, and the efficiency of motion drops globally in response.

Coordination among team-mates engaged in a joint activity requires setting expectations of future actions. Consider the simple example of Steve and Cathy assembling furniture together, in which a number of screws must be inserted and tightened. Steve might pick up the screwdriver, which achieves the functional objective of readying Steve to tighten screws. In context, the action also implies that Cathy should gather screws for insertion in order to help. Since Steve is cooperative, Cathy knows that once she begins to insert screws, Steve will fulfill his implicit promise to tighten them.

Beyond forecasting actions, team-mates might also try to convey information about their capabilities. Human interactional expectations are broadly governed by a common set of human functional and social capabilities, whereas humans are largely uninformed about a robot’s true capabilities. Therefore, robots will likely find themselves being judged according to the wrong standards. Although humans show patience for robots that fail under the right conditions, a robot that seldom works as expected will likely not remain in use, even if the failure is one of expectations rather than capabilities. Properly setting expectations allows human team-mates to avoid being disappointed by robots [26, 89, 92].

## 2.2 Framework

In this section, we describe a framework for implicit communication, modeled as a single-shot act.

### 2.2.1 Definitions

In coordinated activities, Clark [30] distinguishes among several related concepts. A *joint activity* engages a group of two or more agents in acting together toward a common goal. Examples include a marriage ceremony, a classroom lecture, and a football game. Within the context of a joint activity, participants perform *joint actions*, which continuously unfold over some period of time. A specialization is the *joint act*, which is a one-shot joint action. For example, in the joint activity of playing golf, yelling “fore!” is a joint act by which the speaker warns any listeners of a wayward flying ball (their avoidance response, in contrast, is an individual act, performed without consideration of how it will affect the group). The fact of an act being joint or individual is purely a matter of the mental state of the involved agent(s).

Participation in a joint action may be asymmetric – for example, speech is a joint action involving a speaker and listener. Note that the listener actively participates by comprehending and back-channeling (nodding, saying “uh-huh”, etc.). *Knowledge* comprises information believed by an agent to be true and is collected into a set of *facts*, each with associated confidence. Compulsory asymmetry occurs in a joint act or action when one individual, the *actor*, shares knowledge with one or more *observers*. Thus, an important aspect of the joint

action is to communicate information. Frequently, an actor embeds information implicitly in an otherwise purely functional action as part of the joint activity to perform *implicit communication*.

Any communicative action will be perceived by an observer with a certain level of *surprisal*, which is an encoding of how probable the observer believes the action to be given the context. As Hohwy [65] states, surprisal is a declining function of probability: the higher an observer’s surprisal, the more improbable the observer believes the action to be in the given context; the lower an observer’s surprisal, the more probable the observer believes the action to be in the given context. Common-sense knowledge and a shared understanding of the context allows an actor to gauge how surprising her action will be to an observer, which in turn shapes her choice of action depending on the information she would like to convey. In the remainder of this section, we show that greater surprisal corresponds with a more strongly-conveyed message (i.e. the action is more meaningful).

### 2.2.2 Foundations

The interplay of two sets is at the core of the framework.  $\mathcal{A}$  comprises all possible actions, whereas  $\mathcal{M}$  is composed of all possible facts about the world.

In the course of a joint activity, an agent performs a series of actions (including single-shot acts)  $a^1, a^2, \dots, a^n \in \mathcal{A}$ . Each action accomplishes both functional and communicative goals to varying degrees. Let  $A_f \subseteq \mathcal{A}$  be the set of (possibly many) different ways of accomplishing the functional goal of the action. Thus,  $A_f$  can be thought of as a subgoal of the shared goal of the joint activity.

An agent  $Q$  performs actions in a context  $M^Q$  comprising a set of facts  $m_1, m_2, \dots \in M^Q \subset \mathcal{M}$  that capture information about the individuals' knowledge. Only by leveraging this context can implicit communication occur.  $M^Q$  expresses  $Q$ 's beliefs about the world, including the state history of all agents in the joint activity, the observable scene, properties of objects within it, and common-sense knowledge. An individual fact  $m \in M^Q$  can have an associated confidence, thus allowing facts in  $M^Q$  to be added, removed, or changed following the observation of an action.

$M^Q$  is divided into several components. Knowledge that all participants in an interaction know they all share is public knowledge,  $M_{pub}$ , also called common ground. Other knowledge is not known to be public; agent  $Q$ 's private knowledge is denoted  $M_{priv}^Q$ .  $Q$  may be aware that a subset of the other agents know fact  $m \in M_{priv}^Q$ . It is even possible that every agent in a joint activity privately knows  $m$ . In both cases,  $m \notin M_{pub}$  unless all agents are all aware that  $m$  is shared by all.  $Q$ 's total knowledge  $M^Q$  is equal to  $M_{pub} \cup M_{priv}^Q$ .

Finally, the distribution  $P(a|M)$  describes the likelihood that a specific agent may next perform action  $a$  in the specific context  $M$ . Even if we restrict the scope of  $a$  to actions that accomplish a particular goal, there may be a set of possible actions ( $A_f \subseteq \mathcal{A}$ ) to choose among. Some of these actions will be preferred over others for reasons of efficiency, simplicity, or custom.

Posit that the following *common understandings* are agreed upon by all participants in the joint activity:

- the set of alternative actions  $A_f$  that would accomplish a functional goal
- the common ground context model  $M_{pub}$

- the action distribution  $P(a|M)$  (for plausible  $M \subset \mathcal{M}$ )

### 2.2.3 Implicit Communication Criteria

The goal of agent  $\mathcal{Q}$  is to perform an action  $\hat{a} \in A_f$  that satisfies functional goals while also communicating fact  $\hat{m} \in M_{priv}^{\mathcal{Q}}$ . However, it is not always possible to communicate an arbitrary fact  $\hat{m}$  implicitly, nor is it always possible to communicate implicitly via an action  $\hat{a}$ .

The key idea is for the actor  $\mathcal{Q}$  and observer  $\mathcal{R}$  to leverage the common understandings in order to achieve implicit communication.  $\mathcal{Q}$  selects an action that is surprising to  $\mathcal{R}$ , i.e. perceived by  $\mathcal{R}$  as improbable in the given context. However,  $\mathcal{R}$  does not treat the improbable  $\hat{a}$  as a fluke – rather, it triggers her to seek an explanation in the form of a previously-unknown fact  $\hat{m}$  that resolves the surprise. For  $\mathcal{R}$  to correctly interpret  $\mathcal{Q}$ 's intended meaning, we propose that action  $\hat{a}$  and fact  $\hat{m}$  must meet four *implicit communication criteria*:

1.  $\exists \hat{a}, a' \in A_f : \hat{a} \neq a'$
2.  $P(\hat{a}|M_{pub}) < P(a'|M_{pub}) - \varepsilon$
3.  $P(\hat{a}|M_{pub}) < P(\hat{a}|\hat{m}, M_{pub}) - \varepsilon$
4.  $\forall m \in \mathcal{M} \setminus M_{pub} \cup \{\hat{m}\} : P(m|\hat{a}, M_{pub}) < P(\hat{m}|\hat{a}, M_{pub}) - \varepsilon$

The  $\varepsilon$  term incorporates variation caused by personal preference and noise. The strength of a given implicit communication is measured as the largest possible  $\varepsilon$  satisfying the criteria above. Criteria 1–2 govern the actor's generation of implicit communication, whereas criteria 3–4 govern the observer's ability to correctly interpret the intended meaning. We speak of the fact  $\hat{m}$  as the *meaning*

of the action because it *explains*  $Q$ 's choice of action. We next provide additional insight into each of the criteria.

Criterion 1 requires that there must be at least two feasible, distinct actions that accomplish the functional goal, but preferably there are many more. An example of  $A_f$  that violates this criterion is placing a telephone call. Neglecting timing and caller ID, there is only one way to make somebody's telephone ring, leaving no room for a surprising choice of action.

Criterion 2 triggers the observer to search for an explanation of why the actor chose action  $\hat{a}$  over the more obvious choice,  $a'$ . This criterion fails in situations where there does not exist an action  $\hat{a}$  that is a priori substantially less probable than others. An example situation that violates it is one's first time visiting a clown convention, where normally-improbable actions are expected and hence unsurprising.

Criterion 3 requires that the fact  $\hat{m}$  will be easy for the observer to verify as an explanation of  $\hat{a}$ . That is,  $\hat{a}$  is unsurprising when  $\hat{m}$  is known. A well-known historical violation of this criterion was John Hinckley, Jr.'s attempted assassination of President Ronald Reagan in order to gain the favor of actress Jodie Foster – it is unclear how shooting the president is intended to convey infatuation.

Criterion 4 states that no other inferred meaning  $m$  is equally or more likely than the intended explanation  $\hat{m}$ . There are many example violations of this criterion in the form of hand gestures that take different meanings across cultures and geographies. One case in point is a gesture that variously signifies a Satanic association, infidelity, and a college football team in Texas. All three forms have

famously been used by politicians. Only by understanding each individual actor's  $M_{pub}$  at the time he made the gesture can we disambiguate among the three meanings.

## 2.2.4 Understanding and Generation

Suppose that an agent  $\mathcal{Q}$  hopes to convey some information,  $\hat{m} \in M_{priv}^{\mathcal{Q}}$ , to agent  $\mathcal{R}$  without resorting to disclosing it explicitly.  $\mathcal{Q}$  selects an action  $\hat{a}$  consistent with the implicit communication criteria, and  $\mathcal{R}$  determines  $\hat{a}$  to be an improbable action given what he knows.  $\mathcal{R}$ , believing  $\mathcal{Q}$  to be rational, hypothesizes that there must be some unknown factor  $\hat{m}$  that explains seeing  $\mathcal{Q}$  perform  $\hat{a}$ .  $\mathcal{R}$  thus searches over a set of plausible facts  $M$  and chooses  $\hat{m}$  to be the fact with the highest posterior probability given  $\hat{a}$  and  $M_{pub}$ . Maximizing this probability minimizes the surprisal that resulted from  $\mathcal{Q}$  performing  $\hat{a}$ , which in turn causes  $\hat{a}$  to become increasingly stronger evidence for  $\mathcal{R}$ 's hypothesis [65]. Hence, upon seeing  $\hat{a}$ ,  $\mathcal{R}$  proceeds to infer

$$\hat{m} \leftarrow \operatorname{argmax}_{m \in M} \mathrm{P}(m|\hat{a}, M_{pub}), \quad (2.1)$$

and thus  $\mathcal{R}$  concludes that  $\hat{m} \in M_{priv}^{\mathcal{Q}}$ , i.e.  $\mathcal{Q}$  believes  $\hat{m}$  to be true. Using Bayes' rule, we can re-express (2.1) as

$$\begin{aligned} \hat{m} &\leftarrow \operatorname{argmax}_{m \in M} \frac{\mathrm{P}(\hat{a}|m, M_{pub})\mathrm{P}(m|M_{pub})}{\mathrm{P}(\hat{a}|M_{pub})} \\ &= \operatorname{argmax}_{m \in M} \mathrm{P}(\hat{a}|m, M_{pub})\mathrm{P}(m|M_{pub}). \end{aligned}$$

Note that the prior  $\mathrm{P}(m|M_{pub})$  serves to prevent “conspiracy theories” that would otherwise result when noise gets mistakenly interpreted as signal. That is, the fact being communicated must have a reasonably likely prior probability. For

example, if Bob looks up at the night sky and sees a star twinkling, he is unlikely to attribute it to a UFO, given that the prior probability of discovering intelligent extraterrestrial life is small and that there is a more plausible explanation rooted in turbulence of the atmosphere.

Next, we turn to the generation problem. The structure of the generation problem is identical to understanding, except that we now search over actions instead of facts,

$$\hat{a} \leftarrow \operatorname{argmax}_{a \in A_f} P(\hat{m}|a, M_{pub}). \quad (2.2)$$

Applying Bayes' rule again, we can re-express (2.2) as

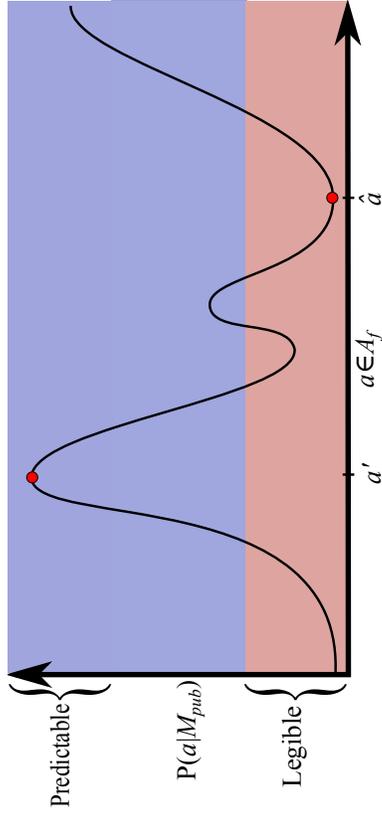
$$\begin{aligned} \hat{a} &\leftarrow \operatorname{argmax}_{a \in A_f} \frac{P(a|\hat{m}, M_{pub})P(\hat{m}|M_{pub})}{P(a|M_{pub})} \\ &= \operatorname{argmax}_{a \in A_f} \frac{P(a|\hat{m}, M_{pub})}{P(a|M_{pub})}. \end{aligned}$$

The resulting expression selects the action for which contributing  $\hat{m}$  to the common ground boosts  $P(a|M_{pub})$  by the greatest amount. See Figure 2.3 for an illustration.

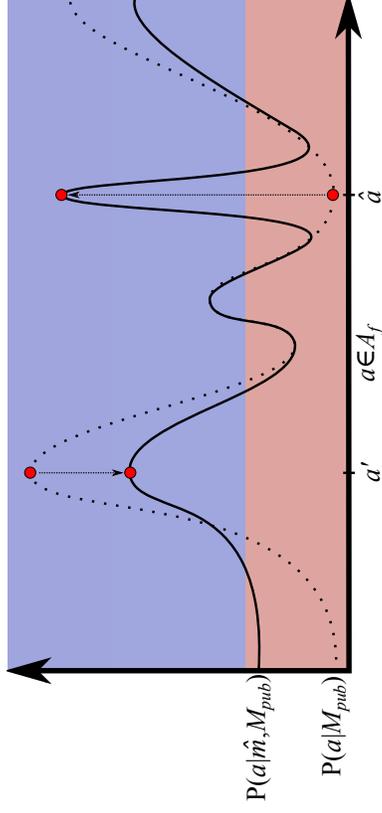
We expand on these ideas and provide examples in Section 2.4, but first we broaden our discussion to include implicit communication occurring over time and in service of joint goals.

## 2.3 Achieving Joint Goals

In a joint activity, rational agents interact with each other and make decisions towards achieving joint goals. These goals could range from completing a collaborative assembly task to smoothly avoiding each other while navigating in a



(a) A likely action such as  $a'$  is termed *predictable*, whereas we say that an unlikely action  $\hat{a}$  is *legible*. See Section 2.4.2.1 for a full discussion of predictability and legibility. Since  $\hat{a}$  is rarely observed in context  $M_{pub}$ , the observer infers that it probably was selected specifically to send a message.



(b) By performing legible action  $\hat{a}$ , an agent implicates the new fact  $\hat{m}$  because knowledge of that fact changes the shape of the distribution, causing  $\hat{a}$  to become a predictable action.

Figure 2.3: These plots illustrate the inference mechanism described in Section 2.2.4 and its effect on  $P(a|M_{pub})$ . Among the set of actions  $a \in A_f$  that accomplish a task, each can be assigned a likelihood of being observed in context. Actions with high likelihoods of  $P(a|M_{pub})$  are deemed *predictable*, low ones *legible*.

hallway. Relying only on implicit communication to achieve joint goals requires the establishment and reinforcement of trust. Implicit communication leverages trust to influence the observer’s belief and converge to a consensus that is beneficial for the accomplishment of a joint goal. In this section we state our model for trust and propose an index for monitoring its evolution in order to inform decision making.

### 2.3.1 Trust

Ordinarily, participants in a joint activity act *rationally* and cooperate to achieve shared goals [61]. This policy forbids deception and supports the assumption that the common understandings (Section 2.2.2) are shared by all participants. Given the great diversity of knowledge and experience among people, however, this assumption is perhaps too strong to apply universally.

In particular, during interactions with strangers, we may be unfamiliar with one another’s judgments regarding  $A_f$ ,  $M_{pub}$ , and  $P(a|M_{pub})$ . If we define trust as confidence in another agent’s future actions [14], then it is natural for one agent to restrict their trust of another based on the limits of common understandings among the individuals, even when all agents behave rationally.

Another obstacle to trust is discrepant beliefs about facts. We allow facts about the beliefs of others to enter  $M_{pub}$ . Thus, it can simultaneously be part of the common ground that  $\mathcal{G}$  believes  $m^{\mathcal{G}}$  and that  $\mathcal{H}$  believes  $m^{\mathcal{H}}$ , even if  $m^{\mathcal{G}}$  and  $m^{\mathcal{H}}$  conflict.  $\mathcal{G}$  and  $\mathcal{H}$  are then free to leverage either of these facts in the generation and understanding of implicit communications between them. Epistemic logic [44] provides tools for representing and analyzing such scenarios. Each

conflicting fact introduces additional uncertainty into the communication process because the observer must infer which fact the actor premised the communication upon. Thus, trust degrades with the number of discrepancies among beliefs within a joint activity. Beyond some limit, implicit communication becomes impossible.

### 2.3.2 Consensus

In a joint activity, agents take actions with functional effects (which contribute to reaching the joint goal) but also with communicative effects. One category of communication, conveying intentions, serves to convey a preference or desire regarding a joint strategy  $S$  for accomplishing the goal. The joint strategy can be thought of as the sequence of subgoals of the joint activity,  $A_f^0, A_f^1, \dots, A_f^n$ , and is drawn from the set of all possible strategies  $\mathcal{S}$ .

A consensus for each subgoal in the joint strategy may unfold gradually or abruptly during the course of the joint activity. As the agents act, the public knowledge  $M_{pub}$  is updated along with the agents' beliefs regarding the emerging strategy  $P(S|M_{pub})$ . Under the assumption of rationality, as formulated in our trust model (Section 2.3.1), a group of competent agents taking actions bearing implicit communication signals will be able to achieve consensus over the joint strategy  $S$ . This essentially means that  $P(S|M_{pub})$  (which we assume is shared by all agents) will converge to a distribution that clearly indicates the emerging joint strategy. The entropy of this distribution is a measure of that convergence.

### 2.3.3 Receptivity

In many joint activities, time and timing are critical attributes of an action. Timing itself often conveys meaning, which we therefore consider as an attribute of an action  $\hat{a} \in A_f$ . Another important aspect of timing is its role in choosing whether (and when) to implicitly communicate. Participants in a joint activity are not equally receptive at all times to certain forms of implicit communication, particularly with regard to consensus over the joint strategy.

When participants in a joint activity lack consensus about a joint strategy, they cannot coordinate effectively to achieve shared goals. Rational agents therefore strive to reach consensus as early in a joint activity as possible in order to maximize coordination efficiency. Consequently, the bulk of implicit communication for consensus should occur towards the beginning of the joint activity. As a joint strategy  $S^*$  emerges and consensus is reached, the agents might favor more predictable, less communicative actions, or they might utilize the implicit communication channel for other purposes. More generally, the implicit consensus formation aspect of joint actions may wax and wane according to the group need. Consequently, there arises the need for monitoring (1) the state of consensus  $P(S|M_{pub})$  and also (2) how receptive the group of agents is to the communicative signals being transmitted.

We formalize this monitoring process by introducing a *Receptivity* score, as

$$Receptivity = - \sum_{S \in \mathcal{S}} P(S|M_{pub}) \log(P(S|M_{pub})) \quad (2.3)$$

which is the information entropy of the distribution over joint strategies, given the common ground,  $P(S|M_{pub})$ . Recall that common ground includes the action history within a joint activity. Intuitively, receptivity measures the willing-

ness of individuals in a group to update their beliefs about the joint strategy, inversely proportionate with clarity. Since  $M_{pub}$  is sequentially updated over time, receptivity reflects the way the agents incorporate observed communicative signals into their own actions. The lower a receptivity score gets, the closer the agents are to a consensus over a joint strategy  $S^*$ . To avoid second-guessing a settled joint strategy, an observer suppresses strategy changes of a larger magnitude than the current receptivity level.

A consequence of a decline in receptivity is that agents can be less expressive when it drops, since other agents will likely ignore the inputs. In a scene with engaged competent agents, receptivity is expected to decrease rapidly, signifying a consensus in the joint activity. This decrease will influence the balance between the functional and communicative aspect of actions taken, shifting the focus of decision making towards the functional component. Beyond some threshold drop in receptivity, agents have become sufficiently certain about the consensus strategy  $S^*$  that they may even ignore their partners using *civil inattention* [81] to reinforce the previously agreed strategy. This behavior involves physically looking away, “so as to express that [one] does not constitute a target of special curiosity or design” [52]. At this point, only a major modification in the strategy will penetrate an agent’s civil inattention.

## 2.4 Case Studies

In lieu of generating new experimental results, which would apply to a single domain and communication modality, we present examples of how implicit communication has been modeled and enforced by several communities in var-

ious collaborative contexts and discuss how their frameworks align with our unifying framework for implicit communication.

## 2.4.1 Implicit Communication through Natural Language

Speech acts are among the richest functional actions in which to embed implicit communication.

### 2.4.1.1 Implicature

In this section, we give a brief background on *conversational implicature*. We seek to draw parallels between implicature and other methods of implicit communication of interest in robotics. Implicature comes from pragmatics, the linguistics subfield that studies the usage of language in context. Basic meaning that is expressed and understood by a speech act is achieved by *entailment* – that is, ideas that logically and unavoidably follow from the words chosen by a speaker.

With implicature, in contrast, the speaker *implicates* (i.e. implies or suggests) an idea without explicitly stating it. It is a frequent phenomenon in English, first described by Grice [56]. Consider this example from Lappin and Fox [91]:

Ann: Do you sell paste?

Bill: I sell rubber cement. (*â*)

*implicature: Bill does not sell paste. (*î*)*

A test for conversational implicature in particular is whether it is *cancelable* – that is, does there exist some phrase that, when appended to the sentence,

cancels the meaning of the implicature? From the above example, a phrase that cancels Bill's implicature is "I sell rubber cement, which is what you really need for your application." An implicature, once canceled, implicitly communicates nothing. The added phrase explains the initial phrase, thus increasing  $P(a|M_{pub})$  and violating implicit communication criterion 2.

When it comes to dialog, people have varied and complex motives for implicating meaning rather than entailing it, including politeness, sophistication, succinctness, and social group cohesion. A detailed consideration of these objectives is beyond the scope of this work.

Grice's *cooperative principle* states, "Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged" [56]. Indeed, the cooperative principle bears more than a passing similarity to the *pedestrian bargain* of Wolfinger [148], which entreats one to behave competently and also to trust others to behave competently. These principles are both forms of the rational actor assumption [61].

A vital component of conversational implicature is provided by the four Gricean Maxims, which describe speech that obeys the cooperative principle. The four maxims are

1. Maxim of Quantity: Make your contribution as informative as is required (but not more so).
2. Maxim of Quality: Make your contribution one that is true.
3. Maxim of Relation: Be relevant.
4. Maxim of Manner: Be perspicuous. Avoid obscurity or ambiguity; be brief

and orderly.

Other maxims have also been proposed, such as “Be polite.” Because adherence to the cooperative principle is assumed, utterances can be interpreted in light of these maxims. A speaker can therefore deliberately flout one of the maxims (an improbable action,  $\hat{a}$ ) in order to convey that he is employing implicature. Returning to the previous example, Ann must apply the following inference steps to conclude that Bill does not carry paste.

- (a) *Contextual premise*: it is mutual, public knowledge that Bill has complete knowledge of the items he sells.
- (b) *Contextual premise*: there is no contextual relationship linking sales of paste and rubber cement (inclusive or exclusive).
- (c) Assume Bill follows the cooperative principle and maxims.
- (d) By (a), Bill can fully resolve Ann’s question, and by (c), he will.
- (e) Only the propositions that Bill does or does not sell paste can completely resolve the question.
- (f) By (b), there is no way to infer from Bill’s answer the proposition that he does sell paste. The cooperative principle forbids obfuscation. Thus, Bill has flouted the maxim of relevance.
- (g) Therefore, we conclude that Bill does not sell paste.

Lines (d)–(g) comprise the narrowing down and resolution of the search for meaning in Equation (2.1).

Conversational implicature is absent when all the maxims are satisfied. One indicates the use of implicature by selecting an action to deliberately flout one of the maxims – in our example, Bill flouts the maxim of Relation.

Sometimes, two maxims conflict and cannot both be satisfied with a single utterance, in which case flouting one or the other maxim is forced. An example of the latter occurs in the following exchange:

Mark: Where is the cat?

Sue: The cat is in the hamper or under the bed. ( $\hat{a}$ )

*implicature: Sue does not know where the cat is.* ( $\hat{m}$ )

Because Sue does not know where the cat is, providing either location alone would violate the maxim of Quality. However, providing both locations conflicts with the maxim of Quantity because the cat is in at most one of the stated locations. Flouting the maxim of Quality would violate implicit communication criterion 2 because either location alone is plausible. Thus, Sue chooses to flout the maxim of Quantity in order to trigger Mark to search for an explanation.

#### 2.4.1.2 Inverse Semantics

Though more direct than conversational implicature, the simpler speech act of entailment is fundamentally described by the same mathematics. Knepper et al. [83] present the *inverse semantics* framework for robots generating natural language help requests. Like most robot speech systems, the framework strives for extremely literal communication. However, it faces a problem of finding pithy, unambiguous means of communicating its needs in an automated assembly scene cluttered with parts that lack unique names. Since words are complex and imperfect containers for meaning, the careful selection of clear language to achieve entailment follows the same rules of generation as described in Section 2.2.4.

The core of inverse semantics is a forward semantics mechanism for understanding natural language, the Generalized Grounding Graph ( $G^3$ ) [129]. This structure takes in natural language expressions  $\lambda$  as inputs and returns their meanings or groundings  $\gamma$  as outputs.

The inverse semantics framework inverts  $G^3$  to perform generation by searching over the space of possible English sentences, sorted from shortest to longest, and inputting each to  $G^3$ . Inverse semantics compares the output of  $G^3$  with the target grounding needed by the help request. The search halts with the first sentence that attains over a threshold confidence match between the two groundings. The expression given for generation,

$$\operatorname{argmax}_{\lambda} P(\gamma|\lambda, \phi, M), \quad (2.4)$$

strongly resembles our own framework’s Equation (2.2). Here,  $\phi$  is a correspondence variable used to indicate the semantic likelihood of a match between  $\lambda$  and  $\gamma$ . Like our model,  $M$  symbolizes the context model in which the meaning is interpreted.

## 2.4.2 Communicative Motion

Besides natural language usage, the robotics community has studied other types of actions. An especially expressive action class for implicit communication is motion.

### 2.4.2.1 Legibility

Let us consider again the joint assembly activity in which Steve and Cathy cooperate to build furniture. Many forms of communicative action arise. One class of actions studied recently by Dragan et al. [41] involves reaching motions. Among parts cluttering a table, Steve has to pick up a particular one. The shape of his reaching trajectory may or may not inform Cathy about Steve’s intent. A direct reaching motion is predictable (high probability  $P(a|M_{pub})$ ) and therefore not communicative. A curved trajectory, in contrast, helps Cathy to identify the target of Steve’s reach before he gets there.

In general, assume that an actor  $Q$  is aiming at reaching a goal  $G^Q$  from a set of goals  $\mathcal{G}$  in front of an observer  $\mathcal{R}$ . The agents share a model  $P(G|\xi)$  that probabilistically attributes a goal  $G \in \mathcal{G}$  to an observed trajectory  $\xi$ . The actor can leverage this knowledge to design his trajectory in a way that indicates his intended goal to the observer. Following the insights of Csibra and Gergely [33] regarding the tendency of humans to interpret observed actions as goal-directed (teleological reasoning), Dragan et al. [41] introduced the *Legibility* score to quantify the intent-expressiveness of a trajectory  $\xi$  with respect to a goal  $G^Q$ :

$$Legibility(\xi) = \frac{\int_0^T P(G^Q|\xi_{0 \rightarrow t})f(t)dt}{\int_0^T f(t)dt} \quad (2.5)$$

where  $T$  is the duration of the trajectory and  $f(t)$  is a function that weights partial trajectories  $\xi_{0 \rightarrow t}$  higher in the beginning and lower later. It should be noted that  $f(t)$  is a proxy for the role of the observer in reducing her receptivity (see Section 2.3.3) as  $Q$ ’s intended goal  $G^Q$  becomes more certain to her. The model  $P(G|\xi)$  scores goals higher if they can be achieved efficiently (with a low energy trajectory  $\xi$ ) and scores goals lower if they require higher energy.

The legibility score is essentially a weighted sum of the probabilities that the observers assign to the actor's intended goal  $G^Q$  throughout the whole trajectory  $\xi$ . Trajectories of higher legibility tend to be more curved towards the intended goal  $G^Q$ , biasing the observers towards predicting the actor's actual goal, while biasing them against predicting other goals. Note that a more curved trajectory is less probable out of context due to the extra energy it expends. As a result, it might be perceived as surprising. This surprise would trigger a search for an explanation, which, in the perceived context, would lead to the conclusion that the actor  $Q$  is aiming at reaching the goal  $G^Q$ .

#### 2.4.2.2 Dynamic Legibility

Consider now the case of a dynamic environment, where the agents are not explicitly collaborating but since the decisions they make are coupled, it is beneficial for everyone to mutually agree on a joint strategy. Assuming again no explicit communication, the only way agents are able to agree on a strategy is to encode their understanding and preferences into their actions.

Social navigation constitutes a representative example of this class of scenarios. Although humans might not often realize that navigation in crowded environments is a collaborative activity, according to sociology studies [148], it is established on implicit cooperation. Pedestrians follow and reinforce the *pedestrian bargain*, a social convention comprising two foundational rules: (1) "pedestrians must behave like competent pedestrians" and (2) "pedestrians must trust that co-present others behave like competent pedestrians". Since the pedestrian bargain serves as a cooperative principle for social navigation, we may formulate a set of maxims for motion that echo the Gricean Maxims of conversational

implicature,

1. Maxim of Efficiency: Be parsimonious.
2. Maxim of Motion: Do not collide with objects or obstruct another agent's motion.
3. Maxim of Manner: Be perspicuous and orderly.

These maxims readily come into conflict where multiple agents are present. Much as in the case of implicature, the actor will choose to deliberately flout one of the maxims – typically the maxim of Efficiency – in order to obey the cooperative principle. It is only by considering the collision-avoidance context that an observer is able to appreciate that by taking an exaggerated trajectory such as  $a_3$  in Figure 2.4, the global welfare is improved, as measured by increased energy efficiency and decreased uncertainty.

Enforcing the pedestrian bargain leads to a consensus over a mutually beneficial joint strategy that allows everyone to comfortably reach their destinations. The agents continuously monitor the progress toward consensus and adjust their decision-making accordingly. Once consensus appears to have been reached, receptivity drops to zero as pedestrians initiate *civil inattention* [52, 81]. Following this mode switch, agents look away from one another as a signal that they have stopped actively avoiding each other and will instead follow their previous planned collision-free path.

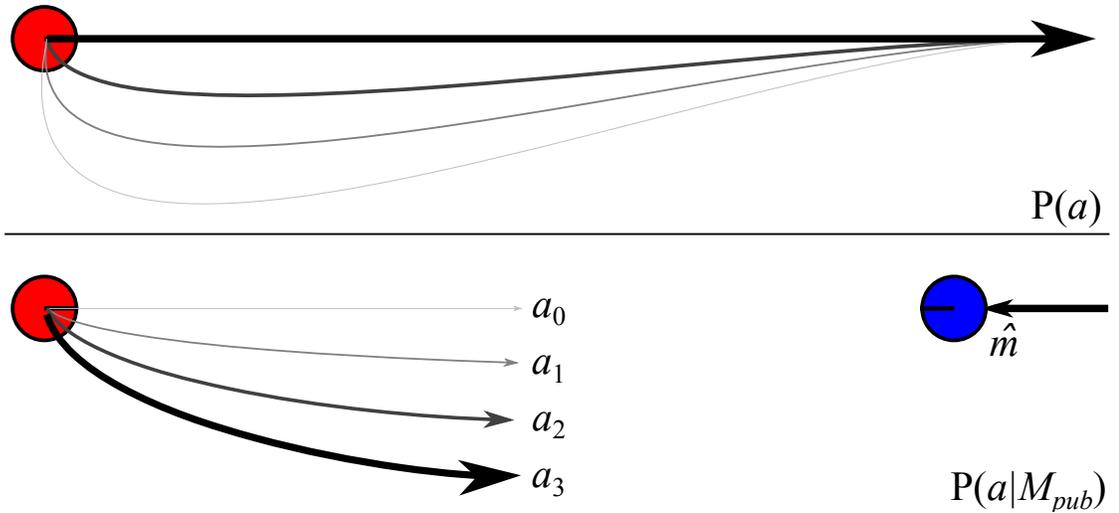


Figure 2.4: The red, navigating agent (human or robot) selects an action  $\hat{a}$ . Out of context (top), the red agent (human or robot) is not avoiding an obstacle, and so the probability of expending needless extra energy is low. In the case of an oncoming blue agent ( $\hat{m}$ ), the likelihood of the oblivious action  $P(a_0|M_{pub})$  is low due to social norms, despite being low energy. Conversely, the normally-improbable act of spending extra energy becomes probable in this context. An observer who sees only the red agent's motion can infer  $\hat{m}$  from observing  $a_3$ .

## 2.5 Other Examples

Teams exchange implicit information in cooperative games when the rules forbid free exchange of information. For example, the bidding conventions of contract bridge allow partners to exchange information about the respective strengths of their hands and arrive at an appropriate contract.

Finally, among married couples, this type of implicit communication eases over time across all modalities (speech, gesture, gaze, etc.) because spouses develop extremely sensitive models of  $P(a|M_{pub})$ , due to familiarity. Remarkably sophisticated notions can be conveyed between spouses by careful action selection in almost any context. We have considerable work remaining before robots can achieve a similar level of understanding of people.

### 2.5.1 Tact

Implicit communication is also the primary tool of tactful communication, as it alleviates the risk of awkwardness due to misunderstandings about what facts the observer already knows. Reflecting on the implicit communication criteria given in Section 2.2.3, an attempted implicit communication of a fact that the observer already knows does not even seem like implicit communication – it would come across as a predictable, functional action. In this case, criterion 3 is clearly violated because  $\hat{m} \in M_{pub}$ , and criterion 2 is probably also violated because  $\hat{a}$  would seem likely.

To offer a concrete example of how speakers leverage implicit communication to achieve tact, consider a married couple discussing dinner plans:

Jack: Remember, my friend Irving is coming for dinner.

*implicatures: Irving is vegetarian; Irving needs to be served a vegetarian meal.*

Kate: Let's make my mother's lasagna recipe.

*implicatures: Kate knows that Irving is vegetarian; Kate's mother's lasagna recipe is vegetarian; the recipe satisfies Irving's need for a vegetarian meal.*

Observe that this exchange can be read at two levels. If both parties are oblivious to the implicature because the sentences are judged predictable, then it is a simple, matter-of-fact dialog.

The statements can also be read as implicature. In both cases, the implicated statements are things that the listener should have already known. Only in the context of the couple's normal conversation can we judge how unusual it is for

Jack to remind Kate about a guest (a fact she may be unlikely to forget), or for Kate to make her mother’s lasagna recipe.

Only if these events are atypical can they truly be regarded as implicit communication. However, they also serve a tactful reminder function, in case Kate forgot about the guest or Jack forgot that Kate’s mother’s lasagna is vegetarian. A failing memory may therefore cause an action to be judged as unusual, in which case the reminder acts as an implicature. Thus, a related virtue of implicit communication is that it allows the observer to maintain the pretense of having already known a fact that they forgot.

## 2.6 Practical Implementation

Inference, both generation and understanding, is implemented as a search over actions and facts, respectively. Techniques are needed to streamline both search problems, due to the intractability of the literal brute force search implied by  $\text{argmax}$  in (2.1)–(2.2). Existing implementations of instances of implicit communication employ AI search-pruning techniques [83, 144] or restrict the action space  $A_f$  in order to narrow the set of options under consideration [41, 107]. In practical terms, the set of feasible actions  $A_f$  is typically hard-coded for a domain, raising the possibility that it mismatches with some human’s expectation. Two people may similarly encounter a mismatch in expectation about  $A_f$ . Interestingly, the machinery described in this work could be used by a robot to infer that an observed human action is intended to accomplish a (surprising) functional goal by leveraging the context, leading to extension of  $A_f$ .

Another challenge is to build  $M_{pub}$ , the common ground model among agents.

A complete model is often both unnecessary (since many facts in the agents' shared knowledge are irrelevant for the joint activity at hand) and infeasible (since the task of modeling the full common ground presents a high cognitive burden). As a result,  $M_{pub}$  need only consist of the facts that are pertinent to the success of the joint activity. For example, in the social navigation of Mavrogiannis and Knepper [98],  $M_{pub}$  might contain an updated belief regarding the destinations and intentions of observed agents.  $M_{pub}$  is therefore instantiated as the mutual understanding that the agents involved intend to participate in the joint activity along with shared knowledge about the kinds of actions that agents will likely take to contribute to the activity [22].

For humans,  $M_{pub}$  does not necessarily include all task-relevant facts at the start of the activity. It is frequently less costly to repair a misunderstanding that results from not sharing a piece of information than to expend the effort required to ground that piece of information through the principle of least collaborative effort [31, 116].  $M_{pub}$  is then updated interactively throughout the course of the joint activity, either when new information about the intents of the agents becomes publicly available or when the agents issue a repair that helps align their own mental models of the situation (and in doing so adds to the common ground) [93]. Machine-interpretable ontologies using tools like RDF and OWL address the general problem of managing and searching  $M_{pub}$ , as exemplified by the KnowRob project of Tenorth and Beetz [130].

Finally, the distribution  $P(a|M)$  is generally best modeled through machine learning. The particular context in which one takes an action affects the probabilities of observing various possible actions, often in complex ways. For example, Knepper et al. [83] employ Tellex's generalized grounding graph ( $G^3$ ) [129].

Based on a conditional random field,  $G^3$  employs a set of *correspondence variables* to valuate the correspondence probability of a given language phrase and grounding concept. These learned relationships capture concepts including objects, actions, and spatial relations.

## 2.7 Discussion

Conversational implicature and legibility, though originating in different domains, are connected by techniques of encoding and decoding meaning using teleological inference [33]. These methods rely heavily on common ground to provide clues about when a message is encoded on an action and what information the message contains. The inference process can be quite complex in real-life situations. Particularly in the case of implicature, many rules must be brought to bear in order to correctly interpret what is being implicated. Several authors [54, 144] show promising early results in modeling a simple form of implicature and performing inference by model inversion.

### 2.7.1 A Call to Action

In the coming years, modeling of implied meaning, including through implicature and legible motion, will become an increasing focus within robotics – not least because humans already use these forms of implicit communication on robots today. Humans are also already interpreting robots’ actions through the lens of implicit communication. Since few robots are cognizant of the implicit meaning of their actions, today’s robots send random signals to humans. By

and large, humans are unable to interpret robot actions in the purely functional manner that they are intended. Thus, the robotics research community must find techniques to efficiently generate and understand implicit communication.

This direction will drive the need for improved modeling of common ground. A major hurdle to performing these inferences on robots in real-world situations is salience; today, the robot must perform a fairly undirected, brute-force search in order to discover which elements of the context are applicable. Humans, in contrast, seem to learn filters and partially pre-compute functions to expedite real-time inference in ambiguous situations. These processes are not yet understood in humans, but they will need to be deployed on robots in order to promote responsive behavior and avoid major misunderstandings.

## **Part II**

# **Planning by Reasoning about Multi-Agent Navigation Strategies**

## CHAPTER 3

### MULTI-AGENT PATH TOPOLOGY IN SUPPORT OF SOCIALLY COMPETENT NAVIGATION PLANNING

This chapter incorporates the information-theoretic point of view exposed in Part I into the design of a motion planner for intent-expressive robot navigation. Specifically, we present a navigation planning framework for dynamic, multi-agent environments, where no explicit communication takes place among agents. Inspired by the collaborative nature of human navigation, our approach encodes the concept of coordination into an agent’s decision making through an inference mechanism about collaborative strategies of collision avoidance. Each such strategy represents a distinct avoidance protocol, prescribing a distinct class of navigation behaviors to agents. We model such classes as equivalence classes of multi-agent path topology, using the formalism of topological braids. This formalism may naturally encode any arbitrarily complex, spatiotemporal, multi-agent behavior, in any environment with any number of agents into a compact representation of dual algebraic and geometric nature. This enables us to construct a probabilistic inference mechanism that predicts the collective strategy of avoidance among multiple agents, based on observation of agents’ past behaviors. We incorporate this mechanism into an online planner that enables an agent to understand a multi-agent scene and determine an action that not only contributes progress towards its destination but also reduction of the uncertainty of other agents regarding the agent’s role in the emerging strategy of avoidance. This is achieved by picking actions that compromise between energy efficiency and compliance with everyone’s inferred avoidance intentions. We evaluate our approach by comparing against a greedy baseline that only maximizes individual efficiency. Simulation results of statistical significance demon-

strate that our planner results in a faster uncertainty decrease that facilitates the decision making process of co-present agents. The algorithm’s performance highlights the importance of topological reasoning in decentralized, multi-agent planning and appears promising for real-world applications in crowded human environments.

### 3.1 Foundations

Consider a set of  $n \geq 2$  agents  $N = \{1, \dots, n\}$  navigating a workspace  $\mathcal{Q} \subset \mathbb{R}^2$ . Denote by  $q_i \in \mathcal{Q}$  the configuration of agent  $i \in N$ . Agent  $i$  starts from an initial configuration  $q_i^s \in \mathcal{Q}$  at time  $t = 0$  and reaches a final configuration  $q_i^d$  at time  $t = T_i$ . The final configuration  $q_i^d$  corresponds to a landmark  $d_i$  from a set of landmarks  $D \subset \mathcal{Q}$  (we assume that  $d_i \neq d_j$  for any two agents  $i, j \in N$ ). The path agent  $i$  follows to reach its destination is a function  $\xi_i : [0, T_i] \rightarrow \mathcal{Q}$ .

The agents do not explicitly exchange any kind of information with each other but are assumed to be acting rationally, which in our context means that (1) they always aim at making progress towards their destinations and (2) they have no motive for acting adversarially against other agents (e.g. blocking their paths or colliding with them). The notion of rationality is in line with the concept of *competence* as described by Wolfinger [148] in his definition of the *Pedestrian Bargain* and also with the concept of *teleological reasoning* that appears to be foundational for human inference [33].

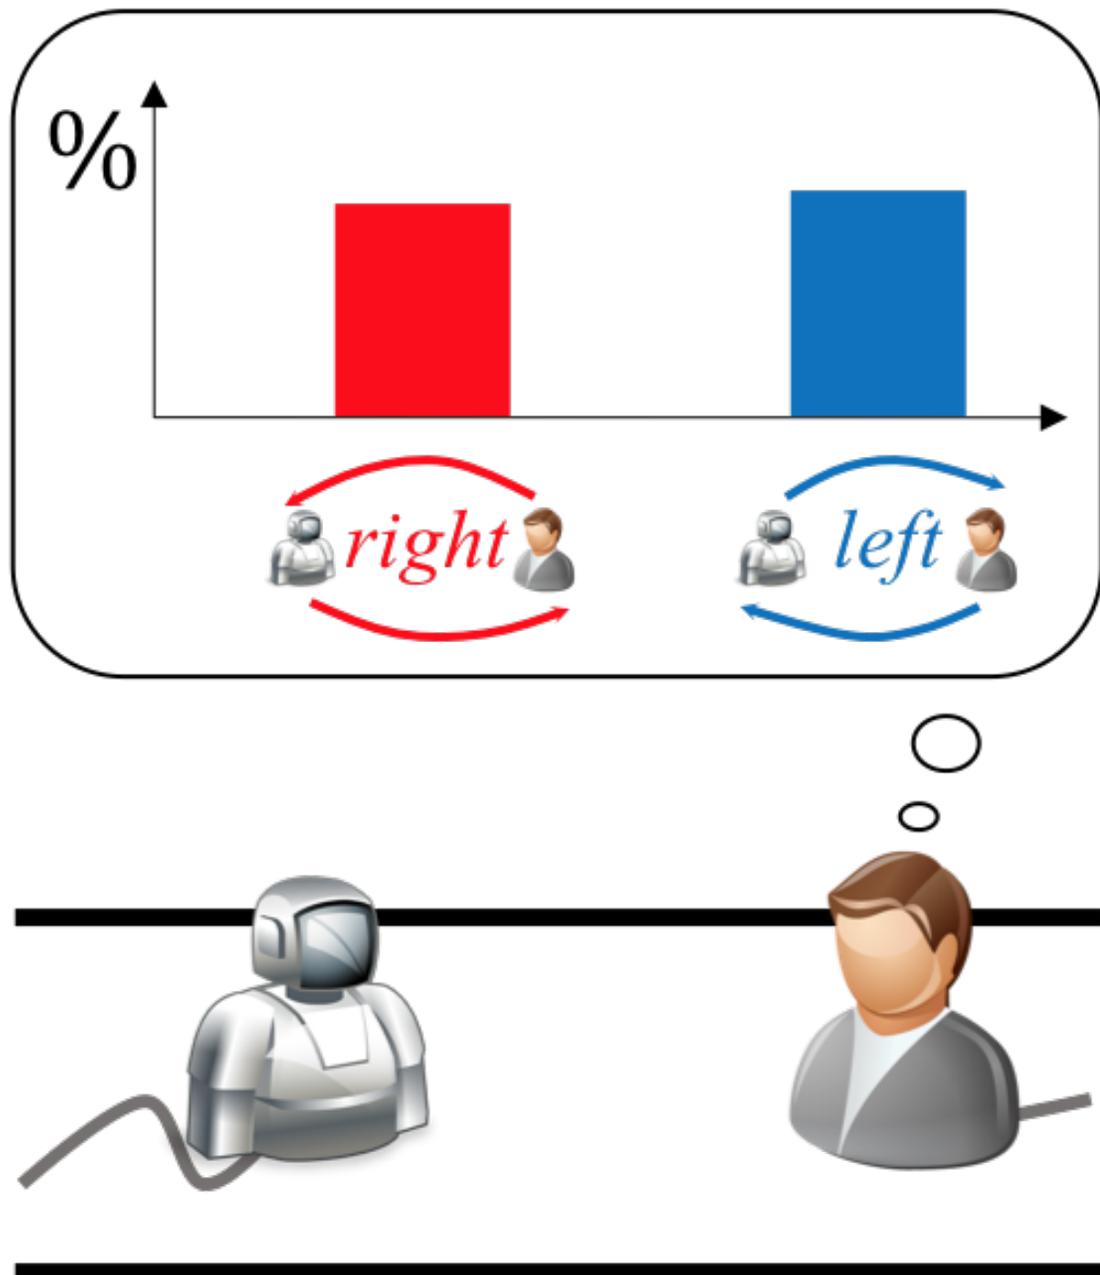


Figure 3.1: A human and a robot are navigating towards opposing directions of a hallway. To avoid collision, they need to agree on an avoidance protocol (passing from the right or left hand side of each other). The jerky behavior of the robot so far and the smooth but non-committal –with respect to a passing side– path of the human agent yield a high-entropy belief distribution over an emerging avoidance protocol from the perspective of both agents. The goal of our planner is to generate a sequence of highly-informative actions that will rapidly reduce the entropy and break a potential livelock or deadlock situation.

### 3.1.1 Game-Theoretic Setup

Inspired by Wolfinger’s observations on the cooperative nature of human navigation, we approach the problem of robotic navigation in multi-agent, dynamic environments as a finitely repeated coordination game of imperfect information and perfect recall. The game is repeated a finite number of rounds  $K$ , which is unknown a priori and corresponds to the round at which the slowest agent reached its destination. At each round  $k \in \{1, \dots, K\}$ , each agent  $i$  decides on an action  $a_i^k$  from a set of available actions (actions that could potentially lead to collisions and actions that violate the agent’s dynamics are excluded)  $\mathcal{A}_i^k$  by minimizing a cost function  $u_i$ . The agents are simultaneously selecting their actions and therefore they have no access to other agents’ plans (imperfect information); we assume however that they maintain a history of all previous rounds (perfect recall). The result of all agents’ decision making at round  $k$  is the strategy profile  $A^k = \{a_1^k, \dots, a_n^k\}$ . The sequence of strategy profiles of all rounds, from the beginning to the end of the game,  $A_1 \dots A_K$ , forms a global *joint strategy*  $\tau$  that the agents engaged in to avoid each other, while making progress towards their destinations.

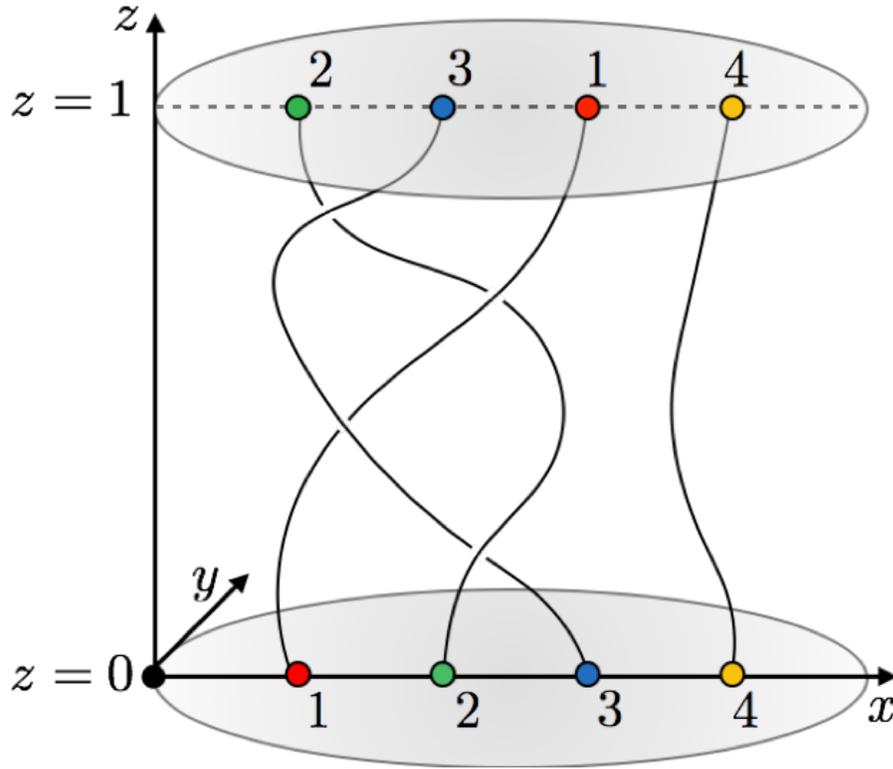
Although the agents do not explicitly coordinate with each other to decide on a joint strategy, their strategy profiles at every round gradually reinforce and contribute to one. Imbuing artificial agents with an understanding of the collective dynamics of a multi-agent scene may allow them to make informed and *socially competent* decisions that contribute to the avoidance of undesired situations, such as hindering others’ paths, deadlocks and livelocks. For this reason, besides personal efficiency, it is important that agents’ cost functions incorporate a model of multi-agent efficiency, reflecting the *social welfare* of the

system of agents.

In this chapter, we develop a topological model of joint strategies, employing the formalism of braids [20], which we use to develop a human-inspired inference mechanism, supported by studies on human action interpretation. Our mechanism provides a principled prediction over the scene evolution that allows agents to take into consideration the effect of their decision making on any observers.

### 3.1.2 A Topological Model of Joint Strategies

Let us collect the state of the system of all agents in a tuple  $Q = (q_1, \dots, q_n) \in \mathcal{Q}^n$ . The system state evolves from a starting configuration  $Q^s = (q_1^s, \dots, q_n^s)$  to a final configuration  $Q^d = (q_1^d, \dots, q_n^d)$ , by following a path  $\Xi : [0, T] \rightarrow \mathcal{Q}^n$ , from the space of system paths  $\mathcal{Z}$ , starting from  $Q^s$  and ending at  $Q^d$ . The system path is a function  $\Xi : [0, T] \rightarrow \mathcal{Q}^n \setminus \Delta$ , where  $\Delta = \{Q = (q_1, q_2, \dots, q_n) \in \mathcal{Q}^n : q_i = q_j \text{ for some } i \neq j \in N\}$  is the set of all system states with agents in collision and  $T = \max_{i \in N} T_i$  (it is assumed that agents remain at their destinations until everyone reaches their own). Naturally  $\Delta$  partitions the space of system paths  $\mathcal{Z}$  into a set of classes of homotopically equivalent system paths. Each such class has distinct topological properties which indicate a distinct joint strategy that the agents followed to reach their destinations. To enumerate such classes of joint strategies but also to characterize topologically the collective behavior of a system of agents, we develop a model of joint strategies using the concept of braids [20]. In the following paragraphs, we provide a primer on braids, establish a correspondence between braids and collective navigation behaviors



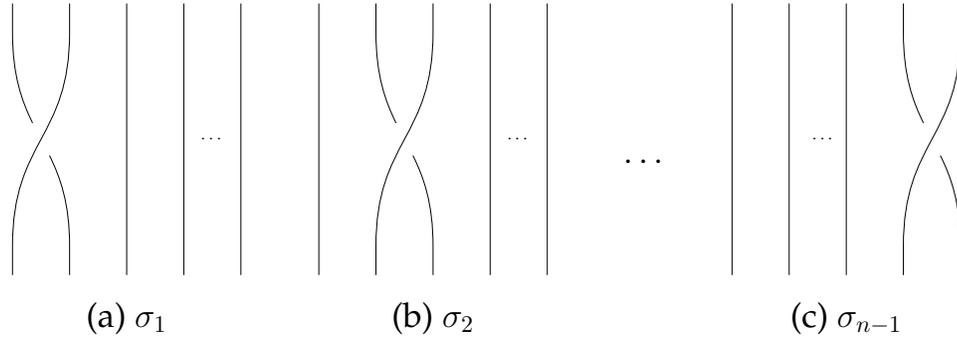
**Figure 3.2:** A set of 4 particles are initially ( $z = 0$ ) arranged along the  $x$ -axis, on the points 1, 2, 3 and 4. Through a sequence of rearrangements, the particles finally ( $z = 1$ ) reach a final arrangement on the points 2, 3, 1 and 4 respectively. The pattern of their trajectories is a *geometric braid*.

and define a topological model of joint navigation strategies.

### 3.1.2.1 Background on Braids

Braids are topological objects with algebraic and geometric presentations. We first introduce them as geometrical entities, following a presentation based on Artin [10] and continue with a discussion of their algebraic presentation and their group formation.

Denote by  $x, y, z$  the cartesian coordinates of a Euclidean space  $\mathbb{R}^2 \times I$ . A braid string is a curve  $X(z) : I \rightarrow \mathbb{R}^2$  that increases monotonically in  $z$ , i.e., has



**Figure 3.3: The generators of the Braid Group  $B_n$ .**

exactly one point of intersection  $X(z) = (x, y)$  with each plane  $z \in I$ . A braid on  $n$ -strings or  $n$ -braid (see Figure 3.2) is a set of  $n$  strings  $X_i(z), i \in N = \{1, \dots, n\}$  for which:

1.  $X_i(z) \neq X_j(z), \text{ for } i \neq j \forall z \in \mathbb{R}$
2.  $X(0) = (i, 0)$  and  $X(1) = (p(i), 0)$ ,

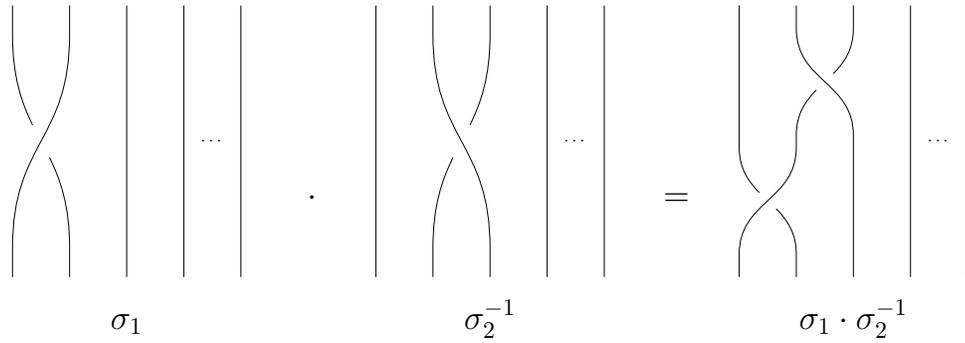
where  $p(i)$  is the image of an element  $i \in N$ , through a permutation  $p : N \rightarrow N$  from the set of permutations of  $N$ ,  $Perm(N)$ , defined as:

$$p = \begin{pmatrix} 1 & 2 & \dots & n \\ p(1) & p(2) & \dots & p(n) \end{pmatrix}. \quad (3.1)$$

This geometric representation of a braid is commonly referred to as a *geometric braid*. More formally, a geometric braid is often represented with a *braid diagram*, a projection of the braid onto the plane  $\mathbb{R} \times 0 \times I$  (see e.g. Figure 3.3).

The set of all braids on  $n$  strings, along with the composition operation, form a group  $B_n$ . The group may be generated from a set of  $n - 1$  elementary braids  $\sigma_1, \sigma_2, \dots, \sigma_{n-1}$  (see Figure 3.3), called the generators of  $B_n$ , that satisfy the following *relations*:

$$\sigma_j \sigma_k = \sigma_k \sigma_j, \quad |j - k| > 1, \quad (3.2)$$



**Figure 3.4:** Example of the Composition operation  $\sigma_1 \cdot \sigma_2^{-1}$  for  $\sigma_1, \sigma_2^{-1} \in B_n$ .

$$\sigma_j \sigma_k \sigma_j = \sigma_k \sigma_j \sigma_k, \quad |j - k| = 1. \quad (3.3)$$

A generator  $\sigma_i$ ,  $i \in \{1, 2, \dots, n - 1\}$  can be described as the crossing pattern that emerges upon exchanging the  $i$ th string (counted from left to right) with the  $(i + 1)$ th string, such that the initially left string passes *over* the initially right one, whereas the inverse element,  $\sigma_i^{-1}$ , implements the same string exchange, with the difference that the left string passes *under*<sup>1</sup> the right (see Figure 3.4). An identity element,  $e$ , is a braid with no string exchanges.

Two braids  $b_1, b_2 \in B_n$  may be composed through the composition operation  $(\cdot)$ , which is algebraically denoted as a product  $b_1 \cdot b_2$ . Geometrically, this composition results in the pattern that emerges upon attaching the lower endpoints of  $b_2$  to the upper endpoints of  $b_1$  and shrinking each braid by a factor of 2, along the  $z$  axis (see Figure 3.3, Figure 3.4). Any braid can be written as a product of generators and generator inverses. This representation is commonly referred to as an *algebraic braid* or a *braid word* (Fig. 3.4).

<sup>1</sup>This specific convention is not particularly important, as long as one is consistent. There are works that use the inverse convention when defining the positive and negative generator exponents. Our selection facilitates the exposition of further concepts in the remainder of the chapter.

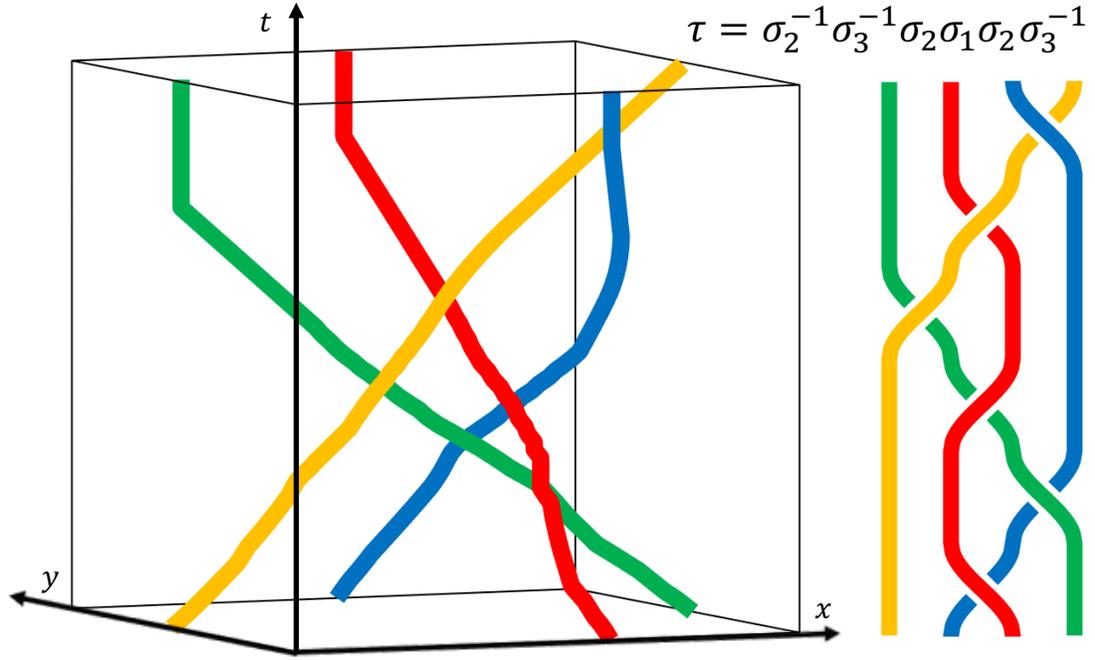


Figure 3.5: A space-time representation of a system path in a workspace with four agents (left) along with its corresponding braid diagram (right) and braid word (top right), defined with respect to the path's  $x$ -projection. The visualization of the braid diagram and the extraction of the braid word was done using BraidLab [132].

### 3.1.2.2 Abstracting Joint Strategies Using Braids

Denote by  $f_x : \mathcal{Q}^n \rightarrow Perm(N)$  a function that takes as input the system state  $Q \in \mathcal{Q}^n$  and outputs a permutation  $p \in Perm(N)$  corresponding to the arrangement of all agents in order of increasing  $x$ -coordinates. As the agents move towards their destinations, they employ navigation strategies – maneuvers to avoid collisions. These contribute to a system path  $\Xi$  which corresponds to a path of permutations  $\pi : [0, T] \rightarrow Perm(N)$  that may be extracted by evaluating  $f_x$  throughout the whole path  $\Xi$ . This path can be represented by a sequence of permutations of minimal length  $\pi^* = (p_0, \dots, p_K)$ , i.e.,  $p_{j-1} \neq p_j, \forall j = \{1, \dots, K\}$

and consecutive waypoints are adjacent transpositions<sup>2</sup>, i.e., permutations that differ by exactly one swap of adjacent elements. Due to continuity, a transition from the  $(j - 1)$ th permutation,  $p_{j-1}$ , to the  $j$ th permutation,  $p_j$ , implies the occurrence of an *event*  $\tau_j$ , which may be described as *the intersection of the  $x$ -projections of the paths of two agents that were adjacent in the permutation  $p_{j-1}$* . The event  $\tau_j$  may be represented as an elementary braid  $\tau_j \in \sigma_i^{\pm 1}, i \in \{1, \dots, n - 1\}$ , where  $i$  corresponds to the index of the leftmost swapping agent in permutation  $p_{j-1}$ . Therefore the whole execution from  $t = 0$  to  $t = T$  may be abstracted into the braid that corresponds to the temporal sequence of events:

$$\tau = \tau_1 \tau_2 \dots \tau_K \in B_n. \quad (3.4)$$

This braid word not only constitutes a topological characterization of the system path (see Figure 3.5 for an example of characterizing a system path as a braid) but it also represents a topological class of system paths that are homotopy-equivalent with the system path in consideration. In the remainder of this chapter, we will be referring to the sequence  $\tau$  as the *joint strategy* or the *entanglement* of the system path. Essentially, we model the space of joint strategies  $\mathcal{T}$  as the braid group, i.e.,  $\mathcal{T} := B_n$ .

**Remark 1.** *In our model, a braid constitutes a two-dimensional abstraction of a three-dimensional pattern of trajectories. Depending on the selection of the projection line, a different braid emerges. Although a change of projection line only changes the braid by conjugation [131], in practice, this implies that a set of non-communicating agents might encode the same joint strategy with a different symbolic representations (braids). However, this does not affect the convergence to a consensus on a mutually acceptable*

---

<sup>2</sup>A transposition can be described as a permutation involving exactly one swap of a pair of elements. An adjacent transposition is a transposition involving an exchange of two adjacent elements. An adjacent transposition implementing an exchange of the elements with order  $j$  and  $j + 1$ , with  $1 \leq j < n - 1$  in a list of  $n$  elements, is commonly denoted as  $\beta_j = [j \ j + 1]$ .

*joint strategy among agents; despite their different representations, they still take actions that contribute towards the same outcome, as will be shown in our simulation results. Therefore, the selection of the projection plane for a planning agent is not important, as long as its action selection process is consistent with it.*

### **3.2 Inference of Collective Behaviors**

An individual agent has no sole control over a specific joint strategy. The joint strategy is an emergent behavior, resulting from the superposition of the individual strategies of all agents. In fact, since agents are not explicitly communicating or coordinating, they cannot have a priori knowledge or a precise sense of the actual joint strategy they are about to follow. However, understanding the dynamics of collective behavior may allow agents to adopt individual navigation strategies that allow others to infer their intentions more clearly, thus facilitating everyone's decision making by reducing uncertainty fast. Judging from our everyday life experience, we may argue that this is the case with humans as well. When humans encounter others in a hallway, they do not exactly know the specific joint strategy they will be following. However, they realize that their decisions are coupled with the decisions of others and are able to reach a consensus regarding an avoidance protocol that is comfortable for everyone.

In this section, we present a probabilistic intention inference mechanism that connects an observed system path with a future system path topology, designed according to the insights of psychology studies on human action interpretation. This design is motivated by our goal of employing our framework on an autonomous social robot that will be navigating in a safe and socially competent

fashion around human pedestrians.

### 3.2.1 Teleological Reasoning in Multi-Agent Navigation

Csibra and Gergely [32, 33] argued that the mechanisms of human action interpretation are *teleological* in nature, i.e., humans tend to interpret observed actions as goal-directed in a given context. Following their insights, we design an inference mechanism of the form  $P(\tau|\Xi_t, M_t)$ , corresponding to a belief over an emerging joint strategy  $\tau \in \mathcal{T}$  given a partial system path  $\Xi_t$  and the state of the context at time  $t$ . The joint strategy represents a collective *goal*, whereas the system path plays the role of the *action*. By context, we refer to publicly available information, such as a model of the static environment (e.g. a map, obstacles, points of interest etc) but also information extracted through processing, e.g. by employing secondary inference mechanisms regarding group formations, identification of reactive agents etc.

### 3.2.2 Inferring Joint Strategies from Context

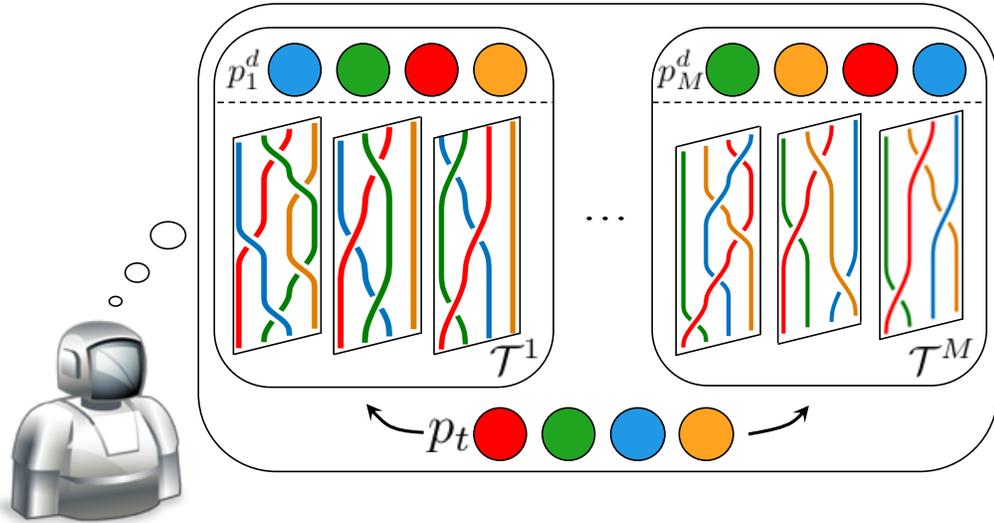
From (3.4), the belief  $P(\tau|\Xi_t, M_t)$  may be expanded as:

$$P(\tau|\Xi_t, M_t) = P(\tau_1, \dots, \tau_K|\Xi_t, M_t), \quad (3.5)$$

which, by applying the chain rule, may be factored as:

$$P(\tau|\Xi_t, M_t) = \prod_{k=1}^K P(\tau_k | \bigcap_{j=1}^{k-1} \tau_j, \Xi_t, M_t). \quad (3.6)$$

This belief quantifies the likelihood of a sequence of events  $\tau_1, \dots, \tau_K$  given observation of agents' past behaviors and the context. Essentially, this corresponds



**Figure 3.6:** Schematic representation of the inference mechanism from the perspective of a robot, navigating in a workspace with 3 other agents. From the perspective of the robot, the system state at time  $t$  corresponds to a permutation  $p_t$  (derived upon projecting on the  $x$ -axis of the robot's body frame), represented graphically with the color permutation at the bottom (the robot order in the permutation is denoted with red color). A set of  $M$  final permutations, taking the robot to its destination are considered and a set of 3 compatible joint strategies-braids are planned for each final permutation. The robot reasons over the set corresponding to the union of all sets of joint strategies  $\mathcal{T} = \bigcup \mathcal{T}^m, m = \{1, \dots, M\}$ .

to predicting the minimal sequence of permutations  $\pi^*$  but also the quality of the physical transitions between consecutive permutation waypoints (passing from the right/left hand side).

A joint strategy describes the avoidance protocol that the agents followed to avoid each other throughout the scene evolution, while navigating from  $Q^s$  to  $Q^d$ . These system path endpoints are not incorporated in the definition of the strategy as geometric entities but rather as the permutations  $p^s = f_x(Q^s)$ ,  $p^d = f_x(Q^d)$ . This design decision reflects the observation that an agent navigating a multi-agent environment does not need to know the precise intended destinations of others to avoid collisions successfully; they just need to understand

their passing preferences/intentions. However, the geometric arrangements of agents' intended final configurations greatly influence the convergence to a joint strategy  $\tau$ . In particular, given the initial permutation  $p^s$ , only a subset  $\mathcal{T} \in B_n$  may lead to  $p^d$ .

Given the importance of the final permutation in the prediction of a joint strategy, we may break the problem of predicting a joint strategy into two separate inference subproblems: (1) a prediction of the final permutation and (2) a prediction of a *compatible* system path *entanglement* – braid word, conditioned on the predicted final permutation. Following this reasoning, eq. (3.6) may be rewritten as the following product:

$$P(\tau|\Xi_t, M_t) = P(p^d, \tau^d|\Xi_t, M_t) \quad (3.7)$$

$$= P(\tau^d|p^d, \Xi_t, M_t)P(p^d|\Xi_t, M_t), \quad (3.8)$$

where  $\tau^d \in \mathcal{T}$  represents a braid that is compatible with the prediction of a final permutation  $p^d$ , given the permutation  $p_t = f_x(Q_t)$  corresponding to the current system state  $Q_t = \Xi(t)$ . Figure 3.6 depicts a graphic representation of the structure of our inference mechanism.

### 3.2.2.1 Inferring the Final Permutation of the System

A planning agent knows with certainty its own destination but has no access to other agents' destinations. Although there is no need to make an inference regarding others' actual destinations, it is important to infer a final permutation  $p^d$  so as to make an informed inference regarding the emerging joint strategy, as discussed in the previous subsection. Under the assumption of rationality and given a model of the world, stored in the context  $M_t$ , the planning agent may

infer the general directions of others through a belief:

$$P(p^d|\Xi_t, M_t) = P(p^d(1), \dots, p^d(n)|\Xi_t, M_t), \quad (3.9)$$

where

$$p^d = \begin{pmatrix} 1 & \dots & n \\ p^d(1) & \dots & p^d(n) \end{pmatrix}. \quad (3.10)$$

For simplicity, let us assume that the planning agent's ID is #1. Then, under the assumption that all agents are moving towards destinations from the known set  $D$ , the agent's index under the permutation  $p^d$ , i.e.,  $p^d(1)$ , is constrained. Leveraging this, the planning agent may conclude to a subset of feasible permutations  $\mathcal{P} \subset Perm(N)$ , by ruling out any incompatible permutations from  $Perm(N)$  as unlikely. Under the assumption that the cardinality  $|D| \geq |N|$ , multiple assignments of destinations to agents may be possible for each permutation in  $\mathcal{P}$ .

More formally, for each compatible permutation  $p_m^d \in Perm(N)$ , with  $m \in \{1, \dots, M\}$  and  $M = (n-1)!$ , we may derive a set of possible destination assignments  $\Delta_m$  that are compatible with (1)  $p_m^d$  and (2) the image of agent #1, through the permutation  $p_m^d$ , i.e.,  $p_m^d(1)$ , corresponds to a final arrangement of all agents to destinations of  $D$ , with agent #1 at its destination  $d_1$ . Essentially, an assignment  $\delta \in \Delta_m$  is an injective function  $\delta : N \rightarrow D$  that maps all agents to a subset of landmarks from  $D$ . Upon marginalizing over all possible  $\delta \in \Delta_m$ , the probability that the final permutation  $p^d = p_m^d$  may be derived as:

$$P(p^d = p_m^d|\Xi_t, M_t) = \sum_{\delta \in \Delta_m} P(p^d = p_m^d|\Xi_t, M_t, \delta) P(\delta|\Xi_t, M_t). \quad (3.11)$$

By definition,  $p^d = p_m^d$ , if we know that agents are going to  $\delta$  and therefore eq.

(3.11) may be simplified as follows:

$$P(p^d = p_m^d | \Xi_t, M_t) = \sum_{\delta \in \Delta_m} P(\delta | \Xi_t, M_t). \quad (3.12)$$

The destination that an agent is aiming for is conditionally independent of the destinations of others, given  $\Xi_t$  and  $M_t$ , therefore we may express  $P(\delta | \Xi_t, M_t)$  as:

$$P(\delta | \Xi_t, M_t) = \prod_{j=2}^n P(q_j^d = \delta(j) | \Xi_t, M_t) \quad (3.13)$$

where we incorporated the fact that agent #1 is certain about its destination. Finally, under the assumption of rationality, we follow an approach similar to Dragan and Srinivasa [40] to model  $P(q_j^d = \delta(j) | \Xi_t, M_t)$  as:

$$P(q_j^d = \delta(j) | \Xi_t, M_t) = \frac{1}{Z} \frac{\exp(-c(\xi_j) - c^*(q_j^s, \delta(j)))}{\exp(-c^*(q_j^s, \delta(j)))}, \quad (3.14)$$

where  $\xi_j$  is the path agent  $j$  has followed so far,  $c$  measures the length of a path,  $c^*$  returns the shortest path between two points and  $Z$  represents a normalizer across the set of landmarks  $D$ .

Algorithm 1 outlines the process of scoring all compatible final permutations from  $Perm(N)$ . Function `Get_Permutation` returns a permutation  $p_l$  of the set  $D$ , corresponding to the arrangement of landmarks in an order of increasing  $x$ -coordinates with respect to the agent's frame. Then, all permutations are accessed and checked for compatibility with the planning agent's destination  $d$  (function `Check_Perm`). In case a permutation is compatible, the set of possible assignments of agents to destinations that are in compliance with the permutation, is extracted with function `Get_Assignments` and then scored (function `Score_Assignments`). Otherwise, the corresponding permutation is assigned a zero score. The scores are finally normalized and returned in the form of a probability distribution  $P$ .

### 3.2.2.2 Inferring the Path Entanglement

The second distribution, over system path entanglements  $P(\tau^d | p^d, \Xi_t, M_t)$  is a harder distribution to approximate. Especially events that take place further than one event ahead may be impossible to be traced back to the decisions that agents have made until time  $t$ . For this reason, in this work, we approximate  $P(\tau^d | p^d, \Xi_t, M_t)$  as:

$$P(\tau^d | p^d, \Xi_t, M_t) = P(\tau_1^d, \dots, \tau_K^d | p^d, \Xi_t, M_t) \quad (3.15)$$

$$\approx \frac{1}{\Gamma} P(\tau_1^d | p^d, \Xi_t, M_t), \quad (3.16)$$

that is, given the state of execution at time  $t$ , as expressed in  $\Xi_t$  and  $M_t$ , the probability of a braid word is set to be approximately proportional to the probability of the next generator, observed after time  $t$ , being equal to the one prescribed by  $\tau_1$ , where  $\Gamma$  is an appropriate normalizer. All entanglements that share the same first generator  $\tau_1$  are assigned the same probability. Finally, in case the current permutation of agents is equal to the predicted final permutation  $p^d$ , it is assumed that the only possible path entanglement is the trivial one, i.e.,  $\tau = e$ , and the rest are assigned zero probabilities.

To model  $P(\tau_1^d | p^d, \Xi_t, M_t)$ , we first encode a generator  $\tau_1^d \in \{\sigma_1, \sigma_1^{-1}, \dots, \sigma_{n-1}, \sigma_{n-1}^{-1}\}$  into a tuple  $g = (swap_g, sign_g)$ , where  $swap_g \in \{1, \dots, n-1\}$  contains the generator's subscript (corresponding to the pair of agents that are exchanging sides) and  $sign_g \in \{-1, +1\}$  contains the generator's superscript (how they are exchanging sides). This model allows us to decompose the generator prediction problem into (1) a prediction of the immediately swapping agents and (2) a

---

**Algorithm 1** Score\_Permutations( $d, \Xi, D, perms$ )

---

**Input:**  $d$  – agent’s intended destination;  $\Xi$  – state history of all agents;  $D$  – list of landmark locations;  $perms$  – list of permutations.

**Output:**  $P$  – prob. distribution over permutations.

```
1:  $p_l \leftarrow Get\_Permutation(D)$ 
2: for  $i = 1 : n!$  do
3:    $compatible \leftarrow Check\_Perm(perms[i], p_l, d)$ 
4:   if  $compatible$  then
5:      $\Delta \leftarrow Get\_Assignments(p_l, perms[i])$ 
6:      $S[i] \leftarrow Score\_Assignments(\Delta, \Xi)$ 
7:   else
8:      $S[i] \leftarrow 0$ 
9:  $P \leftarrow S / sum(S)$ 
10: return  $P$ 
```

---

prediction of the type of their swap:

$$\begin{aligned} P(\tau_1^d = g | p^d, \Xi_t, M_t) &= P(swap_g, sign_g | p^d, \Xi_t, M_t) \\ &= P(sign_g | swap_g, p^d, \Xi_t, M_t) \\ &P(swap_g | p^d, \Xi_t, M_t). \end{aligned} \quad (3.17)$$

Regarding the prediction of the next swap, we employ the following model:

$$P(swap_g | p^d, \Xi_t, M_t) = \frac{1}{H} \prod_{i=1}^{n-1} R_{swap}(i), \quad (3.18)$$

where  $H$  is a normalizer across swaps and  $R_{swap}$  is defined as:

$$R_{swap}(i) = \begin{cases} \frac{1}{1 + \exp(\Delta x_i - \epsilon)}, & i = swap_g \\ \frac{1}{1 + \exp(-(\Delta x_i - \epsilon))}, & i \neq swap_g \end{cases} \quad (3.19)$$

where  $\Delta x_i = x_{i+1} - x_i$  represents the  $x$ -distance between the agents of the pair  $i$  (agents  $p_t(i)$  and  $p_t(i + 1)$ ) in the current permutation  $p_t$ ) and  $\epsilon > 0$  is a constant. Smaller distances indicate swaps corresponding to generators that are exponentially more likely.

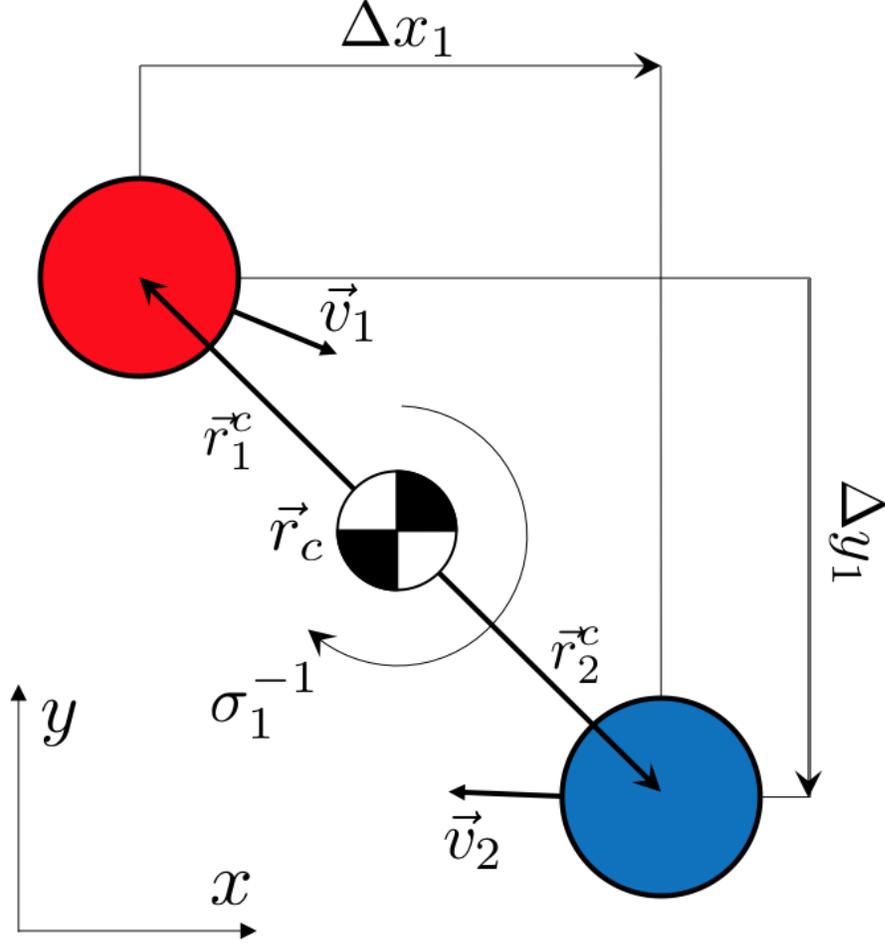


Figure 3.7: Demonstration of the Momentum heuristic for predicting a generator superscript: The  $x$ -projections of the agents' paths are about to cross, forming a  $\sigma_1^{-1}$  generator. The  $z$  component of their angular momentum is negative, indicating a tendency for a counterclockwise rotation, which indicates a negative braid exponent; likewise, in case the  $z$  component of the angular momentum were positive, the emerging exponent would be positive.

Given a prediction of a  $swap_g$ , the next step is to determine the type of swapping, i.e.,  $sign_g$ . In order to model  $P(sign_g | swap_g, p^d, \Xi_t, M_t)$  we employ a heuristic based on the angular momentum of the system of the agents in consideration, which, assuming unit masses, may be defined as:

$$\vec{L}(swap_g) = \vec{r}_i^c \times \vec{v}_i + \vec{r}_{i+1}^c \times \vec{v}_{i+1}, \quad (3.20)$$

where  $\vec{r}_i^c, \vec{r}_{i+1}^c$  are respectively the positions of the currently right and left agents, defined with respect to their current center of mass  $\vec{r}_c = (\vec{r}_i^c + \vec{r}_{i+1}^c)/2$ , whereas  $\vec{v}_i, \vec{v}_{i+1}$  are their respective velocities. For two masses moving on the same plane, the angular momentum is a vector, normal to the plane, along the direction of rotation. If the masses are about to rotate counterclockwise, with respect to an axis of reference, the  $z$ -component of the momentum,  $L_z$ , is positive, and negative in case the masses are about to rotate clockwise. Based on this observation, we model  $P(\text{sign}_g | \text{swap}_g, p^d, \Xi_t, M_t)$  as:

$$P(\text{sign}_g | \text{swap}_g, p^d, \Xi_t, M_t) = \begin{cases} \frac{1}{\Theta} \frac{1}{1 + \exp(-\text{sign}_g L_z(\text{swap}_g))}, & \text{if } |\vec{v}_i| + |\vec{v}_{i+1}| > 0 \\ \frac{1}{\Theta} \frac{1}{1 + \exp(-\text{sign}_g \Delta y_i)}, & \text{otherwise} \end{cases}, \quad (3.21)$$

where  $\Theta$  is an appropriate normalizer. The more positive  $L_z$  is,  $P(\text{sign}_g = +1 | \text{swap}_g, p^d, \Xi_t, M_t)$  gets exponentially closer to 1, whereas in the opposite case,  $P(\text{sign}_g = -1 | \text{swap}_g, p^d, \Xi_t, M_t)$  gets closer to 1. In case the velocities of both agents are currently zero, the corresponding scores only depend on their distance along the  $y$ -axis,  $\Delta y_i = y_{i+1} - y_i$ . Figure 3.7 demonstrates schematically the concept of momentum (for nonzero velocities) and how it is used to predict  $\text{sign}_g$ .

**Remark 2.** *It should be noted that the model of inference presented in this section constitutes is an extended version of the one presented by Mavrogiannis and Knepper [97], as it may handle a) uncertainty over destinations, b) redundancy of destinations (case with no unique mapping from a permutation to the set of destinations) and c) incorporates a novel heuristic for predicting the exponent of braid generators. However, it should also be noted that this distribution is a simplified approximation that cannot*

*guarantee robust performance and generalization. We are using it to provide a proof of our concept. Recent work of ours [105] presented a data-driven framework for directly learning to predict future trajectory topologies from simulated demonstrations of challenging multi-agent scenarios.*

### 3.3 Decision Making

In multi-agent environments, where there is no explicit communication among agents, uncertainty regarding everyone’s actions is typically high, which complicates decision making. Humans usually overcome such a complication by communicating implicitly, mostly through motion. Doing so is made possible through inference mechanisms that allow them to read the intentions of others and select socially compliant actions that reduce uncertainty. This enables them to reach a consensus over an avoidance protocol that serves everyone and ensures comfort, while making progress towards their destinations. The superposition of these considerations represents what – to our interpretation of the pedestrian bargain [148] – constitutes socially *competent* behavior in a pedestrian context.

To generate socially competent behaviors, we design a cost-based policy. Our cost function enables an artificial agent to take actions that not only contribute progress towards its destination but also towards a consensus over a joint strategy that appears to be mutually beneficial for everyone in the scene. Regarding the first specification, a distance-based efficiency cost is employed, whereas for the second one, the entropy of the distribution over joint strategies  $P(\tau|\Xi, M)$  is used. The distribution allows the planning agent to estimate the

long-term effects of an action in consideration and how it might influence the decision making of others. The reduction of the entropy of the distribution may allow the agent to select an action that reduces the uncertainty for everyone.

In the following subsections, we describe in detail our decision making framework.

### 3.3.1 Modeling Agents' Cost Functions

We model the interests of an agent  $i$  with a cost function  $u_i : \mathcal{A}_i \rightarrow \mathbb{R}$  that maps an action  $a_i \in \mathcal{A}_i$  to a real number. We design this cost to comprise two terms: (1)  $E_i$ , which represents the agent's personal *Efficiency* and (2)  $C_i$ , which represents the state of *Consensus* over a joint strategy among agents, from the perspective of agent  $i$ , upon taking an action  $a_i \in \mathcal{A}_i$ :

$$u_i(a_i) = \lambda E_i(a_i) + (1 - \lambda)C_i(a_i). \quad (3.22)$$

We define the personal efficiency term  $E_i$ , to be the length of the shortest path to the agent's destination, whereas  $C_i$  is modeled as the Information Entropy of the belief distribution over joint strategies  $P(\tau|\Xi, M)$ , from the perspective of agent  $i$ , i.e.,

$$C_i(a_i) = - \sum_{\tau \in \mathcal{T}} P(\tau|\Xi^+, M) \log_2 P(\tau|\Xi^+, M), \quad (3.23)$$

where  $\Xi^+$  denotes the system path so far,  $\Xi$ , augmented with the action in consideration  $a_i$ . Finally,  $\lambda$  is a weighting factor, expressing the compromise between efficiency and consensus. Formally, the decision making policy may be described as a minimization of eq. (3.22):

$$a_i^* = \arg \min_{a_i \in \mathcal{A}_i} u_i(a_i). \quad (3.24)$$

Note that the cost function  $u_i$  plays the role of a utility function, with the difference that lower values are better.

Overall, this policy enables an agent to make decisions that not only contribute progress towards its destination but also towards a mutually beneficial consensus over a scene outcome. The faster such a consensus is established, the lower the uncertainty will be for all agents throughout the remainder of the execution. The Efficiency term represents agents' intention of reaching their destinations by spending low energy and is in line with the principle of rational action as highlighted in the definitions of the pedestrian bargain [148] and the teleological reasoning [32]. The Consensus term scores the current state of the global consensus among agents regarding the joint strategy to be followed and therefore, it directly incorporates a form of social understanding into the agent's decision making policy. The lower the entropy, the lower the uncertainty regarding the emerging joint strategy. Thus, by consistently picking actions that contribute to entropy reduction, an agent communicates its intention of complying with a subset of scene outcomes that appear to be preferable by everyone according to the model  $P(\tau|\Xi, M)$ . As a result, the agents are expected to reach a consensus over  $\tau$  easier and faster, avoiding ambiguous situations such as livelocks or deadlocks and reach their destinations with lower planning effort.

### 3.3.2 Planning Joint Strategies

In practice, making use of the distribution  $P(\tau|\Xi, M)$  requires determining a set of candidate joint strategies  $\mathcal{T}$ . The braid group  $B_n$  is countably infinite; in

practice, however, only a subset of joint strategies (braids) are meaningful under the context of a scene  $M_t$  and given observations of agents' past behaviors  $\Xi_t$ . In particular, as discussed in Sec. 3.2, given any compatible final permutation  $p_m^d \in \mathcal{P}$ , only a subset  $\mathcal{T}^m \subset B_n$  may be achievable from the current permutation  $p_t = f_x(Q_t)$ . Consequently, there arises the problem of planning a set of joint strategies  $\mathcal{T} = \bigcup \mathcal{T}^m, m = \{1, \dots, M\}$ , compatible with the set of different final permutations in consideration from  $\mathcal{P}$ .

Planning a joint strategy that transitions the system from a permutation corresponding to the current state of the system,  $p_t$ , to a permutation corresponding to the final state of the system,  $p_m^d$ , may be decomposed into the following sub-problems: (1) planning a path of permutations,  $\pi^m = (p_t, \dots, p_m^d)$  connecting  $p_t$  with  $p_m^d$  through a sequence of adjacent transpositions, and (2) assigning compatible elementary braids to the transitions between consecutive permutations. Each transition may be implemented in 2 different ways, i.e., by a compatible generator or its inverse. For example, a path  $\pi$  of length  $l_\pi$  (the number of transitions required to reach  $p_i^d$  from  $p^t$ , through  $\pi$ ) may be implemented by  $2^{l_\pi}$  different braids.

Assuming that we have concluded to a set  $\mathcal{P} \in Perm(N)$  of potential final permutations, for each one of them  $p_m^d \in \mathcal{P}$ , we plan a set of paths of adjacent transpositions,  $\mathcal{P}^m$ , based on which we generate a batch of joint strategies  $\mathcal{T}^m$ . These are used to form the final set of joint strategies  $\mathcal{T} = \bigcup \mathcal{T}^m, m = \{1, \dots, M\}$ .

In this chapter, we convert the problem of planning a topological joint strategy into a search in a graph of permutations. In the following sections, we describe the construction of the graph and the planning procedure.

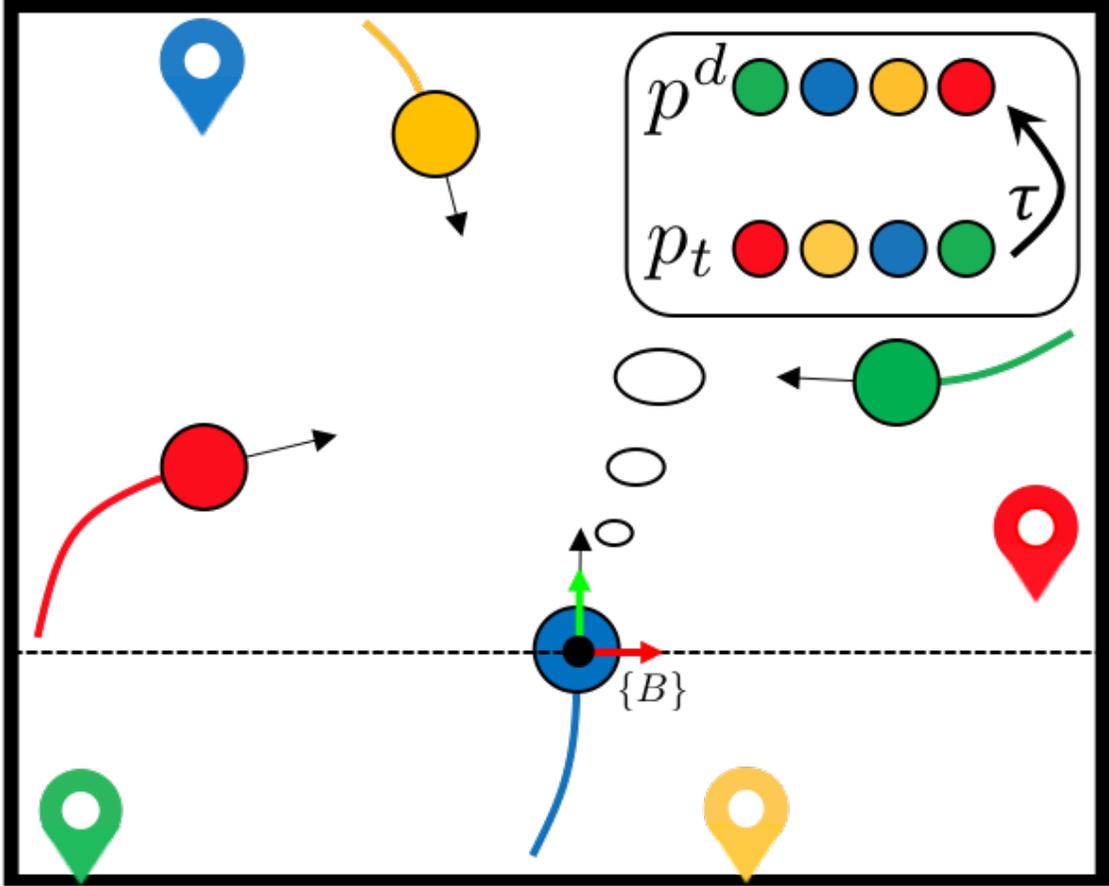


Figure 3.8: A multi-agent scene from the perspective of the planning agent (blue color). At time  $t$ , the agent arranges all agents in the scene in an order of increasing  $x$ -coordinates with respect to the  $x$  axis of its body frame  $\{B\}$  and derives a corresponding permutation  $p_t$ . Based on observation of all agents' past trajectories (solid lines) and given knowledge of existing landmarks in the scene, the blue agent makes a prediction of everyone's destination (colored pointers) and derives a corresponding final permutation  $p^d$ . Transitioning from  $p_t$  to  $p^d$  may be implemented with a joint strategy  $\tau \in \mathcal{T}$ .

### 3.3.3 Permutation Graph Search

The set of all permutations on  $N$ ,  $Perm(N)$ , along with the composition operation, form the symmetric group  $S_n$ .  $S_n$  is a group of order  $n!$ , that can be generated by the set of adjacent transpositions  $\beta_j = \begin{bmatrix} j & j+1 \end{bmatrix}$ , for  $1 \leq j < n-1$ .

We make use of the symmetric group to construct a graph  $G = (V, E)$ , where

$V = Perm(N)$  and any pair of nodes  $\nu_a, \nu_b \in V$  is only connected iff  $\exists \beta_{ab} \in S_n$  that permutes  $\nu_a$  into  $\nu_b$ . The graph  $G$  may be represented as a  $(n - 1)$ -dimensional polytope, embedded in a  $n$ -dimensional space, which is commonly referred to as a *permutohedron* [152]. Fig. 3.9 depicts a permutohedron of order 4, along with example paths and indications of braid transitions.

Planning a path from a permutation  $p^a$  to a permutation  $p^b$ , corresponding respectively to the vertices  $\nu_a, \nu_b \in V$ , is equivalent to finding a path of vertices-permutations that connect them. Figure 3.8 illustrates the concept of planning a joint strategy. At planning time  $t$ , the agents have already followed trajectories  $\Xi_t$ . The planning agent (blue color) has predicted that they are aiming at reaching the destinations denoted by pointers of corresponding colors. Transitioning from the current system configuration to the predicted final system configuration corresponds to transitioning from the current permutation  $p_t$  to a permutation  $p^d$ , both defined with respect to the dashed line, parallel to the x-axis of the agent's body frame  $\{B\}$ .

### 3.3.4 Online Algorithm

In this section, we describe our algorithm design for online navigation planning that makes use of the components detailed in the previous sections. The algorithm compromises between making progress towards the agent's destination and being respectful of everyone's intentions, as inferred by their past behaviors. In case the planning agent does not observe other agents on its way to its destination, it switches to efficiency optimization.

The SCN algorithm (alg. 2) is our online algorithm for *socially competent*

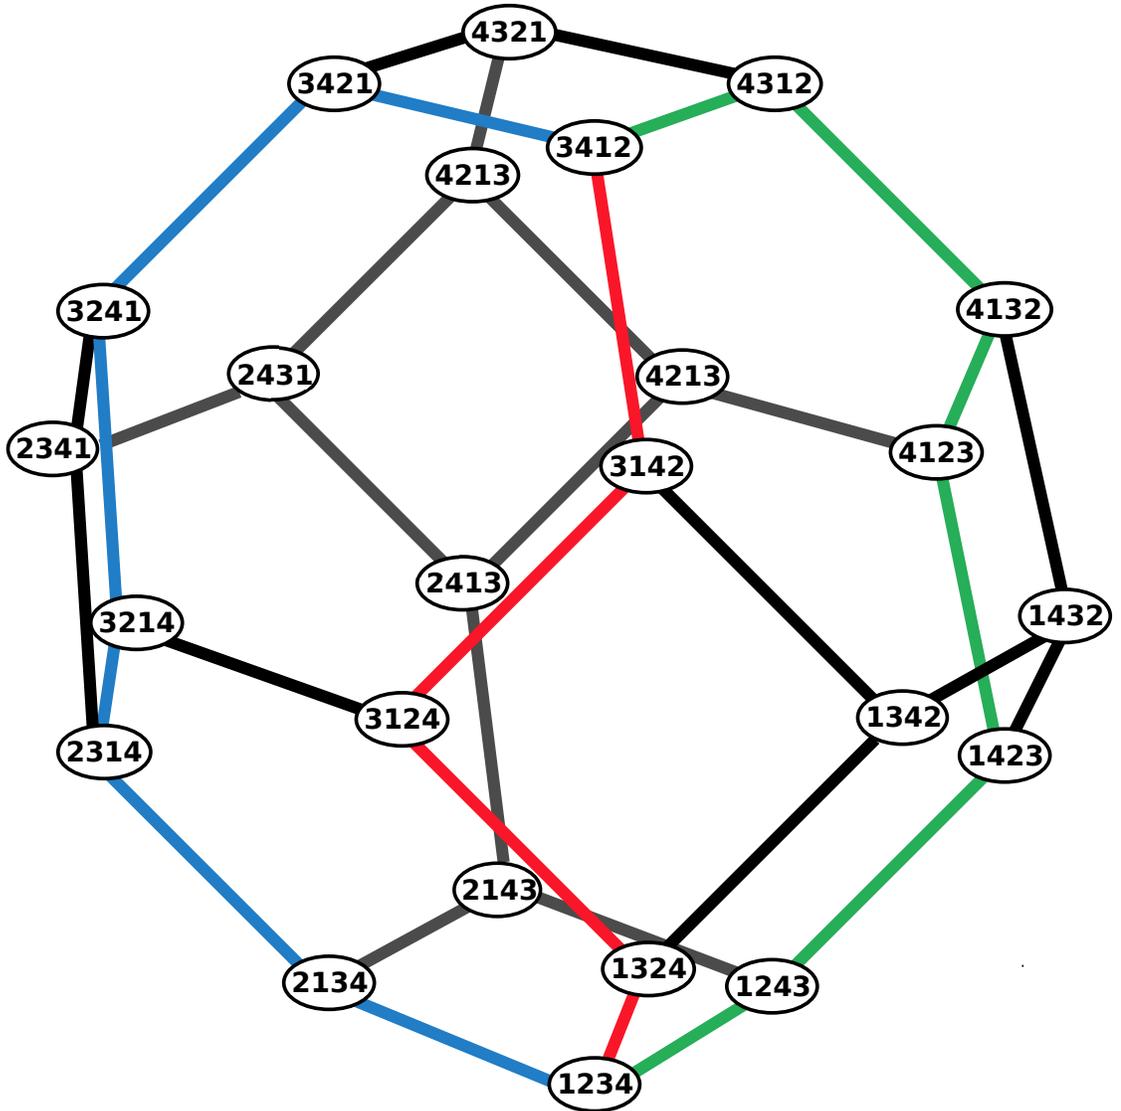


Figure 3.9: A permutahedron of order 4 for a scene with four agents. Three alternative paths implementing the transition from the permutation 1234 to the permutation 3412 are depicted in different colors. Each path consists of a sequence of transitions, each of which can be implemented topologically with a braid generator or its inverse.

*navigation.* At every replanning cycle, the function `Get_Reactive_Agents` first returns a list of agents,  $R$ , which the planning agent should be avoiding collisions with. Then, function `Collision_Check` generates a set of implementable, collision-free actions  $\mathcal{A}$ . If  $R \neq \emptyset$ , the algorithm proceeds by planning a set of joint strategies  $\mathcal{T}$  (function `Get_Strategies`) and then picking the minimizer of (3.22) (function `Cost_Optim`). In case  $R = \emptyset$ , the the function `Efficiency_Optim` computes the most efficient action. The algorithm terminates when the boolean `AtGoal` becomes true, indicating that the planning agent has reached its destination.

### 3.3.4.1 A Discussion of Complexity

The most computationally expensive part of the proposed algorithm is the computation of the set of joint strategies  $\mathcal{T}$ . Obtaining  $\mathcal{T}$  involves determining the set of all compatible final **permutations**  $\mathcal{P} = \{p_1^d, \dots, p_M^d\}$  that correspond to final system states with the planning agent at its destination. This set has cardinality  $|\mathcal{P}| = M = (n - 1)!$  and may be computed in time  $O(n!)$ . For each permutation  $p_m^d$ , we are computing:

1. A set of  $\mathcal{K}$  **permutation paths**  $\Pi^m$  that connect the current system permutation  $p_t$  with  $p_m^d \in \mathcal{P}$ . For computing this path set, we employ an algorithm for finding  $\mathcal{K}$ -shortest paths. Such algorithms typically make use of a shortest path algorithm, e.g. Dijkstra's [37]; therefore their complexity depends on the number of calls to the shortest path algorithm. Katoh's algorithm [75] appears to be the most efficient among them, with a runtime complexity of  $O(\mathcal{K}(|E| + |V| \log |V|))$ , where  $|V|$  and  $|E|$  represent the number of nodes and edges in the graph respectively, which for a permu-

tohedron of order  $n$  are equal to  $|V| = n!$  and  $|E| = n!(n-1)$ . For a constant  $\mathcal{K}$ , this computation runs in time  $O(n!n \log n)$ .

2. A set of **braids**  $\mathcal{T}^m$ , consistent with each permutation path  $\pi^m \in \Pi^m$ . For each permutation path in  $\Pi^m$ , we derive  $2^l$  different braids (where  $l$  is the number of edges in the path), by taking all possible permutations of consistent generator assignments on the permutation path edges. For a constant  $\mathcal{K}$  and considering a maximum path length  $\frac{n(n-1)}{2}$ , this computation runs in time  $O(2^{\frac{n(n-1)}{2}})$  (worst case complexity).

---

**Algorithm 2**  $\text{SCN}(D, Q, d, \Xi, \text{At\_Goal}, M)$

---

**Input:**  $D$  – list of landmarks;  $Q$  – system state;  $d$  – planning agent’s intended destination;  $\Xi$  – state history of all agents;  $\text{At\_Goal}$  – boolean variable signifying arrival at agent’s destination;  $M$  – context.

**Output:**  $a$  – action selected for execution

- 1: **while**  $\neg \text{AtGoal}$  **do**
  - 2:      $R \leftarrow \text{Get\_Reactive\_Agents}(\Xi)$
  - 3:      $\mathcal{A} \leftarrow \text{Collision\_Check}(\Xi, M, R)$
  - 4:     **if**  $R \neq \emptyset$  **then**
  - 5:          $\mathcal{T} \leftarrow \text{Get\_Strategies}(d, Q, D, R)$
  - 6:          $a \leftarrow \text{Cost\_Optim}(\mathcal{A}, \mathcal{T}, d, \Xi, M)$
  - 7:     **else**
  - 8:          $a \leftarrow \text{Efficiency\_Optim}(\mathcal{A}, d, M)$
  - 9: **return**  $a$
- 

The dominant term in the expression of the overall worst case complexity for computing the set of possible braids  $\mathcal{T}^m$  for one permutation  $p_m^d$  is the exponential term  $O(n!)$ . This implies that the complexity of the present algorithm does not scale well with large numbers of agents. Furthermore, considering that the aforementioned computation needs to run for all possible permutations in  $\mathcal{P}$ , the complexity becomes  $O((n!)^2)$ . However, we argue that for our purposes, i.e., eventual deployment on a social robot navigating in real world human environments, the following considerations may enable us to restrict ourselves to low  $n$  and  $|\mathcal{P}|$  and  $\mathcal{K}$ :

- (a) A real robot has limited sensing capabilities, usually corresponding to a local radius of a few meters. Therefore, even in crowded environments, the surrounding agents will be considered as they enter the sensing radius and not universally.
- (b) The rationality assumption for agents, which is supported by studies on human inference [33] and human navigation [148]. Rational agents aim at avoiding undesired divergences from the direction pointing towards their destinations [109]. This implies that (i) humans end up following short permutation paths and (ii) a good final permutation prediction can be done by simply projecting forward agents' current velocity on the boundary of the robot's sensing radius (the robot does not need to know exactly where others are going). These observations motivate low  $\mathcal{K}$  and low  $|\mathcal{P}|$  respectively.
- (c) The rationality assumption also allows us to assume that once two agents pass each other, they stop reacting. For this reason, at replanning time, we restrict ourselves to considering only the number of agents that are ahead and are assumed to be observing the planning agent. Therefore, as the execution progresses and the robot approaches its destination, the number of reactive agents is expected to drop, allowing the robot to switch to efficient –and less computationally intense– execution.
- (d) Our algorithm is based on frequent replanning. Plans have a short horizon but are made with a global reasoning (over joint strategies). The short planning horizon has been shown to be in compliance with human locomotion according to Carton et al. [24], who presented evidence that humans employ a shorter planning horizon as they navigate complex environments, to avoid collisions that could emerge from unexpected distur-

bances. The global planning horizon ensures that the motion of the robot will be consistent throughout the whole sequence of consecutive planning cycles.

For reference, in a game with 5 agents, a replanning cycle of SCN that generates 3 permutation paths per permutation runs at an average of  $\sim 185ms$ , with the worst case being  $\sim 402ms$  in a non-optimized MatLab implementation on a MacBook Pro of 2015 with an Intel Core i7 processor of 2.5 GHz, running macOS Sierra. These times appear to be encouraging for real-time execution on a mobile robot platform, upon the transfer to a faster language and the appropriate code optimizations.

## 3.4 Evaluation

In this section, we present simulation results, demonstrating the benefits of our approach. Subsection 3.4.1 describes the experimental setup and provides implementation details, whereas subsection 3.4.2 presents results extracted by testing our algorithm under different settings.

### 3.4.1 Setup

We consider a setup where  $n$  agents navigate a discretized square workspace, partitioned into a set of  $N_t^2$  tiles, where  $N_t$  is the number of tiles per side (Figure 3.10). Each agent  $i \in \{1, \dots, n\}$  starts from an initial tile  $q_i$  and moves towards a final tile  $d_i$ . The game is played in rounds until all agents reach their

destinations.

In order to assess our approach and demonstrate its benefits, we consider challenging game scenarios that reinforce intense encounters among agents. We do so by positioning agents on the sides of the workspace and having them navigate towards opposing sides. Each game scenario is sampled at random from the space of scenarios of size  $N_t^n \times N_t^n$ , corresponding to the number of distinct assignments of agents to initial and final configurations. At every round, the players simultaneously pick an action, which corresponds to a neighboring, unoccupied square. Forward, backward, left, right and diagonal, collision-free transitions are allowed. Since at planning time each agent has no access to the plans of others, in order to ensure collision avoidance, transitioning to a square that is adjacent to a square currently occupied by another agent is not allowed. Depending on the number of agents and the size of the workspace, this setup might result in deadlocks. In our evaluation, executions that result in deadlocks are discarded. However, note that the purpose of our evaluation is to study how agents behave when they have **multiple actions** available. We examine how the adoption of different strategies in action selection may affect the evolution of the game qualitatively and quantitatively.

---

**Algorithm 3** GREEDY( $Q, d, \Xi, AtGoal, M$ )

---

**Input:**  $Q$  – system state;  $d$  – agent’s destination;  $\Xi$  – state history of all agents;  $At\_Goal$  – boolean variable signifying arrival at agent’s destination;  $M$  – context.

**Output:**  $a$  – action selected for execution

- 1: **while**  $\neg AtGoal$  **do**
  - 2:      $\mathcal{A} \leftarrow Collision\_Check(Q, M)$
  - 3:      $a \leftarrow Efficiency\_Optim(\mathcal{A}, Q, d, M)$
  - 4: **return**  $a$
- 

To demonstrate the importance of incorporating a topological understand-

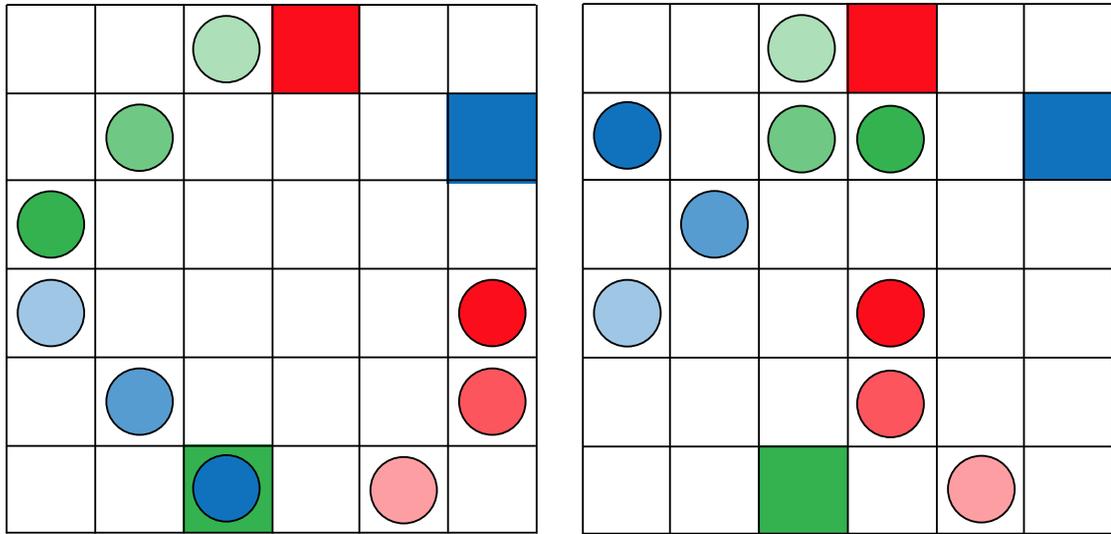
ing into agents' decision making process, we compare the performance of our algorithm (SCN) against a *greedy* baseline (see algorithm 3) that plans actions, greedily seeking to maximize its efficiency (the progress to the agent's destination) at every round. The GREEDY algorithm makes use of the same collision checking function as SCN. Their main difference lies in how they select an action when multiple of them are available. Considering a homogeneous setup ( $n$  agents running SCN versus  $n$  agents running GREEDY), we show that explicitly reasoning about the emerging joint strategy at planning time, benefits everyone in the scene, as it leads to a faster uncertainty decrease that simplifies everyone's decision making. Note that despite the different braid convention that each agent is making, they still manage to converge to a mutually beneficial consensus on a joint strategy of avoidance much faster than GREEDY agents. Qualitatively, our algorithm leads to less ambiguous system configurations, which result in higher average progress per round and lower average time to destination.

### 3.4.2 Simulation Results

In this section, we present the behavior that our algorithm generates and demonstrate its benefits by comparing its performance with the performance of the GREEDY baseline.

#### 3.4.2.1 Qualitative Behavior

Figure 3.10 depicts partial executions, after 2 rounds, of a scenario from a game involving 3 agents. Figure 3.10a depicts the play of agents running SCN, whereas



(a) SCN: By the end of round 2, the agents have reached a system configuration corresponding to a clear consensus over a joint strategy.

(b) GREEDY: By the end of round 2, the agents have reached a system configuration that is about to lead to conflicting encounters.

Figure 3.10: A game with 3 agents in a  $6 \times 6$  workspace. Figure 3.10a and Figure 3.10b depict partial executions of the same scenario (same start and end positions for all agents) with SCN and GREEDY respectively. The current system state is denoted with non-transparent system circles, whereas faded configurations correspond to configurations of past time steps.

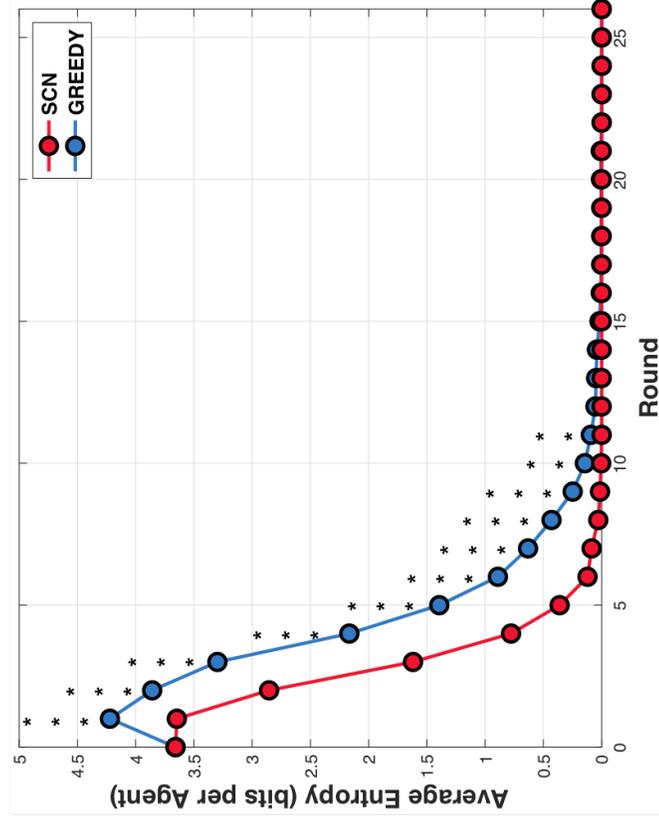
3.10b shows the corresponding play of the agents running GREEDY. It can be observed that the agents running our algorithm have led the game to a configuration that is more beneficial for everyone, as all of their encounters are essentially resolved by the end of the second round. This was achieved by planning informative actions that rapidly led to a significant entropy decrease and accelerated convergence to a consensus over a joint strategy. On the contrary, the agents running the baseline, having initialized their game by focusing on efficiency, are now reacting suboptimally to their constrained action spaces. A video presentation, demonstrating the concept of our approach as well as our system in action can be found at <https://youtu.be/ge9fRI4eav4>.

### 3.4.2.2 Performance

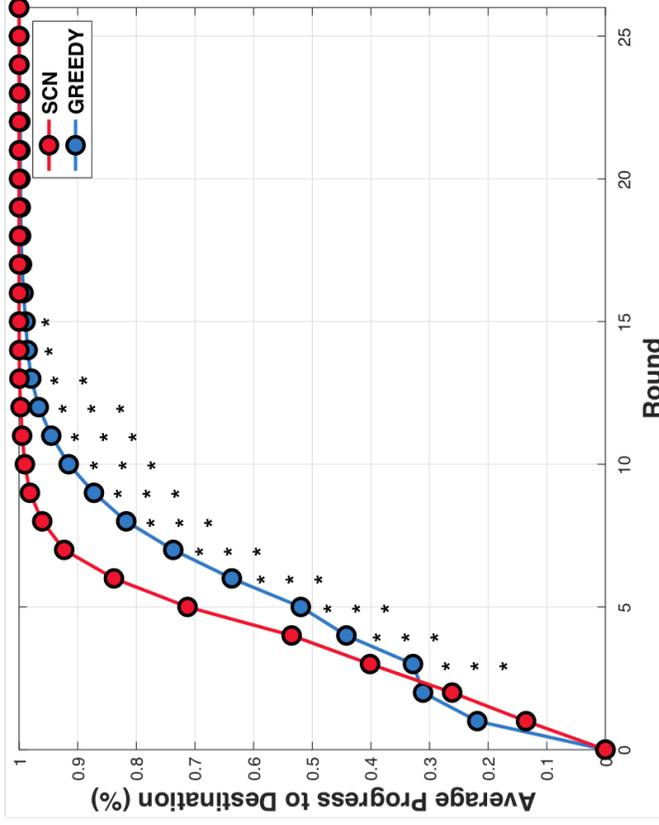
Fig. 3.11 depicts comparative performance diagrams, derived upon running 200 randomly sampled game scenarios involving 3 agents navigating a small workspace of size  $6 \times 6$ . The quality of agents' decision making is illustrated in the profiles of average entropy and average progress to destination, depicted in Figure 3.11a and Figure 3.11b respectively. It can be observed that the systems of agents running SCN achieve faster entropy reduction and higher average progress towards destinations, compared to the systems running the baseline, with statistical significance noted in the diagrams. Figure 3.13a depicts comparative plots of average time to destination (left) and average time to "get free", which corresponds to the time an agent first has full control over the scene evolution, i.e., the first time when no other agents are ahead.

Similar comments can be made for the case of four agents, navigating a workspace of the same size,  $6 \times 6$ . Figure 3.12 presents comparative performance diagrams for entropy and progress to destination, whereas Figure 3.13b depicts a comparison of average time to destination and average time to get free.

For the simulated examples presented, each agent models joint strategies as braids, defined with respect to a projection line that is parallel to its starting side. For the case of the GREEDY agents, entropy was evaluated by employing the inference mechanism of the agents running SCN, i.e., we computed what their belief would be if they had access to the inference mechanism of SCN agents. To ensure proper comparison, the same set of braids was considered at each time step for the same agent in both setups. The weighting factor  $\lambda$  was set to 0.2, as it was found experimentally to lead to a desired compromise between

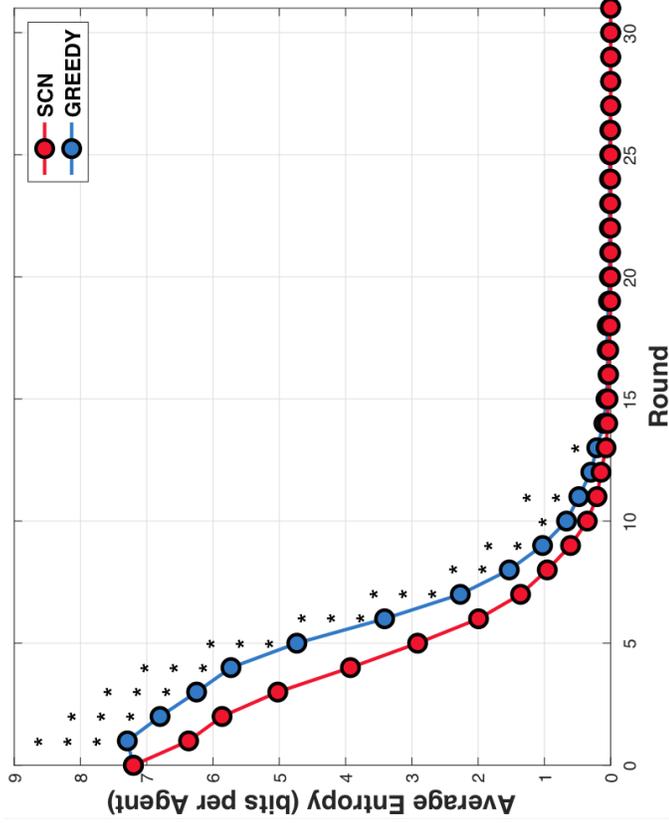


(a) Average Entropy profile per agent per experiment over 200 scenarios involving 3 agents. On average, SCN agents reached a consensus over a joint strategy faster.

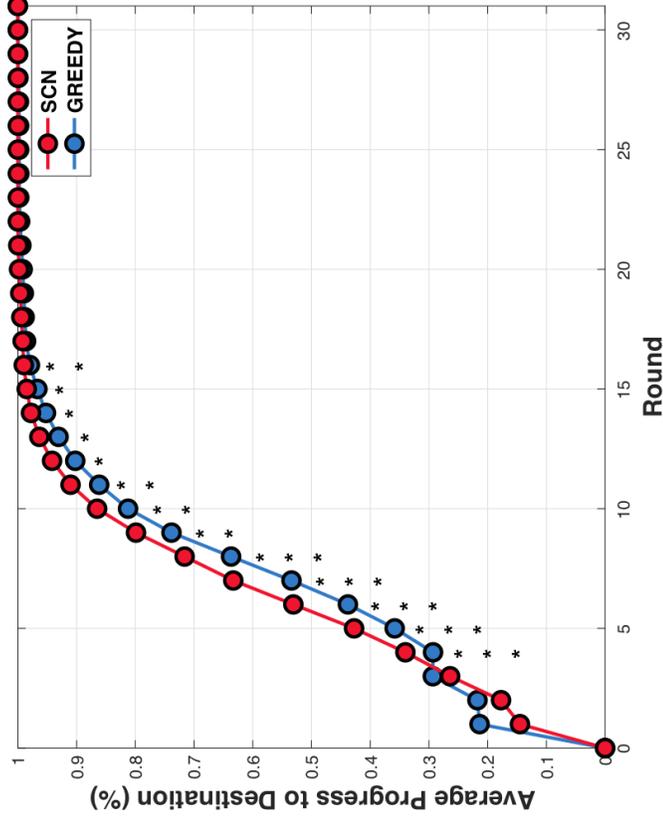


(b) Average Efficiency profile per agent per experiment over 200 scenarios involving 3 agents. On average, SCN agents made greater progress towards their destinations per round.

Figure 3.11: Comparative diagrams, generated upon running 200 experiments with 3 agents in a workspace of size  $6 \times 6$ . The red curves correspond to agents running SCN and the blue curves to agents running GREEDY. The compromise between *Efficiency* and *Consensus* was set to  $\lambda = 0.2$  and the number of paths per permutation to 3. Student's t-tests were performed on all rounds to determine the statistical significance of the profiles difference. The symbols \*, \*\* and \*\*\* denote rounds on which the difference in the performance between SCN and GREEDY was found to be significant to a degree described by p-values  $< 0.05$ ,  $0.01$  and  $0.001$  respectively, according to a paired Student's t-test. Due to space constraints, the significance symbols were place vertically.

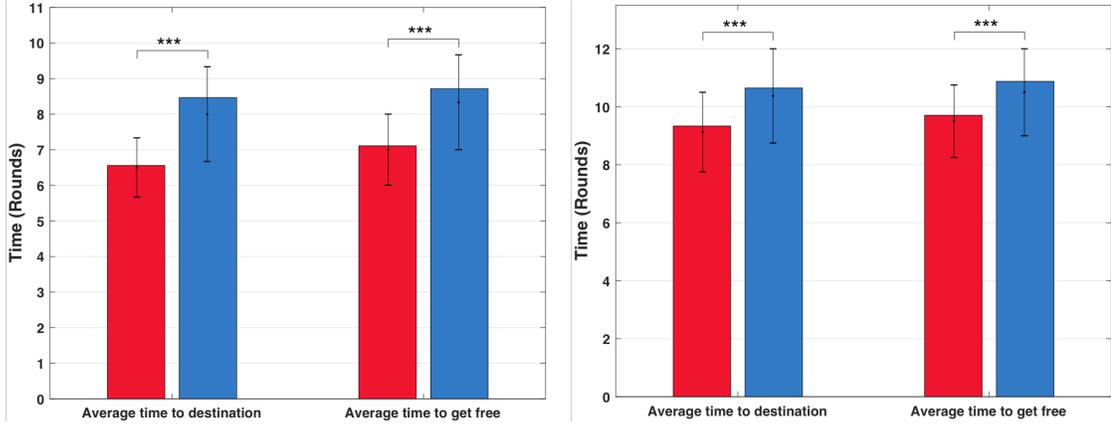


(a) Average Entropy profile per agent per experiment over 200 scenarios involving 4 agents. On average, SCN agents reached a consensus over a joint strategy faster.



(b) Average Efficiency profile per agent per experiment over 200 scenarios involving 4 agents. On average, SCN agents made greater progress towards their destinations per round.

Figure 3.12: Comparative diagrams, generated upon running 200 experiments with 4 agents in a workspace of size  $6 \times 6$ . The red curves correspond to agents running SCN and the blue curves to agents running GREEDY. The compromise between *Efficiency* and *Consensus* was set to  $\lambda = 0.2$  and the number of paths per permutation to 3. Student's t-tests were performed on all rounds to determine the statistical significance of the profiles difference. The symbols \*, \*\* and \* \* \* denote rounds on which the difference in the performance between SCN and GREEDY was found to be significant to a degree described by p-values  $< 0.05$ ,  $0.01$  and  $0.001$  respectively, according to a paired Student's t-test. Due to space constraints, the significance symbols were place vertically.



(a) Average timings per agent over 200 scenarios involving 3 agents. On average, the SCN agents (red color bars) reached their destination faster and managed to “get free” faster than the GREEDY agents (blue color bars).

(b) Average timings per agent over 200 scenarios involving 4 agents. On average, the SCN agents (red color bars) reached their destination faster and managed to “get free” faster than the GREEDY agents (blue color bars).

**Figure 3.13: Average Time to destination and average time to get free, i.e., reach a configuration at which no agents are ahead, generated after running 200 experiments involving 3 (Figure 3.13a) and 4 (Figure 3.13b) agents, in a workspace of size  $6 \times 6$ . Red bars correspond to agents running SCN and blue bars to agents running GREEDY. The error bars indicate 25-75 percentiles. For these experiments, the compromise between efficiency and consensus was set to  $\lambda = 0.2$  and the number of paths per permutation to 3. The \*\*\* symbol denotes a highly significant timing difference, according to Student’s t-test (p-value  $< 0.001$ ).**

efficiency and consensus. For deriving multiple candidate paths in the permutation graph, we use the algorithm of Yen [149] for finding  $K$  shortest paths.

### 3.5 Discussion

We considered the problem of decision making in a navigation scenario involving multiple rational and cooperative agents that do not explicitly communicate with each other. In such a scenario, the uncertainty over the exact strategies of other agents make it hard for an agent to predict their behaviors and thus to

make safe and socially compliant decisions over its own actions. These settings may be found in a variety of real world application, such as robotic navigation in crowded human environments.

To address this problem, we presented an online planning framework, inspired by the insights of recent studies on the cooperative nature of pedestrian behavior [148] and the goal-directed inference of humans [33]. Our framework explicitly incorporates the concept of cooperation by modeling multi-agent collective behaviors as topological global joint strategies, using the formalism of braids [20]. Our topological model forms the basis of an inference mechanism that associates observed behaviors with future collective topologies. In the decision making stage, each agent decides on an action that corresponds to a compromise between its personal efficiency (progress towards destination) and a form of joint efficiency (the status of a consensus on a joint strategy of avoidance). To clearly showcase the benefits of our decision making concept, we deliberately studied a simplified version of the real world problem by considering an abstract, discrete setup, involving artificial agents playing a cooperative game. Extensive trials over randomly-generated, challenging scenarios demonstrated the benefits of reasoning about joint strategies over a baseline that greedily prescribed actions of high efficiency. Our algorithm was shown to lead to a faster decrease of uncertainty regarding the scene evolution, which resulted in efficiency increase and lower execution times with high statistical significance.

## CHAPTER 4

### **SOCIALLY COMPETENT NAVIGATION PLANNING USING A LEARNED MODEL FOR PREDICTING THE MULTI-AGENT PATH TOPOLOGY**

The SCN planner presented in chapter 3 makes use of a probabilistic inference mechanism that outputs a prediction of the emerging multi-agent trajectory topology that agents are about to follow given observations of their past trajectories. We constructed an analytical model of this probability distribution and used it to demonstrate the benefits of reasoning about joint navigation strategies in navigation scenarios with multiple non-communicating agents in a discretized workspace.

In this chapter, we extend the capabilities of SCN through the incorporation of a learned inference mechanism. We derive an approximation of the mechanism by training a neural network on synthetic data comprising trajectories, generated by running randomly generated, simulated multi-agent scenarios. We present the performance of the network in predicting the multi-agent path topology and simulation results, extracted by executing a series of challenging multi-agent scenarios, involving both homogeneous and heterogeneous agents. Most notably, we show that the placement of 1 SCN agent in a workspace where multiple other efficiency-driven agents (running the Social Force Model [62]) also navigate towards their destinations results in modifying the behavior of the latter. Specifically, we observe earlier conflict resolutions, faster consensus among agents and accelerated uncertainty increase, which we interpret as ben-

efits resulting from the socially competent behavior of the SCN agent.

## 4.1 Socially Competent Navigation

According to Wolfinger [148], the social order of human navigation relies on a high-level protocol, comprising two simple rules: (1) *people must behave like competent pedestrians* and (2) *people must trust copresent others to behave like competent pedestrians*. Although Wolfinger did not explicitly define competence, from the examples included in his work, we may deduce that he refers to a notion of *Social Competence*. The concept of Social Competence has been extensively studied in the field of Psychology from different perspectives and for different scenarios (for an extensive review see [111]). In multi-agent navigation, we may define social competence as:

*The ability of an agent to perceive the context<sup>1</sup>, analyze it and pick an action that appears to be compatible with it, according to a pattern of behavior that the agent assumes observing agents expect from him/her by having observed and analyzed the context themselves.*

According to Csibra and Gergely [33], humans tend to attribute *goals* to observed *actions* in a given context. Therefore, socially competent navigation behaviors should be indicative of agents' intentions and compatible with the context. In other words, socially competent agents should be cognizant of the fact that their behaviors implicitly communicate their intentions to any observing

---

<sup>1</sup>By context, we refer to information that is publicly available (e.g. the map), information that may be directly acquirable through sensing (e.g. agent trajectories) and information that may be acquired through standard inference processes (e.g. agent groups).

agents. The importance of implicit communication for human-robot interaction applications has lately been increasingly appreciated [40, 84, 137].

## 4.2 Inferring System Path Topologies from Context

Let us denote by  $M_t$  the *context* of the scene at time  $t \in [0, 1]$ . By context, we refer to information that is either publicly available (e.g. the map of the scene, points of interest, etc.), or directly acquirable through sensing (e.g. agents' state history) or indirectly acquirable through processing (e.g. agents' current arrangement  $p \in Perm(N)$ , inference about agents' destinations, their corresponding final ordering  $p_m$ , agents' groupings, etc.) during the time frame  $[0, t]$ . Assuming that by time  $t \in [0, 1]$ , a sequence of  $k$  events<sup>2</sup>  $\tau_1, \dots, \tau_k$  have already occurred, a model of the form  $P(\tau_{k+1}, \dots, \tau_K | M_t)$  describes the probability of a future system path topology  $\tau = \tau_{k+1} \dots \tau_K \in B_n$  given the context  $M_t$ . For simplicity, when referring to a prediction over future system path topologies, we will be using  $\tau$  to refer to the sequence  $\tau_{k+1} \dots \tau_K$ .

The actions that agents select at each time step become part of the context, as they constitute information that may be directly acquirable by all agents through sensing. Therefore, having an understanding of what collective behaviors  $\tau$  may be compatible with the context  $M_t$ , may allow an agent to contribute to it by executing actions that appear to be in compliance with the emerging collective behavior. In particular, an agent that is considering executing an action from a set of actions  $\mathcal{A}$  may be able to understand how each action  $a \in \mathcal{A}$  may reshape the belief of any observers, by simulating this action and computing

---

<sup>2</sup>For a definition of events, the reader may refer to chapter 3

$P(\tau|M_t, a)$ . Using the chain rule of probability, this distribution may be factorized as:

$$P(\tau|M_t, a) = P(\tau_{k+1}, \tau_{k+2}, \dots, \tau_K|M_t, a) \quad (4.1)$$

$$= P(\tau_K|M_t, a, \tau_{k+1} \dots \tau_{K-1}) \quad (4.2)$$

$$\dots P(\tau_{k+1}|M_t, a)$$

Taking into consideration this distribution at the planning stage may enable an intelligent agent to make decisions that contribute to the context towards what appears to be the more likely or appropriate collective behavior, with respect to the current status of the context,  $M_t$ . In our scope, this is what corresponds to *socially competent* behavior.

### 4.3 Learning Collective Navigation Behaviors

Little variations in agents' decision making and perception mechanisms may lead to significantly different correlations between the context and the topology. Thus, estimating the distribution  $P(\tau|M_t, a)$  realistically with a closed-form model and without introducing over-simplifying assumptions is not a trivial task. For this reason, we adopt a data-driven approach to extract a model of the inference mechanism  $P(\tau|M_t, a)$  from demonstrations of multi-agent navigation. We do so by training a model of the transition probabilities introduced in eq. (4.2). To the best of our knowledge, most publicly available pedestrian datasets either do not contain a sufficiently large volume of sufficiently diverse behaviors or are not in a format compatible to our setup. To overcome these complications, we generate a synthetic dataset of system paths, using the *Social Force* (SF) model [62]. In the following subsections, we describe the process of

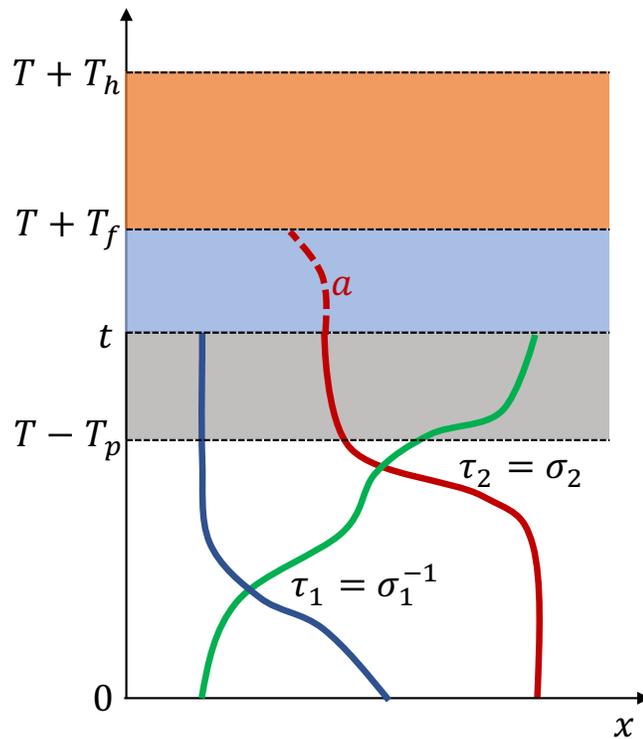


Figure 4.1: Context and time flow around a planning step.

generating our dataset and detail our learning setup and architecture.

### 4.3.1 Generating a Dataset of Diverse Collective Behaviors

#### 4.3.1.1 The Social Force Model

At its simplest form, the core of the Social Force model [62] is a dynamic artificial potential field, constructed by assigning repulsive potentials to agents, workspace boundaries and obstacles and attractive potentials to points of interest or destinations. Each agent is thus subjected to a resultant force that attracts it towards its destination and away from other agents, workspace boundaries or obstacles. Although the Social Force model may produce realistically looking

pedestrian flows in simulation, it lacks a predictive component, which renders it as impractical for real-time robotics applications. However, for training purposes, the model enables us to generate a noise-free set of sufficiently diverse collective behaviors.

#### 4.3.1.2 Experimental Setup

We randomly place a fixed number of agents  $n$  on the circumference of a circular or rectangular workspace. The agents are assigned destinations that enforce intense encounters (lying in the opposite side of the workspace) and move towards them by running individual instances of the social force model. The model parameters, as well as agents' initial positions and destinations are varied across experiments according to gaussian distributions. Experiments on the circular workspace loosely simulate pedestrian crossings in free areas such as atriums or parking lots, whereas experiments on the rectangular workspace resemble crossings in hallways. Each experiment is recorded as a waypoint representation of the system path with fixed time parametrization  $dt$ . The whole dataset is stored in a 4-dimensional tensor  $X$  of size  $T_{max} \times N_{dof} \times n \times N_{exp}$ , where  $T_{max}$  is the maximum number of time steps to destination taken by any agent in the dataset,  $N_{dof}$  is the size of agents' state ( $N_{dof} = 2$  as we do not consider orientation) and  $N_{exp}$  is the total number of experiments.

#### 4.3.2 Learning Setup

We split the acquired dataset into a set of  $N_x$  training examples; each example  $i$  is described by a feature tuple  $\langle M_T^i, a_T^i \rangle$ , where  $M_T^i$  is the context at time step

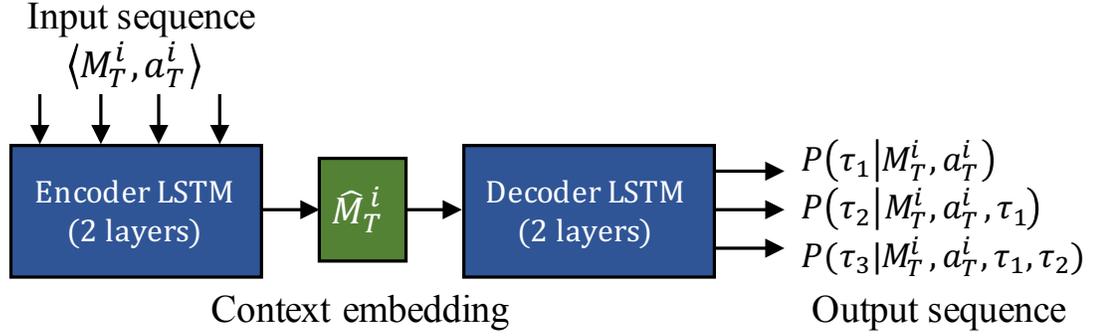


Figure 4.2: Learning Architecture.

$T \in \{1, \dots, T_{max}\}$  and  $a_T^i$  is the action that agent  $i$  executed at that time step, both expressed with respect to frame  $F_i$  centered at the starting position of agent  $i$ , with y-axis pointing towards its destination. We consider the context  $M_T^i$  to be the system path of the time frame  $(T - T_p, T]$ , i.e., we make the assumption that the previous  $T_p$  time steps fully capture the context at time  $T$ . Similarly, we consider the action  $a_T^i$  to be the path that agent  $i$  followed in the frame  $(T, T + T_f]$ . Each example is labeled after the braid word  $\tau \in B_n$  corresponding to the projection of the system path in the horizon  $(T, T + T_h]$  onto the x-axis of frame  $F_i$ . Figure 4.1 demonstrates the time flow around a training example.

### 4.3.3 Learning Architecture

Using the aforementioned setup, the goal of our learning algorithm is to extract models of the conditional probabilities of eq. (4.2), i.e.,  $P(\tau_1 | M_T, a)$ , ...,  $P(\tau_K | M_T, a, \tau_1 \dots \tau_{K-1})$ , so that given an action  $a \in \mathcal{A}$  and a system path topology  $\tau$  of maximum braid length  $K$ , we can compute the probability  $P(\tau | M_T, a)$ . Essentially we need to produce the probability of an output sequence (braid word) given an input sequence (system path). This problem is essentially equivalent

Dataset	$X_{C,2}$	$X_{C,3}$	$X_{C,4}$	$X_{R,2}$	$X_{R,3}$	$X_{R,4}$
Training set	900k	1.6M	2.3M	800k	1.3M	1.8M
Test set	200k	400k	500k	200k	300k	400k

**Table 4.1: Generated dataset sizes (number of examples)**

to a language translation task. Tasks of this form are effectively handled by sequence to sequence neural network models (see e.g. [127]). For this reason, we employ a sequence to sequence encoder-decoder learning architecture. The input sequence  $\langle M_T^i, a_T^i \rangle$  is fed to an encoder Recurrent Neural Network (RNN), which produces an embedding vector  $\widehat{M}_T^i$  that captures the expected future system path topology. The embedding vector is then fed to a decoder RNN that outputs estimates of  $P(\tau_1|M_t^i, a_t^i), P(\tau_2|M_t^i, a_t^i, \tau_1), \dots, P(\tau_K|M_t^i, a_t^i, \tau_1, \dots, \tau_{K-1})$ . For the encoder and decoder RNNs we employ the Long Short-Term Memory (LSTM) architecture [64] due to its effectiveness in capturing long-term sequence dependencies. A schematic representation of our architecture is depicted in Figure 4.2.

### 4.3.4 Implementation Details and Performance

We trained our models on 6 datasets of 100,000 experiments each, labeled  $X_{C,2}, X_{C,3}, X_{C,4}, X_{R,2}, X_{R,3}, X_{R,4}$  where the subscript denotes the number of agents involved ( $n = 2, 3, 4$ ) and the type of workspace considered ( $C$  for circular,  $R$  for rectangular). From each dataset, we used 80,000 experiments for training and the rest for validation. As model parameters, we selected:  $K = 3, T_p = 4, T_f = 3, T_h = 15$ . Using this parametrization, we split the datasets into training examples and test examples as shown in table 4.1. The examples were labeled as braids by using the Braidlab package [132]. As architecture parameters, we

Dataset	$\tau$	$\tau_1$	$\tau_2$	$\tau_3$
$X_{C,2}$	0.93	0.93	1.00	1.00
$X_{R,2}$	0.99	0.99	1.00	1.00
$X_{C,3}$	0.71	0.77	0.93	0.98
$X_{R,3}$	0.78	0.88	0.89	0.97
$X_{C,4}$	0.45	0.62	0.74	0.88
$X_{R,4}$	0.56	0.81	0.78	0.81

**Table 4.2: Braid prediction accuracies for the whole topology  $\tau$  and next, second and third events  $\tau_1, \tau_2, \tau_3$  respectively.**

set both the encoder/decoder to use 2 LSTM layers and a hidden/cell state size of 80. The total number of trainable parameters is 216773. We train using a Dropout [125] of  $p = 0.3$  after each layer. Our models were trained using the LSTM implementation of PyTorch [2]. We used the RMSProp [135] algorithm, considering  $\alpha = 0.99$ , a batch size of 10,000 and an adaptive learning rate schedule, starting from  $LR = 0.001$  and decreasing by a factor of 0.5 if no training loss improvement was observed after 3 epochs until it reached 0.00001. Every mini-batch was constructed with a proportional representation from each dataset, shuffling after every epoch.

The performance of our model in predicting future braid words is presented in tables 4.2 and 4.3. Specifically, table 4.2 contains the accuracies of our models in predicting future topologies in total and per event for each per dataset, whereas table 4.3 contains the accuracies of our models per generator for each dataset, compared with a Prior baseline (each generator is assigned a probability equal to its frequency in the dataset) and Random Guessing (uniform probability for all generators). The accuracies for later time steps improve, because when the future topology  $\tau$  contains less than 3 braid generators, all subsequent generators are trivially identity elements  $e$  i.e. no further crossings occur.

Dataset	$e$	$\sigma_1$	$\sigma_1^{-1}$	$\sigma_2$	$\sigma_2^{-1}$	$\sigma_3$	$\sigma_3^{-1}$
$X_{C,2}$	0.98	0.64	0.76				
$X_{R,2}$	1.00	0.70	0.78				
Prior baseline	0.92	0.05	0.03				
Random guessing	0.33	0.33	0.33				
$X_{C,3}$	0.86	0.60	0.70	0.57	0.70		
$X_{R,3}$	0.90	0.87	0.89	0.85	0.87		
Prior baseline	0.47	0.13	0.13	0.14	0.12		
Random guessing	0.2	0.2	0.2	0.2	0.2		
$X_{C,4}$	0.77	0.51	0.59	0.55	0.61	0.44	0.53
$X_{R,4}$	0.94	0.78	0.81	0.62	0.78	0.78	0.81
Prior baseline	0.29	0.10	0.12	0.12	0.13	0.13	0.12
Random guessing	0.14	0.14	0.14	0.14	0.14	0.14	0.14

**Table 4.3: Per-permutation prediction accuracies for the next braid generator  $\tau_1$ , compared against random guessing and guessing with probability proportional to the prior distribution (frequency of the generator).**

## 4.4 Socially Competent Motion Generation

Our goal is to enable an autonomous agent to exhibit socially competent behavior in a multi-agent setting. From our perspective, this is equivalent to selecting actions that (1) are considered as appropriate within the state of the context  $M_t$ , (2) respect the personal space and the motion plans of others and (3) contribute progress towards the planning agent’s destination.

### 4.4.1 Decision Making Policy

The decision making policy considered in this work is built on the one presented in chapter 3 [97, 104]. We encapsulate the specifications for efficiency and social compliance in a cost function  $\mathcal{C} : \mathcal{A} \rightarrow \mathbb{R}$ , defined as:

$$\mathcal{C}(a) = \lambda\mathcal{E}(a) + (1 - \lambda)\mathcal{H}(a) \tag{4.3}$$

where  $\mathcal{E} : \mathcal{A} \rightarrow \mathbb{R}$  quantifies the *Efficiency* of an action  $a \in \mathcal{A}$ ,  $\mathcal{H} : \mathcal{A} \rightarrow \mathbb{R}$  quantifies the expected state of *Consensus* among agents over the emerging system path topology, upon executing the action in consideration, and  $\lambda$  is a weighting factor. We model Efficiency as the Euclidean distance between the position of the agent upon the execution of the action  $a$  and its destination. Consensus is modeled as the Information Entropy of the distribution over system path topologies  $P(\tau|M_t, a)$ :

$$\mathcal{H}(a) = - \sum_{\tau \in \mathcal{T}} P(\tau|M_t, a) \log P(\tau|M_t, a). \quad (4.4)$$

The higher the consensus cost, the more uncertain the evolution of the scene looks, as, from the definition of the Information Entropy, more outcomes (topologies) will be more likely.

Thus, apart from making progress towards its destination, a socially competent agent has an incentive to actively reduce the uncertainty, by acting according to the context, in a way that reinforces everyone's belief regarding the emerging topology of the system path. The decision making policy for the socially competent agent can be formulated as:

$$a^* = \arg \min_{a \in \mathcal{A}} \mathcal{C}(a), \quad (4.5)$$

where  $a^* \in \mathcal{A}$  is the action that contributes the maximal decrease of  $\mathcal{C}$  in a given context, expressing an optimal compromise between progress to destination and consensus reinforcement, according to the weighting factor  $\lambda$ . Figure 4.3 illustrates an example of reasoning about the future system path topology at planning time.

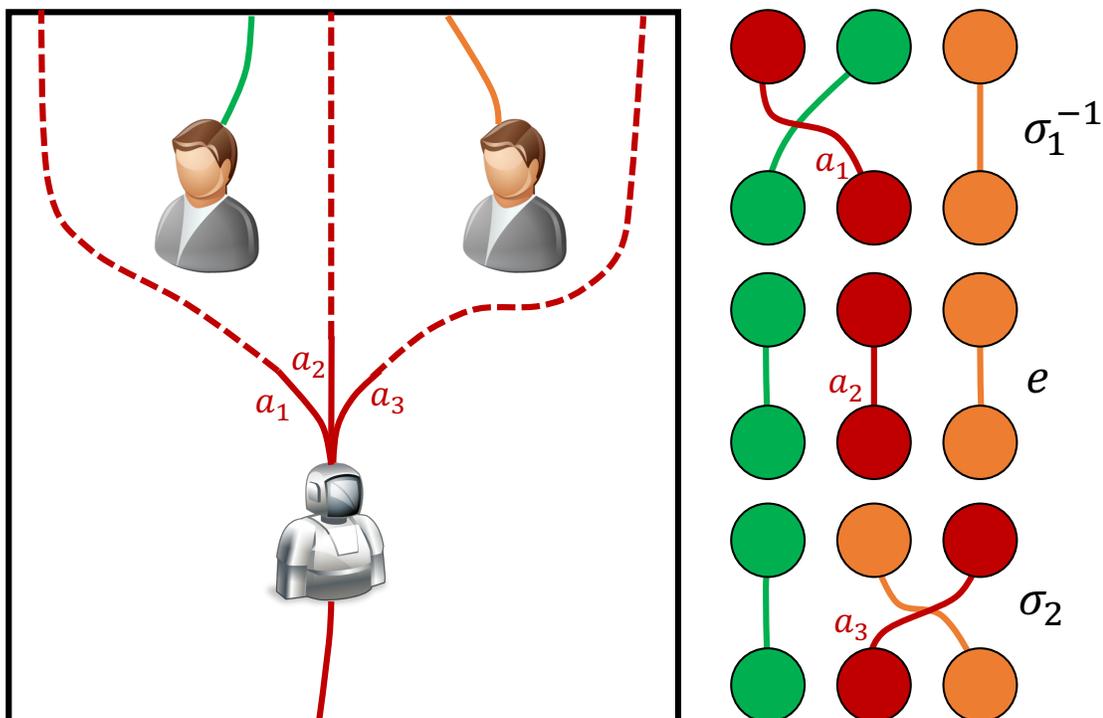


Figure 4.3: The robot is reasoning about different actions leading to qualitatively distinct scene evolutions (left), corresponding to distinct system path topologies (right).

#### 4.4.2 Generating a Set of System Path Topologies

In a scene with  $n$  agents, infinitely many, arbitrarily complex braids could be mathematically possible. However not all of them are likely to emerge. For computational and practical reasons, the planning agent concludes to a set  $\mathcal{T} \subset B_n$  of likely topologies. To do so, similarly to chapter 3, the agent maintains a permutohedron graph, comprising nodes-permutations and edges-elementary braids (see Figure 3.9). At planning time, the agent determines the permutation with respect to the  $x$ -axis of its body frame that corresponds to the current system state  $Q$  and derives the set of all possible future braids words of a given length.

### 4.4.3 Generating a Set of Actions

The planning agent is assumed to have access to a set of actions  $\mathcal{A}$ , comprising trajectories of a fixed number of time steps that are executable by its dynamics (in this paper, we do not incorporate dynamics, assuming that an agent may move towards any direction). The action set is generated offline by considering a set of time-parametrized paths. At planning time, the planner rejects the subset of  $\mathcal{A}$  that is likely to lead to collisions with the environment or other agents. The actions in the collision-free set  $\mathcal{A}_{cf}$  are evaluated with respect to the cost  $\mathcal{C}$  and the best one is executed. This approach is inspired by the works of Green and Kelly [55] and Knepper et al. [82], which provide efficient algorithms and a deeper intuition on path sampling and collision checking.

### 4.4.4 Online Algorithm

Algorithm 4 presents our algorithm for Socially Competent Navigation (SCN). The Function `UpdateContext` incorporates the current system state  $Q$  to the context  $M_t$ . Next, the function `CollisionChecking` checks the action set for collisions and returns a collision-free subset  $\mathcal{A}_{cf} \subseteq \mathcal{A}$ . Subsequently, the function `GetTopologies` derives a set of likely topologies  $\mathcal{T}$ . Then, the function `ScoreTopologies` evaluates every topology in  $\mathcal{T}$  given each action  $a \in \mathcal{A}_{cf}$  and the context  $M_t$  by using our learned model  $P(\tau|M_t, a)$  and returns a corresponding matrix of probabilities  $P$ . Finally, function `MinimizeUtilityCost` evaluates all actions in  $\mathcal{A}_{cf}$  with respect to the utility cost  $\mathcal{C}$  and returns the action  $a^*$  that both contributes the best compromise between progress to destination and communication of compliance with the most likely system path

topology at the given time. The algorithm runs until the agent reaches its destination, i.e., until the boolean variable *AtGoal* becomes 1.

---

**Algorithm 4** SCN( $Q, \mathcal{A}, M_t, d, map, G, AtGoal, a$ )

---

**Input:**  $Q$  – current system state;  $\mathcal{A}$  – action set;  $d$  – agent’s destination region;  $map$ ;  $G$  – Permutohedron;  $AtGoal$  – boolean variable signifying arrival at agent’s destination;  $M_t$  – context

**Output:**  $a^*$  – action selected for execution

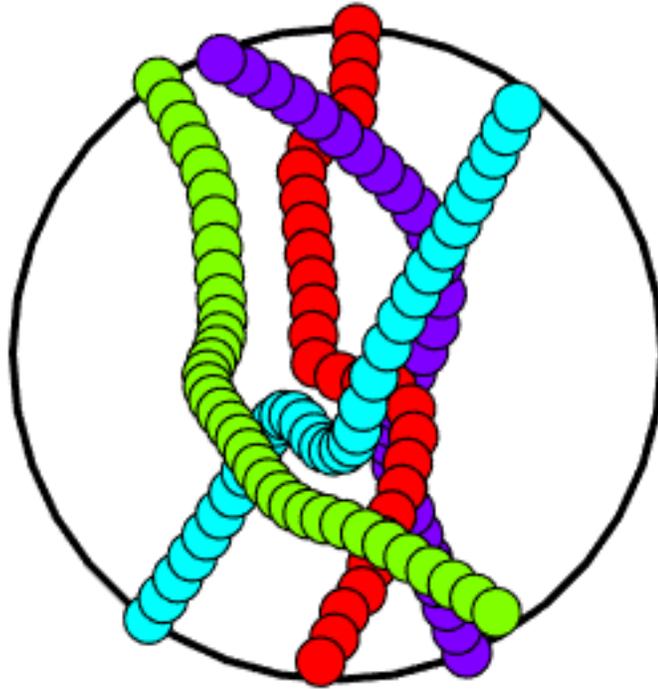
- 1: **while**  $\neg AtGoal$  **do**
  - 2:      $M_t \leftarrow UpdateContext(Q, M_t)$
  - 3:      $\mathcal{A}_{cf} \leftarrow CollisionChecking(\mathcal{A}, M_t, map)$
  - 4:      $\mathcal{T} \leftarrow GetTopologies(Q, G)$
  - 5:      $P \leftarrow ScoreTopologies(\mathcal{T}, \mathcal{A}_{cf}, M_t)$
  - 6:      $a^* \leftarrow MinimizeUtilityCost(P, \mathcal{A}_{cf}, d, M_t)$
  - 7: **return**  $a^*$
- 

## 4.5 Evaluation

In our experimental evaluation, we aim to confirm (1) that an agent running SCN behaves as a socially competent pedestrian, contributing to other agents’ certainty over future topologies and (2) that the presence of a socially competent agent improves the legibility of other agents’ behavior.

### 4.5.1 Experimental Setup

We test our planning algorithm in 50 simulated scenarios involving 3 and 4 agents navigating a circular workspace. Each scenario is defined as a tuple  $(Q_s, Q_d)$ , where  $Q_s$  was defined by placing each agent uniformly at random on the circumference of the workspace (see Figure 4.4) and  $Q_d$  corresponds to points diametrically opposed to  $Q_s$ . The scenarios were deliberately designed



**Figure 4.4:** Swept volumes of 4 agents navigating a circular workspace. The red agent runs SCN whereas the rest of the agents run a separate instance of the social force model.

to reinforce *intense* agent encounters.

For each scenario, we conducted 4 different experiments: (1) experiment  $4SF$ , where all agents are running an individual instance of the social force model, (2) experiment  $1SCN \vee 3SF$ , where agent #1 is running our SCN algorithm and the other 3 agents are running the social force, (3) experiment  $3SF$ , where we remove agent 1 and only simulate the remaining 3 agents using the social force model and (4) experiment  $1SCN \vee 3SF - 1SCN$  that is essentially a playback of  $1SCN \vee 3SF$  with the trajectory of agent 1 (SCN) excluded (treated as invisible). Note that for the first 5 time steps, the SCN agents switch to SF in order to collect enough history to bootstrap the learning algorithm

For the evaluation stage, we train a separate evaluation model of the form

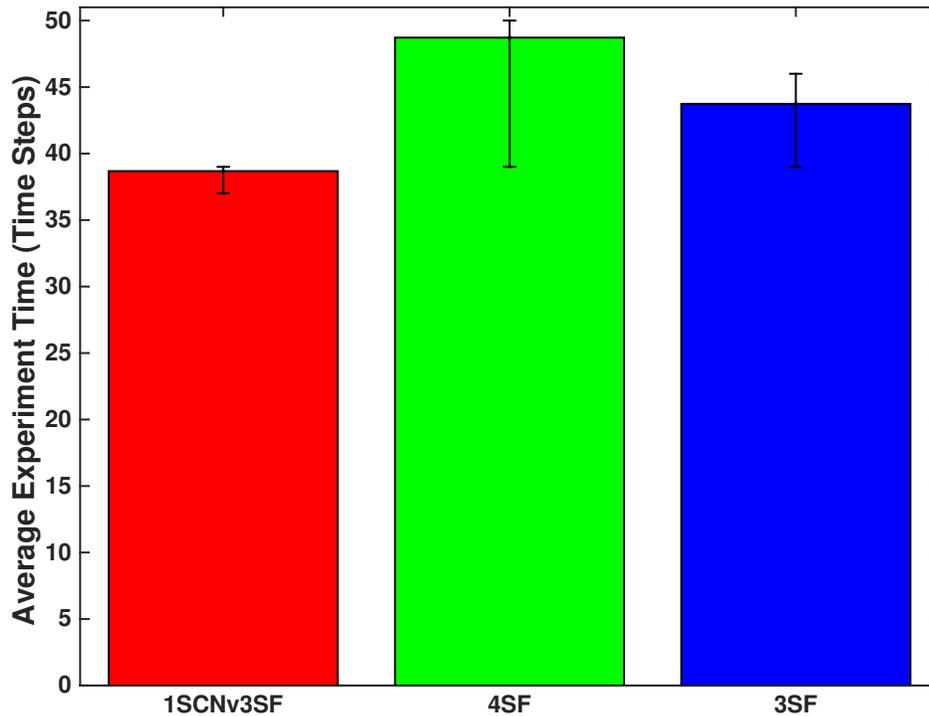
$P(\tau|M_T^i)$  that predicts future topologies  $\tau$  from the perspective of agent  $i$  based on the context  $M_T^i$  at time  $T$ . Both  $M_T^i$  and  $\tau$  are expressed with respect to frame  $F_i$  for agent  $i$  (see section 4.3.2). We use exactly the same model architecture and training procedure as detailed in section 4.3, but with a shorter input sequence that excludes the action  $a$ . This model achieves similar accuracy to  $P(\tau|M_T, a)$  (see Table 4.3).

The action set for the agent running SCN comprised a collection of 31 time-parametrized ( $dt = 0.2sec$ , speed  $1.2 m/sec$ ) straight line path segments of 3 waypoints each, covering  $\pi$  rads, whereas the weighting factor  $\lambda$  was set to  $\lambda = 0.6$ . For reference, the parameters for the agents running instances of  $SF$  were selected as  $v_{max} = 1.7m/sec$ ,  $c = 1$ ,  $\phi = 100^\circ$ ,  $R = 0.2m$ ,  $\sigma = 0.5m$ ,  $\tau_a = 0.4s$ ,  $V_{a\beta}^0 = 20m^2/sec^2$ ,  $U_{aB} = 10m^2/sec^2$ ,  $v_a^0 = 1.6m/sec$  and we kept the same time parametrization  $dt = 0.2sec$ .

## 4.5.2 Results

To demonstrate the benefits of the SCN algorithm for multi-agent navigation scenarios, we measure its effects in the behavior of other agents. More specifically, in each setup, we record (a) the time to destination (Figure 4.5) and (b) the evolution of the entropy of the distribution  $P(\tau|M_T^i)$  (Figure 4.6), both averaged over the same three SF agents (agents #2, #3, #4) per experiment.

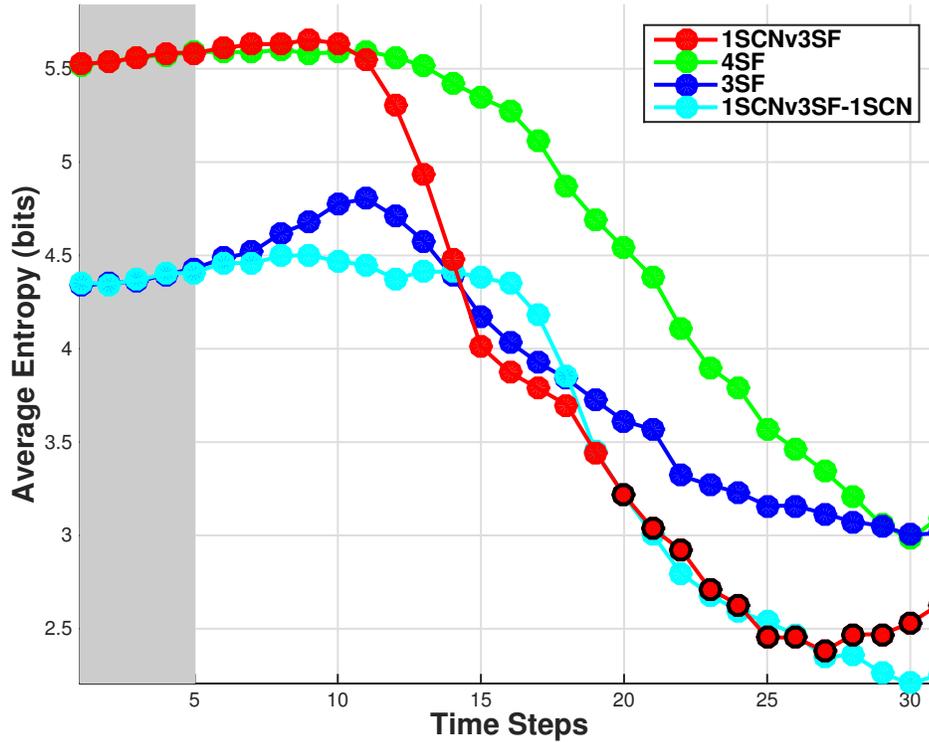
Figure 4.5 shows that the 1SCN $\nabla$ 3SF setup (Figure 4.4 depicts an example execution) achieved the fastest average time to destination, as a result of SCN’s consistently competent behavior. Student’s t-test yields a p-value  $< 0.001$  indicating a highly significant improvement, compared to 4SF and 3SF (Figure 4.5).



**Figure 4.5:** Average time to destination per experiment per agent (excluding agent #1). The error bars indicate 25th and 75th percentiles over 50 experiments. 1SCNv3SF is shown to terminate faster than both 4SF and 3SF, (according to a Student’s t-test with p-value < 0.001).

Note that the time to destination for the SCN agent itself was excluded from this test, indicating that the *three SF agents become more efficient thanks to the presence of one SCN agent.*

In Figure 4.6, at time zero, we see that the entropy of a uniform distribution in  $B_4$  (the braid group with  $n = 4$  strands) is naturally higher than the ones in  $B_3$  simply because of the higher baseline penalty for the probability mass being spread over more discrete possibilities. As we noted earlier, during the first five time steps, SCN agents in 1SCNv3SF run SF in order to collect enough data (grey box in Figure 4.6). Henceforth, they switch to SCN, which results in



**Figure 4.6:** Entropy profiles averaged across experiments and agents (excluding agent #1). The black circles indicate timesteps where the entropy measured for 1SCNv3SF is lower than for 3SF with statistical significance (Student’s t-test, with p-value  $\leq 0.022$ ). The gray area in the plot corresponds to the first time frame  $T_p$ , during which the agent running SCN was moving efficiently and observing the context.

a precipitous drop of the average entropy that continues until it drops below both baselines. In particular, in the time frame [25, 31] the entropy in 1SCNv3SF drops significantly below the entropy of 3SF, according to a Student’s t-test, with p-value  $\leq 0.022$ .

One could object that the SCN agent is an integral part of the braid, and it is therefore unsurprising that a socially competent agent reduces the system entropy. To measure the effect on the entropy of the other three agents alone, we introduce one additional baseline, 1SCNv3SF-1SCN. This result shows the

entropy of the system path for the three SF agents *after removing the SCN agent*. This result shows that the reduced entropy is not due to the direct contribution of SCN alone. Rather, the three SF agents are *themselves* behaving in a more orderly fashion in the presence of the SCN agent. This result suggests that in acting in a socially-competent manner, *SCN increases the social competence of SF agents as well*. For clarity, note that the graph in Figure 4.6 terminates before all of the agents have had an opportunity to quiesce at their goals. At that point, all entropies converge to zero. However, the benefit of social competence in terms of reduced time and confusion is achieved long before.

## 4.6 Discussion

We presented a planning framework for navigation in crowded environments. The foundation of our approach is a topological representation of the collective behavior of a set of agents, based on braids [20]. This representation forms the basis for the design of an inference mechanism that predicts the topology of the future trajectories of a set of agents, given the context of a scene. A model of this mechanism was extracted in a data-driven fashion by employing a deep neural network architecture on synthetic data generated through the use of the Social Force model [62]. The inference mechanism serves as a means of understanding how the agent’s behaviors might affect the observing others. This enables it to select behaviors that are *socially competent*, i.e., constitute the best response to the context. We conducted a set of simulated experiments that provided us with statistically significant evidence suggesting that our framework results in collective behaviors that simplify the planning problem for everyone in the scene. This is reflected in the behavior of other agents: systems of agents containing

an agent running our algorithm achieved significantly reduced average time to destination and were able to get a clear topological understanding of the scene evolution significantly faster, as shown in the average entropy profiles of the agents *not* running our model. Future work involves learning a model of system path topology prediction from human pedestrian data, experimental evaluation on a social robot platform and a study to assess its interactions with humans.

## **Part III**

**Planning by Reasoning about**

**Pairwise Collision Avoidance**

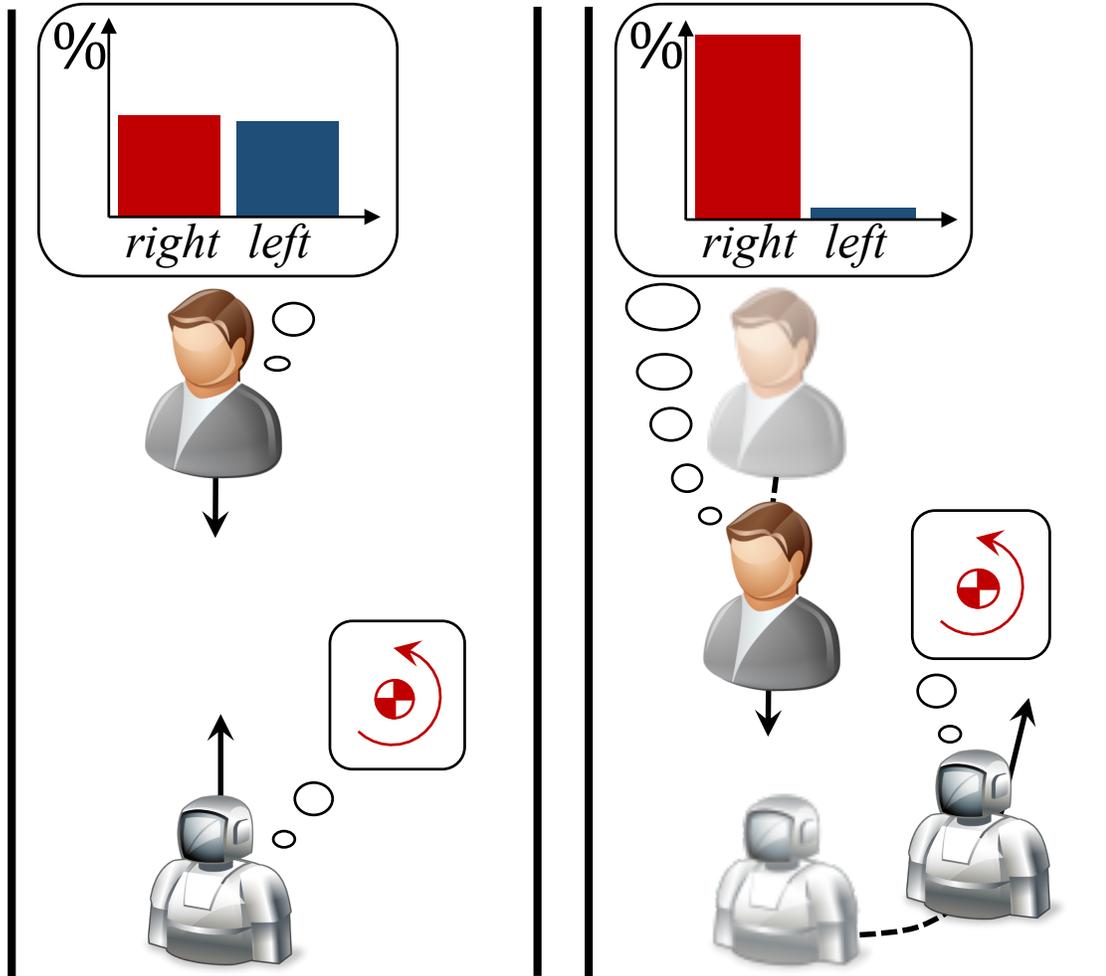
**Intentions**

## CHAPTER 5

### SOCIAL MOMENTUM: A FRAMEWORK FOR LEGIBLE NAVIGATION IN DYNAMIC MULTI-AGENT ENVIRONMENTS

The topological model introduced in Chapter 3 provides a navigation planner with significant expressive power, as the robot may enumerate, consider and anticipate any possible multi-agent navigation strategy. The high computational cost associated with considering large sets of outcomes highlights the need for a practical strategy of pruning. Chapters 3 and 4, introduced two different alternatives for pruning the infinite set of outcomes in consideration: one that makes use a short-horizon lookahead, and one that makes use of a learned inference mechanism, respectively. However, these methods still lack of ways to incorporate long-term intentions into robot actions but also to generalize across different domains.

In this chapter, we present a motion planner that makes use of an additional, physics-inspired way of pruning. Instead of considering a large set of possible multi-agent navigation outcomes, this new planner directly determines the most likely flow of the system and acts in a way that clearly communicates the robot's compliance with it. More specifically, our framework estimates the most likely intended avoidance protocols of others based on their past behaviors, superimposes them, and generates an expressive and socially compliant robot action that reinforces the expectations of others regarding these avoidance protocols (see Figure 5.1 for a graphic illustration of the main concept underlying our approach). This action facilitates inference and decision making for everyone, as illustrated in the simplified topological pattern of agents' trajectories. Extensive simulations demonstrate that our framework consistently achieves



**Figure 5.1: A human and a robot move towards opposing sides of a hallway. The initial configurations of the two agents make it hard for the human to predict the emerging avoidance strategy (“right” or “left”). The robot detects a slight advantage towards the “right” strategy and acts towards amplifying it and thus facilitating the inference of the human regarding the emerging (right) strategy of avoidance.**

significantly lower topological complexity, compared against common benchmark approaches in multi-agent collision avoidance. The significance of this result for real world applications is demonstrated by a user study that reveals statistical evidence suggesting that multi-agent trajectories of lower topological complexity tend to facilitate inference for observers.

More specifically, we contribute: (1) a novel topological consideration of leg-

ibility that captures the complexity of multi-agent dynamics; (2) a novel inference mechanism that reduces the trajectory prediction problem (continuous, infinite number of possible solutions) to the easier problem of topology prediction (discrete, bounded number of candidate solutions); (3) a cost-based planner, inspired by the cooperative nature of human navigation [148] and motivated by the goal-driven nature of human inference [33], that generates motion towards simplifying inference and planning for observers; (4) the introduction of a tool from low-dimensional topology for assessing the complexity of multi-agent trajectories and multi-agent motion planning; (5) extensive simulations demonstrating the topological and geometrical efficiency of our planner, compared against benchmarks in the area of multi-agent collision avoidance; (6) statistical evidence extracted from an online user study with human participants, demonstrating that executions of greater topological efficiency tend to be more *legible*, i.e., allow early and correct inference of agents' behaviors.

This work constitutes a new step in our investigations of the use of topological methods and tools for modeling multi-agent interactions in navigation and planning socially competent robot behaviors in multi-agent domains. Our past works made explicit use of topological braids to plan socially competent robot behaviors [97, 105]. This work makes use of topological braids as an analytical tool for assessing the quality of multi-agent planning.

## 5.1 Multi-Agent Trajectory Analysis

In a scene where multiple agents navigate towards their destinations while avoiding collisions with each other, their decision making over time may be represented as a geometrical pattern, formed by the spatiotemporal superposi-

tion of their trajectories. The topological properties of this pattern are indicative of the interactions that the agents have with each other throughout the course of the scene. In particular, the way agents avoid one another, by passing on the *left, right, front* or *back* of others results in an “entanglement” of their trajectories over time. Depending on the navigation strategies that agents follow, the complexity of this entanglement may range.

*We hypothesize that the complexity of the entanglement of the trajectories of multiple agents that navigate simultaneously a workspace is related to the planning effort they spend.*

In particular, we aim to show that legible behaviors in multi-agent navigation result in trajectory entanglements of low complexity. This will allow us to employ as a measure of topological complexity as a proxy for assessing the legibility of multi-agent behaviors. Based on the work of Carton et al. [25], who showed that legible behaviors in navigation reduce the required planning effort of navigating agents, we aim to show that trajectories of low topological complexity require low planning effort for agents.

In the following subsections, we present a tool for evaluating the topological complexity of multi-agent trajectories. This tool makes use of the representation of multi-agent navigation behaviors as topological braids, introduced in chapter 3.

### 5.1.1 Complexity of Braids

We are interested in quantifying the topological complexity of a multi-agent trajectory  $\Xi$ . Assuming that such a trajectory may be abstracted into a braid word  $\tau$ , with the method discussed in the previous subsection, an intuitive measure of complexity for  $\Xi$  would be the braid word *length*  $l(\tau)$ , defined as the number of generators that form it. In general, topologically complex braids correspond to braid words of larger *length* and equivalently, longer words indicate a higher topological complexity (see Figure 3.5 for intuition on the relationship between braid length and topological complexity). However, the amount of *entanglement* induced by distinct braids of same length may vary. For example, consider the braids  $b_1$  and  $b_2$  shown in Figure 5.2. Although  $l(b_1) = l(b_2)$ , it may be observed qualitatively that the entanglement of  $b_2$  is more intense (or less trivial) than that of  $b_1$ , due to the action of its third strand. This observation is an intuitive indication that the braid word length cannot be used unambiguously as a universal complexity index.

To overcome this degeneracy, Dynnikov and Wiest [42] introduce the braid *Complexity Index*. This index quantifies the amount of entanglement induced on a canonical *curve diagram* (Figure 5.3a) upon its application on it. An  $n$ -braid  $b \in B_n$  may be *applied* on a canonical  $n$ -curve diagram by sequentially enforcing the braid's topological pattern on the diagram's fields (e.g. Figure 5.3b, Figure 5.3c, Figure 5.3d). Intuitively, the canonical curve diagram can be thought of as a heterogeneous mixture, comprising  $n$  clearly separated substances. The application of a braid  $b$  can be thought of as the enforcement to the substances of a mixing pattern that follows the entanglement described by the braid  $b$ . In other words, the strings of the braid are matched to the substances of the curve



$E$  through the action of  $\beta$ . This is only defined up to isotopies fixing  $\partial D^2$  and  $P_n$ .

### 5.1.1.2 The Complexity Index

The *norm* of a curve diagram  $D$  is defined as the number of intersections of  $D$  with the real axis and denoted as:

$$\|D\| = \#(D \cap \mathbb{R}). \quad (5.1)$$

The *Complexity Index* of a braid  $\beta \in B_n$  is then defined through the use of its corresponding curve diagram  $D = \beta \cdot E$ , as:

$$c(\beta) = \log_2(\|\beta \cdot E\|) - \log_2(\|E\|). \quad (5.2)$$

This expression is equal to the logarithm of the *gain* of intersections with the  $x$ -axis, upon the application of a braid. Looking at Figure 5.3, it may be verified that the higher the number of intersections with the real axis, the higher the intensity the entanglement of the corresponding braid.

## 5.1.2 Complexity of Multi-Agent Planning

The *Complexity Index* quantifies the amount of entanglement that a braid induces to the canonical curve diagram. In this work, we employ braids as a representation of the topological pattern that a collection of trajectories forms. Thus, the *Complexity Index* may be used as a measure of the complexity of the braid corresponding to a specific execution of a multi-agent scenario. More importantly, it may serve as a characterization of the complexity of the solution to the mo-

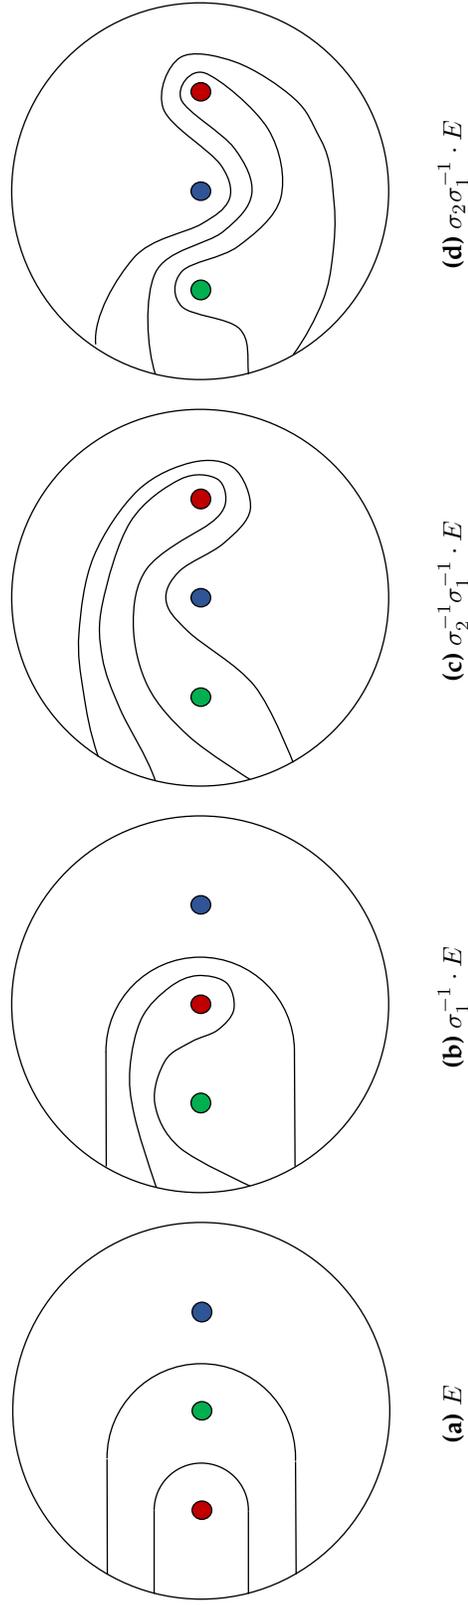


Figure 5.3: Curve Diagrams of braids from  $B_3$  on a disk with 3 punctures. From left to right, curve diagrams of: a trivial braid  $\beta_a = e$  (Figure 5.3a),  $\beta_b = \sigma_1^{-1}$  (Figure 5.3b),  $\beta_c = \sigma_2^{-1} \sigma_1^{-1}$  (Figure 5.3c) and  $\beta_d = \sigma_2 \sigma_1^{-1}$  (Figure 5.3d), with complexities  $c(\beta_a) = 0, c(\beta_b) = 1, c(\beta_c) = 1.5850, c(\beta_d) = 2$  respectively.

tion planning problem of transitioning safely from  $Q^s$  to  $Q^d$  that the navigating agents converged to, without explicitly communicating with each other.

The setup of our problem is quite simple: agents move between two resting positions (from start to goal). However, the lack of explicit communication among them may result in complicated collision avoidance maneuvers, yielding undesired oscillatory behaviors but also potentially undesirably long and entangled paths for one or more agents. The former problem, commonly referred to as the “reciprocal dance” problem [47], has been widely addressed in the literature (see e.g. Trautman et al. [138]). The latter problem though, to the best of our knowledge, has not been modeled or approached appropriately. In this work, we address it by (1) using the *Complexity Index* as a proxy to quantify the complexity of multi-agent planning, (2) proposing a planning framework that explicitly incorporates this understanding into the decision making stage, towards generating *legible* behaviors that reinforce plans of low trajectory entanglement and (3) investigating the effect of trajectory entanglement to human inference through a user study.

## 5.2 The Social Momentum Planning Framework

In this section, we present a planner that enables agents to contribute to trajectory patterns of low topological entanglement, as they navigate towards their destination. Our planner is based on a cost function that detects the intentions of other agents over pairwise collision avoidance protocols (e.g. right or left) and favors actions that are in compliance with them. The cost function, named the *Social Momentum* cost, is defined as the weighted sum of the magnitudes of

the pairwise angular momenta between the planning agent and all others. The optimization of the *Social Momentum* cost results in motion that tends to reinforce the current momenta by locally maximizing their magnitudes along their current directions, which corresponds to reinforcement of the currently established pairwise collision avoidance protocols between the planning agent and others. Our *Social Momentum* (SM) planning algorithm compromises between the *Social Momentum* cost and an *Efficiency* cost that drives the planning agent towards its destination. Throughout consecutive time steps, this policy results in a behavior that appears to be consistently compliant with the agent’s past behavior and with the preferences of others over avoidance strategies. Effective communication of the agents’ intended avoidance strategies results in behaviors that are easy to read and thus enable agents to implicitly cooperate efficiently to avoid each other, which leads to avoiding redundant trajectory entanglement.

### 5.2.1 Angular Momentum for Collision Avoidance

Consider two agents  $A$  and  $B$  moving towards opposing sides of a hallway, as shown in Figure 5.1. The geometry of the shared space renders agents’ decisions coupled. In order to reach their destinations in a collision-free and socially acceptable fashion, they need to (1) agree on a passing side (right or left) and (2) respect the personal space [58] of each other by maintaining a comfortable minimum distance. To quantify how well the agents are doing with respect to both of these specifications, we construct an analogy with the physical quantity of Angular Momentum. Assuming unit masses for the two agents, the angular momentum of their system with respect to its center of mass  $C$ , may be defined

as:

$$L^{AB} = r_A^C \times v_A + r_B^C \times v_B \quad (5.3)$$

where

$$r_A^C = q_A - r_C, \quad r_B^C = q_B - r_C \quad (5.4)$$

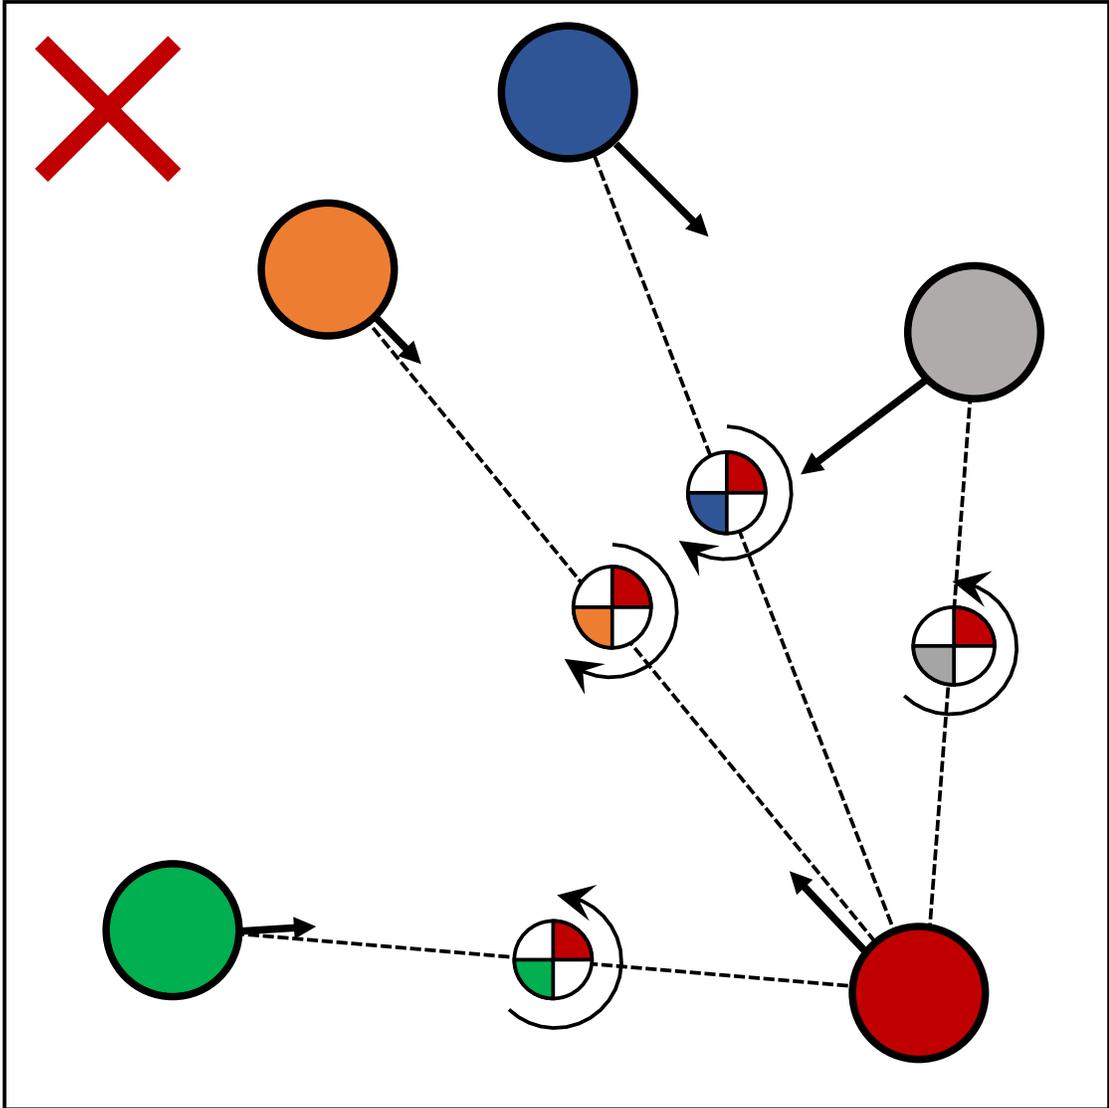
are agents' positions, defined with respect to their center of mass

$$r_C = (q_A + q_B) / 2. \quad (5.5)$$

For a system of agents navigating on the horizontal plane, the angular momentum is a vector perpendicular to the workspace, pointing along the positive direction of the  $z$ -axis for counterclockwise agent rotations and along the negative direction of the  $z$ -axis for clockwise rotations, thus encoding the right and left passings respectively. Its magnitude depends on the distance between the agents and also on the angle of their velocities, with larger distances and antiparallel velocities scoring higher. As a result, the angular momentum may be used (1) as a tool to monitor an emerging avoidance protocol (right/left passing) but also (2) as a tool to generate easily interpretable avoidance maneuvers in compliance with the preferences of the other agent and in consistency with previous behaviors of the agents.

## 5.2.2 Social Momentum for Legible Collision Avoidance

In a crowded multi-agent workspace, an agent interacts with multiple others at the same time, in the sense that every action taken broadcasts signals of intentions or preferences over avoidance strategies. Our framework enables an agent



**Figure 5.4: Social Momentum:** The planning agent (red color) is moving towards the red target X, while complying with its pairwise momenta with all other agents.

to read these preferences, associate them with its own, and act competently towards simplifying everyone’s decision making. To this end, we introduce a novel index, comprising a weighted sum of the magnitudes of all pairwise momenta between the planning agent and the set of all other agents  $N_i = N \setminus \{i\}$ . Higher values indicate a higher certainty over the emerging pairwise avoidance protocols between the agent and all others. We call this cost the *Social Momen-*

tum cost and formally define it for agent  $i$  as a real function  $\mathcal{L} : \mathcal{A} \rightarrow \mathbb{R}$  over the agent's action space  $\mathcal{A}$ , as follows:

$$\mathcal{L}(a) = \begin{cases} \sum_{j \in N_i} w_j \|\hat{L}^{ij}(a)\|, & \text{if } \text{sign} \left( (L^{ij})^T \hat{L}^{ij}(a) \right) > 0, \forall j \in N_i \\ 0, & \text{otherwise} \end{cases} \quad (5.6)$$

where  $\hat{L}^{ij}(a)$  denotes the expected pairwise momentum between agents  $i$  and  $j$ , upon agent  $i$  taking an action in consideration,  $a \in \mathcal{A}$  and agent  $j$  moving with its current velocity,  $L^{ij}$  is their current momentum and  $w_j \in \mathbb{R}$  is a weight, computed as the inverse of the distance between agents  $i$  and  $j$ . The quantity  $\text{sign} \left( (L^{ij})^T \hat{L}^{ij}(a) \right)$  indicates whether the expected evolution of the pairwise momentum between agents  $i$  and  $j$  is in compliance with their current momentum  $L^{ij}$ . A positive sign corresponds to an action that preserves the current momentum sign and thus the currently preferred pairwise avoidance protocol. A negative sign indicates inversion of the established pairwise avoidance protocol, which is undesired. For this reason, an action that results to inversion of a pairwise momentum is assigned a score of zero. Note that the only non-zero components of all pairwise momenta are their  $z$ -components, since we assume that the workspace is a horizontal plane.

---

**Algorithm 5**  $\text{SM}(Q, \mathcal{A}, d, \text{map}, \text{AtGoal}, a)$

---

**Input:**  $Q$  – current system state;  $\mathcal{A}$  – action set;  $d$  – agent's destination region;  $\text{map}$ ;  $\text{AtGoal}$  – boolean variable signifying arrival at agent's destination

**Output:**  $a$  – action selected for execution

- 1: **while**  $\neg \text{AtGoal}$  **do**
  - 2:      $\mathcal{A}_{cf} \leftarrow \text{Collision\_Checking}(\mathcal{R}, Q, \mathcal{A}, \text{map})$
  - 3:      $\mathcal{R} \leftarrow \text{Get\_Reactive\_Agents}(Q)$
  - 4:     **if**  $\mathcal{R} \neq \emptyset$  **then**
  - 5:          $a \leftarrow \text{Get\_Legible\_Action}(Q, \mathcal{A}_{cf}, \lambda, d)$
  - 6:     **else**
  - 7:          $a \rightarrow \text{Get\_Efficient\_Action}(\mathcal{A}_{cf}, d)$
  - 8: **return**  $a$
-

### 5.2.3 Decision Making

In this section, we present the *Social Momentum* (SM) algorithm, a cost-based planning algorithm, built around the *Social Momentum* heuristic. The algorithm is based on frequent replanning; at every planning cycle, it picks an action that corresponds to the optimal compromise between progress to agent’s destination and legible avoidance of others. We formalize this decision making strategy into the following optimization scheme:

$$a^* = \operatorname{argmax}_{a \in \mathcal{A}} \{\lambda \mathcal{E}(a) + (1 - \lambda) \mathcal{L}(a)\}, \quad (5.7)$$

where  $\lambda \in \mathbb{R}$  is a parameter accounting for proper scaling and weighting of the two quantities. We model the progress function  $\mathcal{E} : \mathcal{A} \rightarrow \mathbb{R}$  to be the inverse of the length of the unobstructed line to destination. The action space  $\mathcal{A}$  comprises a pre-sampled set of actions of finite duration that are executable by the agent.

Algorithm 5 describes the SM algorithm in pseudocode format. At each replanning cycle, the function `Collision_Checking` checks for collisions with other agents or bounds and returns a set  $\mathcal{A}_{cf} \subseteq \mathcal{A}$  of collision-free actions. Then function `Get_Reactive` determines the subset of agents  $\mathcal{R}$  to which the planning agent should be reacting: only agents that lie in front of the planning agent are considered (see Figure 5.4). In case  $\mathcal{R} \neq \emptyset$ , the planning agent determines a legible action  $a$  by compromising between *Progress* to destination and *Social Momentum* (function `Get_Legible_Action`); otherwise, the algorithm switches to progress maximization mode (function `Get_Efficient_Action`). Termination occurs once the algorithm reaches a desired distance to destination.

## 5.3 Evaluation

In this section, we evaluate our approach by investigating the following two hypotheses:

1. “The Social Momentum Framework produces multi-agent trajectories of significantly simpler topological entanglement than existing approaches of multi-agent planning”.
2. “In multi-agent domains, multi-agent trajectories of simple topological entanglement are more legible”.

Confirmation of these two hypotheses provides evidence that the Social Momentum framework results in legible navigation in multi-agent environments. To test hypothesis (1), we perform an extensive simulated evaluation. To test hypothesis (2), we conduct a study in which we ask users to predict the evolution of simulated multi-agent scenarios from partial observation.

### 5.3.1 Simulations

We evaluate our planner in simulation by comparing against common benchmarks in multi-agent planning. Specifically, we consider the *Social Force* (SF) model [62] and the *Optimal Reciprocal Collision Avoidance* (ORCA) framework [142]. It should be noted that these frameworks were designed to produce fast and realistic simulations of multi-agent navigation scenarios and not to produce legible behaviors. However, they still constitute relevant baselines due to (a) their wide dissemination and existence of ready implementations, (b) their

proven performance in various types of scenarios, (c) the fact that they constitute common benchmarks and thus common works of reference, (d) their ability to handle any number of agents.

### 5.3.1.1 Experimental Setup

We consider multi-agent scenarios involving sets of homogeneous agents navigating a circular workspace. Each scenario is generated as follows: (1) the workspace circumference is partitioned to  $n$  arcs of equal length, (2) each arc is assigned to an agent, (3) the agent is placed at a random, collision-free position on the arc assigned to it and (4) each agent is assigned a destination that is antipodal to its starting location and lies on the workspace circumference. The workspace considered is a circle with a diameter of  $5m$ , whereas the agents were discs of diameter  $0.6m$ . This setup was specifically selected as it reinforces the emergence of challenging agent encounters. The occurrence of such scenarios highlights the value of intent-expressiveness as a feature that reduces uncertainty by reinforcing implicit coordination. We consider 4 different classes of scenarios, each corresponding to a different number of agents, ranging from 3 to 6. For each class, we generate 200 intense multi-agent scenarios at random and execute them with each of the planners considered. Note that each of the planners considered can be tuned to yield qualitatively different behaviors. In order to ensure a fair comparison, we assumed similar tunings with respect to sensitivity to obstacles and kept them constant across all trials.

### 5.3.1.2 Quality Measures

To evaluate the quality of an execution, we consider two different criteria: (1) the braid *Complexity Index*, which serves as a measure of the topological efficiency of the execution and (2) the *Path Irregularity* index [57], defined as the total amount of unnecessary rotation (divergence of agent’s heading from its direction to its destination) per unit path length, averaged per agent, which serves as a measure of the geometrical efficiency of the execution. *Complexity Index* computations were implemented using the BraidLab package [132], assuming a projection onto the global  $x$ - $t$  coordinate plane.

In the evaluation of the topological entanglement, we also include a theoretical *Lower Bound* baseline, which returns a topological path of minimal topological entanglement that executes the scenario in consideration. This baseline may be described as follows: (1) find the minimal path of transpositions,  $\pi^*$  that connects the initial permutation of the system,  $p^s = f_x(Q^s)$ , with the final permutation of the system,  $p^d = f_x(Q^d)$  and (2) derive a sequence of generator transitions for all consecutive waypoints in  $\pi^*$ , i.e., a braid  $\beta^* \in B_n$  that yields the lowest *Complexity Index*  $c(\beta^*)$  for the scenario. This baseline can be thought of as an ideal case of perfect communication and compliance (or centralized planning).

### 5.3.1.3 Performance Comparison

Figure 5.5 depicts the average *Complexity Index* for each planner and class of scenarios considered. The *Complexity Index* of SF and ORCA appears to be consistently rising with the number of agents. In contrast, SM exhibits a slower rise;

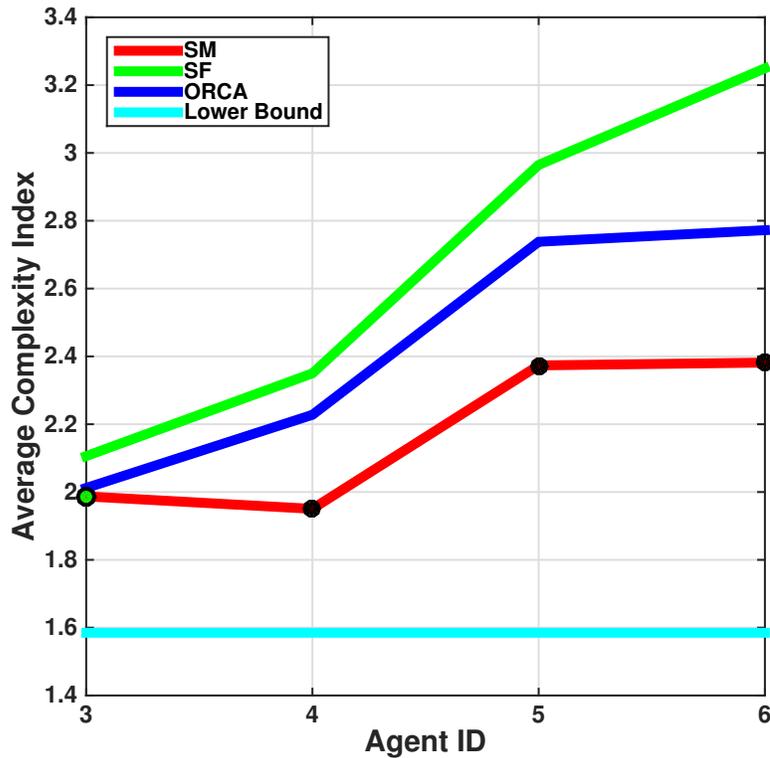


Figure 5.5: Average *Complexity Index* of trajectories generated by executing 200 scenarios with 3, 4, 5 and 6 agents with the *Social Momentum* (SM), *Social Force* (SF) and *Optimal Reciprocal Collision Avoidance* (ORCA) models. A theoretical lower bound baseline is also included for reference. Datapoints marked black correspond to significantly lower average *Complexity* of SM than both SF and ORCA, whereas the datapoint marked green indicates significantly lower average *Complexity* of SM than SF, according to paired Student’s T-test. Test statistics can be found at table 5.1.

the transitions between 3 and 4 agents and between 5 and 6 agents are done with almost constant complexity, with the only rise taking place in the transition between 3 and 4 agents. Overall, SM achieves consistently lower topological entanglement with statistical significance, except from the case of 3 agents, where the scenarios are not geometrically challenging to yield significantly diverse behaviors. Detailed statistics of paired t-tests conducted for the SM-SF and SM-ORCA pairs are reported in table 5.1. Despite this result, the theoretical

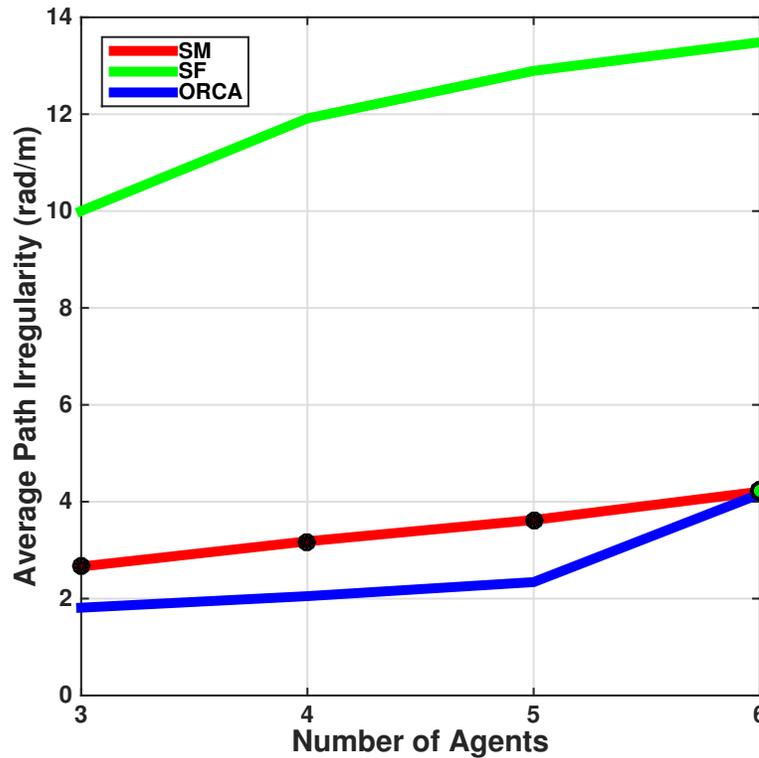


Figure 5.6: Average *Path Irregularity* per agent extracted by executing 200 scenarios with 3, 4, 5 and 6 agents. Datapoints marked black denote different irregularity of SM than both SF and ORCA, whereas the datapoint marked green indicates lower irregularity of SM than SF, according to paired Student's T-tests. Test statistics can be found at table 5.1.

Lower Bound consistently outperforms all planners, providing an illustrative demonstration of their suboptimality in terms of *topological efficiency* which reflects the price of no explicit communication in multi-agent planning. Note that the constant *Complexity Index* value of 1.5850 that the Lower Bound achieves is an artifact of the symmetry of the considered scenarios (agents traveling to antipodal points in the workspace).

Figure 5.6 depicts the average *Path Irregularity* per agent, for each planner and class of scenarios considered. Although for all planners the path irregularity rises with the number of agents, each planner performs differently. The

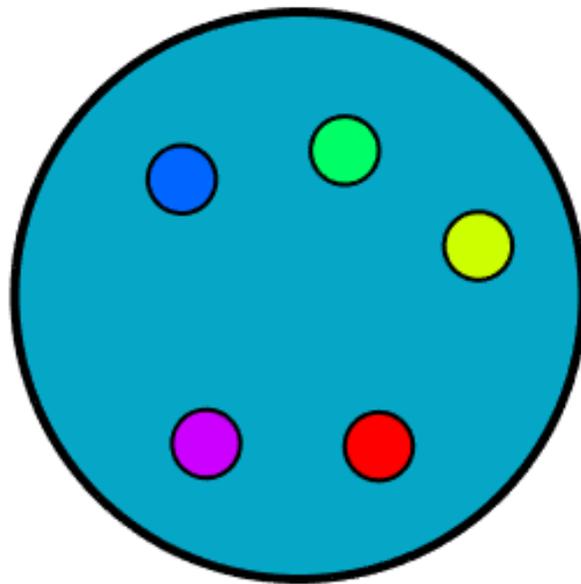
t-Tests		<i>Complexity Index</i>		<i>Path Irregularity</i>	
Agents	<i>Pair</i>	<i>t-value</i>	<i>p-value</i>	<i>t-value</i>	<i>p-value</i>
3	SM-SF	-2.497	0.013	-26.397	< 0.001
	SM-ORCA	-0.593	0.553	9.197	< 0.001
4	SM-SF	-7.963	< 0.001	-34.514	< 0.001
	SM-ORCA	-5.740	< 0.001	17.336	< 0.001
5	SM-SF	-9.424	< 0.001	-41.400	< 0.001
	SM-ORCA	-5.395	< 0.001	7.934	< 0.001
6	SM-SF	-11.561	< 0.001	-51.430	< 0.001
	SM-ORCA	-5.250	< 0.001	0.152	0.879

**Table 5.1: Statistics of paired t-tests between SM and SF, ORCA for different agent numbers. We considered  $N - 1$  degrees of freedom, where  $N = 200$  is the number of scenarios per class.**

different performance of each planner is indicative of the distinct philosophy with which they have been designed. SF, lacking predictive mechanisms yields significantly more irregular paths than SM and ORCA. ORCA achieves consistently the lowest path irregularity, as a result of its geometrically optimal behavior, which in practice results in minimal divergence from the unobstructed line connecting an agent with its destination at any time. SM performs slightly worse ORCA, as a result of its different consideration of collision avoidance as a rotation; SM agents diverge from their shortest paths more often to convey intent. For the case of 6 agents the geometric complexity of the scenarios is too intense even for ORCA which performs almost equally to SM.

### 5.3.2 User Study

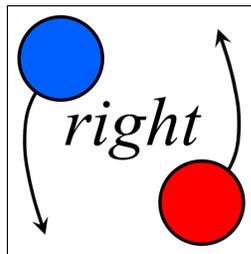
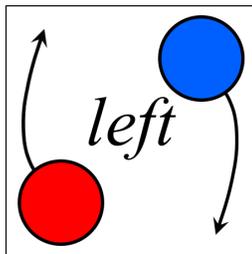
We conducted a user study, in which we asked users to watch a series of videos (shown from a top view) of simulated executions of scenarios involving 5 agents navigating a circular workspace. For each video, users were asked to predict



+2 points!

✓ Correct!

**You answered in 1.818 seconds! That's faster than the median time!**



**Score: 6 points**

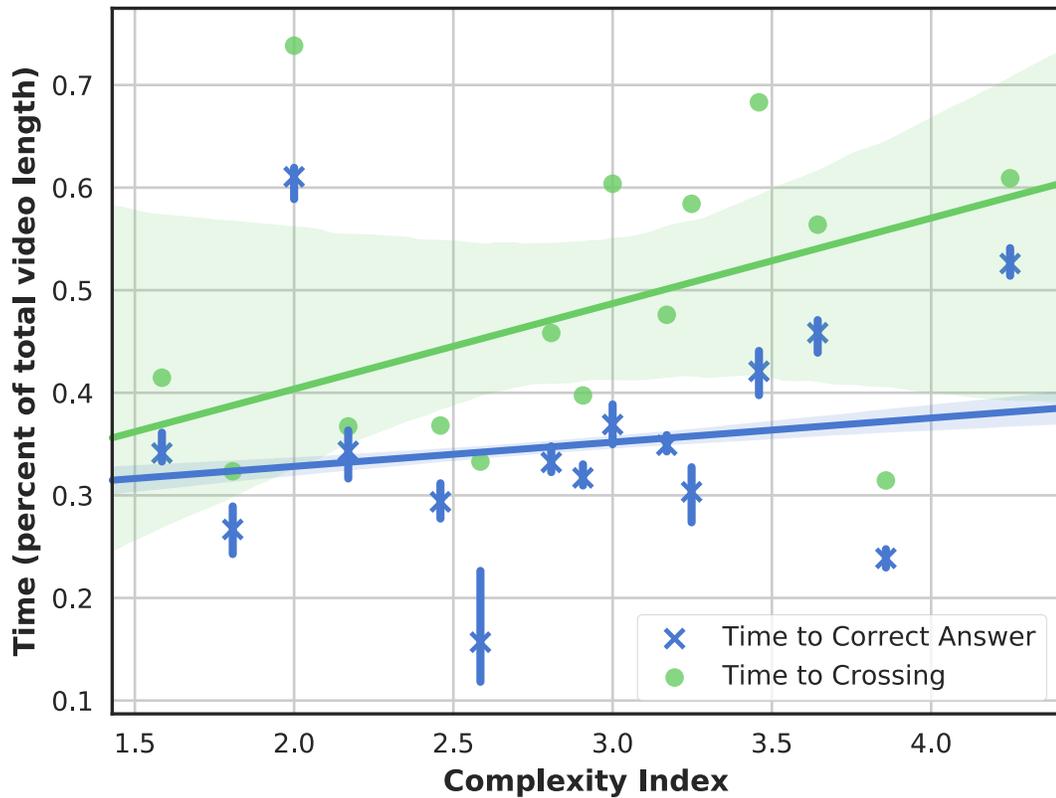
**Figure 5.7: Study interface:** A video of a scenario execution is shown and users predict how the red agent is going to avoid the blue agent by pressing the corresponding button at the bottom. The display of user's score and performance statistics aim to incentivize fast and accurate responses.

the way two agents were going to avoid each other (right or left side). Speed and correctness (the basis of the legibility definition) were incentivized through a scoring system that awarded points for quick and accurate answers and deducted points for wrong or slow responses (Figure 5.7 depicts the study interface). The study used a total of 15 videos, with duration ranging from 6.3 to 15.7 seconds, corresponding to scenarios of *Complexity Index* uniformly ranging from 1.585 to 4.250. More than 180 users, recruited from the social media platforms of Reddit and Facebook, contributed a total of 2704 video views and clicks. An analysis of the collected dataset is presented in Figure 5.8.

The blue trend shows the relation between the *Complexity Index* and the median time of correctly predicting the side on which one agent will pass another. We fit a linear model to the data using iteratively reweighted least squares, shown in Figure 5.8 as a blue line with a 95% confidence interval. The effect of the *Complexity Index* on click time is positive, with a slope of 0.0236, and significant (Student's t-test,  $t = 5.60$ ,  $p < 0.001$ ). In other words, as the topological entanglement intensifies, users take more time to accurately predict the side of passing, i.e., more complex scenarios are less legible. We verified that the rate of incorrect answers for a video is not correlated with the *Complexity Index* of that video via computation of a Pearson correlation coefficient ( $r^2 = 0.1017$ ,  $p = 0.7185$ ).

The green trend shows the relation between the *Complexity Index* and the time of passing between the two agents. We fit a linear model to the data, shown as the green line with a 95% confidence interval. The trend is positive (slope of 0.0833) and nearly significant ( $t = 1.93$ ,  $p = 0.0538$ ). Increased *Complexity Index* correlates positively with increased time of passing, and thus with longer, less efficient interactions. For each video, most users were able to correctly predict the passing side before the passing occurred. We see a trend towards predicting the passing with greater lead times for greater *Complexity Index*. However, we argue that the measure of legibility that matters is not lead time but lag time after navigation begins before the user is able to predict the correct passing.

Our simulated evaluation confirmed our first hypothesis by revealing statistical evidence that SM achieves consistently lower trajectory entanglement than other representative multi-agent collision avoidance approaches. Furthermore, our user study confirmed our second hypothesis by demonstrating a positive



**Figure 5.8: Relation between the *Complexity Index* and (a) time until two specific agents pass each other (green points/line) and (b) median time until users give a correct prediction of the passing (blue crosses/line). Times are normalized to the total length of the relevant video.**

correlation between trajectory entanglement and the time taken for a correct prediction of a passing between two agents. In other words, *executions of lower trajectory entanglement are more legible*. From these, we may assert that the Social Momentum framework appears to produce legible behaviors in multi-agent environments. This feature is of particular importance for robots navigating crowded human environments, where no explicit communication takes place among agents and no formal rules are guiding traffic, such as pedestrian environments.

## 5.4 Discussion

We presented a planning framework for legible motion generation in multi-agent environments. We approached legibility from a topological perspective and introduced the concept of *legible avoidance* as a maneuver that clearly indicates the way an agent plans to avoid another (e.g. *right* or *left* side). Based on this idea, we designed the *Social Momentum* planning framework that enables agents to generate intent-expressive and socially compliant behaviors in multi-agent environments. Statistical evidence, extracted from extensive simulations and from a user study with human participants demonstrated the ability of our framework to produce legible behaviors in multi-agent environments. This result is particularly important for operation in human environments with no explicit communication among agents, such as pedestrian environments. Ongoing work involves validating our approach by conducting experiments on an autonomous social robot, navigating academic hallways.

## CHAPTER 6

### EFFECTS OF DISTINCT ROBOT NAVIGATION STRATEGIES ON HUMAN BEHAVIOR IN A CROWDED ENVIRONMENT

State-of-the-art social robot navigation algorithms often lack a thorough experimental validation in human environments: simulated evaluations are often conducted under unrealistically strong assumptions that prohibit deployment in real world environments; experimental demonstrations that are limited in sample size do not provide adequate evidence regarding the user experience and the robot behavior; field studies may suffer from the noise imposed by uncontrollable factors from the environment; controlled lab experiments often fail to properly enforce challenging interaction settings. In this chapter, we contribute a first step towards addressing the outlined gaps in the literature. We present an original experiment, designed to test the implicit interaction between a mobile robot and a group of navigating human participants, under challenging settings in a controlled lab environment. We conducted a large-scale, within-subjects design study with 105 participants, exposed to three different conditions, corresponding to three distinct navigation strategies, executed by a telepresence robot (two autonomous, one teleoperated). We analyzed observed human and robot trajectories, under close interaction settings and participants' impressions regarding the robot's behavior. Key findings, extracted from a comparative statistical analysis include: (1) evidence that human acceleration is lower when navigating around an autonomous robot compared to a teleoperated one; (2) the lack of evidence to support the conventional expectation that teleoperation would be humans' preferred strategy. To the best of our knowledge, our study is unique in terms of goals, settings, thoroughness of evaluation



**Figure 6.1: This study examines the performance of autonomous robot navigation algorithms in crowded pedestrian spaces.**

and sample size.

State-of-the-art autonomous navigation frameworks have achieved impressive benchmarks in simulation and to exhibit competent behaviors in experimental demonstrations, field studies, and lab experiments. However, their validation is often not sufficiently rigorous and in-depth. Simulated evaluations are inevitably conducted under strong assumptions on the type of the environment,

the context, and the type of behavior exhibited by other agents; thus addressing the *reality gap* problem is not a trivial extension. Experimental demonstrations contribute a significant step towards deployment to the real world but lack a significant sample size of repeated interaction with human users and thus statistical power. Large-scale field studies are an important step in the validation process of any robotic system as they may provide evidence of robust performance under challenging settings. Nonetheless, the noise induced by the frequently massive complexity of a real-world environment may prohibit the extraction of concrete conclusions about the performance of the robot and the user experience. Lab studies may definitely isolate the system from external variables and enable rigorous testing of the desired conditions. However, designing an experiment that will isolate the desired nontrivial interaction between a target system and human participants is a not an easy task and it may often be observed that lab experiments with mobile robots do not test a challenging type of setting. Finally, ensuring the repeatability of the performance of an autonomous robotic system, exposed to close interaction with humans is also not a trivial task and often requires frequent maintenance and high costs. Thus, to approach interesting research questions without the complication of exhaustively testing the autonomy of the platform, a significant amount of research considers only Wizard of Oz experiments [118]. While the findings of such experiments are often of great significance regarding the human-robot interaction, they inevitably leave

a gap in the validation of the autonomy itself.

## 6.1 Contributions

In this chapter, we contribute a step towards addressing the outlined gap in the validation process of social navigation planning algorithms. We present an original experiment design, constructed to enforce naturally a series of challenging implicit interactions between a mobile robot and a group of human participants that navigate in a shared workspace in a controlled lab environment. A simple background scenario serves as a driving force towards interesting, nontrivial interactions but also as a way to cognitively load and distract human subjects from the goal of the experiment. We conducted an extensive, large-scale ( $N = 105$ ), within-subjects user study in which we recorded 945 minutes of interaction between a mobile robot and human subjects (3 at a time). Participants are exposed to a set of three distinct robot navigation strategies, executed by a telepresence robot platform. We considered two autonomous navigation strategies and a teleoperated condition in which a human teleoperates the robot.

We collected human and robot trajectory data, recorded by an overhead motion capture system, and responses to a questionnaire designed to assess participants' impressions of the robot's intelligence, social compliance, and safety. We performed comparative statistical analyses on the collected datasets and report the results. Key findings, extracted by focusing on the close interactions (distance  $< 1m$ ) between the robot and human subjects include the following: (1) human accelerations are significantly lower around an autonomous robot executing Social Momentum [107] than around a teleoperated one; (2) contrary to our expectations, we found no evidence to support the hypothesis that hu-

mans prefer the human-teleoperated navigation strategy—in fact humans did not distinguish between conditions in their ratings; (3) teleoperated motion that follows the same high-level rules as autonomy results in lower topological complexity [42] than autonomy, an observation potentially reflecting the more global character of human decision making for navigation tasks.

## 6.2 User Study

In this chapter we present an IRB-approved (approval code: 1805008009) lab study, focused on the evaluation of a set of distinct robot navigation algorithms with respect to social compliance. The lab environment allows us to have significant control over variables that can interfere with the experimental setting. We leverage this level of control to enforce challenging navigation behaviors in a natural fashion through the design of an original experiment scenario and task.

We enforce a setting of implicit, nonverbal social engagement among agents, similar to the type of interaction among walking pedestrians so that we can study phenomena involving collaborative collision avoidance processes, as observed by Wolfinger [148]. Furthermore, we construct a moderately crowded scene that balances close interactions with space for the robot to showcase its distinct navigation strategies (see Figure 6.1). We also ensure the emergence of nontrivial interactions, involving challenging collision-avoidance maneuvers between participants and the robot through the definition of rules. Moreover, we motivate natural walking behaviors by not disclosing the real purpose of the study until the debriefing process and by increasing participants' cognitive load through the background scenario and task. Finally, we consider three conditions

and keep the total duration less than thirty minutes to facilitate recruiting and minimize potential effects resulting from participants' fatigue.

### **6.2.1 General Experiment Procedure**

Our study is organized into a set of experiment sessions. In each session, three different human subjects participate in a set of three experiment trials. Before the first trial, participants are asked to give written consent to confirm their participation and optionally to be video recorded. A member of our research team delivers the instructions and answers questions. During each trial, participants repeatedly visit a set of stations inside a rectangular workspace of area  $16m^2$  (see Figure 6.3), driven by a fictional scenario. In parallel, a mobile robot (a Suitable Technologies Beam Pro, equipped with a quad core i7 processor laptop from 2017), shown in Figure 6.2, also moves between the stations. During each trial, the human and robot trajectories are tracked and recorded through an overhead motion capture system of six high-accuracy ( $< 1$  mm), high-fidelity (frequency 180 Hz) cameras and videotaped if participants gave consent. Real-time tracking was enabled through the use of construction helmets (Figure 6.2) with reflective markers. After each trial, participants are asked to fill in a questionnaire, containing questions about their impressions from their interactions with the robot. At the end, participants are asked to provide basic demographic data and information regarding their prior experience with user studies and robotics technology. Participants are then debriefed, compensated and dismissed.

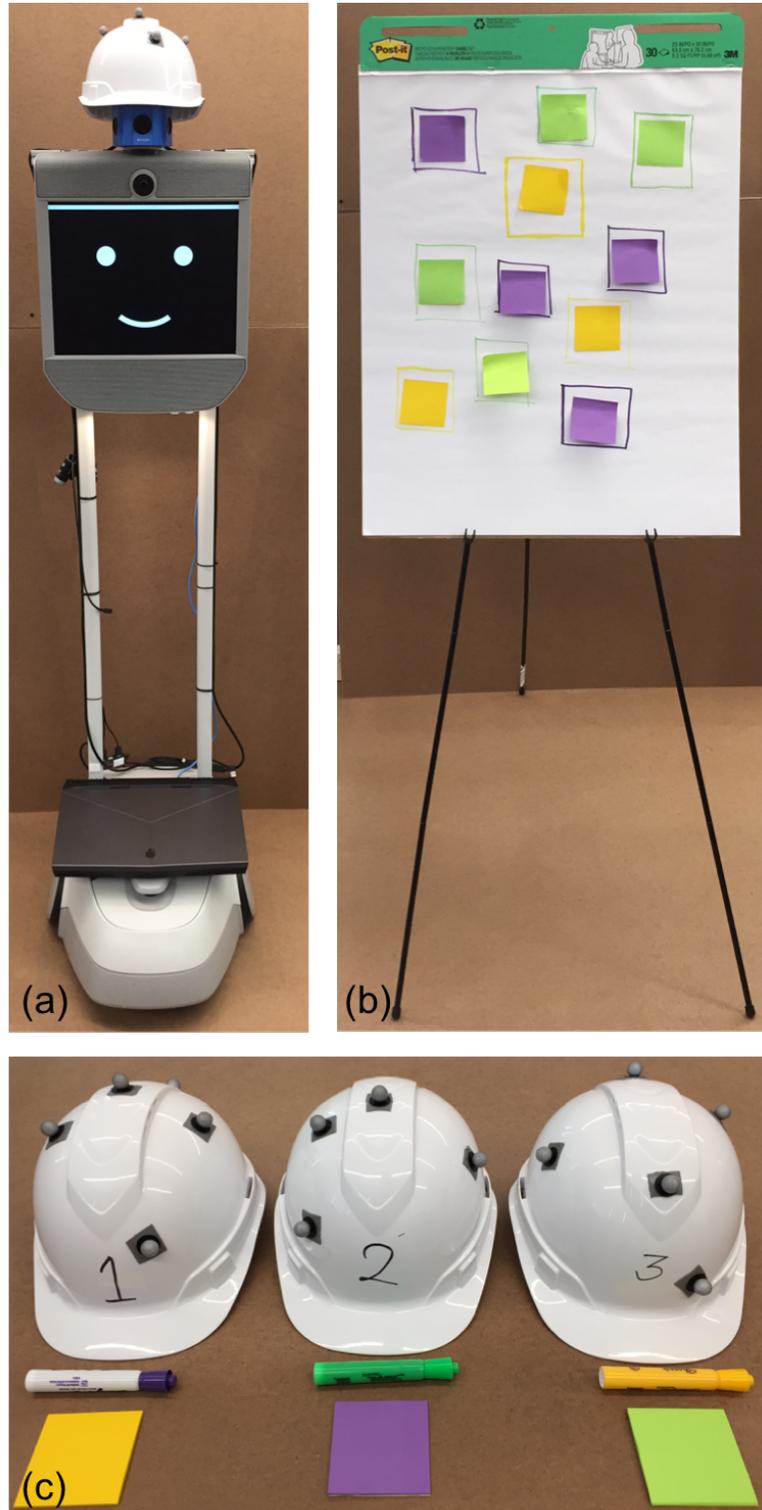


Figure 6.2: Study apparatus: (a) Beam Pro robot [1]; (b) an easel, representing a machine with sticky notes, representing completed maintenance tasks; (c) tracking helmets, sticky notes and markers distributed to participants.

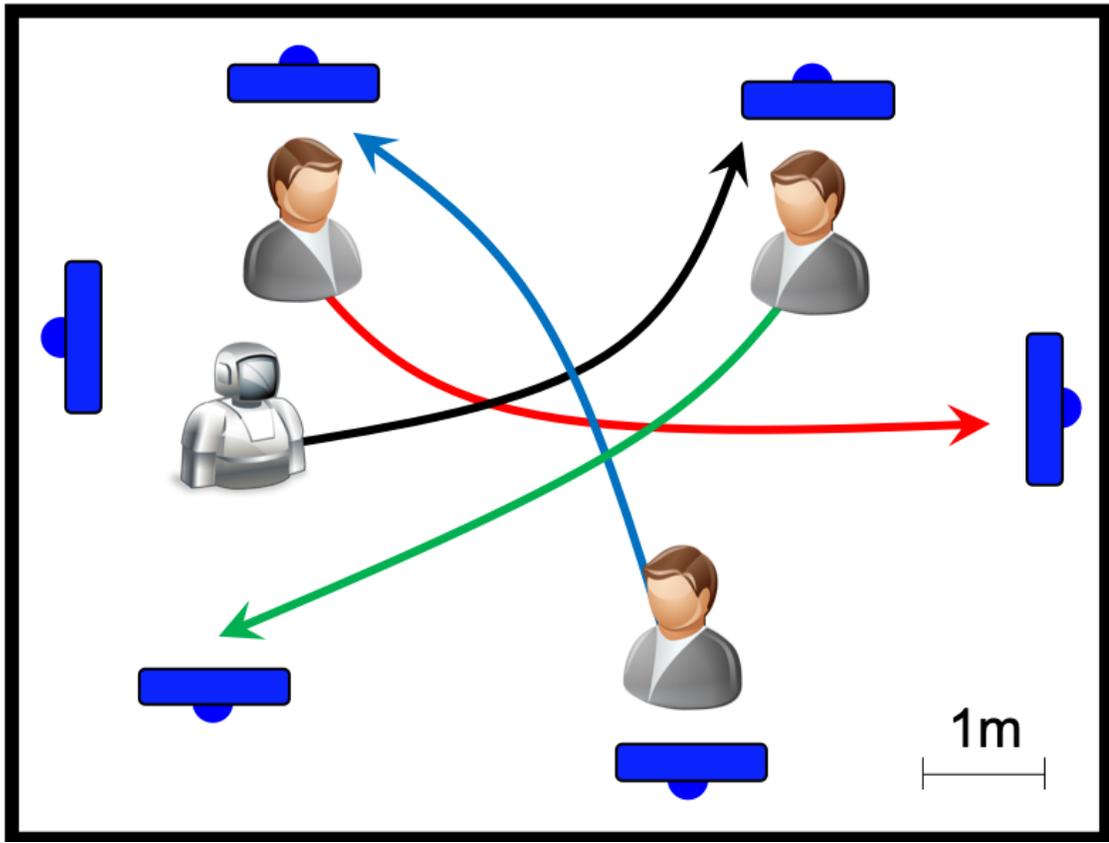


Figure 6.3: Top view of the workspace along with example human and robot trajectories, corresponding to transitions between easels (blue objects).

## 6.2.2 Background Scenario and Task Description

Participants are asked to imagine that they are workers in a factory (the factory setting helps justify the tracking helmets) and the robot is a supervisor. The factory environment (lab workspace) contains six machines, represented as easels, spread around the workspace, as shown in Figure 6.3. Each worker is given a distinctly colored marker and a contrasting, distinctly colored set of sticky notes. The duty of a worker is to **perform** maintenance tasks to machines and **assign** tasks for other workers to perform. Assigning a task is done by drawing a square on the pad of an easel, whereas performing a task is done by posting a sticky note inside a square drawn on an easel pad. Participants are asked to

perform only tasks represented with squares of color that matches the color of their sticky notes. Figure 6.2 depicts the main units of equipment used in this study.

### **6.2.3 Trial Description**

Before the start of each trial, participants are randomly positioned next to different machines and the robot is placed in the middle of the workspace, as shown in Figure 6.3. A trial is organized into a set of maintenance cycles, initiated by a gong sound, played by the robot. Each time the gong is played, participants are instructed to leave their machines towards a non-adjacent machine of their choosing. Each time participants reach a new machine, they are instructed to perform up to one pre-assigned task (if one exists) and assign a new task. At the same time, the robot is navigating in the workspace by following the same rules of transitioning between stations, i.e. it only moves to a randomly picked, non-adjacent machine when the gong sound is played. For synchronization purposes, the gong sound is played when the robot is ready to move towards its next machine. Each trial lasts exactly three minutes, during which an ambient factory sound track is played.

### **6.2.4 Conditions**

All participants were exposed to the same three conditions (within-subjects design), each corresponding to a different navigation strategy, executed by the robot. To account for potential ordering effects (i.e., due to fatigue, frustra-

tion, learning), the condition order was methodically varied and approximately equally spread across all sessions. The selected set of navigation strategies consists of *Optimal Reciprocal Collision Avoidance* (ORCA) [142], *Social Momentum* (SM) [107]) and teleoperation (TE). These strategies<sup>1</sup> were mainly selected due to the diversity of decision making principles that they represent, i.e., ORCA is designed to be optimal, SM is inherently intention-aware; TE is designed to appear humanlike. Additional reasons that influenced our selection included: (1) the fact that ORCA constitutes a common benchmark and work of reference for multi-agent simulations (e.g. [28, 50, 86, 107]); (2) the existence of an open source, optimized C++ implementation of ORCA; (3) the ease of implementation of SM; (4) the widespread use of telepresence robot platforms through teleoperation via their navigation interfaces. The complexity of a real-world pedestrian environment would pose a significant challenge to any of these navigation planners. However, we believe that an extensive and comparative evaluation of planners with distinct philosophies provides us with significant insights and experience for the design of the next generation of social navigation planning algorithms. The following paragraphs provide short descriptions of the mechanisms underlying the selected navigation strategies.

*Optimal Reciprocal Collision Avoidance* (commonly referred to as ORCA; in the results section we will be using the codename OR for brevity) [142] is a decentralized navigation planning framework for the generation of smooth, collision-free, natural-looking simulations of multi-agent scenarios. It is an optimization-based approach that determines the velocity of minimal divergence from an agent’s desired velocity that is guaranteed to be collision-free for a desired time

---

<sup>1</sup>Note that the selected robot features differential drive steering; thus to employ the autonomous algorithms considered, we make use of a standard feedback linearization technique [77].

horizon. This approach makes local collision avoidance considerations by incorporating a model of intentions, based on agents' current velocities. It operates however, under the assumption that other agents also run the same decision making mechanism to guarantee safe and smooth behaviors.

*Social Momentum* (SM) [107] is a decentralized, cost-based, navigation planner, designed to generate legible robot motion in multi-agent environments. The cost is a weighted sum of two functions: one representing efficiency and one representing social compliance. At planning time, the robot selects and executes the action that contributes the best compromise between the two costs. This policy results in consistent progress towards an agent's destination while taking into consideration the collision-avoidance intentions and preferences of other agents. Unlike ORCA, SM does not explicitly assume that others run the same policy; instead, it focuses on reading the intentions of others and incorporates this knowledge into its motion planning process.

The *Teleoperation* strategy (TE) was implemented through the official navigation interface provided by the manufacturer [1], using the arrow keys on a standard laptop keyboard. This interface contains two live streams of video, providing the teleoperator with real-time video streams of a forward, wide-angle field of view (top) and a floor view (bottom). Navigation commands may be executed through a simple keyboard's arrow keys (or with a mouse). Selected commands are demonstrated as projected future trajectories on the video streams, providing visual feedback to the user. The teleoperation condition was executed by the same member of our research team across all sessions, from a remote location (outside of the lab). The teleoperator had significant prior experience of the navigation interface for several years. Before collecting data for our final dataset,

we completed a total of 7 rounds of pilot sessions under different variants of the final study setup. Thus by the time we officially started the study, the teleoperator had reached a skill level that qualitatively appeared to be appropriate for the needs of the condition. Although it is hard to precisely quantify the operator's skill level, his experience was in the order of several hours prior to the start of the study and thus we do not believe that his performance evolved over the course of the study as a result of learning.

### **6.2.5 Hypotheses**

Upon experimenting with the three navigation strategies considered (simulations conducted with SM and ORCA, personal teleoperated teleconference sessions with the Beam), we observed very different patterns of decision making. These patterns were interpreted as the result of the different design principles and objectives behind each framework: ORCA was developed to produce efficient, realistic simulations of virtual multi-agent scenarios; SM was designed to generate legible robot motion in dynamic multi-agent environments; TE was based on a navigation interface [1], specifically designed to allow non-expert users to control a robot intuitively. To the best of our knowledge, these strategies have never been tested against each other under challenging, multi-agent, experimental settings. It was unclear how close interaction between the robot and different human participants would affect the motion generated by the different strategies. Furthermore, it was uncertain how humans would react to different behaviors exhibited by the robot and how this interaction would affect overall performance for both humans and the robot. Using the dataset generated by our study, we explore these questions by examining the validity of the

following hypotheses:

**(H1) - Robot Performance:** In close interactions with humans: (a) ORCA generates the most geometrically efficient paths; (b) SM generates the jerkiest paths; (c) TE generates the most energy-efficient paths.

**(H2) - Human performance:** Humans navigating in close proximity with the robot: (a) follow the least jerky paths when the robot runs SM; (b) spend the least energy when the robot runs TE; (c) spend the most energy when the robot runs OR.

**(H3) - Group performance:** Global group (human and robot) behavior under SM results in trajectories of lower Topological Complexity than the other two conditions.

**(H4) - Human Impressions:** Participants consider the behaviors generated by TE as more socially compliant, intelligent and safe than the rest of the strategies.

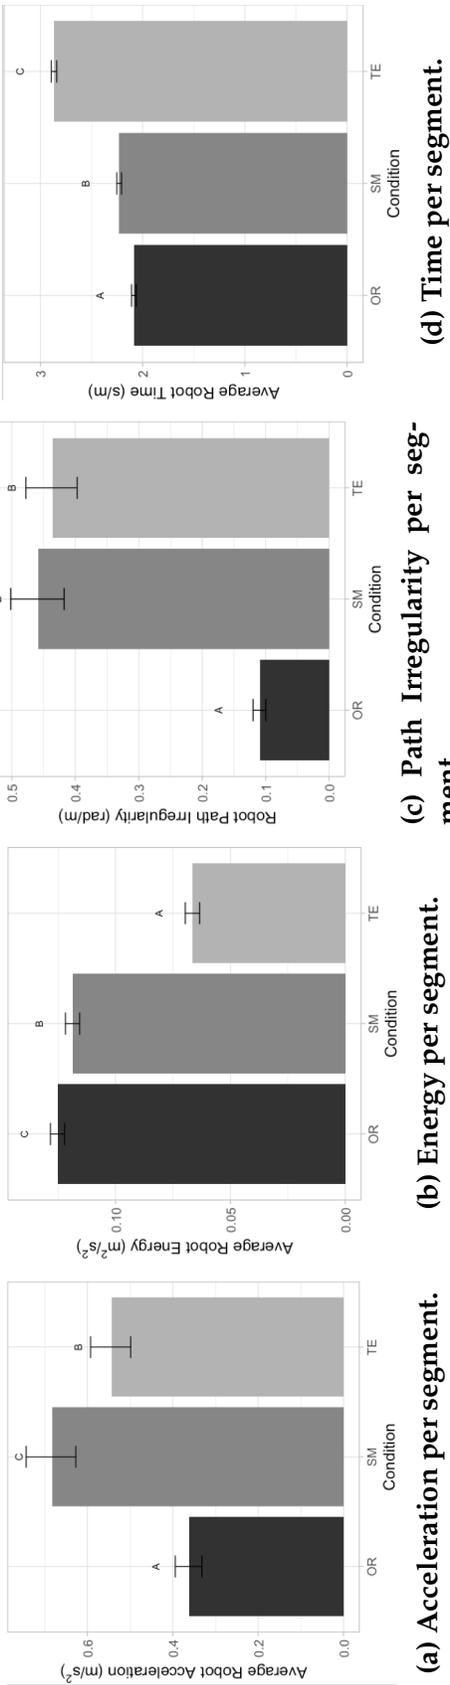
### 6.3 Analysis

We conducted 35 experiment sessions, in which a total of 105 human subjects were exposed to all three conditions. Subjects were recruited from a university population (Cornell University), through a centralized, university-run subject-recruitment website and also through fliers posted across campus. The subjects (59 female, 45 male, 1 unidentified) were 21.45 years old on average ( $SD = 3.19$  years) with their age ranging from 18 to 33 years. About half of them (57) had prior experience of user study participation and they rated their familiarity with

robotics technology with an average of 2.47 ( $SD = 1.27$ ) on a 5-point Likert scale.

We collected a dataset comprising the trajectories of all 105 participants and the robot across all trials. Focusing on dynamic interactions of close proximity, we split this raw dataset into two datasets of trajectory segments: (a) a dataset comprising 1033 robot trajectory segments of close interaction with humans (minimum distance  $d < 1m$ ) and (b) a dataset comprising 1566 human trajectory segments of close interaction with the robot (also, of minimum distance  $d < 1m$ ). We analyze the trajectory dataset using a set of trajectory quality measures from relevant literature [57, 86, 107], computed over fixed timestep intervals (100 timesteps, totaling 0.2 seconds). In particular, we computed: (1) the average *Acceleration* per segment,  $a$ ; (2) the average *Energy* per segment,  $E$ , where energy is defined as the integral of the squared velocity of an agent throughout its trajectory; (3) the minimum *Distance* between the robot and any other humans per segment,  $d$ ; (4) *Path Irregularity* per segment,  $PI$ , measuring the total amount of unnecessary rotation (angle between an agent’s heading and direction to goal) that an agent exhibits per unit path length [57]; (5) *Path Efficiency*,  $\mathcal{E}$ , defined as the ratio of the distance between the endpoints of a segment over the length of the path that the agent actually followed; (6) time spent per unit path length over a segment,  $\tau$ ; (7) *Topological Complexity*,  $TC$  [42, 107], defined as the amount of entanglement among agents’ trajectories throughout a trial (the Braidlab software package [132] was used for these computations).

We also collected a dataset comprising the responses of all 105 participants to a questionnaire, containing Likert-scale style questions, based on the instrument of Bartneck et al. [15] and short response questions.



**Figure 6.4: Expected means and confidence intervals for robot trajectory quality criteria. Quantities labeled with distinct letters (A, B, C) come from significantly different distributions (Tukey's HSD test,  $p < 0.05$ ).**

**Table 6.1: Effect of Navigation Strategy on Robot Behavior**

	Sum Sq	Mean Sq	NumDF	DenDF	F value	Pr(>F)
<i>a</i>	56.09	28.05	2	65.12	58.94	< 0.0001
<i>E</i>	0.7083	0.3541	2	1015	440.1	< 0.0001
<i><math>\mathcal{E}</math></i>	0.05796	0.02898	2	999.1	4.825	0.008213
<i>PI</i>	454.4	227.2	2	1012	355.3	< 0.0001
<i><math>\tau</math></i>	116.5	58.27	2	1016	1056	< 0.0001

### 6.3.1 Effect of Navigation Strategy on Robot Behavior

We model the effect of condition (ORCA, SM, TE) on each one of the trajectory quality measures considered. We use linear mixed-effects regression models, to account for both fixed effects resulting from the conditions but also for random effects resulting from the session and the trial (expected means with confidence intervals are depicted in Figure 6.4).

One-way ANOVA performed on the models demonstrates a significant effect of the condition on all robot trajectory quality criteria at the  $p < 0.05$  level (see table 6.1 for the test statistics and Figure 6.4 for the expected means and confidence intervals for all criteria) and thus, we find that (H1) is confirmed. More specifically, it can be observed that ORCA generates the smoothest motion among all strategies (lowest acceleration, lowest path irregularity, lowest time), which confirms (H1a). This trend was expected as ORCA selects actions that minimize divergence from an agent’s direction to goal and desired speed to ensure collision avoidance for a desired time window. This results in a smoother speed profile than other conditions. SM on the other hand, prioritizes intent-expressiveness by exaggerating its motion to indicate an intended passing-side intention; this results in higher acceleration (due to rotation) and path irregularity, which confirms (H1b). Finally, TE is the most energy-efficient — which

confirms (H1c) — but also the least time-efficient of all strategies. These findings could mainly be attributed to the defensive driving style of the teleoperator and the navigation through arrow keys.

### 6.3.2 Effect of Navigation Strategy on Human Behavior

Similarly to robot trajectory, we model the dependency of the human trajectory quality measures to the condition with linear mixed-effects models, accounting also for random effects of session, trial and helmet per trial. Figure 6.5 depicts the expected means and confidence intervals for the human trajectory quality measures, whereas table 6.2 contains statistics extracted upon performing ANOVA on the models at the  $p < 0.05$  significance level.

**Table 6.2: Effect of Navigation Strategy on Human Behavior**

	Sum Sq	Mean Sq	NumDF	DenDF	F value	Pr(>F)
<i>a</i>	1.415	0.7073	2	250.4	3.888	0.02173
<i>d</i>	0.1075	0.05377	2	231.5	0.5872	0.5567
<i>E</i>	0.112	0.05599	2	253.3	3.449	0.03326
$\mathcal{E}$	0.02977	0.01489	2	68.46	1.959	0.1488
<i>PI</i>	0.5394	0.2697	2	249.4	3.286	0.03904
$\tau$	0.08277	0.04139	2	252.7	2.145	0.1192

Overall, we find that (H2) is confirmed. In particular, we see that humans exposed to the SM condition followed smoother trajectories, of lower acceleration (Figure 6.5a) and path irregularity (Figure 6.5c) than humans exposed to either ORCA or TE, which confirms (H2a). This was in line with our expectations: SM’s intention-aware navigation strategy adapts the robot’s behavior to the preferences of humans, thus facilitating human inference and decision making. Further, it was observed that humans spend the least energy when exposed

to TE, which confirms (H2b). We attribute this finding to the perceived human-likeness of the motion generated by a teleoperated robot: the embodiment of human decision making on a robot platform features humanlike traits that potentially enable a higher level of human comfort. Finally, humans spend the most energy around OR (see Figure 6.5b), which confirms (H2c). This could be perceived as an result of ORCA's more predictable motion (minimal divergence from desired direction). Higher predictability potentially results in higher confidence for participants, which allows them to move faster and thus spend more energy.

### 6.3.3 Effect of Navigation Strategy on Group Behavior

We model the effect of condition on the Topological Complexity of the group trajectory (the set of all agents' trajectories) over a trial, using a linear mixed-effects model (accounting for random effects of session, trial and helmet per trial). Overall, we find that (H3) is rejected. ANOVA performed on the model uncovered a significant variance among conditions ( $F(2, 67.71) = 8.075, p = 0.000716$ , see table 6.3, Figure 6.5d). In particular, it was found that the Topological Complexity of trajectories, generated by groups exposed to TE was significantly lower than both SM and OR. Global group behavior generated in the presence of autonomy was significantly more complex, despite the fact that the human teleoperator was following the same rule for transitioning between machines (random selection of any non-adjacent machine). In other words, autonomous strategies resulted in more intense mixing among all four agents. This could be attributed to the mechanisms underlying human navigation, as the decision making computations under TE were done by the human teleoperator. Lower

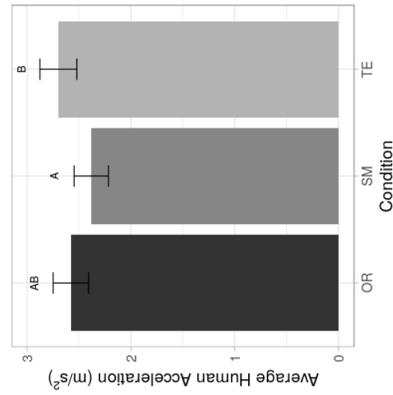
*TC* represents trajectory entanglement which intuitively corresponds to behaviors of passing *around* as opposed to passing *through* others. Thus, this trend could be attributed to the tendency of the human teleoperator to avoid collisions more *globally*, by avoiding any type of encounter with other participants whereas the robot was employing a more local collision avoidance mechanism by sequentially responding to any challenging encounters. This finding is perhaps unsurprising since both autonomous algorithms considered explicitly favor the avoidance of closer collisions over further ones.

**Table 6.3: Effect of Navigation Strategy on Group Behavior**

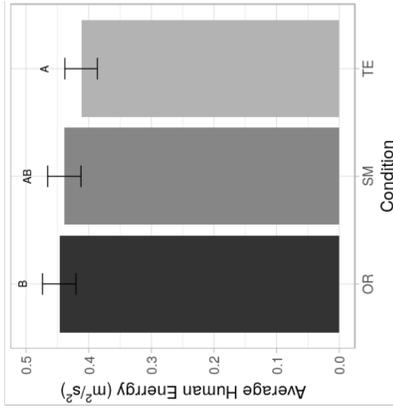
	Sum Sq	Mean Sq	NumDF	DenDF	F value	Pr(>F)
<i>TC</i>	107.3	53.66	2	67.71	8.075	0.000716

### 6.3.4 Effect of Navigation Strategy on Human Impressions

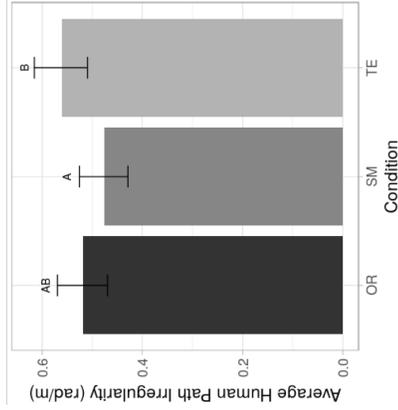
We model the effect of condition to each of the Likert scale questions considered, using a Linear Mixed Effects Regression Model. Table 6.4 contains a list of the questions that were posed to participants (as 5-point Likert scales, with 1 denoting a negative response and 5 denoting a positive response), grouped into three different classes: (a) one referring to the robot’s behavior (orange rows); (b) one referring to participants’ emotional states during the experiment (yellow rows); (c) one referring to participants’ expectations about the future presence of the robot (blue rows). The table also contains the statistics of one-way ANOVA tests, performed to participants’ responses to each question. Significant variance was observed in the responses to the question about the robot’s intelligence ( $F(2, 269.73) = 3.115, p = 0.0460$ ) and it was found that participants rated the intelligence of TE as slightly higher ( $M = 3.29, SE = 0.11$ ) than both



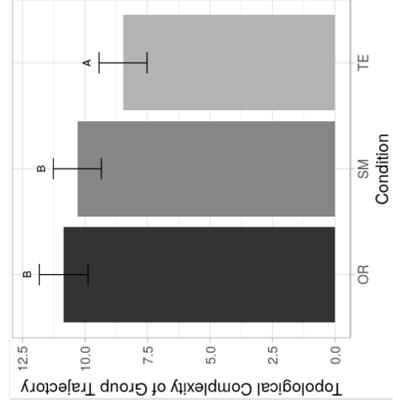
(a) Acceleration per segment.



(b) Energy per segment.



(c) Path Irregularity per segment.



(d) Top. Complexity per trial.

Figure 6.5: Expected means and confidence intervals for human (Figure 6.5a, Figure 6.5b, Figure 6.5c) and group trajectory quality criteria (Figure 6.5d). Quantities labeled with distinct letters (A, B) are significantly different (Tukey's HSD test,  $p < 0.05$ ).

SM ( $M = 3.01$ ,  $SE = 0.11$ ) and OR ( $M = 3.04$ ,  $SE = 0.11$ ). This trend also suggests a potential perception of the humanlikeness of TE from the perspective of the participants, which appears to be in line with the fact that they spent significantly less energy around TE. However, this trend is not reflected in the responses to the rest of the questions. Therefore, we cannot conclusively confirm or reject (H4). It might be the case that the quantitative differences among conditions in terms of the quality criteria were below the precision of human perception.

**Table 6.4: Effect of Nav. Strategy on Human Impressions**

Rating	Sum Sq	Mean Sq	NumDF	DenDF	F value	Pr(>F)
Competent	3.249	1.625	2	269.1	1.875	0.1553
Responsible	0.546	0.2732	2	269.7	0.3047	0.7376
Predictable	0.8769	0.4384	2	270.6	0.3760	0.6870
Compliant	3.599	1.800	2	269.7	2.279	0.1044
Sensible	1.687	0.8435	2	269.7	0.9756	0.3783
Friendly	2.346	1.173	2	269.6	1.299	0.2745
Safe	4.5409	2.270	2	270.0	1.793	0.1684
Pleasant	1.8920	0.9460	2	269.5	1.149	0.3184
Polite	0.9942	0.4971	2	269.8	0.7171	0.4891
Coordinated	2.971	1.485	2	269.4	1.425	0.2423
Intelligent	4.840	2.420	2	269.7	3.115	0.0460*
Trustworthy	0.5773	0.2887	2	269.2	0.4823	0.6179
Socially aware	4.507	2.254	2	268.6	2.131	0.1207
Discreet	4.742	2.371	2	270.2	2.238	0.1087
Relaxed	2.603	1.301	2	269.1	1.210	0.2998
Calm	2.405	1.202	2	269.7	1.123	0.3268
Tranquil	0.1069	0.0535	2	269.4	0.0575	0.9441
Noticeable	3.139	1.570	2	270.3	1.341	0.2633
Predict (future)	0.4635	0.2318	2	270.5	0.1937	0.8241
Bump (future)	7.572	3.786	2	270.2	2.682	0.0702

## 6.4 Discussion

We presented a within-subjects user study design for the experimental evaluation of mobile robot navigation strategies in a controlled lab environment. Our

experiments involved the navigation of a mobile robot in a workspace shared with three human participants, under challenging settings of implicit interaction, emulating aspects of pedestrian navigation. We conducted a total of 35 experiment sessions in which 105 human participants were exposed to the same set of conditions corresponding to three different navigation strategies executed by the robot. We analyzed the collected dataset through the use of objective measures (trajectory analysis) and subjective measures (questionnaires asking for ratings of participants' impressions of robot's intelligence, safety and personality). We found statistical evidence that humans follow less jerky and irregular paths when navigating around one autonomous navigation condition [107] than around a teleoperated robot. Furthermore, contrary to our expectations, humans did not discriminate between conditions, according to their responses to our questionnaire. Finally, we presented evidence that human decision making, as captured in the teleoperated condition, had a more global character than the autonomous strategies. We plan to investigate this finding further in future work.

### **6.4.1 Limitations**

Our study encompasses some limitations generally inherent to any HRI study, and some specific to our scenario. First, a controlled lab environment cannot emulate the complexity of a real-world pedestrian environment and no background scenario or task could give rise to perfectly natural human walking behaviors. Furthermore, humans lack models of interaction with robotic technology, which inevitably affects their behavior around a mobile robot. Even the robot's appearance, structure and dynamics could attract attention and distract

participants from the task. Moreover, the selection of the navigation strategies inevitably impacts the generalizability of the results. Either of the autonomy conditions could struggle with erratic human behavior and specifically with human motion that is suboptimal with respect to intent and flexibility [136] whereas the teleoperator's performance may vary across individuals, experience, skill level, driving style, etc. Finally, the sample of participants, mostly coming from the undergraduate population of a university introduces another confound.

#### **6.4.2 Broader Impact**

Despite its many limitations, this study is unique in terms of its goals, settings, thoroughness of evaluation and sample size. As stated in the introduction, this study was motivated by an observed gap in the literature: we believe that the validation of social navigation algorithms requires a more thorough process. The stage of a controlled lab evaluation is an indispensable part of the validation process and should not be discounted before deploying a robot to the field. The field of social robot navigation could benefit greatly from extensive in-lab validation of additional algorithms, under various interaction settings. We hope that this study will constitute a paradigm for such future studies in terms of its design and scope and a reference for informing the design of future algorithms, within the field of navigation and beyond.

## **Part IV**

# **Planning with Topologically Robust Multi-Agent Trajectory Prediction**

CHAPTER 7  
MULTI-AGENT TRAJECTORY PREDICTION AND GENERATION WITH  
TOPOLOGICAL INVARIANTS ENFORCED BY HAMILTONIAN  
DYNAMICS

Several real-world, multi-agent navigation domains, such as pedestrian or street environments, prohibit the use of explicit communication among agents. Thus, rational agents need to employ mechanisms for predicting the behaviors of others to plan collision-free paths. Although predicting the behaviors of others in a detailed way is challenging, the assumption of rationality, combined with the constraints imposed to agents by the environment bounds results in the definition of a set of qualitatively distinct, global planning alternatives in the joint trajectory space of all agents, corresponding to different strategies of joint collision avoidance. Being cognizant of this structure, agents could anticipate different classes of unfolding multi-agent dynamics. This could enable them to execute actions of global outlook and consistency, which could allow for consistently expressive, smooth and efficient motion, even in the face of unexpected events, such as the appearance of agents with heterogeneous policies or agents with changing intentions. This ability is of particular importance for operation in human-populated domains, where agents continuously read and react to the rapidly changing environment. Important domains with these properties include automobile traffic in parking lots and pedestrian traffic in communal spaces like hallways.

In this chapter, we present a planning framework for decentralized navigation planning, based on online multi-agent trajectory prediction. Unlike SCN from part II, which reasoned abstractly about the emerging trajectory topology

and SM from part III, which reasoned in a distributed fashion about pairwise collision avoidance maneuvers, in this chapter we introduce a planner that generates Cartesian representations of future, topologically distinct multi-agent behaviors at planning time in order to plan smooth, collision-free actions of global consistency and outlook. The main benefit of constructing Cartesian trajectory representations is the ability to evaluate potential futures with respect to their likelihood through the incorporation of cost functions related to energy, acceleration, efficiency, social compliance etc. We show that this architecture allows for rapid rapid adjustment to unexpected events, such as the appearance of heterogeneous agents with no collision avoidance mechanisms or agents with changing intentions.

Our framework is based on the method of Berger [18], which allows for braiding multi-particle trajectories into desired topological patterns. Our prediction mechanism grows trajectories from agents' initial configurations to their predicted destinations in topologically distinct ways. This allows us to introduce desired global properties to the trajectory, in contrast to typical trajectory optimization methods (e.g. [153]), which act on a trajectory locally. Based on this mechanism, we present an online algorithm that generates a diverse set of distinctly entangled multi-agent trajectories and evaluates them with respect to their quality and likelihood, to select an action that best adapts to other agents' impending behaviors and preferences. This allows for rapid adjustment to the changing environment and facilitates robustness to unexpected events such as

the emergence of heterogeneous agents or agents with changing intentions.

## 7.1 Foundations

We consider the planning problem for an autonomous agent that navigates towards its destination in a known environment where other agents are also navigating towards their respective destinations. The agent aims at reaching its destination by following an efficient, smooth, collision-free trajectory and assumes that others share similar objectives, although it has no knowledge of their specific policies. The agent is not explicitly communicating with others and thus has no knowledge of their intended destinations or planned paths or policies but is able to perfectly observe their motion. Assuming that others share similar abilities and objectives, the agent may form a belief about how they are going to move in the future so that it can plan a safe and efficient path towards its destination.

In this chapter, we present an approach inspired by the point vortex problem [7–9] from fluid dynamics. We design a planning framework built around the observation that the collision avoidance process for agents navigating on a plane resembles the dynamics of interacting point vortices in two dimensions. Treating agents as point masses subjected to vortex dynamics allows us to synthesize multi-agent trajectories with desired topological properties. At planning time, this technique enables a planning agent to construct several qualitatively distinct predictions about the future behavior of the system. This allows for an informed action selection towards facilitating a rapid and robust adjustment to the changing environment. Since the predictions are made with a global outlook, this strategy results in a consistently smooth and intent-expressive behav-

ior, even in the face of unexpected events such as agents changing intentions or violating the assumption of rationality. Our approach is based on the method of Berger [18, 19] for generating braided trajectories of multi-particle systems from topological invariants [19]. In this section, we introduce some preliminaries about point vortex flows, review the method of Berger [18] and present the key components of our approach.

### 7.1.1 Hamiltonian Motion for Multi-Particle Trajectory Braiding

A dynamical system whose evolution is described by Hamilton's equations is called a Hamiltonian system. Under the Hamiltonian formalism, the state of a system is completely described by a set of variables corresponding to the generalized coordinates of the system's degrees of freedom and their conjugate momenta. Hamilton's equations relate the evolution of an energy function, called the *Hamiltonian*, to the evolution of the coordinates and momenta for all degrees of freedom of the system. In particular, denoting by  $q_j$  and  $p_j$ ,  $j \in M = \{1, \dots, m\}$  the generalized coordinate and conjugate momentum of the  $i$ th degree of freedom of a Hamiltonian system respectively, its evolution is given by:

$$\dot{q}_j = \frac{\partial H}{\partial p_j}, \quad \dot{p}_j = -\frac{\partial H}{\partial q_j}, \quad \frac{dH}{dt} = 0, \quad j = 1, \dots, m, \quad (7.1)$$

where the dot notation indicates time derivatives,  $H$  denotes the *Hamiltonian* of the system (which is preserved), defined as its total energy, i.e., the sum of the total kinetic and potential energy of all degrees of freedom of the system.

Let us now combine the coordinates and momenta for each degree of freedom into a complex coordinate  $z_j = q_j + ip_j$ ,  $j \in M$ . We define an analytic function

$$F(z_1, \dots, z_m) = \Psi(z_1, \dots, z_m) + iH(z_1, \dots, z_m), \quad (7.2)$$

where  $\Psi : \mathbb{C}^m \rightarrow \mathbb{R}$  and  $H : \mathbb{C}^m \rightarrow \mathbb{R}$ . Berger [18] showed that the Hamiltonian flow (7.1) results in motion  $\dot{z}_j$ ,  $j \in N = \{1, \dots, n\}$ , that follows the Wirtinger derivative of  $\Psi$  with respect to  $z_j$ . Therefore, the collective Hamiltonian motion of all degrees of freedom follows the gradient of  $\Psi$  and points towards its direction of maximum increase. Berger [18] used this finding to generate braided trajectory patterns for systems of two and three particles. In particular, he replaced  $\Psi$  with *Topological Invariants* [19] towards forcing the system to evolve along the growth of the topological invariant.

### 7.1.2 Topological Invariants of Particle Trajectories

Consider a set of  $n$  particles, following trajectories  $\xi_i : [0, T] \rightarrow \mathbb{R}^2$ ,  $i \in N$ , from time  $t = 0$  to time  $t = T$  and let us collect these trajectories into a system trajectory  $\Xi = (\xi_1, \dots, \xi_n)$ . A *topological invariant* over  $\Xi$ , may be defined as a function  $\Psi : \Xi \rightarrow \mathbb{R}$  that maps the system trajectory to a real number that characterizes the spatiotemporal topology of the system dynamics. For any distorted, topology-preserving trajectory  $\tilde{\Xi} \neq \Xi$  with the same endpoints  $\tilde{\Xi}(0) = \Xi(0)$ ,  $\tilde{\Xi}(T) = \Xi(T)$ , for which  $\xi_i(t) \neq \xi_j(t)$ ,  $\forall t \in (0, T)$  and  $i \neq j \in N$ , a topological invariant is preserved, i.e.,  $\Psi(\tilde{\Xi}) = \Psi(\Xi)$ .

### 7.1.2.1 The Winding Number

The so called *Winding Number* is a topological invariant of particular interest for our problem. Consider a curve  $\gamma : [0, t] \rightarrow \mathbb{C} \setminus \{0\}$ . The complex winding number of the curve  $\gamma$ , from time 0 to time  $t$  is defined as:

$$\lambda(t) = \frac{1}{2\pi i} \oint_{\gamma} \frac{dz}{z}, \quad (7.3)$$

where  $z \in \mathbb{C}$ . Let us express  $\gamma$  in polar coordinates as  $\gamma(t) = r(t)e^{i\theta(t)}$ , where  $r(t) = \|\gamma(t)\|$  and  $\theta(t) = \angle\gamma(t)$ . Then, through the use of the Cauchy integral formula, (7.3) may be decomposed into:

$$\lambda(t) = \frac{1}{2\pi i} \int_0^t \frac{\dot{r}}{r} dt' + \frac{1}{2\pi} \int_0^t \dot{\theta} dt' \quad (7.4)$$

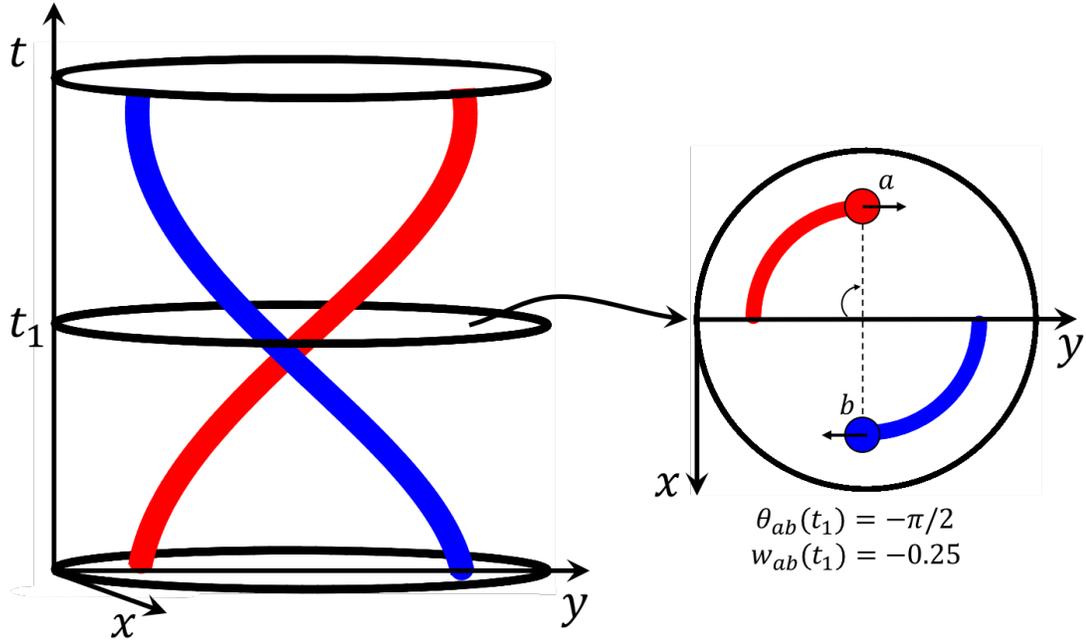
and computing the integrals yields:

$$\lambda(t) = \frac{1}{2\pi i} \log \left( \frac{r(t)}{r(0)} \right) + \frac{1}{2\pi} (\theta(t) - \theta(0)). \quad (7.5)$$

The real part of this integral,  $w = \text{Re}(\lambda)$ , is a topological invariant, counting the number of times the curve  $\gamma$  encircled the origin in the time interval  $[0, t]$ . In other words, fixing the endpoints of the curve, any topology-preserving deformations are mapped to the same value of the winding number. For closed curves, the imaginary part of the winding number is zero. In the following section, considering open curves (evolving trajectories), we describe how it can be used to enforce Hamiltonian motion to interacting particles.

### 7.1.3 Two-Particle Vortex Motion

In this section we put the pieces together to demonstrate a motivating example from fluid dynamics that constitutes the computational basis of our approach.



**Figure 7.1: Spacetime plot of the trajectories of two agents, navigating in a circular workspace (left) and projection of their trajectories until time  $t_1$ , onto the  $xy$  plane, along with the definition of their pairwise winding angle and winding number (right).**

Consider a system of two particles, placed initially at positions  $a = (a_x, a_y) \in \mathbb{R}^2$  and  $b = (b_x, b_y) \in \mathbb{R}^2$  with respect to a fixed coordinate system and assume that a vortex<sup>1</sup> lies between them. Point vortex motion prescribes that the  $x$  and  $y$  coordinates are conjugate to each other (e.g. the conjugate momentum to  $a_x$  is  $a_y$ ) [7]. Let us define the function  $\gamma_{ab}$  from the previous section to track the quantity  $a - b$ , i.e., let us set  $\gamma_{ab}(t) = r_{ab}(t)e^{i\theta_{ab}(t)}$ , where  $r_{ab} = \|a - b\|$  and  $\theta_{ab}(t) = \angle\gamma_{ab}(t)$ .

Assuming unit vorticity, the Hamiltonian for this system may be written as:

$$H = -\frac{1}{2\pi} \log r_{ab}. \quad (7.6)$$

<sup>1</sup>A vortex is a region in a fluid in which the flow revolves around an axis line.

Similarly to Sec. 7.1.2.1, we may define the complex winding number of  $\gamma_{ab}$  as:

$$\lambda_{ab}(t) = \frac{1}{2\pi i} \log(r_{ab}) + \frac{1}{2\pi} (\theta_{ab}(t) - \theta_{ab}(0)), \quad (7.7)$$

and let us set its real part to a dedicated variable

$$w_{ab}(t) = \frac{1}{2\pi} (\theta_{ab}(t) - \theta_{ab}(0)), \quad (7.8)$$

denoting the pairwise winding number of the two curves (see Figure 7.1 for a graphic representation of the pairwise winding number). We may notice that  $Im(\lambda_{ab}) = H$ . Thus, according to Sec. 7.1.1, the Hamiltonian flow for this system maximizes the growth of the real part  $Re(\lambda_{ab}) = w_{ab}$ . This motion corresponds to the two points rotating about each other at a constant radius, in a counter-clockwise direction. Hamilton's equations for this system may be derived as:

$$(\dot{a}_x, \dot{a}_y) = \left( \frac{\partial H}{\partial a_y}, -\frac{\partial H}{\partial a_x} \right) = \frac{1}{2\pi} \left( -\frac{a_y - b_y}{r_{ab}^2}, \frac{a_x - b_x}{r_{ab}^2} \right), \quad (7.9)$$

$$(\dot{b}_x, \dot{b}_y) = \left( \frac{\partial H}{\partial b_y}, -\frac{\partial H}{\partial b_x} \right) = \frac{1}{2\pi} \left( -\frac{b_y - a_y}{r_{ab}^2}, \frac{b_x - a_x}{r_{ab}^2} \right). \quad (7.10)$$

We may control the directionality of the rotation by switching the signs in the right hand side of eqs. (7.9) and (7.10).

#### 7.1.4 Two-Agent Collision Avoidance as Vortex Motion

Treating agents as particles, we may use the outlined method of Berger [18] as a mechanism for generating two-agent, collision-avoidance maneuvers of desired topological specification. Given a winding number  $w_{ab}$ , by multiplying the right hand sides of eqs. (7.9) and (7.10) with  $sign(w_{ab})$ , we have a planning rule that

allows us to grow trajectories for  $a$  and  $b$  that follow the direction indicated by  $w_{ab}$ , with  $\text{sign}(w_{ab}) > 0$  and  $\text{sign}(w_{ab}) < 0$  corresponding to right and left hand side collision avoidance respectively. In a two-agent scene, this may serve as a prediction of the emerging joint behavior. In a scene with high uncertainty, where no agent has committed to a passing side, this mechanism allows a planning agent to anticipate both outcomes. This is useful as it allows the agent to either enforce its own preference or adapt to the preference of the other agent. In the following section, we show how we use this method for synthesizing trajectories of complex topological specifications in environments with multiple agents.

### 7.1.5 Multi-Agent Trajectory Generation from Topological Specifications

Consider the problem of centralized trajectory planning for driving  $n$  agents from their initial positions  $S = (s_1, \dots, s_n) \in \mathbb{R}^{2n}$  to their destinations  $D = (d_1, \dots, d_n) \in \mathbb{R}^{2n}$  in a collision-free fashion and while following a global topological specification  $w$ , prescribing passing sides to agents. We model  $w$  as a tuple of pairwise winding numbers  $w = (w_{12}, w_{13}, \dots)$  from the set of such tuples  $\mathcal{W}$ . Assuming that each agent passes each other exactly once on its way to its destination (agents do not loop around others), the magnitude of  $w_{ij}$ ,  $i \neq j \in N$  is not important, so we will be using  $w_{ij}$  to refer to  $\text{sign}(w_{ij})$ . The cardinality of the set of possible specifications is  $|\mathcal{W}| = 2^{\binom{n}{2}}$ , corresponding to all possible combinations of passing sides for all agents. It should be noted that although all combinations in  $\mathcal{W}$  are topologically possible, in practice, only a subset of them

are meaningful and likely given agents' state history and under the assumption of rationality. Sec. 7.1.6 addresses the problem of evaluating the likelihood and the feasibility of a topological specification.

We now describe a policy  $\pi : \mathbb{R}^{2n} \times \mathcal{W} \rightarrow \mathbb{R}^{2n}$  that can be sequentially iterated to produce a multi-agent trajectory that satisfies a topological specification  $w$ . The policy (referred from now on as HTTG, standing for Hamiltonian Topological Trajectory Generation) prescribes an action  $u_i \in \mathbb{R}^2$  to every agent  $i \in N$ , synthesized from a weighted consideration of all pairwise collision avoidance reactions between the agent and all others, towards meeting the pairwise specifications contained in  $w$ . The policy is executed repeatedly until all agents reach their destinations. It may be formulated for agent  $i$  as follows:

$$u_i = \nu_i \cdot k (u_{att}^i + u_{rep}^i), \quad (7.11)$$

where  $\nu_i \in \mathbb{R}$  is an agent's desired speed,  $u_{att}^i, u_{rep}^i$  are potentials attracting the agent towards its destination and repulsing it from others respectively and  $k \in \mathbb{R}$  is a normalization constant. The potential

$$u_{att}^i = k_{att}(d_i - q_i) \quad (7.12)$$

attracts the agent from its current state  $q_i$  towards its destination  $d_i$  with  $k_{att}$  being an importance weight. The potential

$$u_{rep}^i = k_{rep} \sum_{j \neq i}^N c_{ij} w_{ij} v_j^i, \quad (7.13)$$

repulses agent  $i$  from each other agent  $j \in N, j \neq i$ , through the velocity  $v_j^i$ , derived from eqs. (7.9) and (7.10), with a degree of consideration equivalent to the criticality of their pairwise collision avoidance, expressed by  $c_{ij} \in \mathbb{R}$  (the closer

two agents are, the more critical their avoidance becomes) and along the direction indicated by  $w_{ij}$  whereas  $k_{rep}$  is an importance weight. The choice of the weighting factors  $k_{att}, k_{rep}$  expresses the relative significance between goal attraction and collision avoidance. The criticality term is designed to be a polynomial function of the distance between two agents, activated when the distance becomes lower than a threshold. By sequentially executing the outlined policy, in parallel for all agents, in equal time steps of length  $dt$ , the system of agents is forced to follow the specification  $w$ . Note that this method does not guarantee the attainment of the desired topology. Depending on the number of agents, their initial configurations and intended destinations and the parameter tuning, the method has a varying control over the topological properties of agents' trajectories. Sec. 7.2 explores the performance of the method on scenarios with different numbers of agents.

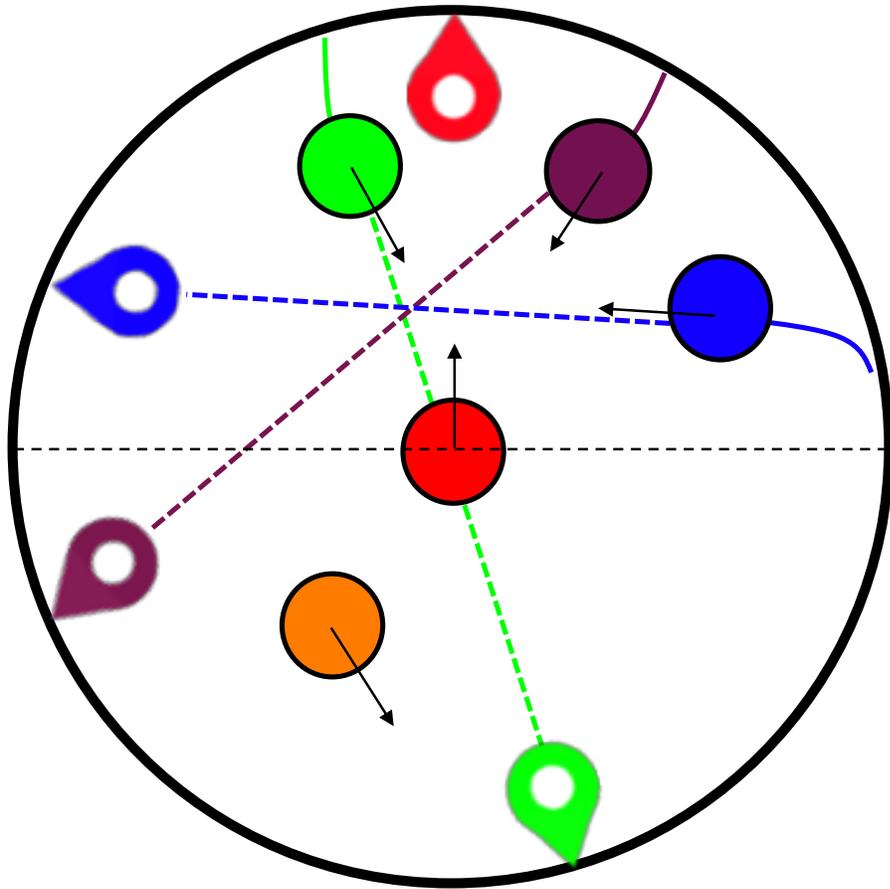
### 7.1.6 TANP: Topologically Adaptive Navigation Planning

In this section, we present an online, decentralized navigation planning algorithm that makes use of the described method for generating online a set of topologically distinct, multi-agent trajectory predictions. The algorithm comprises the following sequence of actions: (1) **predict** the destinations of other agents; (2) **generate** a set of candidate multi-agent trajectories that drive agents from their current locations to their predicted destinations; (3) **evaluate** candidates with respect to a cost function; (4) **execute** the next action assigned to the planning agent from the lowest-cost candidate. In the following subsections, we describe the main components of the algorithm and provide a detailed presentation of it in pseudocode format (see Alg. 6).

### 7.1.6.1 Destination Prediction

In sec. 7.1.5, it was assumed that the planning policy has access to the destinations of other agents. In the settings we are considering (no explicit communication among agents), this is not the case. Thus, the planning agent needs to make a prediction about the destinations of others in order to use the policy. However, in practice, an agent only interacts with others for as long as they lie within its sensing range, which for current robotic systems is quite limited. During this amount of time, other agents' observed behaviors may or may not be revealing about their specific destination. And in fact, detailed predictions of agents' destinations may not be sufficiently informative regarding agents' future behaviors; in crowded environments, the collision avoidance process is a more significant influence over agents' behaviors. For this reason, we take a more practical approach, focusing on coarse predictions of agents' future locations.

In particular, we assume that an agent's sensing range has the shape of a disk of radius  $R$ , centered at the agent's position,  $q_i$ . Any agent lying outside of this disk is not perceived by the agent whereas any agents lying behind the robot are ignored at the planning stage. For each one of the perceived and actively considered agents, we approximate their intended direction of motion by fitting a line to their recent, observed trajectory and projecting their current velocity on it. We then propagate their current speed along this direction until it intersects the boundary of the sensing disk. For our planning algorithm, that point of intersection is considered to be that agent's destination (see Figure 7.2). This prediction is expected to be a coarse approximation of where an agent is heading. However, since our algorithm runs in replanning cycles, this



**Figure 7.2: The destination prediction mechanism. The red agent makes destination predictions for all agents, lying within its circular sensing disk and in front of it.**

approximation provides a sufficient amount of detail for the trajectory prediction mechanism of sec. 7.1.5. This mechanism makes use of the assumption that agents act rationally, i.e., agents' behaviors are driven by an incentive of making progress towards their destinations. Finally, alternative methods of filtering could be employed to provide more accurate destination prediction; however, this is not our focus and as our will be shown in Sec. 7.2, this simplistic model may yield the desired performance.

### 7.1.6.2 Outcome Evaluation

The set  $\mathcal{W}$  contains symbolic representations of topologically distinct *outcomes* for the system of all considered agents. Naturally, a significant question that arises is: which outcome should the planning agent trust and follow? We approach this problem with the following sequence of computations: (1) we first evaluate an outcome with respect to its likelihood; (2) we then generate trajectory representations for the set of the  $K$  most likely outcomes  $\mathcal{W}_K \subset \mathcal{W}$ , using the policy presented in Sec. 7.1.5; (3) finally, we evaluate these  $K$  best outcomes with respect to the physical properties of their trajectory representations.

**Probability of an outcome:** An outcome is initially encoded symbolically as a tuple  $\mathbf{w}$  that prescribes how agents avoid each other throughout the course of the scene. From a topological perspective, these symbols are independent of each other; any motion is allowed even if it is not efficient. However, from a real-world point of view, the collision-avoidance strategies that agents employ to avoid one-another are coupled and modeling the complex probabilistic relationships among them is a challenging problem. We are interested in finding a way to bias our search towards the outcomes that are more likely to occur. We do so by using the following expression:

$$P(\mathbf{w}|\Xi_{past}) = P(w_{12}, w_{13}, \dots |\Xi_{past}) \propto \frac{1}{Z} \prod_{ij} P(w_{ij}|\Xi_{past}), \quad (7.14)$$

where  $\Xi_{past}$  denotes agents' past trajectories and  $Z$  is a normalization constant across all  $\mathbf{w} \in \mathcal{W}$ . This expression was derived by factorizing  $P(w_{12}, \dots |\Xi_{past})$  using the product rule and then substituting each factor with its Bayes' rule expression. Similarly to our past work [104], we model  $P(w_{ij}|\Xi_{past})$  by employing the physical quantity of angular momentum. For two agents  $i, j$ , navigating on a plane, their angular momentum  $L^{ij}$  lies along the  $z$  axis. Notice that the sign

of the  $z$ -component of the momentum,  $L_z^{ij}$  is an indicator of agents' passing side and thus of the winding number of their trajectories  $w_{ij}$ , with  $L_z^{ij} > 0$  indicating the emergence of  $w_{ij} > 0$  (right hand side collision avoidance) and  $L_z^{ij} < 0$  indicating the emergence of  $w_{ij} < 0$  (left hand side collision avoidance). We incorporate the momentum as a heuristic in a sigmoid model as follows:

$$P(w_{ij}|\Xi_{past}) = \frac{1}{1 + \exp(-w_{ij}L_z^{ij})}. \quad (7.15)$$

The greater  $|L_z^{ij}|$  is, the greater the mutual intention or preference of agents  $i$  and  $j$  over a collision avoidance along the direction of  $L_z^{ij}$  is.

**Trajectory Evaluation:** We evaluate a trajectory representation  $\Xi_w$  of an outcome  $w$  by computing its total energy  $\mathcal{E} : \mathcal{Z}^n \rightarrow \mathbb{R}$ , its required immediate acceleration  $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathbb{R}$  and its safety cost  $\mathcal{S} : \mathcal{Z}^n \rightarrow \mathbb{R}$ . The Energy measure (sum of squared speeds throughout the trajectory) gives an estimate of the efficiency of an outcome whereas the acceleration measure is indicative of the aggressiveness of the maneuvers required to comply with an outcome. We model the Safety cost as  $\mathcal{S}(\Xi) = \exp(-d_{min})$ , where  $d_{min} \in \mathbb{R}$  is the minimum distance between any pair of agents in a trajectory  $\Xi$ . Note that other cost functions could be used to incorporate different considerations such as social comfort (see e.g. Sisbot et al. [123]).

### 7.1.6.3 Decision Making

We first rank outcomes at a symbolic level through the use of the probability distribution, presented in Sec. 7.1.6.2 and determine the set of the  $K$  most likely

outcomes  $\mathcal{W}_K$ . Then, we determine the outcome of lowest cost:

$$C(\Xi) = \alpha_e \mathcal{E} + \alpha_a \mathcal{A} + \alpha_s \mathcal{S}, \quad (7.16)$$

where  $\alpha_e, \alpha_a, \alpha_s$  are importance weights and finally extract the optimal outcome through the following optimization scheme:

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}_K} C(\Xi_{\mathbf{w}}). \quad (7.17)$$

The planning agent executes the next action assigned to it from the trajectory of lowest cost  $\Xi_{\mathbf{w}^*}$ . Figure 7.3 depicts a graphic representation of the planning scheme.

#### 7.1.6.4 Pseudocode

Alg. 6 summarizes the described algorithm (Topologically Adaptive Navigation Planning – TANP) in pseudocode format. The algorithm runs in replanning cycles for as long as the boolean variable *AtGoal* is set to *False*, indicating that the agent has not reached its destination yet. At every cycle, the agent first determines a set of *reactive agents*, i.e., agents that lie within the robot’s sensing disk and to the front of the robot’s heading (function `Get_Reactive_Agents`). Then, function `Predict_Destinations` outputs predictions for the destinations of the reactive agents and `Get_Outcomes` returns a set of topological representations for outcomes that could emerge in the remainder of the execution. Function `Get_Outcome_Probability` returns the probability for each of the outcomes considered and function `Get_Best_Outcomes` returns the  $K$  best outcomes. Function `HTTG` executes the HTTG policy and generates trajectory representations for these outcomes and function `Score_Trajectory`

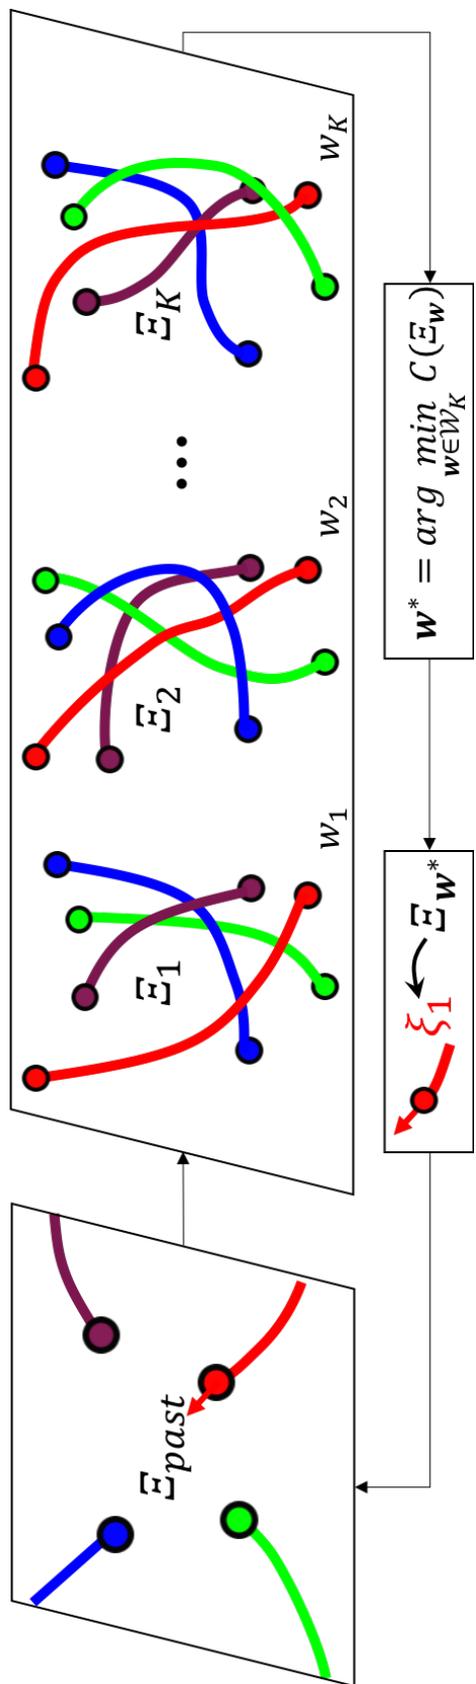


Figure 7.3: Illustration of the replanning scheme. At every replanning cycle, the planning agent generates a set of diverse (topologically distinct) predictions about the joint future behavior of all agents, evaluates them with respect to a cost function  $C$  and executes the action assigned to it from the prediction of lowest cost.

evaluates them with respect to the cost function considered. Finally, function `Get_Best_Next_Action` returns the next action for the planning agent from the trajectory of lowest cost and function `Execute_Action` executes that action. The distance between the resulting agent state and its destination is compared to the predefined threshold  $\epsilon$  and the flag `AtGoal` is updated to `True` in case the agent is sufficiently close to its destination.

---

**Algorithm 6** TANP( $q, d, \Xi$ )

---

**Input:**  $q$  – agent’s current state;  $d$  – agent’s intended destination;  $\Xi_{past}$  – state history of all agents;  $K$  – Number of outcomes to consider;  $\epsilon$  – desired distance-to-goal threshold.

```

1: AtGoal  $\leftarrow$  False
2: while  $\neg$ AtGoal do
3:    $\mathcal{R} \leftarrow$  Get_Reactive_Agents( $\Xi_{past}$ )
4:    $D \leftarrow$  Predict_Destinations( $\Xi_{past}, \mathcal{R}$ )
5:    $\mathcal{W} \leftarrow$  Get_Outcomes( $\mathcal{R}$ )
6:    $P \leftarrow$  Get_Outcome_Probability( $\mathcal{W}, \Xi_{past}$ )
7:    $\mathcal{W}_K \leftarrow$  Get_BestOutcomes( $P, \mathcal{W}, K$ )
8:    $\mathcal{Z} \leftarrow \emptyset$ 
9:   for all  $w \in \mathcal{W}_K$  do
10:     $\Xi_{pred} \leftarrow$  HTTG( $\Xi_{past}, w, D$ )
11:     $\mathcal{Z} \leftarrow \{\mathcal{Z}, \Xi_{pred}\}$ 
12:    $C \leftarrow$  Score_Trajectories( $\mathcal{Z}$ )
13:    $u \leftarrow$  Get_Best_NextAction( $\mathcal{Z}, C$ )
14:    $q \leftarrow$  Execute_Action( $u$ )
15:   if  $\|q - d\| < \epsilon$  then
16:     AtGoal  $\leftarrow$  True
17: return None

```

---

### 7.1.7 Complexity and Practical Considerations

The most computationally intense component of our algorithm is the estimation of the outcome probabilities. For  $n$  agents, this computation runs in time  $O(2^{n^2})$  –the rest of the computations run in polynomial time. In practice, a replanning

cycle of TANP on a scenario involving 4 agents and thus the evaluation of 64 topological classes with  $K = 5$ , runs at an average of  $42ms$ , with the worst case being  $203ms$  in a non-optimized MatLab implementation on a MacBook Pro of 2015 with an Intel Core i7 processor of 2.5 GHz, running macOS High Sierra. Transfer to a faster language and optimization of parts of the code could help vastly improve performance.

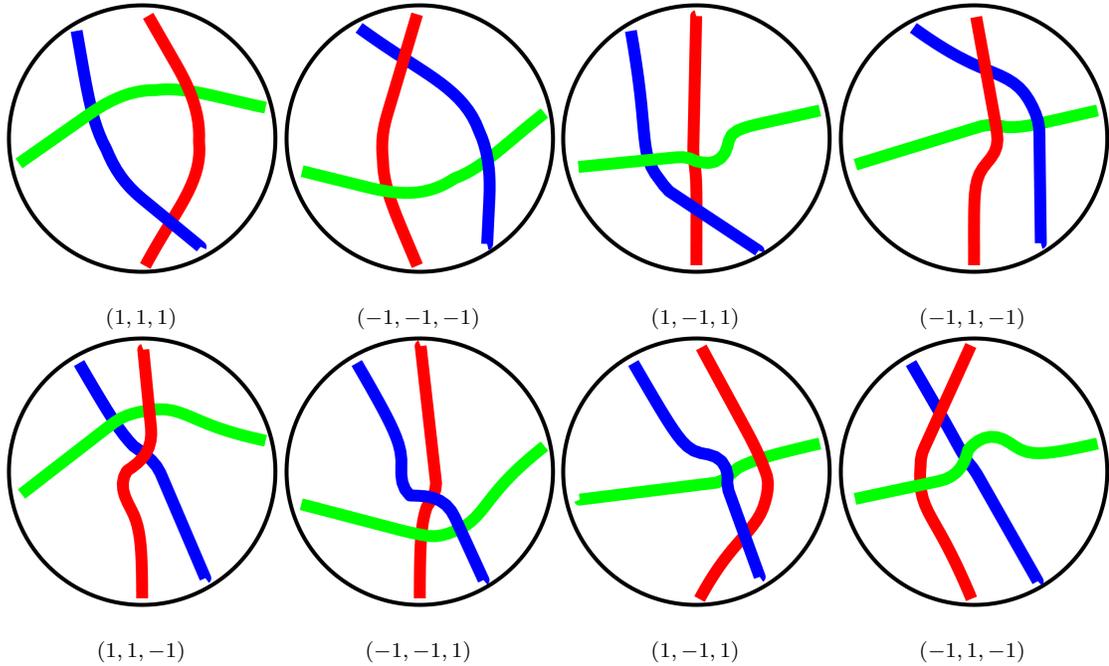
Under the current design, scaling to large  $n$  is not possible. However, for a mobile robot application, we argue that it is also not practical. The sensing limitations would prohibit the emergence of a large number of agents. Even if more agents are sensed, pruning them to the subset of directly reactive agents is a motivated and human-inspired way of reducing the load. Future work involves the design of an online data-driven topology-classification mechanism that would enable agents to directly estimate the most likely candidates, without brute-forcing their evaluation.

## 7.2 Results

In this section, we present the performance of the offline planner in generating trajectories of desired topological properties and the behaviors generated by the online algorithm in different types of scenarios.

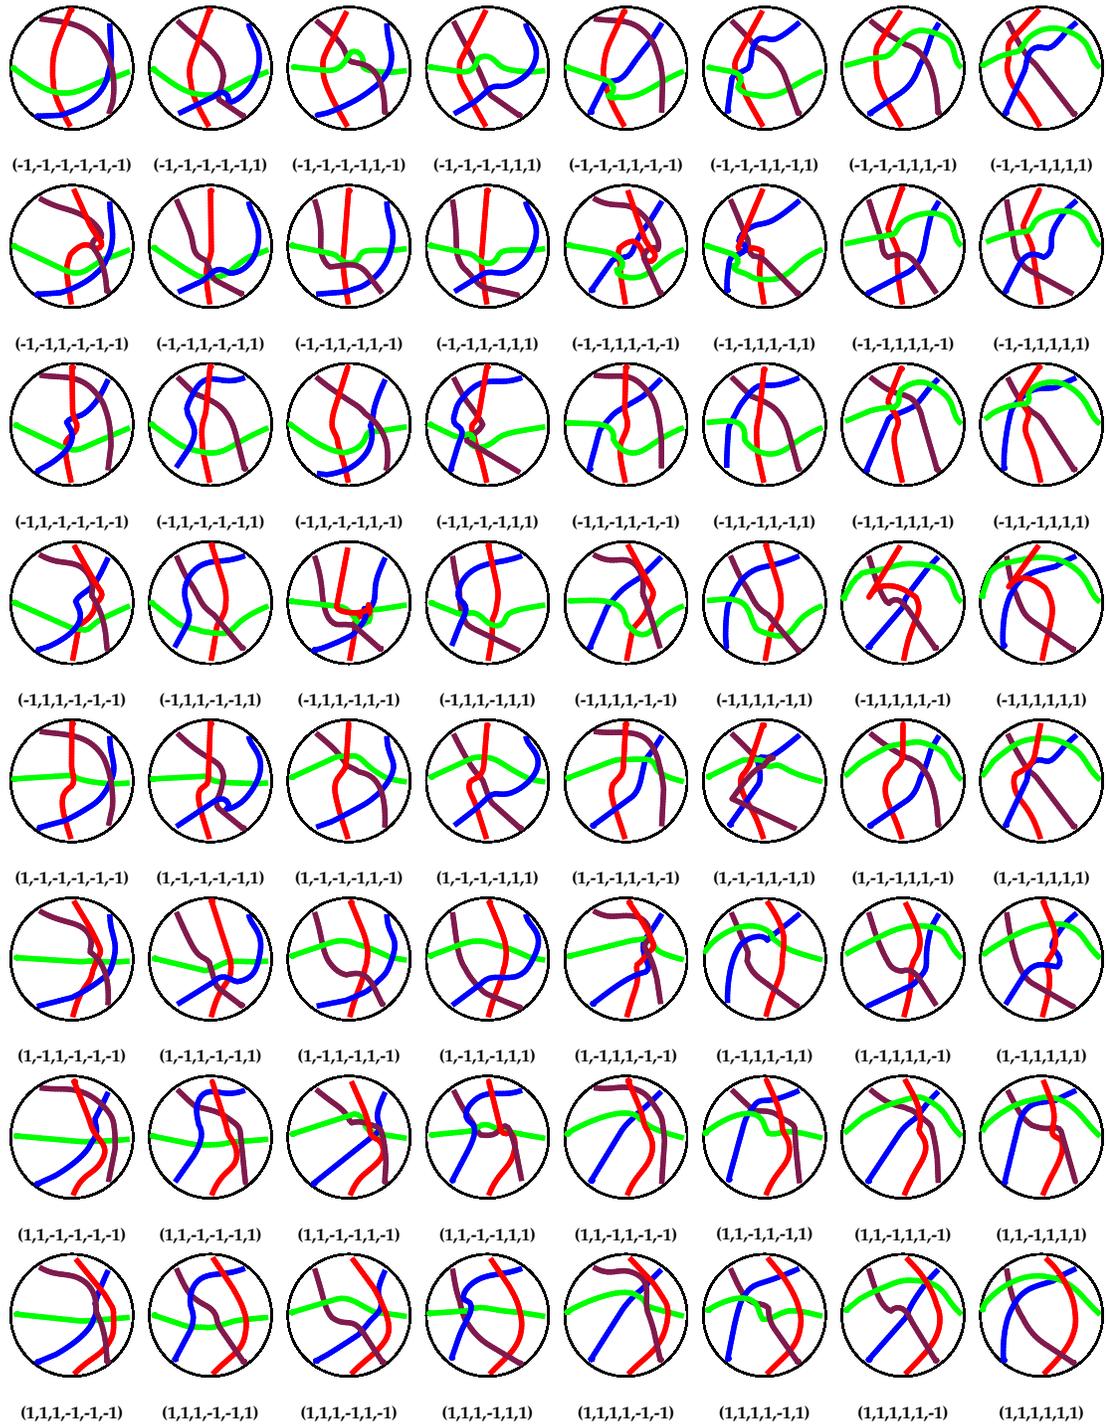
### 7.2.1 Offline Performance

We demonstrate the performance of the offline motion planner in generating topologically distinct, multi-agent navigation trajectories. We consider 4 dif-



**Figure 7.4: Top view of trajectories generated by executing the same 3-agent scenario with all possible topological specifications. The subcaptions denote the topology tuple that was used as a specification for each execution.**

ferent conditions, corresponding to different numbers of agents (2,3, 4 and 5 agents), navigating in a circular workspace of radius  $2.5m$  (agents are represented as disks of radius  $0.3m$ ). For each condition  $n \in \{2, 3, 4, 5\}$ , we randomly generate 100 distinct scenarios, by assigning agents initial and final locations that lead to challenging multi-agent encounters, requiring competent collision avoidance maneuvers. We execute each scenario,  $2^{\binom{n}{2}}$  times, each with a distinct topological specification. We measure the success rate of the planner in generating the desired topology under all conditions considered and report it in Table 7.1 (a trial is considered successful if the planner was able to produce all of the distinct topologies). The planner parameters were kept constant across conditions and scenarios. It can be observed that the planner performance drops as the number of agents  $n$  increases. The method becomes more sensitive to parameter tuning, as the effects of the chaotic nature of the vortex problem



**Figure 7.5: Top view of trajectories generated by executing the same 4-agent scenario with all possible topological specifications. The subcaptions denote the topology tuple that was used as a specification for each execution.**

Condition	2 agents	3 agents	4 agents	5 agents
Number of Outcomes	2	6	64	1024
Success (%)	1	99.75	89.70	65.48

**Table 7.1: Success rate of HTTG in generating the desired, topologically distinct executions for each of the 100 scenarios consider per condition.**

[8] become more significant. In appendix A, Figure 7.4 and Figure 7.5 depict the trajectories generated by following all possible topological specifications on an example 3-agent and scenario and an example 4-agent scenario respectively. Finally, Figure 7.7 shows examples of how the outlined trajectory generation mechanism may be used for online prediction in scenarios involving two, three and four agents.

## 7.2.2 Comparison with Trajectory Optimization

To the best of our knowledge, this is the first work that addresses the problem of generating trajectories for multiple agents in distinct topologies, specified a priori. Conceptually similar, the work of Rösmann et al. [119] considered the problem of generating multiple, topologically distinct trajectories for a single agent in a static environment with obstacles. However, optimizing multiple trajectories together and accounting for topological constraints while ensuring trajectory smoothness is a challenging problem. Typical gradient-based methods (e.g. [153]) act on the trajectory locally, with costs comprising several objectives; thus the gradient action could lose sight of the global, topological specification in favor of a different, local cost improvement. Furthermore, a differentiable cost function that would quantify the progress towards a desired topological outcome is hard to hand-design and we were not able to find any functions of that form in the literature. Our method constitutes a principled alternative

Planner	Success (%)	No. of iterations	Time (s)
CHOMP	78.80	80.3325	0.1291
HTTG	98.40	86.8862	0.0048

**Table 7.2: Success rates and computation times of HTTG and CHOMP over 500 randomly generated 2-agent scenarios.**

to trajectory optimization for this problem. Instead of locally reshaping a set of trajectories according to the gradients on the trajectory waypoints to attain local optima, our method grows the trajectories from initial conditions with a policy that has global knowledge of the desired trajectory topology. Similar to gradient-based optimization techniques, our method cannot guarantee the attainment of global optima. However, the physical encoding of the topological specification into the planning mechanism results in satisfactory performance for a class of problems.

To illustrate the difficulty of automatically synthesizing multi-agent trajectories of desired topological specifications through trajectory optimization techniques, we consider a simple case study, in which we compare the performance of HTTG with the performance of CHOMP [153]. We randomly generate 500 different scenarios involving 2 agents navigating towards opposing sides of a circular workspace (workspace has 5m diameter, starting positions are uniformly distributed along the circumference, speed normally distributed between  $0.3m/s$  and  $1.5m/s$  for each agent). For each scenario, we randomly sample a passing side that agents should pass one another. To encode the objective of respecting a passing side to CHOMP, we construct a cost functional  $\mathcal{F}_{top} = \frac{1}{2}(w_{ab} - w_{des})^2$  which approaches zero as the winding number of agents' trajectories  $w_{ab}$  approaches the winding number corresponding to the desired passing side,  $w_{des}$ . Table 7.2 illustrates the performance of the two approaches, which is measured with respect to success rate and computation time (non-optimized MatLab im-

plementation on a MacBook Pro of 2015 with an Intel Core i7 processor of 2.5 GHz, running macOS High Sierra). For CHOMP, a trial is considered successful if it generates trajectories of the desired topology within 500 iterations whereas for HTTPG, a trial is considered successful if the desired topology is achieved once the agents reach their destinations. It can be observed that HTTPG dominates with a success rate of 98.40% (corresponding to 492/500 successful trials). The computation time is comparable in terms of iterations but HTTPG requires almost two orders of magnitude less time in seconds. The benefits provided by HTTPG in terms of success rate and computation time make the consideration of multiple trajectory topologies at planning time a more practical strategy.

### 7.2.3 Online Performance

In order to demonstrate the benefits of our online algorithm 6, we perform a simulation study comprising a series of experiments on the circular workspace considered in the previous sections (diameter 5m). We consider 9 different experiment configurations, each corresponding to a different group of navigating agents. In particular, we consider groups of 2, 3 and 4 agents, navigating under three different conditions: (a) a homogeneous condition –all agents run the same planner; (b) a heterogeneous condition in which one agent runs our planner and others are moving straight to their goals without avoiding collisions; (c) a heterogeneous condition in which one agent runs our planner and others are changing intentions over a destination twice, without avoiding collisions. Note that the two latter cases are particularly challenging for decentralized planners, as a typical assumption they rely heavily on is homogeneity. For reference, we perform the same experiments with the Optimal Reciprocal Collision Avoid-

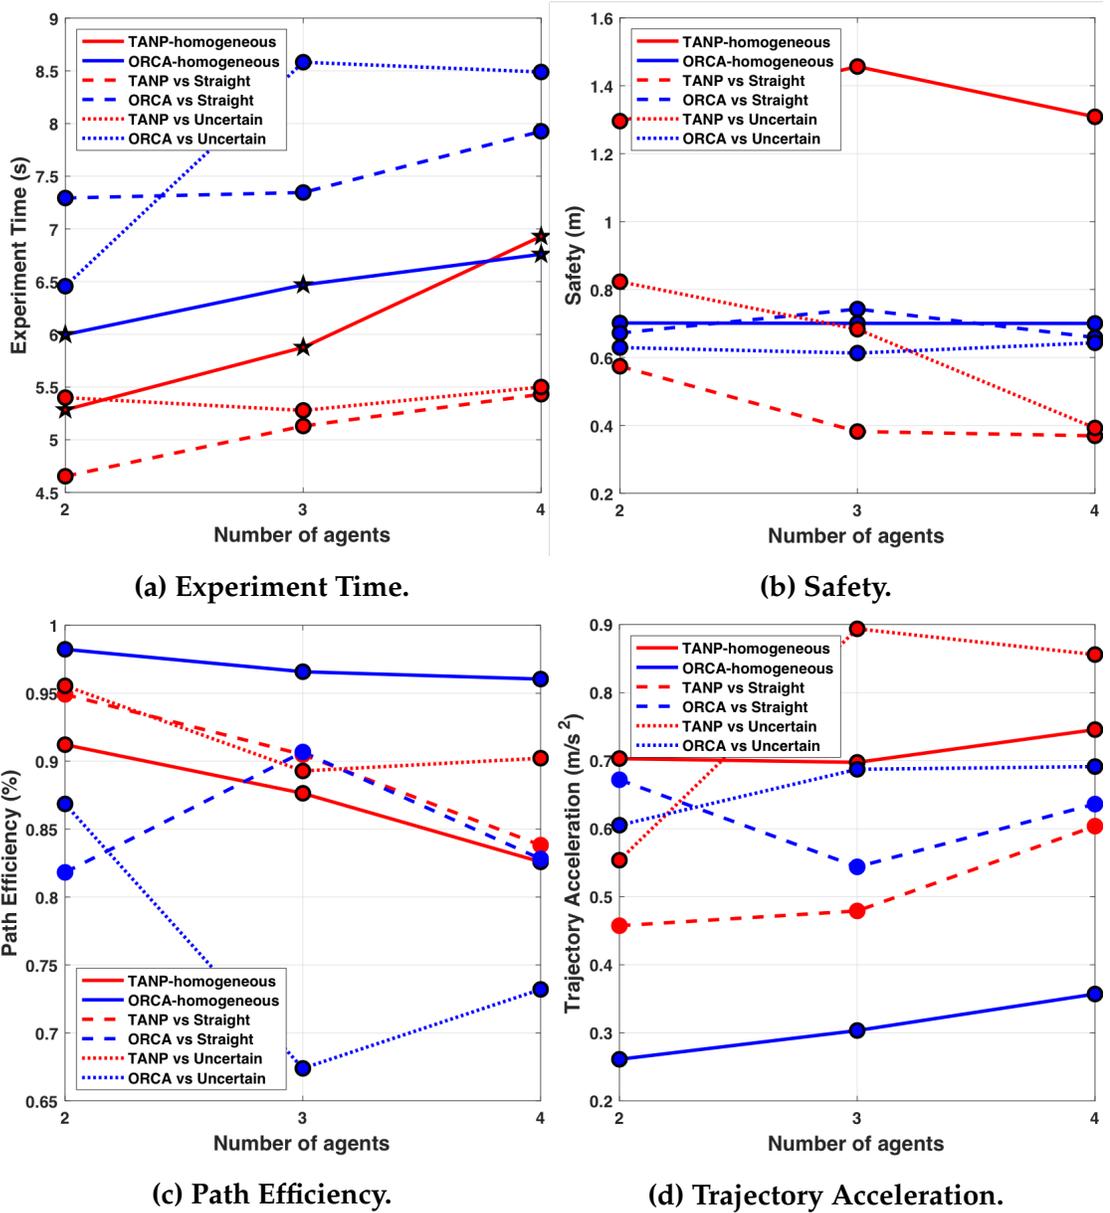
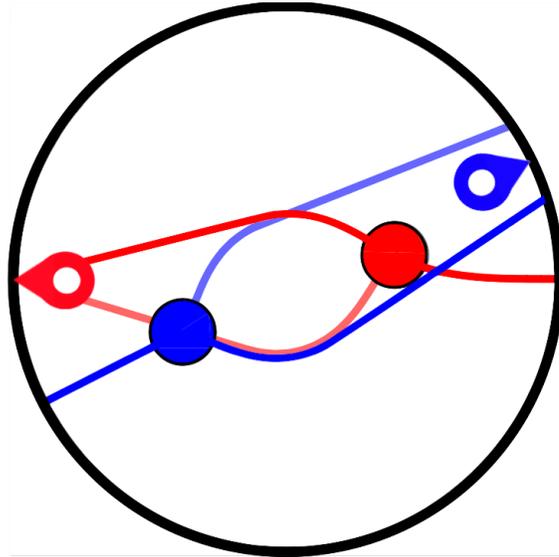
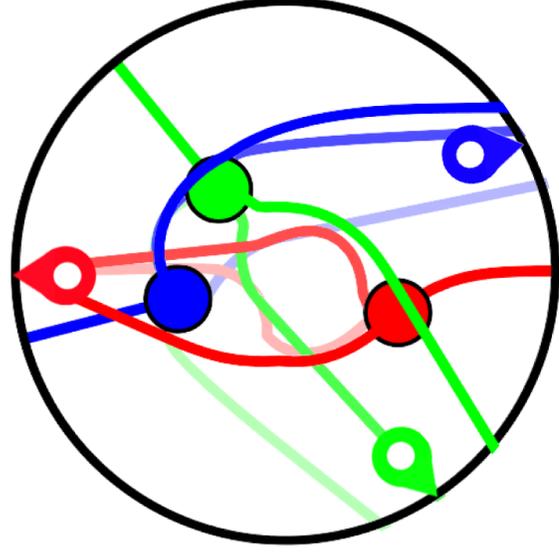


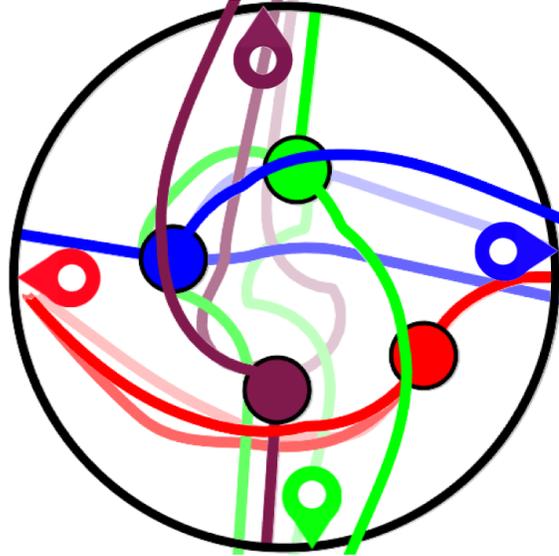
Figure 7.6: Trajectory Quality for all experiment configurations considered. For group size, the same 200 randomly generated scenarios are executed under each of the conditions considered with both planners. For each condition and measure, we perform a paired Student's t-test to compare the populations yielded by TANP and ORCA. Points with black circular boundaries indicate rejection of the null hypothesis with  $p$ -value  $< 0.001$  whereas points with star boundaries indicate rejection of the null hypothesis with  $p$ -value  $< 0.05$ .



(a) Two agents.



(b) Three agents.



(c) Four agents.

Figure 7.7: Overlaid predictions made by a TANP agent (red color) as it navigates towards the red landmark in environments with 2, 3 and 4 agents.

ance (ORCA) [142] (clearance and speed parameters tuned similarly to ensure a fair comparison). We quantify the performance of the planners with respect to four aspects of trajectory quality: (1) Experiment time, measured as the amount of time that the last agent to reach its destination took; (2) Safety, measured as the minimum distance between any two agents for the homogeneous condition and as the minimum distance between a TANP/ORCA agent and any other agent for the heterogeneous conditions; (3) Path Efficiency, measured as the ratio between the length of the optimal path to goal and the length of the path a TANP/ORCA agent followed (averaged over agents in the homogeneous case); (4) Trajectory Acceleration, measured as the average acceleration per time step per TANP/ORCA agent throughout the experiment.

Figure 7.6 depicts the performance of TANP and ORCA under each of the configurations considered. For each configuration, each planner executed the same set of 200, randomly generated scenarios. Overall, TANP exhibits the best time-efficiency for almost all configurations (Figure 7.6a). When executed under homogeneous settings, TANP establishes a consistently high clearance from others, which results in a drop in terms of path efficiency (Figure 7.6c) and a high acceleration per time step (Figure 7.6d). The increased time efficiency of TANP could be attributed to the implicit consensus that is reached through the consideration of joint strategies of collision avoidance. The price the TANP agents pay is increased accelerations and generally lower path efficiency. On the other hand, ORCA is consistently slower but stably safe across all conditions. Under the homogeneous condition, it achieves the highest path efficiency and lowest acceleration, which was expected by its optimality-driven design. This efficiency advantage fades under the heterogeneous conditions, in contrast to TANP, which demonstrates a more balanced behavior.

### 7.3 Discussion

We presented an online planning framework for the generation of adaptive robot motion in dynamic environments where multiple other agents navigate by executing generally different policies. Our framework is based on an offline planner that generates a diverse set of multi-agent trajectory predictions. Each prediction prescribes a different, collision-avoiding behavior to the planning agent. The agent selects the prediction of lowest cost and executes the first action from it. This architecture enables an agent to make local decisions with a global outlook that allows for anticipation of any upcoming agent interactions and rapid adjustment to them. Simulated examples demonstrate the performance of the offline and online parts of our framework. Future work involves (a) evaluating our algorithm in environments of more complex geometry, (b) reducing its computational load by designing a mechanism that efficiently reuses past trajectories and (c) performing real-world experiments on a robot platform navigating in human environments.

## **Part V**

# **Discussion**

## CHAPTER 8

### CONCLUSION

In this section, we summarize the main results extracted within this thesis, comment on some limitations and provide directions for future work. Finally, we close with a note, highlighting the importance of benchmarking standards for the validation of social robot navigation systems and more broadly, any robotic systems that aspire to operate in close proximity or in collaboration with humans.

#### 8.1 Summary

Deploying a fully autonomous mobile robotic system in a crowded human environment, such as a pedestrian scene requires a multitude of system components, including: (a) a map or an online mapping mechanism; (b) a localization mechanism; (c) a decision-making component, involving a mechanism for motion planning. This thesis considered the latter one, specifically focusing on the design of planning algorithms for the generation of robot motion that is intent-expressive, intention-aware and socially compliant, towards improving human comfort and lowering human cognitive load. The emphasis on intent expressiveness and intention awareness was motivated by the insights of studies on human action interpretation, suggesting that human inference is goal-directed [12, 32, 33]. The robotics community has already embraced these insights in research conducted within the area of human-robot interaction resulting in intriguing findings highlighting the importance of legibility or readability for applications such as collaborative manipulation [40] and navigation in the proximity of human subjects [25, 87].

Inspired by these works, the present thesis proposes the incorporation of models of the complex multi-agent dynamics into the robot’s decision-making process. Observing the topological structure of multi-agent collision avoidance behaviors, the proposed models share a common topological foundation, expressed through a series of mathematical constructions such as braids [97, 104], the physical quantity of angular momentum [107], or the winding number [101]. These constructions provide a systematic way of enumerating, labeling and analyzing multi-agent navigation behaviors but also allow for the design of computational mechanisms for online inference and planning, through the use of analytical and computational tools from low-dimensional topology [20, 42], fluid dynamics [8, 131] or generally Hamiltonian dynamics [18]. In practice, these mechanisms allow for a combined treatment of the problems of motion prediction and motion generation, by incorporating models of inference into the process of planning. This architecture enables the robot to anticipate the communicative effects of its actions on the inference processes of nearby humans and select actions that clearly communicate its intention of compliance with the perceived human navigation intentions or preferences.

The benefits of our proposed decision-making schemes can be traced in the accelerated uncertainty reduction [97, 104] of artificial agents as well as their simplified patterns of spatiotemporal, multi-agent trajectory entanglement [107] in simulated scenarios, under challenging interaction settings. Moreover, evidence extracted from the responses of more than 180 human subjects in an online, video-based user study suggests that humans perceive the behaviors generated by one of our algorithms as more legible than a set of state-of-the-art baselines [107]. Finally, evidence extracted from a large-scale (105 human participants) experimental lab study [108] suggests that human participants tend

to follow trajectories of lower acceleration when navigating around a robot running our Social Momentum algorithm [107].

## 8.2 Assumptions, Evaluation and Limitations

The core of this thesis are the technical Parts II, III and IV, presenting three different approaches for planning socially competent robot navigation. Although all of them are united under the same main theme, each one considers a slightly different setup, focusing on a different aspect of socially competent navigation. For this reason, different performance criteria were employed to validate each of them. In this section, we summarize the general assumptions and design decisions made across all parts and then focus on each part separately, explaining the distinct goals set and the evaluation process followed, in a comprehensive and comparative fashion. We also mention some limitations of the considered evaluation.

All frameworks share the same basic setup involving a set of multiple agents navigating from an initial position to a final configuration on a shared workspace. Agents are assumed to be able to generate velocities towards any direction dictated by the planning algorithm, i.e., we do not examine the effect of constraints, imposed by system dynamics. Furthermore, the agents are assumed to be acting rationally in the sense that they always aim at making progress towards their destinations. They also generally intend to avoid collisions with other agents that they encounter on their way. However, they are constrained to not explicitly exchange any information regarding their intended goals or planned paths with each other. Thus, in order to avoid collisions, they rely on the assumption

that others also employ collision-avoidance mechanisms and on their incentive to avoid collisions.

Our simulated evaluations considered an obstacle-free circular workspace and agents were set to move between antipodal points of its circumference. This was done to showcase the ability of our model to capture the multi-agent interactions without the obstruction of obstacles or the influence of the environment geometry. Again, this was done intentionally to remove the introduction of additional noise to our evaluations. However, this setup is a simplification that is rarely found in the real world.

Finally, a central idea in the whole thesis is that a prediction of the emerging topological pattern of a multi-agent trajectory in space-time could allow an artificial agent to understand the implications of its navigation behavior on the decisions and inference of other agents. It was hypothesized that this type of prediction (of qualitative nature) could allow for socially competent navigation in crowded environments. The proposed frameworks explore and investigate the validity, potential and implications of this idea for multi-agent navigation scenarios under a variety of different settings.

### **8.2.1 Reasoning about Multi-Agent Navigation Strategies**

Part II introduces braid groups [20], as a data structure that can represent a multi-agent navigation behavior. The expressiveness of this representation allows us to enumerate the set of all possible multi-agent navigation strategies that a set of agents could follow to navigate towards their destinations in a collision-free fashion. We leverage this feature to design a belief distribution

that enables a robot to reason about the uncertainty of the emerging topology of collision avoidance in a principled way. This belief informs the motion planning; in fact, at every time step the robot compromises between reducing the uncertainty of that belief with progress towards its destination. We evaluate this planning architecture in simulated scenarios involving multiple heterogeneous and homogeneous agents navigating discrete and continuous domains.

One of the benefits from the introduction of braids as a representation for multi-agent navigation is its reciprocity. A multi-agent trajectory embedded in the three dimensions of space and time may be abstracted into a topological braid upon projection onto a selected projection plane that includes time. This abstraction specifies a temporal sequence of trajectory crossings that correspond to pairwise passings between agents. These crossings are labeled by symbols (braid generators) which are put into a temporal order to form a braid word. This transition from 3 to 2 dimensions results in loss of information. More specifically, depending on the selection of this plane, a different braid (and thus, braid word) may emerge. Internally, as long as an agent is consistent throughout consequent time steps with a projection plane, this will not pose any representational issues. However, a set of agents that use different projection planes, may represent the same braid with different symbols (braid words). It should be noted though that these braids are homotopy-equivalent, i.e., one can be deformed to another through an operation involving continuous trajectory deformation without any trajectory intersections. In fact, in Part II, we demonstrated evidence that a set of agents employing different projection planes in their braid representations are indeed capable of achieving a rapid decrease of their uncertainty (entropy reduction) and converge to a state of consensus over a joint navigation strategy fast. We also demonstrated that even

heterogeneous agents, agnostic of the braid representation and lacking a belief over braids are capable of leveraging the prosocial behavior of another agent (that informs its path by reasoning about braid uncertainty) and benefiting by reaching their destinations faster.

Our evaluation is focusing on the spread of information and the uncertainty reduction. We chose this information-theoretic point of view as we wanted to evaluate the ability of our framework to enable a set of non-communicating agents to implicitly coordinate. We model uncertainty as the information entropy of the belief distribution over braids. We compute the entropy of the belief distribution for each agent and compute its average across all agents per time step. We report the convergence of this quantity over time, until all agents reach their destinations, and show that it decreases rapidly when agents run our algorithm. We also examine the implication of this observation for the task performance by measuring the time agents take to reach their destinations. We show that indeed when uncertainty drops faster (under our algorithm), agents arrive at their destinations faster.

## 8.2.2 Reasoning about Pairwise Collision Avoidance Intentions

Part III introduces the Social Momentum (SM) planning framework, a planner that attempts to retain the benefits of reasoning about joint strategies with a lower computational load. SM does so by ignoring joint strategies that do not include the robot, only considering the set of pairwise collision-avoidance maneuvers between the robot and all other agents. The planner generates an action that compromises between progress of the planning agent towards its destina-

tion and compliance with the preferences and intentions of other agents. This is implemented through the use of a heuristic based on the physical quantity of angular momentum. The planning agent repeatedly computes the pairwise angular momenta involving itself and each of the other agents and acts towards reinforcing their directions, while making progress towards its destination.

The SM framework was designed to generate legible navigation behaviors in a multi-agent environment. We focused on evaluating the performance of our framework in achieving this goal. In the absence of established metrics to do so, we introduced the topological quantity of the Topological Complexity Index as a proxy for assessing the legibility of multi-agent navigation behaviors. This metric quantifies the amount of spatiotemporal entanglement of a multi-agent trajectory, which is indicative of the complexity of the behavior that agents followed throughout the execution. We hypothesized that this quantity is related to the legibility of the trajectory followed by agents. Our findings from an online user study with more than 180 participants revealed evidence that indeed low complexity implies high legibility. Under the light of this finding, we conducted extensive simulated experiments involving homogeneous agents executing multi-agent navigation scenarios, comparing the performance of SM with two baselines: the Social Force (SF) model [62] and the Optimal Reciprocal Collision Avoidance (ORCA) framework [142]. It was found that indeed SM results in executions of low topological complexity, which we interpreted as evidence of the potential of our framework to generate legible behaviors in multi-agent environments. As a reference, and to ground the described finding to existing work, we also conducted a comparative analysis across baselines, with respect to the measure of Path Irregularity [57], a metric that characterizes the geometric complexity of a path by quantifying the amount of unnecessary

rotation that an agent follows as a result of its collision-avoidance maneuvers. It was found that the price that SM pays for increased legibility was significantly higher irregularity than ORCA, although still significantly lower than SF. It should be noted that in this part we did not consider the case of heterogeneous agents, i.e., we did not explore how our framework would perform in the presence of agents running different algorithms. This is a simplification in conflict with our scope of deploying our platform in a real-world environment; human navigation is vastly different than SM and still not well-understood or modeled.

We also explored the implications of our findings for real-world scenarios by conducting a lab study involving groups of humans navigating next to a navigating robot in a constrained workspace. We designed a task that enforced interesting and challenging mixing between the robot and human participants and considered a set of three different conditions corresponding to different robot navigation strategies: (a) SM; (b) ORCA; (c) Teleoperation. We collected human trajectories and responses of the participants to a questionnaire including questions rating the robot's intelligence, social compliance and safety and participants' emotional state while navigating next to the robot. We evaluated the collected human and robot trajectories by employing a set of physically grounded measures, including acceleration, energy, time to destination, path efficiency. Interesting findings include: (1) the confirmation of the expected planner traits in the presence of humans (ORCA was still the most efficient, SM had high accelerations, Teleoperation was the most cautious); (2) evidence that humans follow low-acceleration paths next to SM; (3) evidence that the topological complexity of the group (human-robot) behavior is significantly lower than the complexity of the autonomous conditions; (4) no significant variance of human ratings

across conditions.

We interpreted the finding of low-acceleration paths in the presence of a robot running SM as evidence of increased human comfort next to our framework. Low acceleration implies fewer switches between different speed levels, which we considered as an indication of smoother and more comfortable navigation. Further, we interpreted the low-complexity group behaviors in the presence of a teleoperated robot as the result of a significantly different planning horizon between the teleoperation and the two autonomous conditions. Specifically, we interpreted the low-complexity behaviors as the result of long-horizon planning from the part of the human teleoperator. The human anticipated multi-agent encounters more globally than the autonomous algorithms and took action earlier, avoiding complex interactions more effectively. Finally, we interpreted the low variance in human ratings across conditions as an indication that all navigation strategies were similarly perceived from the perspective of participants. This finding rejected our hypothesis and conventional expectation that teleoperated navigation would be perceived as more humanlike or intelligent by human participants (ANOVA performed on all questions did not demonstrate significant variance across conditions), which is an indication that the autonomous algorithms performed competitively enough next to the teleoperation.

All in all, it should be noted that more research should be done to establish our interpretations more concretely. From the correspondence between Topological Complexity and Legibility to the use of physically grounded metrics to quantify the performance of the planners, additional studies are necessary to prove the validity of such measures for the analysis of human and robot nav-

igation. Furthermore, additional research should be conducted to determine an appropriate set of questionnaires that shed light to humans' perceptions of competent robot navigation.

### 8.2.3 Topologically Robust Multi-Agent Trajectory Prediction

Part IV introduced a framework for on-line trajectory generation from topological specifications. The framework makes use of the same machinery as Social Momentum (Part III), focusing on the topology of the pairwise collision avoidance relationships among agents. However, in this part, we explore the foundations of this machinery. We observe that angular momentum –the core of the Social Momentum cost function– is directly related to the topological invariant of the winding number. In fact, the gradient of the winding number of a two-agent trajectory contributes velocities that follow the direction of maximum increase of the angular momentum (assuming constant speeds). We make use of this property to braid trajectories of desired winding, specified in a symbolic form comprising signs of pairwise winding numbers among agents. We demonstrate the robustness of our framework by reporting its performance in generating trajectories of desired topologies involving 2,3 and 4 agents. Our approach outperforms a trajectory-optimization baseline in terms of success rate and computation time. The proposed framework is also used for on-line motion planning in multi-agent navigation scenarios. We design a planning algorithm that iteratively generates a set of topologically distinct multi-agent trajectory predictions, essentially combining the processes of prediction and generation into a single, joint prediction step. We evaluate this algorithm in challenging multi-agent navigation scenarios with homogeneous and heterogeneous agents.

We show how this mechanism of joint trajectory prediction allows the planning agent to rapidly adapt to the behaviors of agents with changing intentions or agents with no collision-avoidance mechanisms. We report the performance of our framework in a comparative simulation study against the Social Force [62] and Optimal Reciprocal Collision Avoidance [142] frameworks. Through this analysis, we observe the benefits of our approach including high time efficiency, low acceleration and safe clearance from other agents.

In terms of our offline trajectory generation framework, as shown in our results section, our framework may fail when tasked with the generation of highly complex multi-agent behaviors. This can be due to the consideration of infeasible specifications and to the sensitivity of the approach to parameter tuning. As the number of agents and the complexity of the desired topological pattern increase, the sensitivity of the system to the parameters defining the balance between goal-direction and collision-avoidance increases. Local velocity decisions could have significant effects of global horizon to the resulting joint behavior. We showed that for systems with up to 4 agents, the success rate of our framework is always higher than 89%. However, it would be interesting and important to explore the stability properties of the approach and find out exactly under which conditions a failure (defined as the violation of a desired topological specification) arises. For systems with 2 and 3 agents, we see almost total stability. For systems with 4 agents the performance drops about 10% and in the case of 5 agents, we see a significant drop to 65%.

Similarly to the previous section, we must acknowledge that the selection of the evaluation criteria (energy, acceleration, time, clearance) could be debated and does not necessarily paint the complete picture of our framework. Again,

the criteria selected were physically grounded and well-motivated, but they inevitably leave several elements out. For example, we observed qualitatively that our framework may quickly adapt to the behaviors of heterogeneous and almost-adversarial agents. Of course, such adaptation requires motion of high acceleration, as reported in our evaluation. However, it was not clear how we could measure the property of adaptation in a quantitative and objective manner.

### **8.3 Future Work**

After the completion of the research presented in this thesis, a number of questions still remain unanswered. In this section, we list and expand on some of the most significant ones.

#### **8.3.1 The Role of Environment Geometry**

A strong assumption made throughout this thesis concerns the geometry of the environment boundary. We deliberately restricted ourselves to environments with a circular boundary or no boundary at all, as we wanted to focus on the topic of multi-agent interactions in an open space. The introduction of geometric constraints imposes specific patterns on the traffic flow that agents follow, enforcing different modes of interaction. Notable examples include collision avoidance on T-intersections or four-way intersections of hallways. Judging from our everyday-life experience, the mechanisms that we employ to avoid collisions on such challenging encounters are unclear and often determined on

the fly by trial-and-error. Understanding how to resolve such encounters or designing a framework for learning how to resolve them poses several interesting research questions.

### 8.3.2 Scalability

How many nearby pedestrians should the robot take into consideration when deciding how to navigate a crowded scene? How can we determine the number of human agents that are reactive and receptive to the communicative signals of the robot's actions? How do humans decide on these numbers? These are questions that we did not address in this thesis and for which we have not been able to find relevant quantitative evidence in the literature. This thesis presented evidence that incorporating an understanding of the unfolding multi-agent dynamics has beneficial properties for the whole system of agents, including improved coordination [104], high legibility [107] and low acceleration [108]. However, throughout all of the experiments conducted to acquire these results, we only considered systems of at most 5 agents. The theoretical properties of the topological models introduced in this thesis allow for representing any scenario with any number of agents. However, the computational load required to determine a navigation strategy increases exponentially with the number of agents (as shown in Chapter II). Our everyday-life experience suggests that it is not necessary to explicitly consider everyone around us when we plan our path in a crowded environment. Besides, the limited on-board sensing and perceptual capabilities of current robotic systems would not allow for the direct consideration of high numbers of agents.

### 8.3.3 Learning Models of Multi-Agent Interactions

In Chapter 4, we presented a learning framework for estimating the emerging multi-agent path topology that a set of agents engage in to resolve a collision. The presented model was trained on synthetic data, generated by simulations of multi-agent scenarios involving three and four agents. These numbers of agents allowed us to show the value of reasoning about multi-agent interactions. However, depending on the number of agents in the scene, we had to switch to a different model, trained on a dataset comprising examples involving the same number of agents. An interesting direction for future work could be the construction of a single model that can handle different numbers of agents by directing its attention to the subsets of agents that are more relevant at each time step.

### 8.3.4 Multi-Modal Legibility

In this thesis, we only considered the modality of path shape as a tool for intent-expressiveness. This decision was motivated by the fact that the path shape constitutes perhaps the most basic modality that all mobile robots –independently of their design– possess at the minimum. Clearly, humans make use of a variety of additional modalities to communicate clearly their intentions and objectives to others, while navigating. These include body language, eye gaze and less frequently, even arm gestures and verbal communication, and combinations of them. Understanding the mechanisms of multi-modal implicit communication in humans and transferring them to robots, equipped with the aforementioned modalities pose ambitious and interesting research questions. Related to that, it

would be interesting to explore the preferences of humans regarding the choice of signaling in artificial agents and study the effectiveness of different modality combinations.

### **8.3.5 Mechanisms Underlying Human Decision Making**

In the lab study presented in Chapter 6, it was shown that the Topological Complexity of the mixed group (robot and human participants) tends to be significantly lower when the robot is teleoperated, compared to when the robot is running one of the two autonomous conditions. We interpreted this finding as an indication of intrinsic differences in the decision-making mechanisms of humans and robot motion planners. Specifically, we conjecture that human decision-making mechanisms for navigation feature a more global horizon than robot motion planners. This results in the selection of actions that proactively avoid collisions by avoiding the regions of rich interaction. Of course, our study was limited in the number of planners considered (2 autonomous conditions) and in the number of teleoperators (only one across all experiments). We believe that it would be interesting to investigate this hypothesis further through a series of experiments. The consideration of additional motion planners and the recruitment of a diverse set of teleoperators would be direct ways of extending the findings of our study. However, new study designs could also help focus more directly on the investigation of this question. Overall, this investigation could be considered as a part of a more general research direction of studying the structure of human decision making, in navigation and beyond.

### 8.3.6 Topological Representations for Robotics

A central foundation of the work presented in this thesis comes from the field of low-dimensional Topology, the branch of Topology that focuses on topological spaces of up to four dimensions. More specifically, observing that the entanglement of the trajectories of multiple navigating agents over time carries information about their joint collision avoidance strategies, we sought a construction that would capture it. We found that the formalism of topological braids [20] provided a formal theory for enumerating all possible classes of entanglement and thus all possible classes of joint collision avoidance strategies that could emerge in a multi-agent scene. The use of this representation enabled us to directly construct an inference mechanism for predicting the unfolding multi-agent dynamics of interaction in a crowded scene. It also provided us with computational tools and metrics that were used for analysis and motion generation.

Braid groups constitute just an example of the many constructions in the field of Topology that could find use in Robotics. One such example is the use of annular braids [76]. Annular braids could be thought of as sets of strings entangled around an annulus and can be labeled through projection onto a virtual cylinder centered on the annulus. A potential benefit of this representation is that, unlike the standard braid formalism, it frees us from the decision of a projection plane (see Chapter II), which could potentially result to ambiguity between consecutive time steps. A potential problem with annular braids is that they have not been studied as extensively as the standard braid formalism. Another example is the extension to problems in 3 dimensions with the use of knots [3]. A knot is an embedding of a circle in  $\mathbb{R}^3$ . Knots form a group and

could be used as multi-agent behavior primitives for three-dimensional phenomena, in a similar fashion as braids were used for multi-agent navigation on the plane. This could allow for extensions to interesting problems in air-traffic control or multi-agent systems of quadrotors.

Overall, the main general benefit for the employment of topological representations for robotics applications arises from their powerful abstraction capabilities that allow for the encoding of complex behaviors into symbols. This naturally favors the development of methodologies for inference and symbolic planning that are of particular interest for robotics applications. One drawback of topological representations is that they naturally result in data compression and loss of information. For example, in order to label a multi-agent trajectory as a braid, we need to project it onto a selected plane. Once we do so, we lose information such as the clearance between agents or their velocity components that are orthogonal to the projection plane. This property may or may not be problematic, depending on the application. Caution however must be taken when transitioning from geometric representations to topological ones and vice versa. This observation could also motivate further research on the development of novel representations that result in desired levels of abstraction while retaining a sufficient amount of information for a given task.

## **8.4 The Need for Benchmarking**

This thesis makes a unique contribution in the area of social robot navigation. First, it contains a comprehensive presentation of the complete development process, from the conception of abstract mathematical models of multi-agent

navigation to the design of motion planning algorithms. Then, it presents an extensive validation process, ranging from simulations of challenging multi-agent scenarios to lab experiments with human participants.

Despite the wide interest in the development of social navigation frameworks, the majority of existing works contains simulated evaluations, under strong assumptions on the scenarios, environments and agent behaviors considered. Relatively fewer works contain qualitative evidence collected from limited in sample size experimental demonstrations whereas even fewer present findings from large-scale field studies, which tend to suffer from the noise induced by the uncontrolled variables of the real world. Particularly rare are the lab studies that make a rigorous validation of the fundamental properties of proposed frameworks. The design of such studies is challenging; isolating the exact elements that one desires to test requires a carefully selected setup, script and performance from the part of the experimenters. Furthermore, there is a considerable cost associated with acquiring the required equipment and compensating human subjects for their participation.

However, we argue that social navigation frameworks and more broadly, any human-robot interaction framework or technology that aspires to generate socially compliant behaviors and integrate seamlessly in a human environment requires an extensive and in-depth validation process that should verify the properties claimed by its developers. The importance of extensive in-lab evaluation for the transition from simulation to field deployment should not be overlooked. We hope that this thesis will constitute a paradigm and a source of inspiration for future studies to come. The field is under extensive growth and there is a multitude of topics to be explored to help develop a fully autonomous

and socially competent robot pedestrian.

The long process of conducting the research presented in this thesis uncovered a series of complications associated with the process of validation. First, very few works have released official open-source implementations of their proposed frameworks (e.g. the Optimal Reciprocal Collision Avoidance framework [142]). Some works are relatively easy to implement from scratch (e.g. the Social Force Model [62]) and thus also relatively directly accessible. However, others—including a significant portion of the state-of-the-art—are not easy to implement, as a result of the lack of public access to their datasets (e.g. [86, 138]). This renders any comparison with them practically impossible, especially considering the lack of existing well-formatted, noise-bounded, open-source datasets. In contrast to other fields (e.g. Computer Vision), in which a multitude of well-formatted and openly accessible datasets exist to train on, in the area of robot navigation, to the best of our knowledge, no such datasets exist.

This situation calls for a more active engagement of the research community in the area of social navigation. The establishment of a concrete set of benchmarking standards appears to be a significant step that the field needs to make in order to ensure the quantification of the progress and the definition of future goals. For example, an extension of the protocol proposed by Sprunk et al. [124] could include a set of validated instruments and tasks that focus on the evaluation of social compliance and human comfort.

## Bibliography

- [1] Suitable technologies, Inc. <https://suitabletech.com/>. Accessed: 2018-09-13.
- [2] Sam Gross Adam Paszke and Soumith Chintala. PyTorch. <https://github.com/pytorch/pytorch>. Accessed: [01/28/2017].
- [3] Colin C. Adams. *The Knot Book*. W.H. Freeman, 1994.
- [4] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In *Proceedings of the 2016 IEEE International Conference on Computer Vision and Pattern Recognition (CVPR '16)*, 2016.
- [5] Rachid Alami, Aurélie Clodic, Vincent Montreuil, Emrah Akin Sisbot, and Raja Chatila. Task planning for human-robot interaction. In *Proceedings of the 2005 joint conference on Smart objects and ambient intelligence: innovative context-aware services: usages and technologies*, pages 81–85. ACM, 2005.
- [6] P. Althaus, H. Ishiguro, T. Kanda, T. Miyashita, and H. I. Christensen. Navigation for human-robot interaction tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1894–1900 Vol.2, 2004.
- [7] Hassan Aref. Point vortex dynamics: A classical mathematics playground. *Journal of Mathematical Physics*, 48(6):065401, 2007.
- [8] Hassan Aref, Scott W. Jones, S. Mofina, and Irek Zawadzki. Vortices, kinematics and chaos. *Physica D: Nonlinear Phenomena*, 37(1):423 – 440, 1989.

- [9] Hassan Aref, Nicholas Rott, and Hans Thomann. Gröbli’s solution of the three-vortex problem. *Annual Review of Fluid Mechanics*, 24(1):1–21, 1992.
- [10] E. Artin. Theory of braids. *Annals of Mathematics*, 48(1):pp. 101–126, 1947.
- [11] E. Artin. Braids and permutations. *Annals of Mathematics. Second Series*, 48:643–649, 1947. ISSN 0003-486X; 1939-8980/e.
- [12] C. L. Baker, J. Tenenbaum, and R. Saxe. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
- [13] Tirthankar Bandyopadhyay, Kok Sung Won, Emilio Frazzoli, David Hsu, Wee Sun Lee, and Daniela Rus. Intention-aware motion planning. In *Proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2012.
- [14] Jack M Barbalet. Social emotions: confidence, trust and loyalty. *International Journal of Sociology and Social Policy*, 16(9/10):75–96, 1996.
- [15] Christoph Bartneck, Elizabeth Croft, and Dana Kulic. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1):71–81, 2009.
- [16] Maren Bennewitz, Wolfram Burgard, Grzegorz Cielniak, and Sebastian Thrun. Learning motion patterns of people for compliant robot motion. *International Journal of Robotics Research*, 24:31–48, 2005.
- [17] Aniket Bera, Tanmay Randhavane, Rohan Prinja, and Dinesh Manocha. Sociosense: Robot navigation amongst pedestrians with social and psychological constraints. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7018–7025, 2017.

- [18] Mitchell A Berger. Hamiltonian dynamics generated by Vassiliev invariants. *Journal of Physics A: Mathematical and General*, 34(7):1363, 2001.
- [19] Mitchell A. Berger. Topological invariants in braid theory. *Letters in Mathematical Physics*, 55(3):181–192, 2001.
- [20] Joan S. Birman. *Braids Links And Mapping Class Groups*. Princeton University Press, 1975.
- [21] Stephane Bonneaud and William H. Warren. An empirically-grounded emergent approach to modeling pedestrian behavior. In *Proceedings of the Conference on Pedestrian and Evacuation Dynamics (PED)*, pages 625–638. Springer International Publishing, 2014.
- [22] Michael E Bratman. Shared cooperative activity. *The philosophical review*, 101(2):327–341, 1992.
- [23] Baptiste Busch, Jonathan Grizou, Manuel Lopes, and Freek Stulp. Learning legible motion from human–robot interactions. *International Journal of Social Robotics*, Mar 2017.
- [24] Daniel Carton, Verena Nitsch, Dominik Meinzer, and Dirk Wollherr. Towards assessing the human trajectory planning horizon. *PLOS ONE*, 11(12):1–39, 12 2016.
- [25] Daniel Carton, Wiktor Olszowy, and Dirk Wollherr. Measuring the effectiveness of readability for mobile robot locomotion. *International Journal of Social Robotics*, 8(5):721–741, 2016. ISSN 1875-4805.
- [26] Elizabeth Cha, Anca D. Dragan, and Siddhartha S. Srinivasa. Perceived robot capability. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 541–548, 2015.

- [27] Konstantinos Charalampous, Ioannis Kostavelis, and Antonios Gasteratos. Recent trends in social aware robot navigation: A survey. *Robotics and Autonomous Systems*, 93:85 – 104, 2017.
- [28] Yu Fan Chen, Michael Everett, Miao Liu, and Jonathan P. How. Socially aware motion planning with deep reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1343–1350, 2017.
- [29] Sanjiban Choudhury. *Adaptive Motion Planning*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, February 2018.
- [30] Herbert H. Clark. *Using Language*. Cambridge University Press, 1996.
- [31] Herbert H. Clark and Deanna Wilkes-Gibbs. Referring as a collaborative process. *Cognition*, 22(1):1–39, 1986.
- [32] Gergely Csibra and György Gergely. The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental Science*, 1(2):255–259, 1998.
- [33] Gergely Csibra and György Gergely. ‘Obsessed with goals’: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, 124(1):60–78, January 2007.
- [34] José Grimaldo Da Silva Filho and Thierry Fraichard. Human Robot Motion: A Shared Effort Approach. In *European Conference on Mobile Robotics*, 2017.
- [35] Nikhil Chavan Dafle, Rachel Holladay, and Alberto Rodriguez. In-hand manipulation via motion cones. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*, 2018.

- [36] Yancy Diaz-Mercado and Magnus Egerstedt. Multirobot mixing via braid groups. *IEEE Transactions on Robotics*, 33(6):1375–1385, 2017.
- [37] Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.
- [38] Minh Binh Do and Subbarao Kambhampati. Planning graph-based heuristics for cost-sensitive temporal planning. In *Proceedings of the International Conference on Artificial Intelligence Planning Systems (AIPS)*, pages 3–12, 2002.
- [39] Mehmet R. Dogar and Siddhartha S. Srinivasa. A planning framework for non-prehensile manipulation under clutter and uncertainty. *Autonomous Robots*, 33(3):217–236, 2012.
- [40] Anca D. Dragan and Siddhartha Srinivasa. Integrating human observer inferences into robot motion planning. *Autonomous Robots*, 37(4):351–368, 2014.
- [41] Anca D. Dragan, Kenton CT Lee, and Siddhartha S. Srinivasa. Legibility and predictability of robot motion. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction (HRI)*, pages 301–308, 2013.
- [42] Ivan Dynnikov and Bert Wiest. On the complexity of braids. *Journal of the European Mathematical Society*, 009(4):801–840, 2007.
- [43] Michael A. Erdmann and Matthew T. Mason. An exploration of sensorless manipulation. *IEEE Journal on Robotics and Automation*, 4(4):369–379, 1988.
- [44] Ronald Fagin, Joseph Y Halpern, Yoram Moses, and Moshe Vardi. *Reasoning about knowledge*. MIT press, 2004.

- [45] Francesco Farina, Daniele Fontanelli, Andrea Garulli, Antonio Giannitrapani, and Domenico Prattichizzo. Walking ahead: The headed social force model. *PLOS ONE*, 12(1):1–23, 2017.
- [46] Sarah Ferguson, Brandon Luders, Robert C. Grande, and Jonathan P. How. Real-time predictive modeling and robust avoidance of pedestrians with uncertain, changing intentions. In *Proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, pages 161–177, Cham, 2015. Springer International Publishing.
- [47] Franck Feurtey. Simulating the collision avoidance behavior of pedestrians. Master’s thesis, University of Tokyo, Tokyo, Japan, 2000.
- [48] Amalia F. Foka and Panos E. Trahanias. Probabilistic autonomous robot navigation in dynamic environments with human motion prediction. *International Journal of Social Robotics*, 2(1):79–94, 2010.
- [49] Robert R. Ghrist. Configuration spaces and braid groups on graphs in robotics, 1999. URL <https://arxiv.org/abs/math/9905023>.
- [50] Julio Godoy, Ioannis Karamouzas, Stephen J. Guy, and Maria Gini. Moving in a crowd: Safe and efficient navigation among heterogeneous agents. *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 294–300, 2016.
- [51] Erving Goffman. *Behavior in Public Places: Notes on the Social Organization of Gatherings*. Free Press, 1966.
- [52] Erving Goffman. *Relations in public: microstudies of the public order*. New York: Basic Books, 1971.

- [53] Matthew C. Gombolay, Ronald Wilcox, and Julie A. Shah. Fast scheduling of multi-robot teams with temporospatial constraints. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*, 2013.
- [54] Noah D. Goodman and Andreas Stuhlmüller. Knowledge and implicature: Modeling language understanding as social cognition. *Topics in cognitive science*, 5(1):173–184, 2013.
- [55] Colin Green and Alonzo Kelly. Toward optimal sampling in the space of paths. In *Proceedings of the International Symposium of Robotics Research (ISRR)*, 2007.
- [56] Herbert P. Grice. Logic and conversation. *Syntax and Semantics*, pages 41–58, 1975.
- [57] Jérôme Guzzi, Alessandro Giusti, Luca M. Gambardella, Guy Theraulaz, and Gianni A. Di Caro. Human-friendly robot navigation in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 423–430, 2013.
- [58] Edward T. Hall. *The Hidden Dimension*. Garden City, NY: Anchor Books, 1990.
- [59] Bradley Hayes and Brian Scassellati. Discovering task constraints through observation and active learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4442–4449, 2014.
- [60] Bradley Hayes and Brian Scassellati. Effective robot teammate behaviors for supporting sequential manipulation tasks. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6374–6380, 2015.

- [61] Peter Hedström and Charlotta Stern. Rational choice and sociology. *The new Palgrave dictionary of economics*, pages 872–877, 2008.
- [62] Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51:4282–4286, 1995.
- [63] Peter Henry, Christian Vollmer, Brian Ferris, and Dieter Fox. Learning to navigate through crowded environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 981–986, 2010.
- [64] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [65] Jakob Hohwy. *The Predictive Mind*. Oxford University Press, 2013.
- [66] Rachel Holladay, Anca D. Dragan, and Siddhartha S. Srinivasa. Legible robot pointing. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2014.
- [67] Serge Hoogendoorn and Piet H. L. Bovy. Simulation of pedestrian flows by optimal control and differential games. *Optimal Control Applications and Methods*, 24(3):153–172, 2003.
- [68] J. Hu, M. Prandini, and S. Sastry. Optimal coordinated motions of multiple agents moving on a plane. *SIAM Journal on Control and Optimization*, 42(2):637–668, 2003.
- [69] Mubbasir Kapadia, Shawn Singh, William Hewlett, Glenn Reinman, and Petros Faloutsos. Parallelized egocentric fields for autonomous navigation. *The Visual Computer*, 28(12):1209–1227, 2012.

- [70] Ioannis Karamouzas, Peter Heil, Pascal van Beek, and Mark H. Overmars. A predictive collision avoidance model for pedestrian simulation. In *Motion in Games*, pages 41–52. Springer Berlin Heidelberg, 2009.
- [71] Ioannis Karamouzas, Brian Skinner, and Stephen J. Guy. Universal power law governing pedestrian interactions. *Physical Review Letters*, 113:238701, Dec 2014.
- [72] David A. Karp, Gregory P. Stone, and William C. Yoels. *Being urban: a social psychological view of city life*. Lexington, MA: Heath, 1977.
- [73] Christian Kassel and Vladimir Turaev. *Braid Groups*, volume 247 of *Graduate Texts in Mathematics*. New York, NY: Springer, 2008.
- [74] Yusuke Kato, Takayuki Kanda, and Hiroshi Ishiguro. May I help you?: Design of human-like polite approaching behavior. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 35–42, 2015.
- [75] Naoki Katoh, Toshihide Ibaraki, and Hisashi Mine. An efficient algorithm for k shortest simple paths. *Networks*, 12(4):411–427, 1982.
- [76] Richard P. Kent and David Peifer. A geometric and algebraic description of annular braid groups. *International Journal of Algebra and Computation*, 12(01n02):85–97, 2002.
- [77] Hassan K. Khalil. *Nonlinear Systems*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition, 1996.
- [78] Beomjoon Kim and Joelle Pineau. Socially adaptive path planning in human environments using inverse reinforcement learning. *International Journal of Social Robotics*, 8(1):51–66, 2016.

- [79] Rachel Kirby. *Social Robot Navigation*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, May 2010.
- [80] Rachel Kirby, Reid Simmons, and Jodi Forlizzi. COMPANION: A constraint-optimizing method for person-acceptable navigation. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 607–612, 2009.
- [81] Ross A. Knepper and Daniela Rus. Pedestrian-inspired sampling-based multi-robot collision avoidance. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 94–100, 2012.
- [82] Ross A. Knepper, Siddhartha S. Srinivasa, and Matthew T. Mason. Toward a deeper understanding of motion alternatives via an equivalence relation on local paths. *The International Journal of Robotics Research*, 31(2):167–186, 2012.
- [83] Ross A. Knepper, Stefanie Tellex, Adrian Li, Nicholas Roy, and Daniela Rus. Recovering from failure by asking for help. *Autonomous Robots*, 39(3):347–362, 2015.
- [84] Ross A. Knepper, Christoforos Mavrogiannis, Julia Proft, and Claire Liang. Implicit communication in a joint action. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI '17)*, pages 283–292, 2017.
- [85] Heather Knight. *Expressive Motion for Low Degree-of-Freedom Robots*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, August 2016.

- [86] Henrik Kretzschmar, Markus Spies, Christoph Sprunk, and Wolfram Burgard. Socially compliant mobile robot navigation via inverse reinforcement learning. *The International Journal of Robotics Research*, 35(11):1289–1307, 2016.
- [87] Thibault Kruse, Patrizia Basili, Stefan Glasauer, and Alexandra Kirsch. Legible robot navigation in the proximity of moving humans. In *Proceedings of the IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*, pages 83–88, 2012.
- [88] Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726–1743, 2013.
- [89] Minae Kwon, Malte F. Jung, and Ross A. Knepper. Human expectations of social robots. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 463–464, 2016.
- [90] Chi-Pang Lam, Chen-Tun Chou, Kuo-Hung Chiang, and Li-Chen Fu. Human-centered robot navigation—towards a harmoniously human–robot coexisting environment. *IEEE Transactions on Robotics*, 27(1): 99–112, 2011.
- [91] Shalom Lappin and Chris Fox. *The handbook of contemporary semantic theory*. John Wiley & Sons, 2015.
- [92] Min Kyung Lee, Sara Kiesler, Jodi Forlizzi, Siddhartha S. Srinivasa, and Paul Rybski. Gracefully mitigating breakdowns in robotic services. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 203–210, 2010.

- [93] Séverin Lemaignan and Pierre Dillenbourg. Mutual modelling in robotics: Inspirations for the next steps. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 303–310, 2015.
- [94] Matthias Luber, Luciano Spinello, Jens Silva, and Kai O. Arras. Socially-aware robot navigation: A learning approach. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 902–907, 2012.
- [95] Wei-Chiu Ma, De-An Huang, Namhoon Lee, and Kris M. Kitani. Forecasting interactive dynamics of pedestrians with fictitious play. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4636–4644, 2017.
- [96] Jim Mainprice and Dmitry Berenson. Human-robot collaborative manipulation planning using early prediction of human motion. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 299–306, 2013.
- [97] Christoforos Mavrogiannis and Ross A. Knepper. Decentralized multi-agent navigation planning with braids. In *Proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.
- [98] Christoforos Mavrogiannis and Ross A. Knepper. Towards socially competent navigation of pedestrian environments. In *Workshop on Social Trust in Autonomous Robots, Robotics: Science and Systems Conference (RSS)*, 2016.
- [99] Christoforos Mavrogiannis and Ross A. Knepper. Interpretation and communication of pedestrian intentions using braid groups. In *Workshop on*

*Intention Recognition in HRI, ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2016.*

- [100] Christoforos Mavrogiannis and Ross A. Knepper. Designing algorithms for socially competent robotic navigation. In *Proceedings of the Companion of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 357–358, 2017.
- [101] Christoforos Mavrogiannis and Ross A. Knepper. Multi-agent trajectory prediction and generation with topological invariants enforced by hamiltonian dynamics. In *Proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2018.
- [102] Christoforos Mavrogiannis and Ross A. Knepper. Decentralized navigation planning using multi-agent trajectory prediction governed by hamiltonian dynamics. In *Workshop on Multi-robot Perception-Driven Control and Planning, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [103] Christoforos Mavrogiannis and Ross A. Knepper. Online multi-agent trajectory generation for adaptive navigation planning. In *Pioneers Workshop, Robotics: Science and Systems Conference (RSS)*, 2018.
- [104] Christoforos Mavrogiannis and Ross A. Knepper. Multi-agent path topology in support of socially competent navigation planning. *The International Journal of Robotics Research*, 2018.
- [105] Christoforos Mavrogiannis, Valts Blukis, and Ross A. Knepper. Socially competent navigation planning by deep learning of multi-agent path

- topologies. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6817–6824, 2017.
- [106] Christoforos Mavrogiannis, Valts Blukis, and Ross A. Knepper. Inferring multi-agent path topologies: Towards socially competent navigation in human environments. In *Workshop on Mathematical Models, Algorithms and Human-Robot Interaction, Robotics: Science and Systems Conference (RSS)*, 2017.
- [107] Christoforos Mavrogiannis, Wil B. Thomason, and Ross A. Knepper. Social momentum: A framework for legible navigation in dynamic multi-agent environments. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 361–369, 2018.
- [108] Christoforos Mavrogiannis, Alena M. Hutchinson, John Macdonald, Patrícia Alves-Oliveira, and Ross A. Knepper. Effects of distinct robotic navigation strategies on human behavior in a crowded environment. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2019.
- [109] Mehdi Moussaïd, Dirk Helbing, and Guy Theraulaz. How simple rules determine pedestrian behavior and crowd disasters. *Proceedings of the National Academy of Sciences*, 108(17):6884–6888, 2011.
- [110] Kunio Murasugi and Bohdan Kurpita. *A Study of Braids*. Mathematics and Its Applications. Kluwer Academic Publishers, 1999.
- [111] Douglas W. Nangle, David J. Hansen, Cynthia A. Erdley, and Peter J. Norton. *Practitioner’s guide to empirically based measures of social skills*. Springer Science & Business Media, 2009.

- [112] Stefanos Nikolaidis, David Hsu, and Siddhartha S. Srinivasa. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, 36(5-7):618–634, 2017.
- [113] Elena Pacchierotti, Henrik I. Christensen, and Patric Jensfelt. Evaluation of passing distance for social robots. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 315–320, 2006.
- [114] Jong Jin Park, Collin Johnson, and Benjamin Kuipers. Robot navigation with model predictive equilibrium point control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4945–4952, 2012.
- [115] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc J. Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 261–268, 2009.
- [116] Martin J. Pickering and Simon Garrod. Toward a mechanistic psychology of dialogue. *Behavioral and brain sciences*, 27(02):169–190, 2004.
- [117] David Richeson. The maypole braid group. <https://divisbyzero.com/2009/05/04/the-maypole-braid-group/>, 2009. Accessed: 2019-03-19.
- [118] Laurel D. Riek. Wizard of oz studies in HRI: a systematic review and new reporting guidelines. *Journal of Human-Robot Interaction*, 1(1):119–136, 2012.

- [119] Christoph Rösmann, Frank Hoffmann, and Torsten Bertram. Integrated online trajectory planning and optimization in distinctive topologies. *Robotics and Autonomous Systems*, 88:142 – 153, 2017.
- [120] Dorsa Sadigh, Shankar Sastry, Sanjit A. Seshia, and Anca D. Dragan. Planning for autonomous cars that leverages effects on human actions. In *Proceedings of the Robotics: Science and Systems (RSS) Conference*, pages 239–248, 2016.
- [121] Paul Scovanner and Marshall F. Tappen. Learning pedestrian dynamics from the real world. In *IEEE International Conference on Computer Vision (CVPR)*, pages 381–388, 2009.
- [122] Masahiro Shiomi, Francesco Zanlungo, Kotaro Hayashi, and Takayuki Kanda. Towards a socially acceptable collision avoidance for a mobile robot navigating among pedestrians using a pedestrian model. *International Journal of Social Robotics*, 6(3):443–455, 2014.
- [123] Emrah Akin Sisbot, Luis Felipe Marin-Urias, Rachid Alami, and Thierry Siméon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23(5):874–883, 2007.
- [124] Christoph Sprunk, Jörg Röwekämper, Gershon Parent, Luciano Spinello, Gian Diego Tipaldi, Wolfram Burgard, and Mihai Jalobeanu. An experimental protocol for benchmarking robotic indoor navigation. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*, pages 487–504, Cham, 2016. Springer International Publishing.
- [125] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural net-

- works from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [126] Kyle Strabala, Min Kyung Lee, Anca Dragan, Jodi Forlizzi, Siddhartha S. Srinivasa, Maya Cakmak, and Vincenzo Micelli. Toward seamless human-robot handovers. *Journal of Human-Robot Interaction*, 2(1):112–132, 2013.
- [127] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [128] Stefanie Tellex, Ross A. Knepper, Adrian Li, Nicholas Roy, and Daniela Rus. Asking for help using inverse semantics. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*, 2014.
- [129] Stefanie A. Tellex, Thomas Kollar, Steven R. Dickerson, Matthew R. Walter, Ashis Banerjee, Seth Teller, and Nicholas Roy. Understanding natural language commands for robotic navigation and mobile manipulation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1507–1514, 2011.
- [130] Moritz Tenorth and Michael Beetz. Knowrob: A knowledge processing infrastructure for cognition-enabled robots. *The International Journal of Robotics Research*, 32(5):566–590, 2013.
- [131] Jean-Luc Thiffeault. Braids of entangled particle trajectories. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(1):017516, 2010.
- [132] Jean-Luc Thiffeault and Marko Budišić. Braidlab: A software package for braids and loops, 2013–2017. URL <http://arXiv.org/abs/1410.0849>. Version 3.2.3.

- [133] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. Computational human-robot interaction. *Foundations and Trends in Robotics*, 4(2-3):105–223, 2016.
- [134] Sebastian Thrun, Michael Beetz, Maren Bennewitz, Wolfram Burgard, Armin B. Cremers, Frank Dellaert, Dieter Fox, Dirk Hähnel, Chuck Rosenberg, Nicholas Roy, Jamieson Schulte, and Dirk Schulz. Probabilistic algorithms and the interactive museum tour-guide robot minerva. *The International Journal of Robotics Research*, 19(11):972–999, 2000.
- [135] T. Tieleman and G. Hinton. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning, 2012.
- [136] Pete Trautman. Sparse interacting gaussian processes: Efficiency and optimality theorems of autonomous crowd navigation. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, pages 327–334, 2017.
- [137] Pete Trautman. Breaking the human-robot deadlock: Surpassing shared control performance limits with sparse human-robot interaction. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*, 2017.
- [138] Peter Trautman, Jeremy Ma, Richard M. Murray, and Andreas Krause. Robot navigation in dense human crowds: Statistical models and experimental studies of human-robot cooperation. *International Journal of Robotics Research*, 34(3):335–356, 2015.
- [139] Adrien Treuille, Seth Cooper, and Zoran Popović. Continuum crowds. *ACM Transactions on Graphics*, 25(3):1160–1168, 2006.

- [140] Xuan-Tung Truong and Dung Ngo Ngo. Toward socially aware robot navigation in dynamic and crowded environments: A proactive social motion model. *IEEE Transactions on Automation Science and Engineering*, 14(4):1743–1760, 2017.
- [141] Vaibhav V. Unhelkar, Claudia Pérez-D’Arpino, Leia Stirling, and Julie A. Shah. Human-robot co-navigation using anticipatory indicators of human walking motion. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 6183–6190, 2015.
- [142] Jur van den Berg, Stephen J. Guy, Ming C. Lin, and Dinesh Manocha. Reciprocal  $n$ -body collision avoidance. In *Proceedings of the International Symposium on Robotics Research (ISRR)*, pages 3–19, 2009.
- [143] Dizan Vasquez, Billy Okal, and Kai O. Arras. Inverse reinforcement learning algorithms and features for robot navigation in crowds: An experimental comparison. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1341–1346, 2014.
- [144] Adam Vogel, Christopher Potts, and Dan Jurafsky. Implicatures and nested beliefs in approximate decentralized-pomdps. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 74–80, 2013.
- [145] William H. Warren. The dynamics of perception and action. *Psychological Review*, 113(2):358–389, 2006.
- [146] Eva Wiese, Agnieszka Wykowska, Jan Zwickel, and Hermann J. Müller. I see what you mean: How attentional selection is shaped by ascribing intentions to others. *PLOS ONE*, 7(9):1–7, 2012.

- [147] Ronald Wilcox, Stefanos Nikolaidis, and Julie Shah. Optimization of temporal dynamics for adaptive human-robot interaction in assembly manufacturing. *Proceedings of the Robotics: Science and Systems Conference (RSS)*, pages 441–448, 2012.
- [148] Nicholas H. Wolfinger. Passing Moments: Some Social Dynamics of Pedestrian Interaction. *Journal of Contemporary Ethnography*, 24(3):323–340, 1995.
- [149] Jin Y. Yen. Finding the k shortest loopless paths in a network. *Management Science*, pages 712–716, 1971.
- [150] Bolei Zhou, Xiaogang Wang, and Xiaoou Tang. Understanding collective crowd behaviors: Learning a mixture model of dynamic pedestrian-agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2871–2878, 2012.
- [151] Brian D. Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, J. Andrew Bagnell, Martial Hebert, Anind K. Dey, and Siddhartha S. Srinivasa. Planning-based prediction for pedestrians. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3931–3936, 2009.
- [152] Günter M. Ziegler. *Lectures on polytopes*. Springer-Verlag, New York, 1995.
- [153] Matthew Zucker, Nathan Ratliff, Anca Dragan, Mihail Pivtoraiko, Matthew Klingensmith, Christopher Dellin, J. Andrew (Drew) Bagnell, and Siddhartha Srinivasa. Chomp: Covariant hamiltonian optimization for motion planning. *International Journal of Robotics Research*, 32(9-10): 1164–1193, 2013.