An Institutional Approach to Normative Distributed Robotics for Mixed Societies of Humans and Robots

THIS IS A TEMPORARY TITLE PAGE
It will be replaced for the final print by a version
provided by the service académique.

Thèse présentée le 11 mars 2020
à la Faculté ENAC
Distributed Intelligent Systems and Algorithms Laboratory
Programme doctoral en Systèmes de Production et Robotique
École Polytechnique Fédérale de Lausanne

pour l'obtention du grade de Docteur dès Sciences
par

Alicja Roelofsen née Wąsik

acceptée sur proposition du jury:

Prof. Alcherio Martinoli, directeur
Prof. Pedro U. Lima, directeur
Prof. Auke Ijspeert, president
Prof. Alexandre Alahi, rapporteur
Prof. Rachid Alami, rapporteur
Prof. Lorenzo Sabattini, rapporteur

EPFL, Lausanne, 2020

Mamie i Tacie,

za cierpliwość, wsparcie i zachętę.

A przede wszystkim za miłość i przyjaźń
aż do końca świata.

Mężowi i dzieciom
zadedykuję następną.

Acknowledgements

I am very grateful for having been given an opportunity to participate in the IST/EPFL Joint Doctoral Initiative, which allowed me to enjoy the advantages of my studies twofold. I had a chance to live in two countries, I was immersed in two different academic environments, met twice as many valuable people, got to hike in Swiss mountains and bike on the beaches of the Atlantic ocean. I was also lucky enough to be supervised by not one, but two excellent Professors, Alcherio Martinoli and Pedro U. Lima. They both provided me with support and guidance, and put their confidence in me, even at times when I was lacking it myself. Most importantly, I would like to express my utmost gratitude for their understanding in the last year of my studies, when they allowed me greater flexibility to meet the challenges of my personal life, trusting that I will be able to tackle both. It meant the world to me – thank you.

I would like to thank the thesis committee – Prof. Rachid Alami, Prof. Lorenzo Sabattini, Prof. Alexandre Alahi and Prof. Auke Ijspeert, who kindly accepted to review this manuscript and provided me with invaluable feedback.

I would like to thank all members of the MOnarCH project for the truly great experience of working with them. In particular, I am grateful to Prof. João Sequeira, Prof. Rodrigo Ventura, Duarte Gameiro, José Carlos Castillo, Víctor González and João Mendes, with whom I had the pleasure to closely collaborate, and Paulo Alvito from IDMind for his indispensable technical support and advice.

I would like to thank my collaborators at the Örebro University, Stevan Tomic, Alessandro Saffiotti and Federico Pecora. It was through all the discussions, close supervision and constructive feedback during the time they hosted me in Örebro when my thesis took the final shape. Special thanks go to my collaborators at EPFL, José Nuno Pereira, Zeynab Talebpour,

David Mansolino and Lorenzo Sarti for for everything that I have learned from them. I would also like to mention two enthusiastic and dedicated Master students whom I had the pleasure to supervise: Emil Bryngelsson and Michiaki Hirayama.
I have many thanks to Denis Rochat and Emmanuel Droz, who were very kind to help me when I needed technical support. And very special thanks to Corinne Farquharson, who is the most organized, dependable and dedicated person I have met during my studies. And a really fun person to talk to about the challenges of parenting.

I cannot fit here the names of all my friends and colleagues at EPFL, IST and in Örebro, who

made me fell so welcome and who let me integrate into their social circles so quickly. Nonetheless, I will mention Duarte, who made the EPFL-IST transitions so much less intimidating. I'm sorry about the biking trip. The strong, confident women in our robotic lab – Maria Boberg, Bahar Haghighat, Faezeh Rabhar and Chiara Ercolani. It was fun to see people being surprised to see us together in a male-dominated profession. Postdocs and senior PhD students who shared with me their outlook at the life during the studies and made the whole less intimidating, Ali Marjovi, Meysam Basiri, Iñaki Navarro, Tiago Veiga, Pedro Miraldo, Jorge Soares, Masoumeh (Iran) Mansouri, Adrian Arfire, Ezequiel Di Mario, Aamir Ahmad, Asif Arain, Chris Evans, Felix Schill, Alexander Bahr and Thomas Lochmatter. And my colleagues, with whom I had so much fun and laugh, Anwar Quraishi, Miloš Vasić, Kağan Erünsal, Cyrill Baumann and Oscar Lima.

My deepest gratitude goes out to my family. Mamo, bez Ciebie nie dałabym rady. Dzięki Tobie miałam siłę by wyjść do pracy, bo wiedziałam ze zostawiam Arunię w najlepszych rękach. Za to samo jestem wdzięczna Tobie, Ciociu Danusiu. Dziekuję że pokochałyscie moje dziecko tak mocno jak własne. Tato, zawsze byłeś i jesteś moim oparciem. Gdy tylko mam wątpliwości, wyobrażam sobie jak Ty byś sobie poradził, i to mi dodaje sił. Uli i Karolowi, za to że są moim wspaniałym rodzeństwem. Ula, dzięki za zmobilizowanie reszty żeby przyjechali na mój ślub. Karol, wilkie dzięki za zmianę opon o czwartej rano. Jesteście fajni.

And finally, I would like to thank my husband Steven for accompanying me during this most challenging period of my life and for the exciting future we have ahead of us. Finally, most loving thanks go to our delightful daughter Ariana for teaching me that love and happiness have no boundaries.

Lausanne, 8 April 2020 Alicja

Abstract

We are entering the era of robots, for everyone, everywhere. Robots work alongside humans in hospitals, museums, at the airports, they provide assistance to elderly and handicapped people. The expectations of the robots being reliable, intelligent and friendly artificial creatures drive the key challenges of the state-of-the-art robotics: technological, practical and social. The latter is of particular importance, as the extent to which the robots will be accepted in our social environments largely depends on whether they will accommodate the norms of our society.

In this thesis, we venture a multidisciplinary effort at the intersection of social studies, economy and robotics. We explore the notion of institution to design social multi-robot behaviors, where institutions allow for interpretation of abstract norms formulated in human language in terms of robot-understandable terminology. In our formalism, institutions are reusable structures that provide abstraction, encapsulation and formalization of generic social norms and allow for governance over miscellaneous robot behaviors and integration of social norms of diverse nature.

The core mechanism of our framework – norm realization, forms the translation layer between the institutional abstraction and a specific system. Norm realization is founded on three key ingredients. First, the representation of norms in a universal form of human language, which introduces a degree of generality necessary to achieve conceptualization, systematization and reusability of norms. Such universal representation is shown to have a high potential to facilitate the integration of existing solutions for diverse applications and social contexts without resorting to the use of heuristics. Second, the development of a clear semantics, necessary for the interpretation of the norms at both abstract and concrete level, and for the implementation of social norms in a plug-and-play manner instead of programming hard-wired social compliance in ad-hoc behaviors. Third, a step-by-step approach for addressing the question of *how* to apply a generic, language-defined norm into robot terminology, making it readily implementable and executable in physical systems under specific constraints. The proposed formalism embraces low-level sophistication to adopt the complexity of continuous multi-robot behaviors, at the same time retaining the desirable high-level properties. We showcase the power of norm realization through a number of norms encompassing a large variety of social aspects, ranging from human comfort achieved through navigational compliance to legibility of robot intentions reflected through gestures, expressions and sounds. We carry out a multi-facet validation of our institutional formalism through three extensive case studies, where we address diverse social contexts involving mixed human-robot teams.

We target the deployment of multi-robot systems in real human-populated environments, which are cluttered, unpredictable, and highly dynamic. To address the practical challenges of such conditions, we develop methods for overcoming the technological shortcomings along three research thrusts. First, we propose an approach to agile formation control called Local Formation Transformation, where the shape of the formation adapts locally and gradually to meet the demands of complex indoor environments. Second, we develop a cooperative localization method called Formation Information Gaussian Mixture Probability Hypothesis Density filter for achieving robustness of collective navigation in case of communication failures by combining data from diverse sources. Third, we adopt existing single-robot social navigation methods within our institutional framework to achieve social awareness in the context of multi-robot systems. With these ingredients, our work is pioneering within the scope of human-aware multi-robot navigation in real, human-populated environments.

Keywords: Normative Robotics; Institutional Robotics; Social Robots; Cooperative Navigation; Human-Robot Interaction; Cooperative Tracking; Adaptive Formation Control; Multi-Robot Systems.

Résumé

Nous entrons dans l'ère des robots, pour tous le monde et partout. Les robots travaillent aux côtés des humains dans les hôpitaux, les musées, les aéroports. Ils donnent assistance aux personnes âgées et aux personnes handicapées. Les attentes pour des robots en matière de fiabilité, intelligence et socialité poussent l'état de l'art de la robotique en termes technologiques, expérimentaux et sociaux. Ce dernier point est particulièrement important, car l'ampleur à laquelle les robots seront acceptés dans notre environnement social dépendera grandement de leur capacités à s'adapter à des normes sociales.

Dans cette thèse, nous entreprenons un effort multidisciplinaire à l'intersection des études sociales, économiques et robotiques. Nous explorons la notion d'institution pour la conception de comportements sociaux multi-robots, où cette notion permette l'interprétation de normes abstraites formulées en langage humain par les robots. Dans notre formalisme, les institutions sont des structures réutilisables qui permettent l'abstraction, l'encapsulation et la formalisation de normes sociales génériques, ainsi que leur integration dans le comportment des robots.

La partie centrale de notre formalisme, la réalisation des normes (*norm realization*), forme la couche traductrice entre les abstractions institutionnelles et un système en particulier. *Norm realization* est composé de trois ingrédients clés. Premièrement, la représentation des normes dans une forme universelle du langage humain, qui introduit la quantité de généralité nécessaire pour atteindre la conceptualisation, la systématisation et la réutilisation des normes. Une telle représentation universelle a un fort potentiel pour faciliter l'intégration de solutions existantes pour applications et contextes sociaux différents sans avoir à recourir à des heuristiques. Deuxièmement, le développement d'une sémantique claire, nécessaire pour l'interprétation des normes, à la fois au niveaux abstrait et concret, et pour l'implémentation de normes sociales interchangeables, à la place de programmer directement les conformités sociales dans des comportements ad-hoc. Troisièmement, une approche pas-à-pas pour répondre à la question de *comment* appliquer une norme générique et définie dans un langage humain dans des termes robotiques, la rendant implémentable et exécutable en systèmes physiques sous des contraintes spécifiques. Le formalisme proposé inclut la sophistication nécessaire à bas niveau pour adopter les complexités des continuités comportementales des systèmes multi-robots, et en même temps garde les propriétés de haut niveau désirées. Nous montrons la puissance de la réalisation des normes à travers plusieurs normes qui englobent une large diversité d'aspects sociaux, allant du confort humain atteint à travers la conformité de navigation, jusqu'à la lisibilité des intentions des robots par des mouvements, expressions et sons. Nous avons conduit une validation multi-facette de notre formalisme institutionnel à travers trois cas d'études extensifs, où nous abordons plusieurs contextes sociaux impliquant des groupes mixtes de robots et d'humains.

Nous ciblons des déploiements de systèmes multi-robots dans des environnements réalistes peuplés d'humains, qui sont encombrés, imprévisibles et hautement dynamiques. Pour aborder les défis pratiques de telles conditions, nous avons développé trois differentes méthodes pour gérer des groupes de robots en mouvement. Premièrement, nous proposons une approche pour un contrôle agile et graduel de la topologie du groupe basée sur une transformation locale de la formation (*Local Formation Transformation*). Deuxièmement, nous avons développé une méthode de localisation coopérative basée sur des filtres Bayesiens (*Formation Information Gaussian Mixture Probability Hypothesis Density filter*) pour atteindre une navigation collective fiable en cas de perte de communication, et cela en combinant les informations de differentes sources. Troisièmement, nous adoptons des méthodes existantes de navigation sociale pour un seul robot dans notre système institutionnel pour atteindre une réalisation sociale de nos systèmes multi-robots. Avec ces ingrédients, notre travail est novateur dans le cadre de la navigation sociale de groupes de robots dans des environnements réalistes peuplés d'humains.

Mots Clés : Robotique Normative ; Robotique Institutionelle ; Robots Sociaux ; Navigation Cooperative ; Intéraction Humain-Robot ; Suivi Coopératif ; Contrôle en Formation Adaptif ; Systèmes Multi-Robots.

Resumo

NESTA tese empreendemos um esforço multidisciplinar, englobando estudos sociais, economia e robótica. Exploramos o conceito de instituição para projetar comportamentos sociais envolvendo vários robôs, permitindo a tradução de normas abstratas formuladas em linguagem humana para uma terminologia capaz de ser compreendida pelos robôs. Na nossa formulação, instituições são estruturas reutilizáveis que proporcionam abstração, encapsulamento e formalização de normas sociais genéricas, permitindo a gestão de diversos comportamentos robóticos e a integração de normas sociais de diversas naturezas.

O mecanismo principal da nossa plataforma – realização de normas, forma a camada de tradução entre a abstração institucional e o sistema em concreto. A formulação proposta abrange uma especificação de baixo nível, de forma a adoptar a complexidade de comportamentos entre vários robôs, retendo ao mesmo tempo, as propriedades de alto nível do sistema. Mostramos a importância da realização de normas utilizando um conjunto de normas que envolvem uma grande variedade de aspetos sociais, que vão do conforto humano, atingido através de navegação ciente de humanos, à leitura das intenções dos robôs através de gestos, expressões e sons. Realizamos uma validação multifacetada do nosso formalismo institucional através de três estudos de casos, onde abordamos vários contextos sociais envolvendo equipas mistas entre robôs e humanos.

A tese tem como objectivo o desenvolvimento de sistemas de múltiplos robôs em ambientes povoados por pessoas, que são desordenados, imprevisíveis, e muito dinâmicos. De forma a abordar os desafios práticos inerentes a essas condições, desenvolvemos métodos para ultrapassar as deficiências técnicas explorando três linhas de investigação. Primeiramente, propomos uma abordagem para controlo flexível de formações, denominada transformação local de formação, onde a geometria da formação se consegue adaptar localmente e gradualmente, de forma a poder suportar as complexidades de ambientes internos. Seguidamente, desenvolvemos um método de localização cooperativa, denominada de filtro de densidade de hipóteses com mistura de gaussianas para informação de formações, que pretende atingir a robustez de uma navegação colectiva no caso de falhas de comunicação, combinando informação de várias fontes. Finalmente, incluímos métodos de navegação social para robôs singulares na nossa plataforma institucional de forma a conseguir implementar em sistemas de vários robôs cientes de normas sociais. Com estes ingredientes, o trabalho é pioneiro no contexto de navegação de vários robôs levando em conta pessoas em ambientes reais e povoados.

Palavras-Chave: Robótica Normativa; Robótica Institucional; Robótica Social; Navegação Cooperativa; Interacção Entre Humanos e Robôs; Seguimento Cooperativo; Controlo de Formações Adaptativas; Sistemas Multi-Robô.

Streszczenie

WKRACZAMY w erę robotów, dla wszystkich i wszędzie. Roboty pracują razem z ludźmi w szpitalach, muzeach i na lotniskach, zapewniają opiekę osobom starszym i niepełnosprawnym. Oczekiwania, że roboty powinny być niezawodne, inteligentne i przyjazne kierują kluczowymi wyzwaniami współczesnej robotyki – mianowicie wyzwaniami technologicznymi, praktycznymi i społecznymi. To ostatnie ma szczególne znaczenie, ponieważ stopień, w jakim roboty będą akceptowane w naszym społeczeństwie, zależy w dużej mierze od tego, czy dostosują się one do ogólno przyjętych norm.

W niniejszej pracy doktorskiej podejmujemy się tematyki łączącej badania społeczne, ekonomię i robotykę. Wykorzystujemy pojęcie *instytucji* w celu określania i budowania społecznych zachowań robotów, które nie tylko współpracują ze sobą nawzajem, lecz także z ludźmi. Instytucje pozwalają nam na interpretację norm społecznych sformułowanych w mało precyzyjnym języku ludzkim oraz przetworzenie ich w terminologie zrozumiałą dla robota. Osiągnięta w ten sposób formalizacja ogólnych norm społecznych zapewnia możliwość ich powtórnego użytku w szerokiej gamie odmiennych sytuacji, pozwala na zarządzanie różnorodnymi zachowaniami robotów oraz integrację norm społecznych o wielorakim charakterze.

Podstawowym mechanizmem naszego formalizmu jest proces *realizacji norm*, który tworzy warstwę pośrednią między instytucjonalną abstrakcją a konkretnym, fizycznym systemem robotów. Realizacja norm opiera się na trzech kluczowych składnikach. Po pierwsze, normy społeczne są przedstawione formie ludzkiego języka, a zatem formie ogólnie zrozumialej i uniwersalnej, niezbędnej do osiągnięcia systematyczności w ustalaniu społecznych zachowań robotów i możliwości ich ponownego zastosowania w odmiennych sytuacjach. Co więcej, uniwersalna reprezentacja norm ma ogromny potencjał ułatwiania integracji istniejących rozwiązań bez uciekania się do heurystyki. Po drugie, realizacja norm opiera się na przejrzystej semantyce, niezbędnej do interpretacji norm społecznych zarówno na poziomie abstrakcyjnym, jak i wprowadzania ich w czyn na poziomie konkretnych zachowań robotów. Dzięki temu normy społeczne są wdrażane na zasadzie podłącz i graj, a nie wprowadzane jako zakodowane na stale zgodności, jak to jest powszechnie stosowane we współczesnej robotyce normatywnej. Po trzecie, realizacja norm to instrukcja *jak*, krok po kroku, przetłumaczyć normy zdefiniowane w języku ludzkim na terminologie zrozumiałą dla robotów, dzięki czemu proces ten jest łatwy do wdrożenia w konkretnych systemach. Proponowany formalizm pozwala osiągnąć precyzję, która jest wymagana by moc osiągnąć pożądane zachowania robotów, przy jednoczesnym zachowaniu korzystnych właściwości. Weryfikacja naszego formalizmu

przeprowadzona jest poprzez trzy obszerne studia przypadków, w ramach których zajmujemy się różnorodnymi kontekstami społecznymi z udziałem grup ludzi i robotów. Poprzez zastosowanie szeregu norm obejmujących wiele różnych aspektów społecznych, takich jak zapewnienie komfortu człowieka dzięki przystosowaniu ruchu robotów, czy osiągnięcie czytelności intencji robotów, odzwierciedlonej poprzez gesty, mimikę i dźwięki, demonstrujemy potencjał przedstawionego formalizmu oraz procesu realizacji norm.
Systemy wielorobotowe, którymi się zajmujemy, weryfikowane są w miejscach często uczęszczanych przez ludzi, w realistycznych sytuacjach, charakteryzujących się złożonością, nieprzewidywalnością i dynamiką. Aby sprostać powyższym wyzwaniom, w niniejszej pracy podejmujemy się badań dotyczących trzech zagadnień, które według naszych przewidywań niezbędne są do wdrożenia systemów wielorobotowych w naturalnych przestrzeniach zajmowanych przez ludzi. Po pierwsze, proponujemy metodę pozwalająca grupie robotów poruszającej się w zwartej formacji na adaptację owej formacji, by móc efektywnie poruszać się we wnętrzach z zawiłymi przeszkodami. W owej metodzie każdy robot indywidualnie i stopniowo modyfikuje kształt formacji w swoim najbliższym otoczeniu, by wspólnie ominąć przeszkodę. Po drugie, opracowujemy metodę pozwalającą na utrzymanie formacji nawet w przypadku zaburzeń systemów komunikacji pomiędzy robotami, poprzez łączenie danych z różnych źródeł, mianowicie z systemów komunikacji, z pomiarów czujników oraz informacji o pożądanym kształcie formacji. Po trzecie, stosujemy istniejące już metody służące do osiągnięcia stosownego poruszania się pojedynczego robota w środowiskach ludzkich w nowym kontekście systemów z wieloma robotami. Dzięki powyższym składnikom nasza praca przoduje w zakresie nawigacji z wieloma robotami w realistycznych miejscach uczęszczanych przez ludzi i przystosowanej społecznie do obecności człowieka.

Słowa Kluczowe: Robotyka Normatywna; Instytucjonalny Formalizm; Roboty Świadome Społecznie; Kooperacyjna Nawigacja; Interakcja Człowiek-Robot; Kooperacyjne Śledzenie; Adaptacja Formacji Robotów; Systemy Wielorobotowe.

Contents

Introduction Part I

It's a dangerous business, Frodo, going out your door.
You step onto the road, and if you don't keep your feet,
there's no knowing where you might be swept off to.

J.R.R. Tolkien, The Lord of the Rings

1 Introduction

It feels like technology and its prevalence in our everyday lives is progressing faster than ever. 30 years ago hardly anyone had access to a computer, while nowadays most of us carry an equivalence in their pocket. Raymond Kurzweil, a computer scientist and futurist, in 2001 wrote "We won't experience 100 years of progress in the 21st century – it will be more like 20,000 years of progress."

After the boom of privately-owned computing technologies – personal computers, laptops and smartphones, now it is the time for interactive, mobile devices to ensue – we are entering the era of robots, for everyone, everywhere. Current trends in robotics slowly shift beyond the primary applications in manufacturing and automation towards our homes, offices and public places. Robots work alongside humans in hospitals, museums, at the airports, they provide assistance to elderly and handicapped people; household robots are available in almost every convenience store and almost every child possesses a simple robotic toy. The numbers are staggering – by 2020, 5.9 million new service robots will be in operation, raising the total to 22.5 million world-wide[1]. Furthermore, the market potential will reach its full capacity once solutions are found for extending their operational range. Irrespective of whether we accept this trend or not, robots are slowly becoming an integral part of our lives and there are no signs of the trend slowing down.

A successful introduction of robots into human environments will rely on the development of systems that are reliable, safe, and easy to use. However, operation "in the wild" is very different from that in a controllable factory environment [1]. Human-populated environments are cluttered, unpredictable, and highly dynamic, and so, the first challenge of deployment of robots in human spaces is that of overcoming the technological shortcomings and limitations related to their reliable operation. Such challenges become further exacerbated when multiple robots are supposed to operate cooperatively.

The second challenge, disputably more consequential, is of social nature. Needless to say that similarly to the Industrial Revolution of the 19th century, which transitioned handcraft

1 International Federation of Robotics, Frankfurt, https://ifr.org/.

to manufacturing, nowadays the viewpoints on the pervasiveness of machines in our lives vary. The perspective of sharing everyday life with robots or any other artificial machines causes many controversies – people are not only scared of losing their jobs to automation, but also worried that technology slowly replaces real human relationships and even that one day intelligent machines could develop consciousness and will take over control. While the fear is partially caused by the science-fiction literature, the wide-spread apprehension towards robotics remains a fact[2].

As the progress of robotics is riding this new wave of uncertainty, efforts towards smooth and natural introduction of robots are of paramount importance. The expectations of the robots being reliable, intelligent and friendly artificial creatures drive the key challenges of the state-of-the-art social robotics, as the extent to which the robots will be accepted largely depends on whether they will accommodate the norms governing our society. To this end, research on social robots proposes increasingly sophisticated methods for human-robot interaction, strongly inspired by well-versed findings on human social behaviors. Socially intelligent robots and other artificial agents are taught to follow social norms, create relationships with humans and invoke social responses. However, the vast majority of research deals with specific challenges or applications. Current developments in social robotics remain within the comfort zone of controllable single-robot single-human settings and address isolated problems, the solutions to which are difficult to generalize or reuse for applications other than what they have been originally proposed for. As of now, only few attempts are focused on developing a holistic theory that would bring the interdisciplinary efforts together.

In human societies, social order is achieved through institutions, while the core of each institution relies on social norms. Thereupon it is only natural that institutional mechanisms inspire solutions for the development of artificial systems for interacting with humans. A number of approaches drawing on the economics paradigms are proposed within the field of Multi-Agent Systems (MAS), with the aim of modeling complex systems. The MAS frameworks, however, are disconnected from physical operations and thus cannot easily be applied to robots. At the other end of the spectrum, robotic approaches to normative behaviors suffer from poor reusability and scalability, as the norms are often designed for, and integrated into, a specific behavior, and not for a general use. Such design principle possibly hinders the progress of social robotics, as methods only target limited scope, and consequently are difficult to compare and reuse for other applications.

The challenges of technological and social nature vary across different applications. Problems to be tackled in the case of a stationary robot engaged in sophisticated conversations with humans are different from those of a mobile platform moving in a human crowd in a safe and social way. The latter case necessitates a number of developments that are yet to be made, in particular in the context of multi-robot systems with real world deployments. A possible reason for such state of affairs is the fact that traditionally multi-robot methods tend to be

2 World Development Report 2019: The Changing Nature of Work. World Bank. 2019. Washington, DC: World Bank. doi:10.1596/978-1-4648-1328-3. License: Creative Commons Attribution CC BY 3.0 IGO

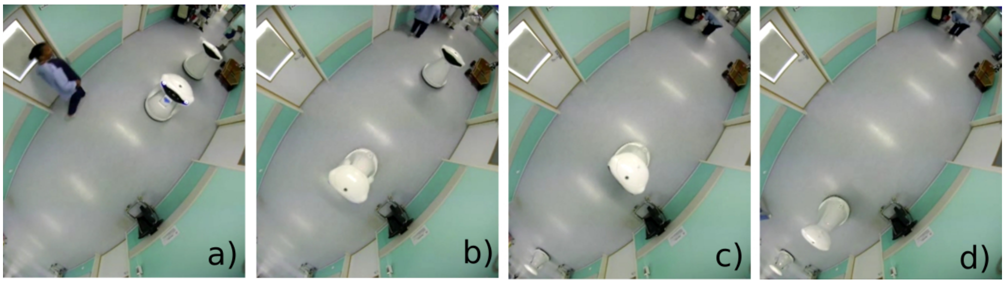

Figure 1.1 – Images from overhead cameras showing a formation of two robots navigating in the corridor of the IPOL ward (a). During execution, the robots change access point (b), which causes the formation to break temporarily (c) before it gets corrected after the robots communicate again (d). Similar problems could be solved if the robots used a combination of communication and tracking based on onboard sensing.

validated through simulation, in controlled physical settings, or with simplifying assumptions. However, once such systems are deployed in real human spaces, the challenges that are yet to be addressed become evident.

1.1 Motivation: Lessons Learned at a Hospital

One can truly understand the challenges of multi-robot navigation in human-populated environments through a real deployment. Within the course of this thesis we had the unique opportunity of performing multi-robot experiments with two robots moving in a formation at the pediatric ward of the *Instituto Português de Oncologia de Lisboa* (IPOL) in Lisbon, Portugal. The lessons learned during that short experimental campaign became the leading motivation of this thesis.

During the experiments, we identified challenges along three main axes: multi-robot navigation, cooperative localization, and social awareness. Details and results of the IPOL experiments are presented in Chapter 6.

I. Multi-Robot Navigation. Human indoor environments are enclosed spaces, constrained by structural building features such as doors and corridors, and cluttered with various appliances. While single-robot navigation methods proved to be successful in such settings, the operation of multiple mobile robots turned out to be far more challenging, even more so in case the robots needed to exhibit spatial coordination. Without any specific application in mind but rather with a two-fold objective, namely that of entertaining the kids hosted at the pediatric ward and that of gathering experience with multiple robots operating in such settings, we implemented a simple navigation experiment involving two robots in a column formation moving back and forth from one end of the ward to the other. Based on this field experiment, we imagined scenarios in which more than two robots had to navigate safely in such cluttered environment while maintaining a certain spatial topology and proposed in Chapter 7 a method for

Figure 1.2 – Social challenges of deploying multi-robot systems in human-populated environments. Images taken during experiments at the IPOL hospital.

adaptively changing the formation shape.

II. Cooperative Localization. Collective navigation methods such as formation control rely on constant exchange of state information among the robots. While wireless communication is generally reliable, it might suffer issues in settings such as those of an hospital, including message losses, delays, and even temporal loss of connection. During the experiments conducted at IPOL, robots temporarily lost their ability to communicate when changing access points of a local WiFi network, which resulted in situations similar to that shown in Figure 1.1, where a formation of two robots broke temporarily, resulting in a less intelligible collective behavior from a human perspective and ultimately giving the onlookers an impression of incompetent robots. To overcome such difficulties, in Chapter 8 we propose a system that incorporates the available communication data, sensory information, and knowledge about the desired formation geometry in a multi-target tracking filter.

III. Social Awareness. The greatest challenge of deploying multiple robots in an instrumental and yet fragile environment such as a hospital is of social nature. On the one hand, robots represent an added value by, for instance, providing entertainment, edutainment, or encouragement for physical interactions; on the other hand, robots can hamper activities of hospital staff and, in the long term, be an annoyance. The formation behavior we have deployed in the hospital did not incorporate any human awareness, and the results presented later in this thesis show that although the robots were well accepted, they also disturbed regular activities of the ward. In Figure 1.2 we show examples of situations where robots were generating disturbances and some discomfort to the surrounding humans while showcasing a technically correct behavior. We believe that if the robots were aware of the mechanisms of human society and complied with its social norms, the disturbances and discomfort mentioned above would have been significantly mitigated. This is the core challenge of this thesis and is addressed in Part III through an institutional framework in which social norms are to be respected by the multi-robot system and incorporated into robot behaviors.

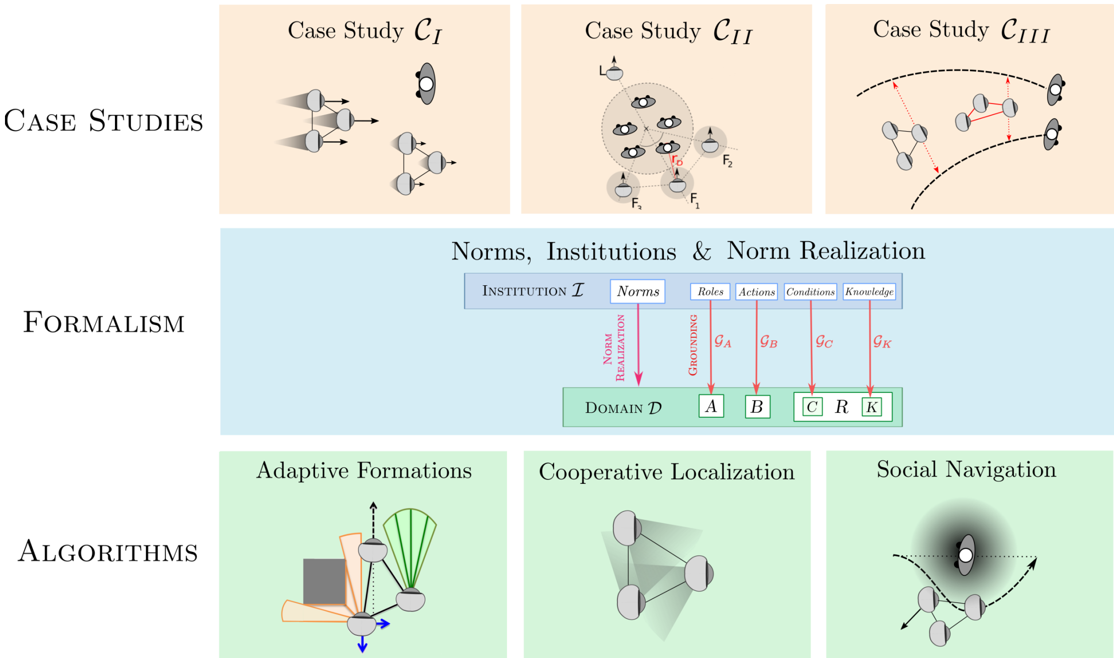

Figure 1.3 – Thesis overview.

1.2 Objectives and Outline

Our objective is to develop a distributed framework for multi-robot systems enabling safe navigation in structured, indoor environments populated by humans. We aim to achieve a socially aware behavior from a human observer perspective, where a multi-robot team navigates around the humans without disturbing their habitual activity and takes into account human comfort and conventions that people abide by. We choose to focus on spatially coordinated behaviors resulting in formations and flocks, as they accentuate the perpetual balance between purely robotic cooperation and human-robot interactions.

In an attempt to address the aforementioned challenges, the focus of this thesis will be on answering the following research questions:

A) What are the distributed, multi-robot algorithms that would result in a reliable and robust navigation in the environments of interest?

B) How can the team of robots respect the conventions and rules of human societies?

C) Which coordination mechanism could deal with both high-level behavior management and low-level, reactive coordination of robots?

D) Can the same mechanism introduce social norms for guiding robots towards socially acceptable collective behaviors and lead to better mutual understanding between humans and robots?

The focus of this thesis resides on the three layers visualized in Figure 1.3, from the bottom to the top: the algorithmic approach towards enabling human-aware multi-robot coordination in real, human-populated environments; the formalism that allows for the introduction of social norms to robot behaviors in a plug-and-play manner; and the demonstration thereof through three distinct case studies. With this modular bottom-up approach, the basic elements related to the continuous control of the multi-robot system are operated upon by the institutional formalism, to result in the normative behaviors demonstrated in the case studies. As such, the normative behaviors are achieved through the institutional formalism and rooted into the algorithms.

Although the algorithmic methods are strongly intertwined into the case studies and linked through the institutional formalism, our approach emphasizes abstraction and encapsulation, therefore allowing for the use of other algorithms at the bottom layer and application to other case studies at the top layer. The structure of this manuscript will generally follow Figure 1.3 in a bottom-up and left-right fashion. Besides the introductory and conclusive parts, this thesis includes the following two core parts:

- **Part II: Cooperative Navigation**

  We focus on the deployment of multi-robot systems engaged in spatially coordinated behaviors in physical environments. In this part, we address two of the main challenges identified in the previous section:

  - **Multi-Robot Navigation.** We build upon state-of-the-art graph-based control laws, where the graph-based framework serves as a tool to formulate rules for rigid formations and a naturally loose flocks. We enhance algorithmic robustness in complex, indoor environments and facilitate dynamic adjustment of the spatially coordinated behavior to the environmental situation. To this end, we propose a method for the real-time adaptation of the formation geometry to accommodate the constraints imposed by structured indoor environments.
  - **Cooperative Localization.** With the purpose of providing reliable robot state estimates to be used for formation control when communication fails, we propose a method that incorporates data from diverse sources to result in an enhanced multi-robot tracking filter.

- **Part III: Institutions and Norms**

  We address the challenge of formalization of human social norms in the context of multi-robot systems as follows:

  - **Social Awareness.** We adopt well-established single-robot methods for social awareness in the context of multi-robot systems.
  - **Institutional Formalism.** We propose a model-based approach for abstraction, encapsulation, and formalization of generic social norms into reusable structures called institutions.

– **Case Studies**: We demonstrate how the proposed abstract representation of institutions allows for the governance over miscellaneous robot behaviors and integration of social constraints of diverse nature.

1.3 Contributions and Publications

To the best of our knowledge, ours is the first effort towards the integration of social norms into continuous robot behaviors through a formalism that allows to abstract the normative layer from its execution in a physical system. Furthermore, up to date, no contributions have reported experimental results about a group of robots robustly moving in structured human-populated environments or engaging in mixed human-robot formations, and only few attempts have been made to expand the state-of-the-art beyond the coordination mechanisms of miniature robots with a low-level of complexity or beyond highly controlled environments. With this thesis we bring the following contributions to the state-of-the-art in robotics.

Adaptive Formation Control. We deploy cooperative robot teams in scenarios with real human participants in typical indoor spaces. To this effect, we developed a Local Formation Transformation (LFT) method for realizing adaptive robot formations in constrained indoor environments that yields *local* and *gradual* change of formation geometry with the level of alteration that accommodates the structure of the environment. Our method stands in contrast with the state-of-the-art approaches on formation control, where the change of formation shape is *global* and *discrete*, and robots are typically deployed in controlled environments with obstacles scattered around a large arena. A relevant publication for this part is:

- A. Wasik, J. N. Pereira, R. Ventura, P. U. Lima, and A. Martinoli, "Graph-based distributed control for adaptive multi-robot patrolling using local formation transformation", in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 1721–1728.

Cooperative Localization for Formation Control. We proposed a method for combining global positioning data exchanged by the robots, information about the formation geometry, and on-board sensory detections in an extension of a multi-target tracking filter, a Gaussian Mixture Probability Hypothesis Density (GM-PHD) filter. Our method called Formation Information GM-PHD (FI-GM-PHD) filter is capable of maintaining the state estimates even when long-duration sensing occlusions or communication losses occur, and it allows for maintaining formations in cluttered environments under high measurement uncertainty and low communication quality. Relevant publications include:

- A. Wasik, P. U. Lima, and A. Martinoli, "A robust localization system for multi-robot formations based on an extension of a Gaussian mixture probability hypothesis density filter", *Autonomous Robots*, vol. 44, 395–414, 2019, DOI:10.1007/s10514-019-09860-5.

- A. Wasik, A. Martinoli, and P. U. Lima, "A robust relative positioning system for multi-robot formations leveraging an extended GM-PHD filter", *Proceedings of the First International Symposium on Multi-Robot and Multi-Agent Systems*, pp. 71–77, 2017.
- A. Wasik, R. Ventura, J. N. Pereira, P. U. Lima, and A. Martinoli, "Lidar-based relative position estimation and tracking for multi-robot systems", in *Robot 2015: Second Iberian Robotics Conference*, Springer International Publishing, 2016, pp. 3–16.

Social Awareness in Multi-Robot Systems. We adopted well-versed findings of single-robot social navigation research in the context of multi-robot systems. In contrast to the state-of-the-art research on multi-robot teams deployed in human-populated environments that either leverage multiple, uncooperative robots or cooperative solutions that fail to consider realistic situations, we deploy socially aware cooperative robot teams in scenarios with human participants in typical indoor spaces, enacting the deployment of spatially coordinated multi-robot systems in shared environments, as well as mixed groups of multiple humans and multiple robots. By the virtue of representing a selection of human-aware methods in the form of social norms that are represented at an abstract level using human language, we take the first step in a process of achieving a comprehensive method for unifying the existing approaches to normative robotic systems.

Institutional Formalism. We performed an interdisciplinary effort towards bridging research findings in economics, multi-agent systems and normative robotics. We developed a model-based approach for abstraction, encapsulation and formalization of generic social norms into reusable structures, called institutions. Our methods draw upon the original principles of Institutional Robotics (IR) [6] [7] and the approaches derived there. Nonetheless, we reach beyond the current state of affairs by targeting continuous, collective, social robot behaviors, venturing the formerly addressed work on IR beyond the swarming principles and discrete planning methods.

- **Norm Realization** We have identified three key elements necessary for interpreting social norms and putting them into practice. First, the representation of norms in a universal manner, in a human-understandable form. Second, the development of a clear semantics, necessary for the interpretation of the norms at both abstract and concrete levels. Third, a step-by-step approach for addressing the question of *how* to apply a given norm, resulting in its concretization into robot-understandable terminology. These three elements are brought together in our main contribution – *norm realization*, a mechanism for translation of high-level, language-defined norms in terms of robot-understandable terminology, making such norms readily implementable onto concrete restrictions of robot behaviors and executable in real physical systems. Our method embraces low-level sophistication to address the complexity of continuous multi-robot behaviors, while and at the same time retaining the desirable high-level properties, including abstraction, encapsulation, and modularity. We showcase the power of norm realization through a number of norms targeted to encompass a large variety of social aspects, ranging from human comfort achieved though navigational compliance, to legibility of robot intentions reflected through gestures, expressions, and sounds.

- **Norm Abstraction** The aforementioned norm realization provides a tool for encoding behavioral specifications according to plug-and-play principles instead of programming social compliance in hard-wired ad-hoc behaviors, as it is done currently. We pioneer in a number of worthwhile properties lacking in state-of-the-art approaches to social and normative robotics, namely high reusability of institutions and norms, which we freely allot over different domains and behaviors without the need of redesigning them, and modularity and scalability, where the encapsulation of norms allows to decouple their operation from the behavior design and where introducing a new norm does not require heuristics on how to merge it with the current solution. Instead of the design being driven by norms, norms are imposed as constraints operating over the parametrization of already existing behaviors.

Relevant publications include:

- A. Wasik, S. Tomic, A. Saffiotti, F. Pecora, A. Martinoli, and P. U. Lima, “Towards norm realization in institutions mediating human-robot societies”, *2018 IEEE/RSJ International Conference On Intelligent Robots And Systems (IROS)*, IEEE International Conference on Intelligent Robots and Systems, pp. 297–304, 2018.
- S. Tomic, A. Wasik, P. U. Lima, A. Martinoli, F. Pecora, and A. Safiotti, “Towards institutions for mixed human-robot societies”, in *Intternational Joint Conference on Autonomous Agents and Multiagent Systems*, 2018, pp. 2216–2217.
- A. Wasik, A. Martinoli, and P. U. Lima, “An institutional robotics approach to the design of socially aware multi-robot behaviors”, *Proceedings of the RO-MAN 2017 Workshop on Towards Intelligent Social Robots: Social Cognitive Systems in Smart Environments*, pp. 2–7, 2017.

Contributions of Predecessors and Collaborators

I would like to give credit to my close collaborators who influenced the final shape of this thesis.

My institutional formalism has been strongly influenced by the work of Stevan Tomic and his supervisors, Professor Alessandro Saffiotti and Professor Federico Pecora, all belonging to the Cognitive Robotic Systems Laboratory of the Centre for the Applied Autonomous Sensor Systems (AASS) at Örebro University. They were the first to abstract the institutions from the concrete systems within the IR framework. In particular, I adapted the high-level definitions of institutions, domain and grounding, and modified them to better capture the complexity of norms for continuous cooperative behaviors. Our close collaboration resulted in two joint publications: first in [9], where my contribution is minor and hinges on the development of the multi-robot behaviors the institutions operate upon, and second in [8], where I propose my major contribution of norm realization.

This thesis includes the work of a master student Michiaki Hirayama, who under my supervision carried out the real robot experiments evaluating the performance of the FI-GM-PHD filter in Chapter 8. A significant part of my experimental setup, including installation of the motion capture system and camera system has been prepared by my colleagues, Zeynab Talebpour, Duarte Dias Emmanuel Droz and Steven Roelofsen. Furthermore I strongly relied on Zeynab's expertise when preparing my participative study.

Finally, I would like to give credit to all my fellow members of the European Multi-Robot Cognitive Systems Operating in Hospitals (MOnarCH) project, the work of whom has laid the instrumental foundations of this thesis, in particular the single-robot navigation functionalities, calibration of the simulated setup and preparation of the software stack I have been relying upon long after the conclusion of the project. The credit goes to Professor Rodrigo Ventura, research and development engineers Lorenzo Sarti and David Mansolino, and post-doctoral researcher José Nuno Pereira. Moreover, José Nuno provided me valuable feedback on my early work on formation control and made the experiments at IPOL possible.

Summary

The research effort presented in this thesis addresses three key aspects, namely the algorithmic methods enabling human-aware multi-robot coordination in real, human-populated environments, the formalism that allows for introduction of social norms to robot behaviors in a plug-and-play manner, and the demonstration through three distinct case studies. Our modular and layered bottom-up approach puts together the basic elements related to the continuous control of the multi-robot behaviors in order to provide the physical baseline for our institutional formalism to operate over, while the case studies validate normative behaviors that are built upon the institutional formalism and rooted into the algorithms.

2 Systems and Methods

THE deployment of mobile, child-sized, complex robots in human-populated environments necessitates the co-existence of multiple systems and components interplaying at a different level of centralization, competency and physicality. In this chapter, we describe the robotic platforms, auxiliary systems and simulation tools used for experimentation and validation of our methods. It should be noted that our approach to normative, multi-robot coordination is independent of the chosen robotic platform or the system architecture, so the choice of the system components does not impose further restrictions on our methods.

2.1 System Overview

An overview of the system is presented in Figure 2.1. Some of the components are used only in a subset of the experiments, depending on their availability in the experimental facility or on whether their presence is required.

At the core of a networked multi-robot system is communication. All robots and external systems are connected to a local wireless network and, as we will describe in Section 2.3.1, by the means of Robot Operating System (ROS) messages, they can broadcast information to all other connected elements.

The main robotic platform used throughout this thesis is the MBot robot - a child-sized, friendly-looking robot equipped with navigation, perception and low-level safety sensors. The navigation of the MBot robots is based on a standard occupancy grid map, serving for both motion planing and self-localization. For spatial coordination, the robots communicate to each other the poses generated through self-localization. In order to minimize the dependence of the system on a single point of failure, we developed an onboard detection and tracking module that complements information obtained via communication with data from onboard sensors. Our multi-robot navigation functionalities make use of two already existing single-robot navigation capabilities of the robot. First, a Fast Marching Method (FMM) is used for

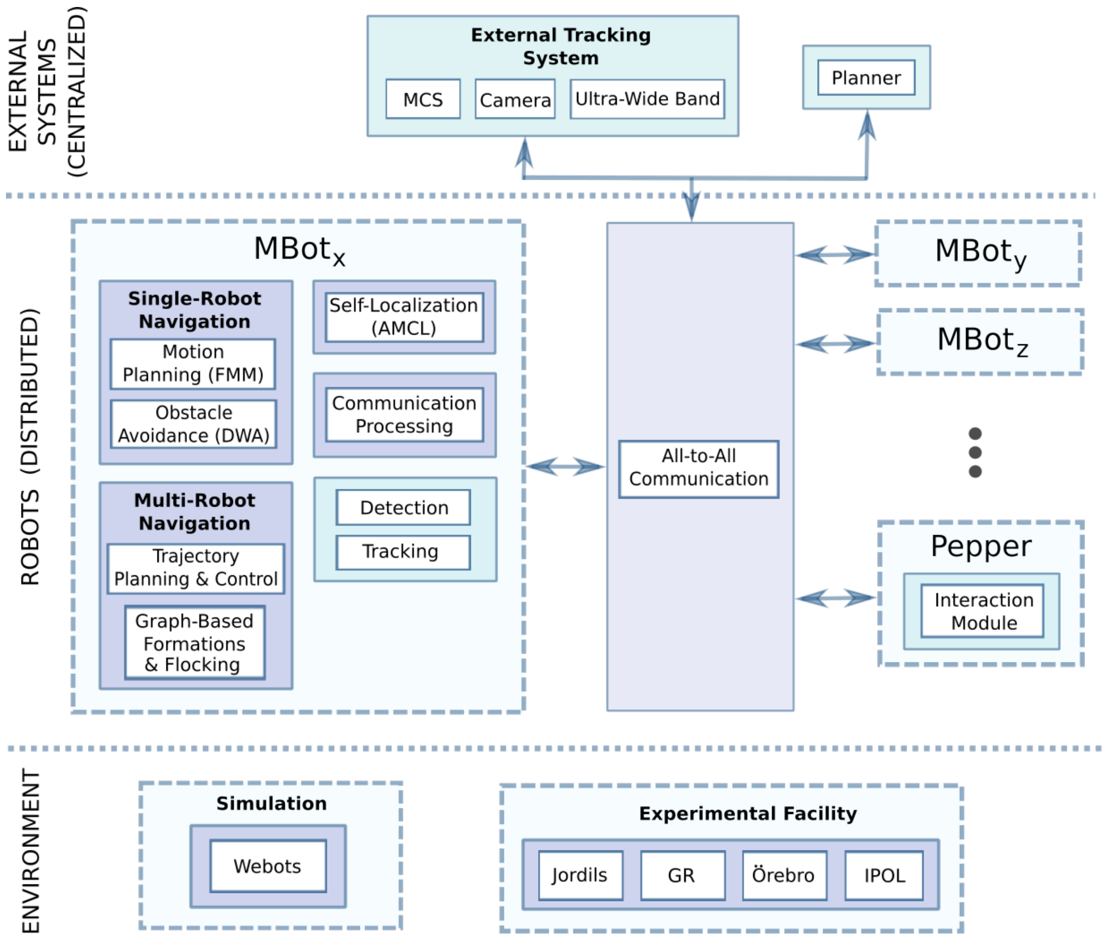

Figure 2.1 – System overview.

trajectory planning and point-to-point navigation. Second, a Dynamic Window Approach (DWA) reactive obstacle avoidance method is always activated in parallel with our multi-robot algorithms, preventing robots to collide with each other and with objects that are detectable by sensors. Additionally, in one set of experiments, we use another robotic platform – a Pepper robot, with its interaction capabilities, but not its navigation.

To obtain reliable pose estimates of the humans during experiments, we use one of the three following systems: an overhead camera, a Motion Capture System (MCS), or an Ultra-Wide Band (UWB) localization system. These systems are connected to ROS and broadcast data that are available to all the robots. Furthermore, data from the MCS provides ground truth localization of the robots, a key information used for experimental evaluation. The modularity of our setup allows for flexibility in terms of adding or removing additional components.

Experimental evaluation is performed in both simulation and reality. For simulation, we use the high-fidelity simulator Webots, which faithfully reproduces the robots with all their sensors and actuators, calibrated to match the reality. Real experiments are carried out in four facilities, and exhaustive evaluation has been realized in two facilities of DISAL (the off-campus Motion

Arena at Rue de Jordils 41 in St. Sulpice, in short Jordils, and the on-campus Robotic Laboratory in the GR building, in short GR) as well as one of the AASS Center of the Örebro University (in short Örebro). We also present observations collected during trials at the IPOL hospital.

2.2 Robotic Platforms

In this section, we describe two robotic platforms used in this work, the MBot and the Pepper. The MBot robot is our main experimental platform and will be used throughout this thesis, while the Pepper robot, available at the Örebro experimental facility, has been used to a limited extent. Pictures of both robots are shown in Figure 2.2.

2.2.1 MBot

The MBot is a child-sized, friendly-looking robot developed within frame of the FP7 European project MOnarCH[1] with the goal of introducing social robots in the pedriatric ward of the IPOL hospital. Shown in Figure 2.2, the MBot has an approximately round footprint of 0.65 m in diameter and a height of 0.98 m. A complete description of the MBot robot can be found in [11].

The robot is equipped with navigation, perception and low-level safety sensors. For mapping, localization and obstacle avoidance, it fuses measurements provided by odometry encoders, IMU sensors and two laser range finders (Hokuyo URG-04LX-UG01) mounted at a height of approximately 13 cm above the ground, one in front and one at the back of the robot, providing a 360° Field of View (FOV) and approximately 4 m sensing distance. An omnidirectional locomotion system with four Mecanum wheels provides a maximum speed of 2.5 m/s and maximum acceleration of 1 m/s^2. The robot has two onboard computers running Linux OS (Ubuntu 12.04), one responsible for navigation, localization, system control and actuation of low-level interaction devices, the second responsible for the control of the interaction functions. Interaction components include arms and head actuated with servo motors, speakers, touch screen and LEDs located at the mouth, eyes, cheeks and at the bottom of the base. Batteries, depending on the usage, give an autonomy of up to 3 hours.

2.2.2 Pepper

Pepper is a 1.2 m high humanoid robot manufactured by SoftBank Robotics [2], designed to interact with people at a personal level, at homes or for education, as well as aid businesses at a professional level as customer help. It is equipped with interaction, motion detection and safety sensors, including 2D and 3D cameras, touch screen, tactile sensors and IR sensors and 6 lasers, each providing 60° field-of-view and used for safety purposes, but not for navigation.

1 MOnarCH, FP7, FP7-ICT-2011-9-601033 (http://monarch-fp7.eu)
2 SoftBank Robotics (https://www.ald.softbankrobotics.com)

Figure 2.2 – Two robotic platforms used in this thesis - the MBot robot, used throughout our work, and the Pepper robot, accompanying MBot in experiments at one of the facilities.

Pepper's interaction functions, such as speech, LEDs or gestures can be controlled using dedicated software or can be scripted in Python. The robot is used in our case study in Chapter 16.

2.3 Software

In our work we heavily rely on two open-source software tools: the Robot Operating System (ROS), which provides the basis upon which we build our software stack, and Webots, a high-fidelity simulator we used for design, development and first-step validation of our methods.

2.3.1 ROS

The ROS[3] is an open-source framework providing a collection of libraries and tools for design, implementation, and execution of robotic applications. The primary benefits of using ROS include abstraction of hardware and low-level device control, simplicity of message-passing between processes and devices, and ready-to-use implementation of commonly-used robot functionalities, making it easy to develop software by connecting existing modules from

3 ROS (http://wiki.ros.org)

various developers. We have been using ROS1, based on a master-slave architecture, with a master node running on each robot. In ROS, a robot application is broken down into a network of concurrent processes called *nodes*, which are loosely coupled using the ROS communication infrastructure, where messages are relayed using *topics*. The same infrastructure allows seamless communication over multiple devices or multiple robots. Nodes, either within one device or distributed over several machines, communicate using a publisher/subscriber architecture, where nodes can subscribe to topics and receive data from other nodes that are publishing.

2.3.2 Simulations

Webots [12] is an open-source, physics-based robot simulator that allows for modeling, programming and simulation of robots. Webots is integrated with ROS, allowing for seamless transition from simulation to reality with no changes in code. In our simulations, the appearance and the devices of the MBot robot are faithfully simulated, with all the sensors and actuators carefully calibrated to match the reality, so the performance of simulated robots, especially at the multi-robot level, is comparable to real robots. Webots also provides simple models of humans, for which we have developed custom motion controllers based on research on crowd behavior. For experimental evaluation, Webots provides access to ground truth positioning information of robots and humans, accommodating the role of a real external tracking system. The real experimental facilities, including GR, IPOL and Jordils are reproduced with care, so that it is possible to carry out the same experimental scenarios in both simulation and reality. Snapshot of simulated humans and MBot, and examples of Webots worlds used in this thesis are shown in Figure 2.3.

2.3.3 Temporal Planner

In the case study of Chapter 16, the temporal course of the experiments (i.e. behavior activation and deactivation) is established offline by a temporal planner. The planner uses constraint satisfaction problem reasoning methods for determining relations between time intervals when the behaviors are executed based on specifications. More details about the planner can be found in [13], where it is used for enforcing temporal norms within the institutional formalism.

2.4 Detection and Tracking

At the heart of a networked multi-robot system lies the ability to communicate. In formation control and other consensus-based methods agents simply *must* know the state information of at least a subset of the neighbors. Furthermore, when we start to consider mixed societies of humans and robots, it becomes necessary to be able to reliably detect and track humans. In this section we describe the methods and tools used for obtaining human and robot positioning

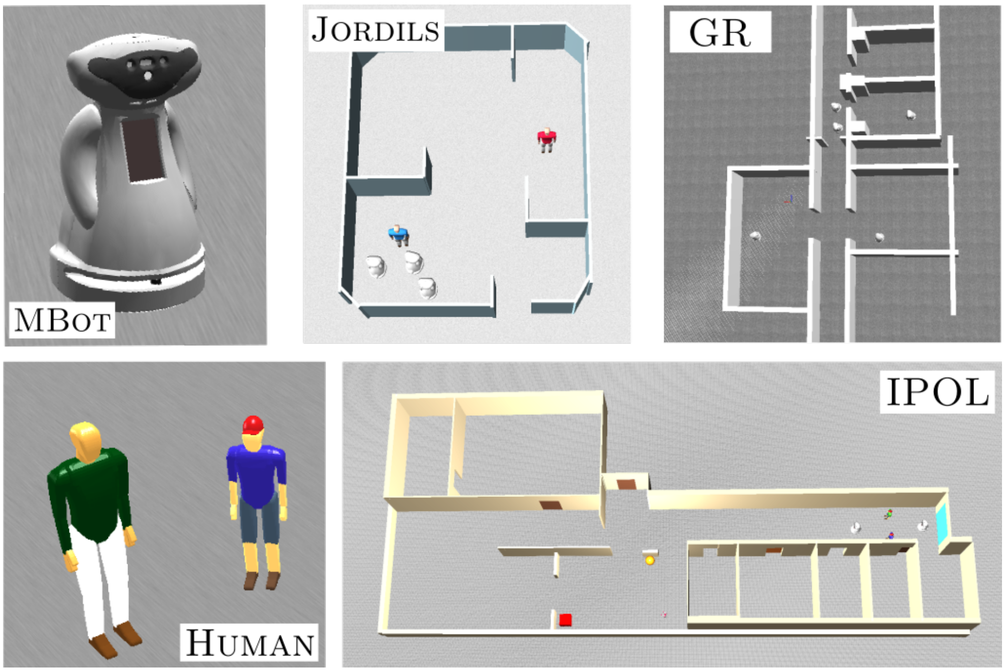

Figure 2.3 – Models of the MBot robot, humans, and selected worlds in the Webots simulator.

information and explain the data flow in our system.

2.4.1 Communication

All robots and external devices connect to a local wireless network and exchange messages using a ROS-based system, where selected topics and their content are automatically synchronized via multicast routing provided by a third-party ROS package[4] running on all connected devices.

2.4.2 Onboard Robot Detection

For robot detection we use an onboard relative localization system, which we describe in detail in [5]. At each time step, the laser range finder point cloud measurements are associated using the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm[5] from scikit-learn[6]. For a robot with a circular model of the base and known diameter, the clusters are subject to circle fitting using a least squares solver, where the center of the circle corresponds to the center of the base of the detected robot. The coordinates are compared against a known occupancy grid map to filter out false positives stemming from known obstacles. The final measurement is the position of the tracked robot. The accuracy of the detections, as

4 ROS package multimaster_fkie (http://wiki.ros.org/multimaster_fkie)
5 In [5] we use sliding-window, nearest-neighbor classification for that purpose.
6 DBSCAN data clustering algorithm, https://scikit-learn.org

determined in [5] is around 20 cm. We use the onboard detection system as the source of measurements for the tracking filter presented in Chapter 8.

2.4.3 Overhead Camera

A GigE color camera from Basler mounted under a ceiling at the GR facility is used for human position tracking in our experiments on adaptive formations in Chapter 7. The tracking method is based on active marker tracking with OpenCV-based color blob detection[7]. In these experiments the tracking information is only used for performance evaluation and is not fed back to the robots, as they rely on onboard sensing to avoid human in a manner that does not consider the social aspect.

2.4.4 Motion Capture System

The Motion Capture System (MCS) available at Jordils facility is a tool developed by Motion Analysis Inc. for tracking 3D movements of multiple targets with millimetric accuracy. In this work, we use a MCS for two purposes: for obtaining ground truth pose of the robots and for human tracking. The MCS leverages a set of 28 cameras anchored to a structure at the ceiling and tracks sets of reflective markers placed rigidly on the tracked objects. In the robot case, markers are placed at the top of robot's head, while for the humans we place the markers on plastic construction caps. The data from the cameras are processed by a dedicated software provided by the MCS manufacturer and results in a pose estimate for each tracked object. The pose estimates of humans are further translated to reference frame of the robots' map and relayed through ROS to the robots.

2.4.5 Ultra-Wide Band Localization

Ultra-Wide Band (UWB) localization is an emerging technology that has proven effective in indoor positioning and tracking. It provides low-cost, low-computation localization in semi-structured environments. However, it also suffers from spatially-varying measurement biases when emitters and receivers operate in Non Line Of Sight (NLOS) conditions, leading to a spatially-varying offset between the actual and the estimated position.

We have adopted the Kio UWB technology from Eliko[8] for human localization in a subset of experiments at the Örebro facility. Kio is a small size, easy to deploy UWB system consisting of tags - movable sensors attached to the tracked person, and anchors placed at known locations, used as references to compute the positions of the tags. UWB radio signals are exchanged between the anchors and the tags, and the tag's position is estimated by performing trilateration on the distances to the four nearest anchors. In our implementation, each tag carried by a person is powered by a power bank and connected to a Raspberry Pi 3 with an

7 OpenCV (https://opencv.org)
8 Eliko https://www.eliko.ee

on-board WiFi adapter[9].

2.5 Navigation Methods

The single-robot functionalities of the MBot robot have been designed and developed as part of the MOnarCH project, but are key for the development of our multi-robot methods. In this section, we describe the localization and the navigation methods of the MBot robot and shortly explain how they are integrated in our algorithms.

2.5.1 Mapping and Localization

For constructing a static occupancy grid map of the environment we use the off-the-shelf ROS package GMapping[10], which implements a Rao-Blackwellized particle filter for SLAM [14], generating an occupancy grid map from laser scan data and odometry. Robot self-localization is performed using another off-the-shelf ROS package AMCL[11], which by combining odometry and laser scan data tracks the pose of a robot within a known map using an adaptive Monte Carlo localization approach [15]. The accuracy of the self-localization system is around 10 cm in the facilities considered in this thesis.

2.5.2 Motion Planning and Obstacle Avoidance

For motion planning, we use a Fast Marching Method (FMM) [16], which provides a local-minima-free potential field that encodes the optimal direction of motion towards the goal. FMM provides a solution to the boundary value problems of the Eikonal equation - equation that models a wave front propagation, starting from the initial hypersurface, and propagating it with a speed specified by a scalar field. FMM is a numerically efficient solution employing discretization of space, therefore it can be directly applied to the occupancy grid map. Planning is followed by a Dynamic Window Approach (DWA) method to perform obstacle avoidance while driving the robot along the optimal direction. This is done by casting the problem of determining the next actuation command into a constrained optimization problem over a discrete set of candidate velocity commands. The set of valid candidates is obtained from an equally spaced grid of points in the actuation space, constrained to respect the speed and acceleration limits of the robot and to guarantee that, for each candidate, the robot can stop at maximum deceleration before hitting a perceived obstacle. In our work, the multi-robot algorithms deliver a velocity reference for each robot. The obstacle avoidance method described above integrates this information so that the DWA method will tend to prefer velocity candidates closer to the velocity reference, while assuring collision-free trajectories.

9 Raspberry Pi (https://www.raspberrypi.org/)
10 ROS package GMapping (http://wiki.ros.org/gmapping)
11 ROS package AMCL (http://wiki.ros.org/amcl)

	DIMENSIONS	ROBOTS	EXTERNAL TRACKING	CONTROLLA-BILITY	STRUCTURED
JORDILS	8 × 10	4	MCS	high	✓
GR	14 × 6	4	-	medium	✓
ÖREBRO	20 × 5	2 MBots, 1 Pepper	UWB	medium	✓
IPOL	23 × 8	2	-	low	✓

Table 2.1 – Characteristics of experimental facilities. *Controllability* of an environment indicates whether the area is specifically dedicated for robotic use (high), or it is a space for everyday use, frequently visited by people that do not explicitly participate in the experiments (low). *Structured* environment is constrained by structural obstacles, such as walls or doors.

2.6 Experimental Facilities

Our methods have been validated in a variety of environments, each posing a different set of challenges. Jordils is a controlled, open arena with walls scattered around, where we enact highly dynamic scenarios with multiple people walking around. The facility that we refer to as GR is a real office and laboratory space in the GR building of the EPFL campus, where we challenge the robot formation to move through multiple doors and kitchen frequently visited by users. The Örebro area at the Örebro University, where we deploy mixed human-robot formations, is a busy section of the AASS center, full of obstacles scattered around. Finally, we perform a number of experiments with a two-robot formation at the IPOL hospital – a highly dynamic, sensitive environment with challenging obstacles and several humans essentially always present in each room. Pictures of the experimental areas are shown in Figure 2.4, and in Table 2.1 we summarize their main characteristics.

2.6.1 Jordils

The experimental area at Jordils is a controlled environment, the structure of which can be build on-the-fly using modular building blocks. This facility is endowed with a MCS (described in Section 2.4.4), which we use to obtain human and robot poses with millimetric accuracy. Experiments performed at Jordils are presented in Chapter 8, Chapter 15 and Chapter 17.

2.6.2 GR

This area can be considered a fairly standard indoor setting cluttered with various appliances, including furniture, lab equipment and structural building features such as doors and columns. Our experiments take place between two rooms connected by a corridor, in an area

Figure 2.4 – Maps and pictures of facilities where we carried out our multi-robot experiments.

frequently visited by people. The challenges posed by this environment inspired our study on transformable formations in Chapter 7.

2.6.3 Örebro

Being in part a robotics lab, in part a corridor, the Örebro arena is cluttered with fragile equipment and frequently crossed by humans. At this facility we use the two available MBot robots and a Pepper robot, and track positions of humans with the Kio UWB system described in Section 2.4.5. We deploy altogether 12 anchors to simultaneously track two humans. Experiments performed in Örebro are presented in Chapter 17.

2.6.4 IPOL

In contrast to the other facilities, IPOL is a sensitive environment, where performing experiments with even one robot might generate severe disturbances to the regular operation of the pediatric ward if not properly designed and at least remotely supervised. Thanks to an incrementally built trust relationship with the hospital staff over long periods of operation with a single robot, we have been afforded the opportunity to deploy two robots for one day. In Chapter 6, we discuss the results of experiments in this highly dynamic, fragile environment, where groups of visitors block robots' paths, children run next to the robots trying to stop them, and busy staff can be easily disturbed by the robots' motion. The lessons learned at IPOL inspired the leading idea of this thesis - to make multi-robot behaviors social. Moreover, the communication problems we encountered at this facility are the leading motivation of our work on multi-robot tracking, presented in Chapter 8.

Summary

Performing experiments with multiple robots in complex, human-populated environments requires a reliable robotic platform with robust localization and navigation tools. If robots are expected to behave socially, they must be able to perceive humans, or at least obtain their poses from external systems. In this chapter, we have presented the key system components upon which the methods developed in our research heavily rely. From algorithmic design in simulation to validation in experimental facilities, all presented tools will accompany our journey through the subsequent chapters.

Part II Cooperative Navigation

"Darwin had no idea..."
"That life is so unbelievably complex," Malcolm said.
"Nobody realizes it.
I mean, a single fertilized egg has a hundred thousand genes, which act in a coordinated way,
switching on and off at specific times, to transform that single cell
into a complete living creature.
That one cell starts to divide, but the subsequent cells are different.
They specialize. Some are nerve. Some are gut. Some are limb.
Each set of cells begins to follow its own program, developing, interacting.
Eventually there are two hundred and fifty different kinds of cells, all developing together, at exactly the right time.
(...)
Week after week, this unimaginably complex development proceeds perfectly – perfectly."

Michael Crichton, The Lost World

3 Introduction

MULTI-ROBOT cooperative systems are becoming increasingly relevant to real-world applications, including search-and-rescue missions, evacuation of humans in emergency situations, or environmental monitoring. The ability to efficiently navigate as a group enables a team of robots to perform activities not possible for single robots, contributing not only to enhanced performances when carrying out specific missions but also bringing additional flexibility and robustness to failures. However, cooperative navigation in complex, structured environments poses many technical challenges, including how to handle spatial constraints when robots are moving as an ensemble (for example, in a formation of large robots). Furthermore, introduction of robots in social environments entails some basic level of human awareness, i.e. the ability to act around humans in a way that appears friendly, social, and safe. Finally, in order to ensure inter-robot positioning robustness, also in case of communication failures, we identify the need for a reliable cooperative localization system.

In this chapter we motivate the use of cooperative approaches for multi-robot navigation and look into the aforementioned challenges, identifying the steps necessary to address them.

3.1 Cooperation and Coordination in Natural and Artificial Systems

Compared to single robots, multi-robot systems can offer a large variety of advantages, including, to name a few, enhanced efficiency and flexibility achieved through concurrent execution of the task, distributed sensing and actuation, capability to specialize, and increased tolerance to robot failure. Moreover, some tasks might not even be possible for a single robot. However, to thrive upon the benefits of multi-robot systems it is not enough to simply send out several robots and let them act on their own - robots must cooperate with each other. From the very onset of multi-robot systems, solutions for effective cooperation have been sought for in nature. Since in this work we are particularly interested in multi-robot navigation, we will now take a closer look at solutions for motion coordination in natural and artificial systems.

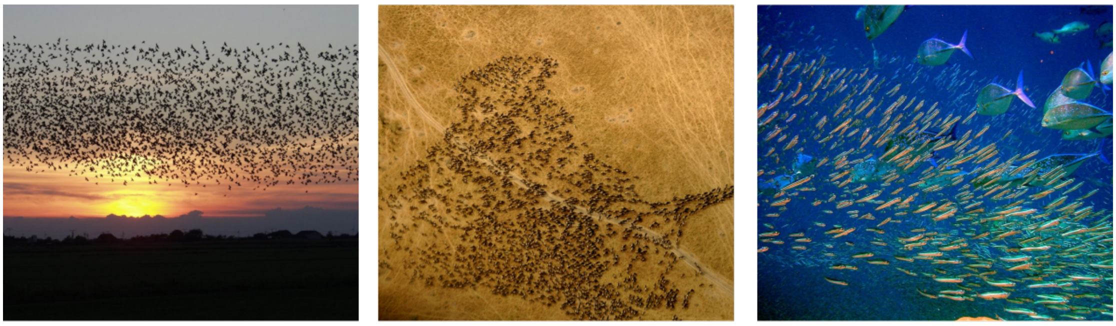

STARLING FLOCK | WILDEBEEST MIGRATION | FISH SCHOOLING

Figure 3.1 – Examples of flocking in the nature.[1]

3.1.1 Inspiration from Animal Flocking

Collective animal behavior is a fascinating phenomenon where large groups of social animals coordinates their movement, resulting in a functional global structure. Such cooperation provides numerous advantages ranging from effective foraging to protection from predators. Such systems are typically leveraging interactions among individuals executed on the basis of purely local information, with no external guidance or central coordination.

Collective movements in animal societies, such as flocking of birds, schooling of fish or herding of land animals (shown in Figure 3.1) are popularly implemented with three simple local rules:

1) Cohesion – remain close to your neighbors

2) Separation – avoid collisions with the neighbors

3) Alignment – adjust heading to that of the neighbors

The three rules above have been originally proposed by Craig W. Reynolds [17] in order to produce realistic computer animations of bird flocks. Similar rules are used for modeling of crowd behaviors [18], where humans move as a result of attraction towards destination and repulsion from other people and obstacles. The flocking ruleset has received a great deal of attention in artificial multi-agent systems [19] [20], because of its straightforward applicability to a variety of virtual and physical platforms.

While the flocking rules applied to multi-robot systems result in robust, flexible motion, the resulting inter-agent spatial coordination remains loose: the overall grouping can be typically influenced by the balance among the three rules and a perceived field guiding the individual robots towards their destination (Reynolds called it the “migration urge”). To control the inter-robot distances with higher accuracy, we will now take a closer look into robot formations.

1 (left) Image by Tommy Hansen, https://en.wikipedia.org/wiki/Collective_animal_behavior (CC BY 3.0)
(middle) Image by T. R. Shankar Raman, https://en.wikipedia.org/wiki/Wildebeest (Public Domain)
(right) Image by Bruno de Giusti, https://en.wikipedia.org/wiki/Shoaling_and_schooling (CC BY-SA 2.5)

3.1.2 Robot Formations

Compared to the flocking rules, formation control methods provide an additional degree of control over inter-robot positions – range and bearing control, which at the team level results in the control over the overall geometric shape of the robot team. Formation control is thus employed where spatial topologies are important – in aerospace and outer space applications (e.g., space-borne optical interferometry [21]), in systematic search missions (e.g., odor source localization [22]), and in surveillance and mapping. The methods for formation control can be loosely classified into a) behavior-based approaches, where simple motion primitives are composed to build more complex patterns [23]; b) virtual structures, which consider the formation as a rigid body with motion of each robot following a point on a moving virtual shape [24]; c) leader-follower strategy, where the leader moves towards a group goal while the followers maintain an offset from the leader to keep their desired place in the formation [25]. Other alternatives include fuzzy systems [26], and artificial potentials [27].

The ability to represent the individual robots and their relative spatial positions in an abstract form makes it possible to ignore technical details of a particular implementation. Graph-theoretic methods provide means for such representation by abstracting away the complex interplay of sensing, actuation and communication of individual robots and offer tools for synthesis and analysis of networked systems [21]. We explore a graph-theoretic framework in our work to achieve cooperative group movement. It is worth noting that the adoption of such framework is well-aligned with our institutional formalism which also aims at abstraction from the physical implementation.

3.2 Challenges of Multi-Robot Systems in Real Environments

Robustness and performance in multi-robot teams is achieved at the cost of behavioral complexity - multi-robot systems are more difficult to synthesize and analyze, as the resulting behavior emerges from a large number of interactions. The challenges of single-robot navigation are exacerbated – while it might be difficult for a single robot to move successfully in a complex environment cluttered with obstacles and frequently visited by humans, a multi-robot system deals with the same problems at much larger scale and with often the additional constraints to maintain some sort of coordination with the teammates. In the following subsections, we will discuss the main challenges of deploying multi-robot systems in real environments.

3.2.1 Formation Control in Complex, Structured Environments

Indoor human environments are typically characterized by a large number of appliances such as furniture or equipment, and structural building features such as corridors, doors, or other wall-restricted spaces. State-of-the-art methods for multi-robot formations yield good performance when tested in laboratory conditions, and the underlying obstacle avoidance is sufficient to deal with open environments with obstacles appearing sporadically. However, strong assumptions regarding sparseness of the obstacles and openness of the environment with lack of structure makes them ill-prepared for real, indoor environments. For a successful navigation of multiple large-sized robots in complex environments additional dexterity is needed - robots not only must be able to avoid obstacles or shrink the shape of the formation, but also tolerate high flexibility in terms of formation geometry when passing through corridors and doors.

3.2.2 Unreliable Communication

As the robots are becoming more present in environments populated with humans, safety, robustness, and reliability of the methods are of highest importance. While many of the multi-robot approaches do not require continuous coordination among the agents, methods such as formation control and other consensus-based algorithms heavily depend on access to the state information of the other robots and require robust and reliable information flow among them.

The simplest solution for the robots is to communicate the required information to each other, which involves constant information exchange between a robot and all of its neighbors in the network [28]. Unfortunately, wireless communications suffers from many problems, including message losses, delays, and even temporal loss of connections. Volatile communications threatens formation stability, while temporary loss of communication during close proximity navigation might even lead to collisions. Our robot formations are deployed in settings which are by far not standard for the majority of the formation methods, namely Global Navigation Satellite System (GNSS)-denied, complex indoor environments populated with humans and obstacles, where a positioning system based on direct inter-robot measurement suffers from long-term occlusions and false detections [2] and the communication can suffer from short-term outage. While communication is often the only possible approach to formation control in dynamic and structured environments full of obstacles, it is necessary to provide a backup solution for securing the formation when the communication fails. An example of such situation is shown in Figure 1.1. There exist well-studied approaches for connectivity maintenance aiming at improving the robustness of the distributed robot network [29][30], however in our setup network failure is location-independent. For this reason, we interpret the problem from a global perspective and explore a multi-target tracking approach to face the issue of unreliable communication.

3.2.3 Human Awareness

Addressing the technical challenges related to reliable navigation in complex environments is a must when robots are to be deployed in human environments. However, it is not sufficient. When moving among humans, robots can cause discomfort by either moving too closely, too quickly, or too loudly. They can be a nuisance, if they get into a way or interfere with human activities. Therefore, research on human-aware navigation aims at maximizing robot acceptance, minimizing annoyance and stress caused by a robot, and making the interaction more comfortable and natural.

According to [31], human-aware navigation studies focus on three aspects: 1) human comfort, with the goal of minimizing human stress and annoyance, 2) naturalness, aiming at making robot predictable and intuitive, and 3) sociability, where robots act according to social norms. To improve human comfort and make a person feel safe, the literature proposes strategies such as maintaining proper distance, speed control or planning to avoid interference. To make robot motion understandable, navigation methods focus on trajectory smoothness and general resemblance to how humans move.

Sociability (or social awareness), refers to the ability of a robot to follow behavioral norms expected by the people with whom the robot is intended to interact [32]. For this, robots need an adaptable model of human social behaviors, ability to recognize social contexts and conventions, and, eventually, bear means to support diversity of user cultural and social backgrounds, different ages, genders, etc. in order to deal with different human abilities and preferences [33]. Moreover, robots should clearly convey their intentions to the onlookers and be proactive in their actions. While the state-of-the-art research is far yet from designing a perfect social robot, there is a rich body of work in human-aware and social navigation providing a wide range of solutions [31] and further insights into the design of a socially acceptable robot.

3.2.4 From Single-Robot to Multi-Robot Social Navigation

While the design of a socially acceptable navigation strategy for a single robot is a challenge by itself, addressing it from a multi-robot perspective boosts the problem complexity. Well established human-aware navigation methods have no equivalence for multi-robot systems, as there is no analysis of the impact of the presence of multiple robots on their acceptance by the humans. Similarly, social behavior of a group of people can differ from that of a single person, for instance, in the management of space [34]. It is out of scope of this thesis to perform extensive user studies necessary for understanding the impact of multi-robot (and multi-human) aspects on social navigation. Our approach is to use well-established single-robot methods in the multi-robot context with an assumption that their core mechanisms are adequate, even if not complete.

Summary

Collective systems are known for their flexibility to adapt to the environment and robustness to individual failures. Multi-robot systems offer a large variety of advantages, including parallelism, scalability, flexibility, and robustness through redundancy and absence of a single point of failure. However, the deployment of a robust cooperative team of robots navigating around humans in a socially aware manner requires overcoming many challenges, including collective navigation in complex environments, development of a no-single-point-of-failure system architecture, and exploration of human-aware single-robot navigation solutions in a multi-robot context. In the next chapter, we will review the state-of-the-art approaches to address these challenges.

4 Related Work

THE deployment of cooperative multi-robot systems navigating in human-populated environments in a socially aware manner requires overcoming a number of challenges of technical and social nature. In the previous chapter, we put an emphasis on three such challenges we have identified during our preliminary experiments at the IPOL hospital. First, the need for robust navigation methods enabling a team of robots to move collectively in complex environments. Second, the development of a no-single-point-of-failure system architecture to provide the robots with a reliable source of information, even in case of communication failure. And finally, the integration of human-aware solutions in a multi-robot context. In this chapter, we review the related work addressing these challenges.

4.1 Formations in Complex Environments

As we emphasized in Section 3.2.1, existing methods for multi-robot formations rely on strong assumptions regarding the nature of the environment. For instance, an attractor dynamic approach in [35] controls the formation topology in situations where no obstacles are present, but breaks the formation during obstacle circumvention, therefore assuming an implicit capacity to split and join a formation upon encountering a single object in a sparse space. Dynamic change of formation in [36] performs well in open outdoor spaces, but may not be adequately reactive indoors. The decision to change the formation shape is taken globally by a leader, which upon detection of an obstacle informs the followers about the new desired formation geometry. Approaches to formation scaling, morphing and rebuilding behaviors in [37] allows for the selection of new formation shapes minimizing the total formation error after completing an obstacle avoidance maneuver. However, all the robots participate in the deliberation of the formation topology; the formation change is global (as opposed to local) and implemented as a discrete switch between two different shapes. Transitions between formation control and other modes of operation such as obstacle avoidance and wall following are tackled in [38], but only for unstructured environments. Other approaches that rely on a modification of the formation geometry for obstacle avoidance include priority-based arbitration among behaviors [39], formation change using a transition matrix [40], scaling [41], or avoiding obstacles while handling the formation as a rigid body [42]. Deformable formations are studied in [43] and in [22], where the exploration of a noisy environmental field drives the shape of the multi-robot formation. A method with potential to create flexible formations proposed by the gaming community [44] represents a crowd of simulated characters as a deformable mesh, with vertices being control points that allow the user to prescribe desired trajectories. To prevent congestions, the characters are reassigned in formation and obstacles difficult to avoid are negotiated by splitting and rejoining, but the high-level control is centralized.

A common assumption in the above approaches is that (a) the experimental area is large, i.e. the formation has enough space to reorganize after negotiating the obstacle, or (b) the obstacles are cluttered in the environment and do not restrict the environment itself, i.e. the experimental settings reflect outdoor, unstructured environments without boundaries.

Perspective

In view of the limitations summarized above, in Chapter 7 we introduce the Local Formation Transformation (LFT) algorithm – an approach to dynamic formation change that is *local* (meaning that the formation is only reshaped in the immediate neighborhood of the robot that initializes the change) and *gradual* (meaning that the formation does not switch topologies but is modulated in its shape to some extent), with the level of shape alteration proportional to the density of obstacles around the robot. We relax the traditional assumptions on experimental settings and present results of cooperative formation patrolling in structured indoor areas characterized by confined spaces, i.e. spaces not large enough to deploy the desired formation.

4.2 Tracking for Formation Control

The problem of tracking for realization of multi-robot formations has been addressed as early as the first approaches to formation control appeared. Among many existing formation control algorithms, the most common ones rely on the pose estimates in the global reference frame [24], on relative positions of the other robots [45], on range only [46] or bearing only information [47]. While acquisition of accurate state of the other team members has been addressed previously using various perception tools, including cameras [35], infrared sensors [45], sonars [48] or laser range finders [5], tracking is significantly simplified in order to be reliable enough for formation control.

Simplification usually casts a multi-target tracking problem to single-robot tracking by providing the robots with unique identification tags (IDs) that can be extracted by the tracking robot. Tracking the multi-target estimates with known IDs is trivial, as long as not all the tracks are lost. To realize an ID-based formation, multiple solutions have been proposed in the literature. In [25], robots perform teammate detection using a combination of a LIDAR and a camera, tracking a uniquely colored marker, which in turn can reveal the identity of the neighboring robot. In [38], robots recognize themselves by extracting color blobs from a camera image. Both approaches aim at localizing a single local leader distinguishable by a marker. In [48], a follower vehicle maintains a formation with two leaders using acoustic ranging. The distinction between the leaders is ensured by an appropriate time-multiplexing scheme of the acoustic relative positioning signals. For tracking multiple quadrotors, [49] use active markers and an on-board camera. The markers provide 3D poses of the robots, which by pulsating at a predefined frequency, provide a unique aircraft ID. ID-dependent graph-based formation is achieved in [45]. Robots use a dedicated infrared range-and-bearing system and exchange messages containing robot IDs. In [50], agents localize using bearing-only measurements but are constrained to move with the motion type that is known by all agents a priori.

The above approaches all share a common element – ID of the neighbor robots can be extracted from the sensing data. When this is not the case and multiple robots are to be tracked, one must either perform data association at the tracking level or assign each estimate to a given role in the formation. Within the context of multi-robot coordination, the problem of role assignment has been addressed previously using potential fields [51], market-based algorithm for task allocation [52], and the Hungarian algorithm in [53] and in [54] for formation initialization. Those works however consider only a static case, i.e. the roles must be assigned only once, and are kept throughout the experiment. Dynamic role assignment, where roles are updated online by the robots as they navigate in the environment, can be performed as part of the tracking algorithm. In [55] state estimates are associated with unique track labels within a Gaussian Mixture Probability Hypothesis Density (GM-PHD) filter. In [56] the IDs of the robots are reconstructed by incorporating odometric data directly in the Probability Hypothesis Density (PHD) filter. Both methods maintain the track-label association, however an explicit assignment of the labels to formation roles would require further manipulation. To overcome the problem of mutual multi-robot localization with ID-less measurements, in [57] robots are allowed to communicate their IDs and the ID-less measurements of the other robots. The robots self-localize in their individual coordinate frames, the relative configuration of which has to be estimated. The proposed method uses probabilistic data association techniques, combinatorial in their nature, together with multiple particle filters, one per each robot.

Up to date, most of the existing research addresses robot localization and multi-agent formation control as two separate problems and only very few studies combine them as an integrated control problem. Velocity and relative position estimation integrated with formation control is studied in [58], but each agent is forced to carry out a specific combined circular and linear motion during the entire process. To control the agents in the formation that are not able to measure the relative positions of their neighbors, [46] devised a method called stop-and-go. A consensus-like relative localization using measurements and local communications in [59] is integrated with leader-follower formation control by combining the proposed relative localization scheme and a Laplacian-based formation control method, allowing to achieve desirable convergence properties. ID data are available to the agents as they communicate locally. The above methods mostly focus on the theoretical aspects of the problem and show only results in simulations with simplistic models. Assumptions on robot motion and availability of additional information make them not suitable for ID-based formations with ID-less tracking data.

Perspective

To ensure robustness of the navigation methods to communications failure, in Chapter 8 we study robot localization and multi-agent formation control as an integrated control problem. We work with ID-based robot formations, where each robot is assigned a role in the formation (target position), but the tracking data does not provide the identity of the robot; it is impossible therefore to utilize single-target tracking methods. For this reason, we introduce our method for incorporating communication data, tracking information, and knowledge about the desired formation geometry in the Gaussian Mixture Probability Hypothesis Density (GM-PHD) filter [60]. Furthermore, we perform an online role assignment, where the estimates are matched with the expected target positions in the formation. Designed to deal with short-term communication outages or low communications throughput, our methods perform well even in obstacle-cluttered complex environments with high measurement uncertainty and sensing occlusions.

4.3 Single-Robot Human-Awareness

Human-aware navigation deals with modeling and respecting norms of human societies, not only to avoid discomfort or animosity but also to improve robot acceptance and integration in human environments. Human-aware navigation introduces elements of research on Human-Robot Interaction (HRI) to robot motion planning, and so, a robot not only takes into account constraints posed by an environment, but also those related to human comfort and social rules. From among many possible types of interaction between robots and humans, that affecting navigation is the most subtle as no direct communication with the humans is considered. In this section, we review a selection of human-aware navigation methods, which will inspire our solutions in the case studies in Part III.

The study in [31] argues that a human-aware robot should possess the following capabilities:

Respect personal spaces. A personal space is an area around a human, intrusion into which can cause discomfort. It can vary across cultures and familiarity groups [61] and depend on personality traits of the individual [62]. Analysis of spatial distances in social and interpersonal situations is captured in the Proxemics Model (PM) [63]. The model has been adapted in robotics to describe virtual spaces around a person that the robot should respect and became the most popular paradigm to describe space management around humans. A multi-human counterpart of proxemics – O-spaces – deals with the management of space around a (static) interacting group of people. The space, being more than a mere addition of individual personal spaces, is a set of concentric circles around the group, which vary depending on poses of the participants [64]. In dynamic situations, particularly in crowds, spatial constraints around humans are popularly represented using potential field methods [65] [66] and cost functions concentrated around human poses. Once assigned to the map of a robot, the repulsive forces influence the robot motion planning [31]. It has been shown in [67] that the performance of the reactive methods can be improved by adding a predictive component in the model of human motion [68].

Respect affordance spaces. The affordance space is another type of virtual space arranged around a human activity (or a potential human activity) that should be avoided by robots to prevent potential interference [69]. An example of such space is an area between a person watching a TV and the TV itself, which should be circumvented to avoid interference. Most commonly, affordance spaces are represented using cost functions, which can be static or dynamic, deterministic or probabilistic [70] and can take multiple shapes depending on the human activities [64] [69].

Approach for explicit interaction. The simplest and most common strategy to identify locations in the environment adequate to stop for interaction is based on proxemics, but other methods, for instance based on potential fields, have been proposed [31]. Various aspects of human approaching behavior are addressed, ranging from pose selection [71] and velocity control to visibility [72] and modulation of gaze direction.

Adjust velocity around humans. Studies on the impact of robot velocity on human comfort [73] suggest that there is a limit to which robot speed is acceptable. If chosen inappropriately, velocity can have a high impact on human perception of a robot – a large robot moving too fast can be threatening, while robot moving too slow might be perceived as incompetent.

Modulate gaze direction. Gaze direction has been found to have an effect on human comfort, but some contradictory results are obtained depending on human gender [64]. It is natural to assume that during interaction a robot should direct its gaze towards the human partner. It is generally recommended to avert gaze when approaching human to avoid conveying a threat [74].

Avoid erratic motions or noises. To achieve smooth motion one must account for various criteria including environment structure, unknown objects, social conventions, and human proximity constraints [75], but, as it has been pointed out in [31], for achieving optimal results it is necessary to predict human movement. Natural, human-like navigation can be achieved with strategies based on models of human motion in crowds [76], but some results suggest that not all behaviors appropriate for humans are suitable for human-friendly robots [77].

Avoid culturally inappropriate behaviors. – In other words, adhere to human social norms, acknowledging personal and cultural preferences [78]. Examples of such norms refer to people ordering in queues, letting people leave a room or an elevator before entering, giving priority to elderly or disabled, walking on the right (or left) side of a corridor [79] [75], apologizing for crossing somebody's personal space, or avoiding to pass through a group of people [75].

Perspective

The methods mentioned above are the tip of an iceberg of a large body of work on human-aware robot navigation. In Chapter 10, we provide more details on the selection of methods used in this thesis, including the proxemics model, cost maps and examples of social norms embodied in our case studies in Part III. By exploring a variety of approaches typically adopted by socially aware robots, we illustrate the power of the institutional framework to represent and abstract the underlying methods, allowing for unification of the development of social-aware, multi-robot behaviors.

4.4 Multi-Robot Human-Aware Navigation

Research on human-aware navigation addresses the question of how to mitigate the inner reserve we have towards the non-living objects. Human lack of comprehension of robot behavior and inability to predict the robots' movements may cause vigilance, especially with people that rarely have a contact with autonomous machines. The issue becomes even more exacerbated when the environment is shared with multiple robots. Although human-aware navigation is a widely studied subject, only few studies on human-aware multi-robot systems exist. Moreover, one has to distinguish cooperative methods, where robots act and move as a team, from approaches, where robots act as individuals. In the latter case the robotic team has a common goal (such as greeting visitors or moving among a crowd to provide information services), but the robots undertake independent actions. For example, in [80] four robots greet visitors in a shopping mall and guide them towards goals allocated in a Wizard-of-Oz manner, therefore no coordination between the robots is required. An obstacle avoidance method with miniature robots in [77] based on the imitation of pedestrian behavior results in trajectories intuitive for the interpretation by humans, but no experiments with humans and robots are performed. Human-operated multi-robot systems are also employed in urban search and rescue missions [81].

The above mentioned methods consider multiple robots moving independently. To the best of our knowledge, only a small number of studies involves cooperative multi-robots systems. In the context of coordinated multi-robot navigation, a large body of research is related to human guidance. Strategies stem from early research on flocking [17] and control of herds [82], also in a multi-robot case [83], they treat the group of humans as dynamic particles [76], dynamic obstacles [84], or assume that humans simply follow the robot [85]. Formation control of aerial vehicles navigating in environments with static and dynamic obstacles – humans – has been shown in [84], but no social aspects are considered. Motion control of groups of formations in [86] is designed with the intention of achieving team's legibility, i.e. ability to convey information such as the state or the task to the user. Experiments are carried out in virtual reality, therefore the physicality, and consequently, the social aspect of the human presence is left out. A preliminary study on human crowd dynamics [66] has shown that in some situations the presence of cooperative robots can improve pedestrian flow, while in others can lead to unexpected detrimental effects. Methods are evaluated only in simulation and with an idealized human model. The above solutions are largely over-simplistic and fail to consider realistic situations. A more realistic study [87] proposes a human-aware prediction and anticipation model that can handle realistic situations and regrouping of people who have left the team. The approach is thoroughly tested in simulation, but no real-world results are presented. An adaptive multi-robot task allocation strategy in [88] deals with cooperative re-planning in presence of humans, where the social aspect and human motion uncertainty are encoded into risk-based bids. Methods are validated in experiments with two robots and up to two humans. Robots cooperate at a task allocation level, but not at the navigation level.

To conclude, although numerous studies attempt to develop human-aware robotic teams, the current models of humans are too simple, the algorithms work under too many assumptions and the experiments are too controlled. In contrast to what is claimed in the literature, we believe that currently roboticists are not yet able to launch groups of robots that could interact and move around groups of humans. From our perspective, the state of research on human-aware navigation is largely mature in the single robot case, but when it comes to studies of cooperative multi-robot systems, only few studies target realistic applications, while the research on social robots navigating as a group handles the presence of a person inappropriately or even naïvely, and solutions are heuristic and difficult to generalize.

Summary

In view of the literature summarized in this chapter, we conclude that robot formations are largely deployed in controlled environments and formation reconfiguration in presence of structured indoor obstacles is a novel concept. Although solutions to cooperative tracking have been employed in the context of formation control, the problem is reduced to that of tracking a single robot by assuming that robot IDs can be sensed. Only very few works attempt to combine formation control and robot localization as an integrated problem, but these approaches are not suitable for ID-based formations with ID-less sensing, as it is our case. Finally, little research has been conducted on cooperative multi-robot teams deployed in human-populated environments: most of the multi-robot systems focus on the deployment of individual, uncooperative robots or cooperative solutions that fail to consider challenging situations from a group navigation perspective .

In the next chapter, we provide the background on graph-theoretical approaches to the formation control and flocking. That chapter lays the background for the remaining parts of this thesis, as all our case studies involve cooperative robot teams. After that we present our method for adaptive formation control in complex environments and present our approach for dealing with the challenge of unreliable communication – an issue we must address before deploying formations in human-populated environments.

5 Cooperative Spatial Coordination

SPATIAL group coordination realized by a team of cooperative robots is characterized by distinct geometrical patterns that result from a significant number of low-level interactions, as opposed to a global specification given by a centralized authority. Graph theory provides means for practical representation of such interactions, and allows for abstraction of technical details and encapsulation of concepts necessary for the specification of complex multi-agent behaviors. In this chapter, we introduce the basic elements of graph theory, which we will then leverage for formation and flock control in the rest of the chapter.

While tight motion coordination achieved with formation control allows for flexibility, at the same time retaining good control of the team spatial configuration, flocking as a bio-inspired solution can result in more natural motion. For this reason we consider the two methods to be complementary and use both of them to showcase the mechanisms of our formalism, in particular the ability to abstract not only from a physical system, but also from the underlying behaviors and behavior specification.

5.1 Introduction to Graph Theory

An undirected graph with N elements is defined as a pair $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_i, i = 1 \dots N\}$ is the *vertex set* and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the *edge set*. Edges represent the flow of information, that is if v_i can observe v_j, then there exists an edge ε_{ij}. In this work if v_i can communicate with v_j then v_j can also communicate with v_i, thus the edges are undirected.

An *incidence matrix* $\mathcal{H} \in \mathbb{Z}^{N \times |\mathcal{E}|}$ describes which edges connect which nodes. It is usually defined on directed graphs, but if (without loss of generality) a random orientation is assigned to the edges of the undirected graph $\mathcal{G}$, the incidence matrix takes the following values:

$$\mathcal{H}_{ik} = \begin{cases} -1 & \text{if } \varepsilon_k = (v_i, v_j) \\ 1 & \text{if } \varepsilon_k = (v_j, v_i) \\ 0 & \text{otherwise} \end{cases} \tag{5.1}$$

where $|\mathcal{E}|$ is the cardinality of the edge set and ε_k is the k^{th} edge of $\mathcal{G}$. In this work, vertices correspond to the robots and edges represent the existence of an information flow among them.

Rendezvous

Consider the robots modeled using a single kinematic integrator, $\dot{p}_i = u_i$, where $p_i \in \mathbb{R}^2$ represents the position of robot R_i on a plane and u_i is a control input. For a team of N robots, a simple averaging strategy commonly used in continuous consensus problems declares that each robot should be attracted to its neighbors with a strength proportional to the weights w_{ij} of the corresponding edges of the graph:

$$\dot{p}_i = \sum_{j\sim i} w_{ij}(p_j - p_i) \tag{5.2}$$

where $j \sim i$ means robot R_i and robot R_j are connected. The above solution to the *rendezvous* problem (all robots converging to the same position) can be rewritten in a standard form of the Laplacian-based feedback control [21]:

$$\dot{p} = -\mathcal{L}p \tag{5.3}$$

The *Laplacian matrix* $\mathcal{L} \in \mathbb{R}^{\mathrm{N x N}}$ defined as $\mathcal{L} = \mathcal{H}\,\mathcal{W}\,\mathcal{H}^T$ for an undirected graph is a symmetric and positive semi-definite matrix [89], with one eigenvalue equal to zero. The zero eigenvalue ensures state convergence $\lim_{t\to\infty} p(t) = v_1 v_1^T p_0$, where $v_1 = \frac{\mathbf{1}}{\sqrt{N}}$ is the normalized eigenvector corresponding to the zero eigenvalue.

A *weighted Laplacian* is a Laplacian with weights assigned to the edges: $\mathcal{L} = \mathcal{H}\,\mathcal{W}\,\mathcal{H}^T$, where $\mathcal{W} = \mathrm{diag}(\{w_k, \forall \varepsilon_k \in \mathcal{E}\}) \in \mathbb{R}^{|\mathcal{E}|\times|\mathcal{E}|}$ is the weight matrix. By controlling the relative weights of the connection edges one can achieve better control over system convergence, as we will see in the next sections.

Formation Control

While $\dot{p} = -\mathcal{L}p$ drives the agents to a rendezvous point, formation control requires an additional bias matrix b, so that:

$$\dot{p} = -\mathcal{L}(p - b) \tag{5.4}$$

defines a desired deviation from a center point of the formation [21].

5.2 Graph-Based Formation Control

The unconstrained holonomic locomotion system of the MBot robots allows us to use the standard Laplacian-based feedback directly as a control input.

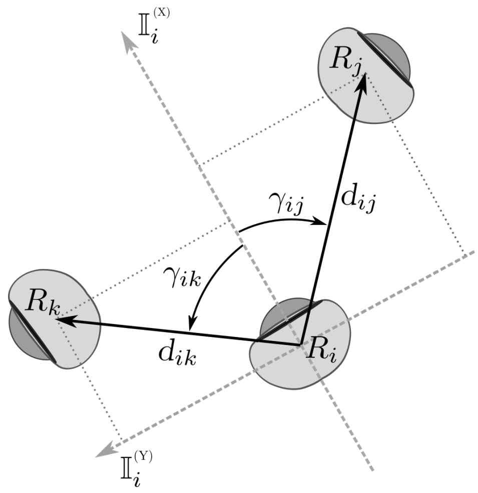

Figure 5.1 – Illustration of the graph-based connectivity information with respect to the robot R_i, where d_{ij} is the Euclidean distance between R_i and its neighbor R_j, and γ_{ij} is the bearing. Dashed lines indicate the local coordinate frame $\mathbb{I}_i$ of R_i.

Leaderless Formations

For a robot R_i with position $p_i = [x_i, y_i]$ and orientation α_i, the formation control presented in Equation 5.4 is achieved as follows:

$$\begin{aligned}\dot{x}_i &= K_u \frac{1}{|\sum_j \mathcal{L}_{ij}|} \sum_{i\sim j} \left[-\mathcal{L}_{ij}(d_{ij}\cos(\gamma_{ij}) - b^x_{ij})\right] \\ \dot{y}_i &= K_u \frac{1}{|\sum_j \mathcal{L}_{ij}|} \sum_{i\sim j} \left[-\mathcal{L}_{ij}(d_{ij}\sin(\gamma_{ij}) - b^y_{ij})\right]\end{aligned} \tag{5.5}$$

where $\mathcal{L}$ is a piecewise time constant Laplacian, d_{ij} and γ_{ij} are the Euclidean relative range and the bearing between the robots R_i and R_j respectively, as shown in Figure 5.1, and K_u is a positive constant.

As mentioned above, a relation $i \sim j$ means robot R_i and robot R_j are connected. If there is a directed edge from R_i to R_j, R_i has access to state information of R_j. In our work, unless stated otherwise[1], the graph is considered fully connected with unidirectional edges, because each robot maintains the information about the spatial states of the others at all times. The state information can be communicated through a standard radio channel (e.g. WiFi). Communication neighborhood set of robot R_i, $\Xi_i = \{j | i \sim j, i \neq j\}$, is a set of agents to which state information R_i has access to, while the term $|\sum_j \mathcal{L}_{ij}|^{-1}$ is used for normalization with respect to the number of neighbors.

The bias vectors $b^x_i, b^y_i \in \mathbb{R}^N$, where N is the number of robots in the formation, define the desired inter-robot distances along the x and y axes of the formation's coordinate frame respectively. Example of a bias for a diamond-shaped formation of five robots is shown in Figure 5.2. Comprehensive demonstration of this controller stability has been presented in [90].

Finally, the robots are omnidirectional, so the heading is decoupled from the velocity control

1 This assumption is relaxed in our work on cooperative tracking methods for formation control presented in Chapter 8.

to match the desired orientation α_D:

$$\dot{\alpha}_i = K_\phi(\alpha_D - \alpha_i) \tag{5.6}$$

where K_ϕ is a positive constant.

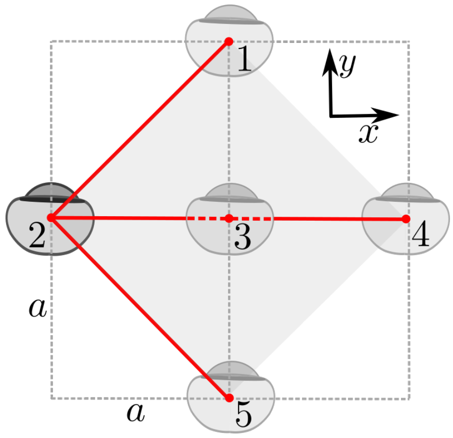

Figure 5.2 – Illustration of the bias matrix b for a diamond-shaped formation of five robots, specified in an absolute coordinate frame. With respect to the robot R_2, the bias vectors along the x and the y axes of the formation take the following form: $b_2^x = [a, 0, a, 2a, a]$ and $b_2^y = [a, 0, 0, 0, -a]$, the index of which corresponds to the robot ID.

5.2.1 Leader-Follower Formations

Leader networks extend the idea behind the graph-based control by allowing some team members – *leaders* (denoted as L) to perform an independent behavior (such as moving towards a targeted destination), while others – *followers* (denoted as F) to converge to the predefined configuration, simultaneously keeping up with the leader. With regard to the underlying graph, the follower-follower edges remain bidirectional, but the leader-follower edges are unidirectional, leaving the leader and entering each follower, indicating information flow from the leader to the followers, but not the opposite.

Specifically, the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ becomes $\tilde{\mathcal{G}} = (\tilde{\mathcal{V}}, \tilde{\mathcal{E}})$, where the vertex set $\tilde{\mathcal{V}}$ holds the *leaders subset* $\mathcal{V}_l$ and the *followers subset* $\mathcal{V}_f$. A partition of the vertex set leads to an analogous subdivision of the edges, $\tilde{\mathcal{E}} = \mathcal{E}_l \cup \mathcal{E}_f \cup \mathcal{E}_{fl}$. Details of how to partition the Laplacian matrix to encompass dissimilarities between leader and follower vertices can be found in [91].

The weight matrix is partitioned so that changing weights on some particular edges triggers a desired group dynamics. The relative strength of the connection edges, given by the magnitude of the associated weight, determines the convergence rate towards the bias. If $w_{ff} < w_{fl}$, the followers have a higher potential to reach the leader first. On the other hand if $w_{ff} > w_{fl}$, the followers are forced to converge to the formation in priority. The relative strength of the edges depends on the desired result and varies across applications, therefore we will provide the implementation details along with the description of our experiments, later on in the manuscript.

Now that the distinction between the leaders and the followers is defined, we would like to point out that, unless stated otherwise, the orientation of the follower robots achieved through

Equation 5.6 is aligned to the orientation of the leader, while the leader's heading is pointed towards the direction of motion of the formation and so, it also defines the coordinate frame of the formation. Such design choice results in a clearer indication of the team's intention as well as a better readability of the individual robot behavior from a human perspective.

The number of leaders in the formations studied in this thesis is dictated by the application requirements. Formations with one leader are used in Chapter 7 and Chapter 15, and formations with multiple (coordinated) leaders in Chapter 17.

5.2.2 Methods for Mixed Formations of Humans and Robots

Humans, even when moving in a formation with robots, do not conform to the rigid configuration specified by the formation bias. Quite the opposite, we do not assume that human motion is controllable. We consider two aspects of mixed teams of humans and robots. First, robots perform social human avoidance, where by generating a repulsive field around the human the robot is prevented from interfering with human personal space. Second, based on an assumption that humans tend to maintain free space around themselves, in our models robot proximity causes repulsive effects on human motion.

Social Human Avoidance

Each human H_h at position p_h generates a repulsive field [92] that modifies the motion of a follower robot F_i by adding a new term Γ_R to the formation consensus in Equation (5.5):

$$\dot{p}_i = \Big(|\sum_j \mathcal{L}_{ij}|\Big)^{-1} \sum_{i\sim j} -\mathcal{L}_i\big[(p_i - p_j) - b_{ij}\big] + \Gamma_R \tag{5.7}$$

$$\Gamma_R \;=\; \sum_{i\sim h} w_{ih}(\|p_h - p_i\|)\cdot\big(p_h - p_i\big) \tag{5.8}$$

where $i \sim h$ means that i has access to state information of h. The weight coefficient w_{ih} has a repulsive effect, assuring that F_i stays away from H_h:

$$w_{ih} = \begin{cases} -\frac{K_o}{(d_{ih}-\Delta_c)^2} - \frac{K_o}{(\Delta_a-\Delta_c)^2} & \text{if } \Delta_c + \epsilon_c < d_{ih} < \Delta_a \\ -\frac{K_o}{\epsilon_c^2} - \frac{K_o}{(\Delta_a-\Delta_c)^2} & \text{if } 0 \le d_{ih} < \Delta_c + \epsilon_c \\ 0 & \text{otherwise} \end{cases} \tag{5.9}$$

The weight changes continuously with the distance $d_{ih} = \|p_h - p_i\|$ between the robot and the human, and is parametrized with a constant gain K_o, range Δ_a where repulsion is a quadratic function of d_{ih}, and an imminent collision range $\Delta_c < \Delta_a$, where the weight w_{ih} grows to a large, finite number. Later on in this chapter we will explain how we leverage existing human behaviors models to determine the parameters in Equation 5.9.

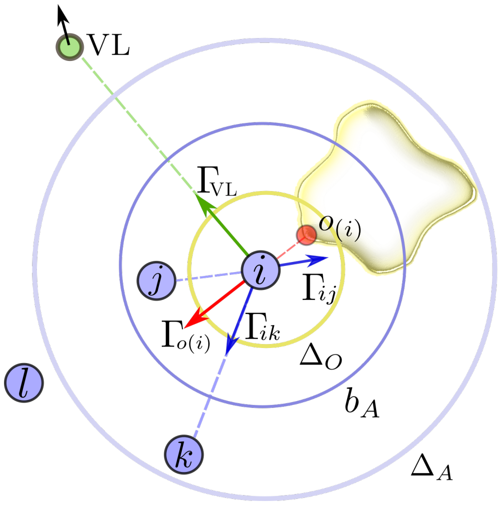

Figure 5.3 – Illustration of the flocking components. We denote the agent-to-agent motion vector of R_i to R_j as $\Gamma_{ij} \in \Gamma_\alpha$, the obstacle repulsion motion vector from point $o_{(i)}$ at obstacle o as $\Gamma_{o(i)} \in \Gamma_\beta$, and the migration motion vector towards VL as $\Gamma_{\mathrm{VL}} \in \Gamma_\gamma$. An obstacle is represented as a virtual β-agent, which is created at the surface of an obstacle at a point closest to R_i to which this β-agent is connected. Finally, b_A is the desired agent-to-agent distance and Δ_A and Δ_O are interaction ranges for agents and for obstacles, respectively.

5.3 Flocking

Flocking is a collective movement of a group of agents with a common destination. As we described in Section 3.1.1, flocking has been first simulated by Reynolds [17] with three basic rules: 1) cohesion, i.e. staying close to the neighbors, 2) separation, i.e. avoidance of the nearby agents, and 3) alignment, i.e. matching velocity with neighbors. However, as pointed out in [19], none of the rules is mathematically stated, so it is unclear when and how the rules are applied. The same study proves that Reynolds rules can be insufficient for flocking, and that under certain conditions they can lead to fragmentation. Although the basic conditions for achieving flocking according to [20] include stabilization of the distances between the agents and similarity of their velocity vectors, [19] point out that the ability to avoid obstacles, split or rejoin the flock as well as the means to migrate to a destination are equally important, but not explicitly stated in the canonical flocking rules.

5.3.1 Components of Flocking

The flocking algorithm we use in this work is following the approach proposed in [19] composed of the following three terms:

1. *Inter-agent distance balancing*, where if too close, the agents (called α-agents) are repelled; if too far, agents are attracted. This term leverages a gradient-based method and includes velocity consensus among agents.

2. *Repulsion from obstacles*, where obstacles are represented as a type of kinematic agents, β-agents, that move on the surface of the obstacles.

3. *Team migration*, namely attraction to a dynamic point in space – a *virtual leader*, represented by a γ-agent, which drives the team through the environment.

The first term embodies all the three rules of Reynolds. In this section we provide a high-level description of the flocking algorithm, while the computational details can be found in [19].

5.3.2 Flocking Algorithm

The set of α-agents $\in \mathcal{V}_\alpha$ (robots), β-agents $\in \mathcal{V}_\beta$ (obstacles) and γ-agents $\in \mathcal{V}_\gamma$ (migration point) make up a *proximity graph* in the form:

$$\mathcal{G}_{\alpha\beta\gamma} = \mathcal{G}_\alpha + \mathcal{G}_\beta + \mathcal{G}_\gamma \tag{5.10}$$

where $\mathcal{G}_\alpha = (\mathcal{V}_\alpha, \mathcal{E}_\alpha)$ is an undirected robot-robot graph, $\mathcal{G}_\beta = (\mathcal{V}_\beta, \mathcal{E}_\beta)$ is a directed bipartite graph where the nodes represent the β-agents simulated at the surface of the obstacles at the points closest to the individual robots and the edges connect the β-agents to the respective α-agents $\mathcal{E}_\beta \subset \mathcal{V}_\alpha \times \mathcal{V}_\beta$, and $\mathcal{G}_\gamma = (\mathcal{V}_\gamma, \mathcal{E}_\gamma)$ is a complete bipartite graph, where the γ-agents $\in \mathcal{V}_\gamma$ are the virtual leaders and the edge set $\mathcal{E}_\gamma$ connects each α-agent to all γ-agents (i.e. each robot has access to state information of all virtual leaders). The resulting graph $\mathcal{G}_{\alpha\beta\gamma}$ can be therefore explicitly expressed as $\mathcal{G}_{\alpha\beta\gamma} = \big(\mathcal{V}_\alpha \cup \mathcal{V}_\beta \cup \mathcal{V}_\gamma,\ \mathcal{E}_\alpha \cup \mathcal{E}_\beta \cup \mathcal{E}_\gamma\big)$. Since $\mathcal{G}_{\alpha\beta\gamma}$ is a proximity graph, edges of that graph are defined between agents situated within interaction range from each other (in case of the γ-agents the range is assumed to be infinite). For a robot R_i, the neighbor set with other robots is defined as $\Xi_{\alpha i} = \{j \in \mathcal{V}_\alpha : \|p_j - p_i\| < \Delta_A\}$ and with obstacles (β-agents) is $\Xi_{\beta i} = \{k \in \mathcal{V}_\beta : \|p_{k(i)} - p_i\| < \Delta_O\}$ where Δ_A and Δ_O are interaction ranges for agents and for obstacles, respectively, and the β-agents are represented as virtual agents, indicated with $(\cdot)_{k(i)}$, created at the surface of an obstacle k at a point closest to R_i. For illustration, see Figure 5.3.

Finally, flock migration is achieved by directing the agents towards a *virtual leader* (VL) - a dynamic point at position p_{VL} that moves with velocity $\dot{p}_{\mathrm{VL}}$. In [19] the γ-agent is *static*, meaning there is only need for a single information exchange between the α-agents and the γ-agent at initialization. In our work, the γ-agent is dynamic (and embodied by a human leader) and connected to α-agents throughout the experiment.

Each robot R_i realizes the flocking behavior using:

$$\dot{p}_i = K_c\big(\Gamma_\alpha + \Gamma_\beta + \Gamma_\gamma\big) \tag{5.11}$$

where K_c is a positive gain, Γ_α drives the neighbors towards a equally-distanced configuration, Γ_β is the obstacle avoidance term and Γ_γ is the migration term:

$$\Gamma_\alpha \;=\; K_{\alpha 1} \sum_{i\sim j} \phi_\alpha(\|p_j - p_i\|_\sigma)\mathbf{n}_{ij} \;+\; K_{\alpha 2} \sum_{i\sim j} a_{ij}(\dot{p}_j - \dot{p}_i) \tag{5.12}$$

$$\Gamma_\beta \;=\; K_{\beta 1} \sum_{i\sim k} \phi_\beta(\|p_{k(i)} - p_j\|_\sigma)\mathbf{n}_{k(i)} \;+\; K_{\beta 2} \sum_{i\sim k} a_{ik}(\dot{p}_{k(i)} - \dot{p}_i) \tag{5.13}$$

$$\Gamma_\gamma \;=\; -K_{\gamma 1}\sigma_1(p_i - p_{\mathrm{MP}}) \;-\; K_{\gamma 2}(\dot{p}_{k(i)} - \dot{p}_{\mathrm{MP}}) \tag{5.14}$$

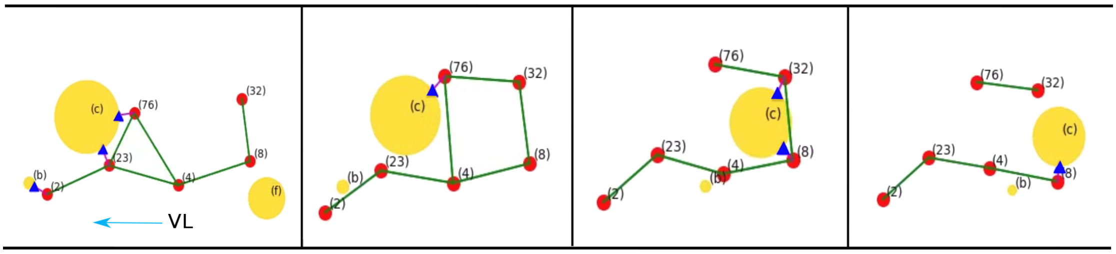

Figure 5.4 – Simple simulation of the flocking algorithm. A group of six holonomic agents (in red) moves in an environment cluttered with obstacles (yellow spheres). The VL is located at the left (from the point of view of the reader) and moves in the direction indicated by the arrow. The numbers and the letters in parenthesis correspond to the IDs of agents, and obstacles, respectively.

The first term in each equation represents cohesion, obstacle avoidance, and migration, respectively. The second term is the velocity matching with respect to the other agents, to the obstacles, and to the dynamic virtual leader, respectively.

Parameters K_α, K_β and K_γ are positive gains. Link $i \sim j$ means that R_i and R_j are within an interaction range and ϕ is a function that represents a smooth pairwise attractive or repulsive potential with a global minimum at the desired agent-agent or agent-obstacle distance of b_A and b_O respectively. To complete the notation, $\|\cdot\|_\sigma$ is a differentiable map $\mathbb{R}^m \to \mathbb{R}_{\geq 0}$, and σ_1 is a gradient defined as $\sigma_1(z) = 1/\sqrt{1+\|z\|^2}$. $\mathbf{n}_{ab}$ is a vector from p_a to p_b, and $\mathcal{A} = [a_{ij}]$ is an adjacency matrix satisfying the property $a_{ij} \neq 0 \Leftrightarrow (i,j) \in \mathcal{E}$.

The flocking motion vector are illustrated in Figure 5.3. For simplicity, we denote the agent-to-agent motion vector of Γ_i to Γ_j as $\Gamma_{ij} \in \Gamma_\alpha$, obstacle repulsion motion vector from point $o_{(i)}$ at an obstacle o as $\Gamma_{o(i)} \in \Gamma_\beta$ and migration motion vector towards VL as $\Gamma_{\mathrm{VL}} \in \Gamma_\gamma$. In Figure 5.4 we show result of applying the above algorithm to a group of six holonomic agents moving in an environment cluttered with obstacles.

5.4 Performance Evaluation

The performance of the formation is analyzed with regard to two criteria to be minimized:

$$\text{FORMATION ERROR} \qquad e_F = \Big(|\sum_j \mathcal{L}_{ij}|\Big)^{-1} \sum_{i\sim j} |(p_i - p_j) - b_{ij}| \tag{5.15}$$

$$\text{ORIENTATION ERROR} \qquad e_O = \Big(|\sum_i \mathcal{L}_{iL}|\Big)^{-1} \sum_{i\sim L} |\alpha_i - \alpha_L| \tag{5.16}$$

The formation error e_F is the average difference between the desired distances and the actual distances between the robots in the formation. The formation orientation error e_O evaluates the orientation control by verifying differences in heading alignment between the leaders and the followers.

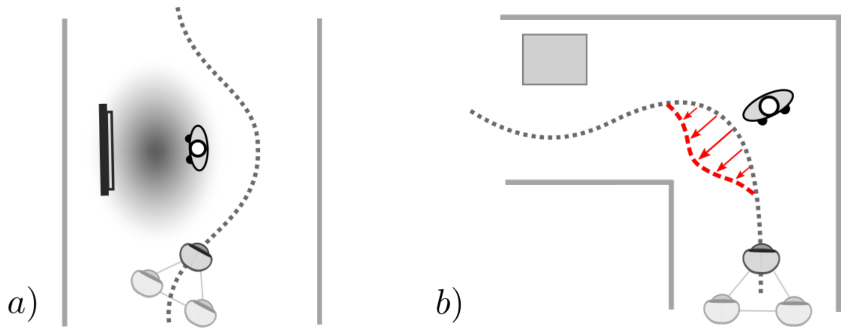

Figure 5.5 – Illustration of selected leader behaviors, a) trajectory re-planning for avoidance of social affordance spaces modeled as cost maps and b) on-line trajectory modification to move sideways when passing next to a human.

To evaluate the flocking behavior, we use the following performance metrics:

$$\text{CONNECTIVITY} \qquad e_C = \left(||\mathcal{V}_\alpha| - 1|\right)^{-1} \operatorname{rank}(\mathcal{L}) \tag{5.17}$$

$$\text{COHESION} \qquad e_R = \max_{i \in \mathcal{V}_\alpha} \|p_i - p_c\| \tag{5.18}$$

$$\text{DEVIATION} \qquad e_E = \left(|b_A^2(|\mathcal{E}_\alpha| + 1)|\right)^{-1} \sum_{i \in \mathcal{V}_\alpha} \sum_{i \sim j} \left(\|p_j - p_i\| - b_A\right) \tag{5.19}$$

where p_c is the center of positions of the robots. Connectivity $e_C \in [0,1]$ is a measure of how well the α-agents (robots) are connected in the graph; bear in mind that the edges are broken once two agents are further than Δ_A apart. The cohesion metric e_R indicates the group dispersion, and the deviation metric e_E measures how well the robots maintain the desired inter-robot distances b_A.

5.5 Encoding the Collective Movement Behaviors

The collective movement behaviors we have discussed in this chapter have a number of parameters that, when manipulated, can yield different results. A specific parameter setting that yields a unique behavior will be referred to as a *behavior modality*. In our formalism, we will explore the behavior parameter space for achieving norm-satisfying results in the case studies presented in Part III. Answering to how parameters should be manipulated is left for the later chapters, while in this section we list the relevant parameters of the formation and flocking behaviors, and those of the individual leader behavior.

5.5.1 Behavior of the Formation Leader

The task of the leader robot is to plan a path in a known environment and guide the team along this path, while making its current position known to all the agents in the team. Given a goal destination, the leader path $\mathcal{P}_L$ is determined using the Fast Marching Method (FMM) method described in Chapter 2.5.2. The leader independently follows the planned path with a desired speed, providing the overall trajectory $\mathcal{T}_L$ for the team. This leader behavior is referred to as MoveOnTrajectory.

The leader has the freedom to dynamically modulate low-level aspects of its behavior such as adjusting the velocity, moving sideways or pausing. The capabilities are represented by two parameters: 1) *TrajectoryShape* and 2) *TrajectorySpeed*.

The *TrajectorySpeed* parameter allows the leader to adjust its speed; for instance, to slow down when navigating close to a human or to wait for the other team members to catch up. Note that although we consider norms which make the leader wait for the followers, in general the leader does not take the followers into account when planning, in particular with regard to spatial constraints – the leader plans the path for itself, while the task of the followers is to be able to follow.

The *TrajectoryShape* parameter allows the leader to re-plan its path taking into account the additional constraints, or modify the original path on-the-fly. The first capability is useful when a robot has to consider social norms such as avoidance of affordance spaces. We model such spaces as cost maps and add to the occupancy map for re-planning with FMM, as illustrated in Figure 5.5a. The second capability serves for normative human avoidance, where a robot is to move sideways when passing next to a human, as illustrated in Figure 5.5b.

The resulting movement of the leader follows a trajectory $\mathcal{T}_L$ being a function of the path $\mathcal{P}_L$ provided by FMM and the speed $\dot{\mathcal{P}}_L$ with which the leader completes that path, both modified as a function of the norms, i.e.

$$\dot{p}_L \leftarrow \mathcal{T}_L = f\left(\mathcal{P}_L, \dot{\mathcal{P}}_L\right) \tag{5.20}$$

5.5.2 Behavior of Formation Followers

The MoveInFormation behavior of the formation follower has the following parameters:

- *Connectivity* - an abstraction of the edge set $\mathcal{C} = \left\{(i,j)|i \sim j, i \neq j, i, j \in \{\, \mathrm{L}, \mathrm{F}, \mathrm{H}\,\}\right\}$, where H is a human
- *FormationShape* - choice of formation topology, e.g. triangle, line, etc., encoded by the formation bias b that fully defines the geometrical configuration of the formation
- *FormationSize*- a collection of the desired inter-agent distances provided by the bias b
- *LaplacianWeights* - given in the form of the weight matrix $\mathcal{W} = \{w_{fl}, w_{ff}\}$
- *RepulsionWeights* $\mathcal{W}_{\mathcal{R}}$ - weights of the motion vectors used for human avoidance
- *ControlGain*, K_u for velocity control and K_ϕ for orientation matching

We decouple *FormationShape* and *FormationSize* to facilitate high-level formation specification. For instance, it is advantageous to be able to specify that the formation should be a *triangle* with *small* inter-robot spacing. Details on computation of the formation repulsion weights we use for social human avoidance are given in Section 5.2.2.

The resulting formation control signal in Equation 5.5 can be represented as follows:

$$\dot{p}_i = f\left(\mathcal{C},\ b,\ \mathcal{W},\ \mathcal{W}_{\mathcal{R}},\ K_u, K_\phi\right) \tag{5.21}$$

5.5.3 Behavior of Flocking Followers

The MoveByFlocking behavior of the formation follower has the following parameters:

- *Connectivity* - abstraction of the edge set $\mathcal{E}_{\alpha\beta\gamma}$, $\mathcal{C} = \{(i,j)|i \sim j, i \neq j, i, j \in \{\mathrm{L}, \mathrm{F}, \mathrm{O}, \mathrm{VL}\}\}$, where O is a set of obstacles
- *DesiredDistance* - desired agent-agent b_A and agent-obstacle b_O distances
- *InteractionRange* - interaction range for agents Δ_A and for obstacles Δ_O
- *ControlGain* - gain of the flocking behavior K_c
- *FlockingTermGain* - gain of the flocking terms K_α, K_β and K_γ

The resulting formation control signal in Equation 5.11 can be represented as follows:

$$\dot{p}_i = f\left(\mathcal{C},\ b_A,\ b_O,\ \Delta_A,\ \Delta_O,\ K_c,\ K_\alpha,\ K_\beta,\ K_\gamma\right) \tag{5.22}$$

5.6 Modeling Human Behaviors

Human modeling is a key mechanism for interpreting human behavior, and adapting the robot's behavior to accommodate users with varying skills, experience, and knowledge [33]. In this thesis we will rely on two models of human behavior, the Proxemics Model (PM) and the Social Forces Model (SFM), which will be integrated with the formation control to result in human-aware robot motion. In particular, we use a PM to express human comfort zones, as it provides simple and intuitive representation of personal space preferences. However, a distance value alone is insufficient to determine the parameters K_o, Δ_a and Δ_c of the robot repulsive motion vector necessary for computing its weight in Equation 5.9. For that purpose we introduce a SFM.

5.6.1 Proxemics

The original PM we use for representation of human comfort spaces characterizes only one aspect of human comfort in only one situation – a distance to establish for explicit interaction when the human stands still [31].

To represent human comfort spaces we use the PM proposed in [63], which classifies the

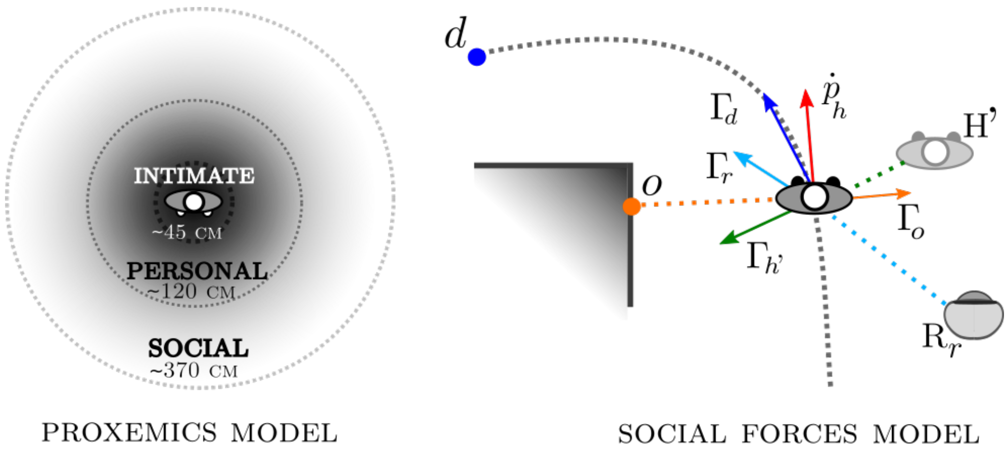

Figure 5.6 – Two basic models of human behavior. PM defines the preferable areas that people maintain between themselves and the others. The SFM represents human motion $\dot{p}_h$ as a result of interaction between forces that drive the person to a destination (Γ_d), and forces that keep away from other people ($\Gamma_{h'}$) and obstacles (Γ_o), including robots (Γ_r).

virtual area around the human into intimate zone (< 45 cm), personal space (< 120 cm) and social space (< 3.7 m) (see Figure 5.6). The radii of human spaces serve to determine the degree of compliance the robot must apply to ensure human comfort. We label the extent of the comfort zones for human H_h with $\Delta_{S,h}$ for social space, $\Delta_{P,h}$ for personal space, and $\Delta_{I,h}$ for intimate space.

Although proxemics has also been investigated in dynamic situations [31], we have chosen to use a well known human model – the SFM [18] to understand the effects that the robot repulsion forces have on human behavior and determine the parameters of such forces.

5.6.2 Social Forces Model

The SFM represents humans as masses subject to effects of virtual gravitational forces. Each individual is drawn towards a destination (or a leading person), while simultaneously being repelled by obstacles and from other agents (see Figure 5.6 for illustration). The SFM has been popularly used to model crowd dynamics [68], and represent human behavior in agent simulation [66], while in robotics it is used to model the effects of robot's presence on human motion [76], by explicitly modeling the robot-repulsion term [66] [67]. It is also used for determining the robot's behavior required for enforcing a desirable human action [93]. We adapt the model from [18], with an additional term for robot effects.

The movement of the human H_h is regulated by the forces exerted by the environment:

$$\dot{p}_h = \Gamma_d + \sum_O \Gamma_o + \sum_H \Gamma_{h'} + \sum_R \vec{F}_r + n \tag{5.23}$$

where $n \sim \mathcal{N}(0, \sigma_f)$ is motion fluctuation, Γ_d is an attractive force driving human towards destination d with the desired velocity $\dot{p}_d$, relaxation time τ_f and strength K_{fd}:

$$\Gamma_d = K_{fd}(\dot{p}_d - \dot{p}_h)\tau_f^{-1} \tag{5.24}$$

Force Γ_o prevents the human from colliding with obstacles $o \in O$, while the repulsive effects from the other humans H' ∈ H and from the robots $\mathrm{R}_r \in \mathrm{R}$ are modeled using $\Gamma_{h'}$ and Γ_r forces respectively:

$$\Gamma_o = K_{fo} e^{-\|p_v - \mathbf{o}\|/\Delta_o} \tag{5.25}$$

$$\Gamma_{h'} = K_{fh} e^{-\gamma_f/\Delta_h} \tag{5.26}$$

$$\Gamma_r = K_{fr} e^{-\|p_h - p_r\|/\Delta_r} \tag{5.27}$$

where K_{fo}, K_{fh} and K_{fr} are constant gains and Δ_o, γ_f, Δ_h and Δ_r specify the areas of influence of the forces. Each force is additionally discounted to account for the perceptive effects. For details on the model and calculating γ_f, please refer to [18]. The resulting potential repulsive resulting from this model is a monotonically decreasing function with equipotential lines in the form of an ellipse oriented towards the direction of motion.

Unfortunately, there is no physical equivalence between the personal distance of proxemics and the parameters of the SFM, so we chose to adapt the parameters based on observations of real pedestrian behavior presented in original work [18], namely $\sigma_f = 0.001$, $\tau_f = 0.5$, $|\dot{p}_d| \sim \mathcal{N}(1.34, 0.26)$, $K_{fd} = 2.0$, $K_{fo} = 2.0$, $K_{fh} = 2.1$, $K_{fr} = 6.5$ and $\Delta_o = 0.1$ m, $\Delta_h = 0.3$ m and $\Delta_r = 0.38$ m. We use the SFM to model human motion in the Webots simulator, from which we determine the parameters K_o, Δ_a and Δ_c of the robot repulsive force empirically in [9]. The resulting parameters are used in the case studies in Part III.

Summary

This chapter has provided the background on collective movement approaches central for this thesis – formation and flocking. While the graph-based representation enables the conceptualization of individual agents and their interactions necessary for achieving coordination, we have introduced another level of abstraction – that of *behavior modality* representing the behavior parametrization, a concept that will be used in our formalism for denoting specifications for norm-satisfying behaviors. Before moving to the formal framework however, in the next chapter we analyze our preliminary experiments at the IPOL hospital to understand the implications of socially-ignorant robot behaviors. Then, we focus on practical challenges of deploying multi-robot systems in real environments.

6 Multi-Robot Deployment in a Hospital

THE multi-robot experiments at the IPOL hospital are the leading motivation for this thesis and reveal the major challenges we attempt to address - the need for robust multi-robot navigation in complex environments, augmented cooperative perception and normativeness of multi-robot behaviors. In this chapter, we closely analyze the experimental results and aim to understand the impact of non-social robots on a place with strict ethical regulations - a hospital environment.

6.1 Experimental Setting

Our experiments have been carried out at a pediatric ward of the IPOL hospital in Lisbon on the 14th Oct 2015. Two MBot robots moving in a leader-follower formation navigated through a child playroom, then through a busy corridor and back. Robots displayed random facial and bodily expressions (such as changes of light, displaying pictures on a touch screen, random head or arms movements), but did not engage in interactions. The formation was controlled by a canonical graph-based law with active obstacle avoidance but no social layer. Over the course of a day we performed 23 experimental runs (i.e. back and forth traversing of the ward) with an overall duration of 74 min.

All experiments were recorded using five cameras: four overhead omnidirectional cameras, two mounted in the playroom and two in the corridor, and one camera providing a side-view of the corridor. We have annotated the videos in post-processing using subjective factors that quantify the influence of the robot presence on the end-users with respect to the following categories:

Motion. Interactions are analyzed with regard to relative motion of humans and robots. We distinguish three subcategories: a) when robots interrupt static humans, b) when robots disturb moving humans and c) when humans join the formation.

Impact. How the presence of the robot team affects the behavior of humans. The sub-

Figure 6.1 – Examples of interactions between the team of robots and humans at the IPOL hospital.

categories include: a) situations when humans take actions to ensure safety (such as removing obstacle from the path of the robots), b) when humans visibly change their intentions, c) when humans engage in a play with the robot, and d) when humans engage in a passive, long-term observation of the robots.

Staff. The last category describes the impact of robot's presence on the efficiency of hospital staff. This includes all the interruptions caused by the robots, judged either as significant or acceptable (when only the motion is disturbed, but no actions are required from staff to continue their work).

6.2 Results

During the experiments, robots navigate autonomously, but were observed from a nearby room for safety reasons. An operator intervention was required four times, 1) when two humans with an Intravenous Device (IV) pole were passing through a door and one of the robots blocked the base of the pole, 2) when a robot navigated towards a child with an IV pole, 3) when the robots were blocked by a wheelchair, and 4) when the robots could not reach a goal point due to a navigation glitch. Note that the base of the IV poles, having a maximal height inferior to that of the Mbot's laser scanning plane, were practically invisible to the robot at the time the experiments were run (a camera-based solution was introduced later in the project).

We have labeled a total of 174 interactions, 19% of which are of long-duration (> 20 s). A summary of the interaction activation rates per second ($\times 10^{-3}$) is presented in Table 6.1.

With regard to the relative motion of humans and robots, the highest number of interactions occurred when robots and humans were passing by each other - a total of 60 interactions of this type were recorded. Slightly less, but equally important are the interactions with static humans (a total of 33). Examples of such interactions are shown in Figure 6.1. Interestingly, there was a relatively high number of instances when a human or multiple humans join a formation, with a total of 21. Most of the cases occurred in the morning – at the time when visitors with children were present. Both children and adult visitors joined the formation, either for amusement and as a part of a play, or to extend the interaction time with the robots. Figure 6.1 shows examples of one human moving side-by-side with robots, and a child and an adult following the robots.

In the category of robot's impact on human behavior it can be clearly seen that the robots noticeably affect the IPOL environment. It can be seen in the videos that the robots were the main attraction for the children and the visitors, who played with the robots 45 times in total over the course of the day (see Figure 6.1). Many people (26 occurrences) engaged in long duration observation of the robots (annotations are based on human gaze and gaze direction that follows the robots). Robots were also found to affect the environment practically – in 57 occurrences humans visibly changed their initial action in order to accommodate the robots, either by giving them a way or actively avoiding. A high number of humans (37 occurrences) took visible actions to assure safety by removing obstacles (e.g., IV poles, cleaning equipment) from the path of the robots or by moving away from the robot in an attempt to prevent potential collisions (see Figure 6.1). Work of hospital staff was only moderately affected by the robots – a total of 15 times the movements of the staff were mildly disturbed, but we noted 9 occurrences when the staff needed to take an action in order to continue work (e.g., move a cart with food away, overtake robots while pushing a bedware cart).

6.3 Discussion

It is difficult to estimate the effects of a multi-robot system on the humans in a hospital based on a short-term exposure. It has been a great challenge of the MOnarCH project to analyze such effects even when the amount of data collected over years with a single robot was large. However, the relatively small set of experiments described in this section can already give some insights into the potential advantages and drawbacks of a long-term deployment of multi-robot systems in sensitive and cluttered environments such as that of a pediatric ward. For example, one can expect continuous excitement from children and newcomers: the appearance of the robots and their coordinated movements are definitively entertaining. However, it should be expected that the patience and leniency of staff and long-term occupants might eventually wear off. The high number of humans taking action to ensure safety indicates that even though humans enjoy the robot presence in general, they are not fully comfortable with it. One can imagine that if the cleaning staff with heavy equipment had to avoid robots on a daily basis, or doctors had to watch out every time they traversed corridor, the robots would have quickly become an annoyance. The primary objective when designing robots

MOTION	Static human interrupted	4.05
	Static group interrupted	3.38
	Moving human interrupted	10.8
	Moving group interrupted	2.7
	Formation with one human	2.7
	Formation with multiple humans	2.03
IMPACT	Human taking action to assure safety	8.33
	Human changing intentions	12.83
	Human playing with the robots	10.13
	Human observing robots	5.85
STAFF	Significant interruption of staff work	2.03
	Acceptable interruption of staff work	3.38

Table 6.1 – Activation rates per second ($\times 10^{-3}$) of human-robot interactions observed during deployment of two MBot robots at the IPOL hospital.

for human environments should be therefore to minimize disturbances for any operation typically carried out in that environment.

Summary

The analysis of results gathered during multi-robot experiments carried out at the IPOL hospital suggests that the robots are a great addition to highly sensitive environments, providing entertainment, and encouraging physical interactions. At the same time, it is apparent that non-social robots are of high disturbance to the regular hospital activities and in the long-term robots might become a nuisance. While solving the technical problems related to navigation and obstacle avoidance can bring improvements, it is obvious that robots need to comply to human social norms in order to be better accepted in our societies. However, if multi-robot systems were to operate autonomously in social environments, they must be aware of the mechanisms of human society. Such key mechanisms by which humans operate are social norms – they regulate all human interactions, behaviors, and provide order and predictability in the society. If the robots were to follow simple norms, such as "wait at the door until the human passed" or "move out of way of hospital staff", their acceptance could be tremendously improved.

7 Adaptive Formations

As we have pointed out in Chapter 3, the state-of-the-art approaches to multi-robot navigation rely on fairly restrictive assumptions regarding the experimental settings, namely that (a) the experimental area is large, i.e. the formation has enough space to reorganize after negotiating obstacles, or (b) the obstacles are cluttered in the environment and are not surrounding it (i.e., no enclosure). Therefore, such canonical experimental settings reflect outdoor, unstructured environments without boundaries.

In view of these limiting assumptions, we introduce the Local Formation Transformation (LFT) algorithm – an approach to dynamic formation change that is *local* (meaning that the formation is only reshaped in the immediate neighborhood of the robot that initializes the change) and *gradual* (meaning that the formation does not switch topologies but is modulated in its shape to some extent), with the level of shape alteration proportional to the density of obstacles around the robot. Although locally adaptive formations are receiving attention in the robotics community [94] [27] [43], to the best of our knowledge, LFT is the first formation change method validated on real robots that allows any robot to alter locally and independently the formation shape without imposing a global reconfiguration.

In this chapter, we explain how the LFT method extends the graph-based formation control approach and validate the algorithm through formations consisting of three real MBot robots in structured indoor areas of increasing complexity.

7.1 Overview

The LFT method allows the follower robots to autonomously modify the formation in reaction to the environment as a function of the available area, topology of that area and density of obstacles. Intuitively, the LFT algorithm works as follows. We define a variable η_i that reflects a local density of obstacles with respect to the follower F_i and drives its gradual and local formation change. Two extreme circumstances can be distinguished: a) when there are no obstacles around the follower, $\eta_i = 0$ and the robot should remain in the defined place

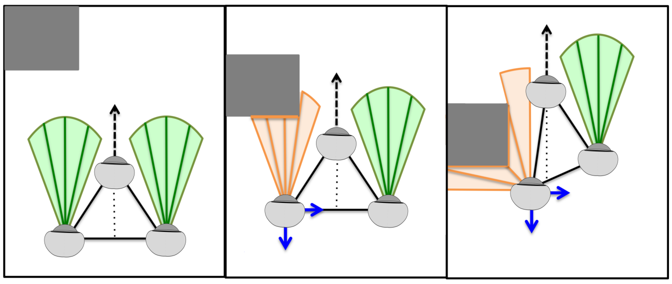

Figure 7.1 – Illustration of the Local Formation Transformation algorithm. (Left) The default formation. (Middle) The follower on the left detects an obstacle and starts gradual alignment behind the leader. (Right) The follower on the left remains partially aligned behind the leader until the obstacle is cleared.

in the formation; b) in case of maximal density of obstacles, $\eta_i = 1$, the safest path for the follower is to be aligned within a parametrized tolerance margin behind the leader. Under such circumstances we can assume its path to be collision free. In particular, this is the case for sufficiently static environments[1] when the leader's heading is constrained to be tangential to its trajectory by a motion planning module.

The variable η_i varies continuously between 0 and 1, and thus the follower gradually wavers between its desired place in the formation and a safe place behind the leader. The estimated obstacle density $\eta_i \in [0,1]$ serves as a sole indicator of how the follower should modify the original formation matrix to navigate around the obstacles. It is computed based on a constrained virtual sensor FOV characterized by a sensing range and a sensing angle. The FOV covers an area to be traversed in the near future, with an exception of when $\eta_i \to 1$ and the robot queries the area around the desired place in the formation to return there when it is safe to do so. The LFT method is illustrated in Figure 7.1.

7.2 Algorithm

Because of the continuous adaptation of the follower F_i to maintain its desired relative position with respect to the leader, the resulting formation change can be formalized as a transformation function $\psi_i : [0,1] \to \mathbb{R}^{N_i \times D}$ in a D-dimensional space as follows:

$$S_{LFT} = \left\langle \psi_i(0),\ \psi_i(1),\ \psi_i : \eta_i \to \mathbb{R}^{N_i \times D} \right\rangle \tag{7.1}$$

where $N_i = |\Xi_i|$ is the number of neighbors of robot F_i. That is, given two geometries $\psi_i(0)$ and $\psi_i(1)$, we design a smooth, continuous function of class C^1 that maps the density of obstacles perceived by the robot F_i in the environment to the modification of the formation geometry $\psi_i \in [0,1]$. Globally, the geometry $\psi(0)$ specifies the original shape of the formation and is assumed to be determined by a higher level controller or external user. The purpose of $\psi(1)$ is to define a virtual column formation shape and specify an allocation of the robots

1 With the exception of when a dynamic obstacle cuts the course of the follower; in such situation collision is prevented by the means of the DWA obstacle avoidance described in Section 2.5.2.

within such structure. Using $\psi(1)$ guarantees that each follower has a dedicated place behind the leader in the extreme case when $\forall \mathrm{F}_i \in \mathrm{F},\ \eta_i \approx 1$ and all followers must fall behind the leader. An illustrative example of the transformation function ψ is shown in Figure 7.2.

Remark 1 (Generality). The specification of the LFT algorithm requires that ψ is able to transform the shape of the formation from $\psi(0)$ to $\psi(1)$. Consider the specification of the formation in Equation 7.1. If the transformation function is designed with the constraint:

$$\forall \mathrm{F}_i, \mathrm{F}_j \in \mathrm{F},\ \| \psi_i(\eta_i) - \psi_j(\eta_j) \| \geq \epsilon \tag{7.2}$$

where ϵ is the minimal allowed distance between any two robots (e.g., the robot diameter plus a safety margin), then collisions among formation members will not occur. The algorithm is generalizable to N robots and different formation topologies, because it is possible to design $\psi(1)$ and $\psi : [0,1] \to \mathbb{R}^{N\times D}$ so that the transformation between two formation shapes is collision-free. ∎

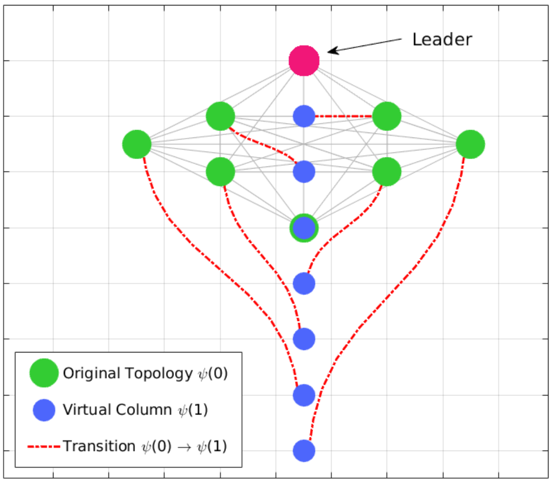

Figure 7.2 – Illustration of the transition function ψ. The desired formation of the robots is $\psi(0)$ (green), upon detection of obstacles they transition to a column topology $\psi(1)$ (blue).

7.3 Implementation

While the abstract definition of LFT in Section 7.2 does not restrict the algorithm to a specific formation control law, here we present its realization within the graph-based framework. In particular, the transition function $\psi_i(\eta) : \eta_i \to \mathbb{R}^{N_i \times 2}$ corresponds to the bias matrix $b_i \in \mathbb{R}^{N_i \times 2}$ as follows:

$$\tilde{b}_{iL} = \psi_i(\eta_i) : [0,1] \to \mathbb{R}^{N_i \times 2} \tag{7.3}$$

The bias to the leader is changed on two axes simultaneously: by modifying b^x_{iL}, the follower F_i aligns behind the leader L, while by modifying b^y_{iL}, the follower maintains a closer or further distance from the leader. The bias b_{ij} to all local neighbors of F_i is changed accordingly. The functions for changing the x and the y components of b_{iL} are chosen as follows:

$$\begin{aligned} \psi_i^x(\eta_i) &= b_{iL}^x\big(1-(1+e^{-K_1(\eta_i-K_2)})^{-1}\big) \\ \psi_i^y(\eta_i) &= \operatorname{sign}(b_{iL}^y)\,\big((\psi_i(0)-\operatorname{abs}(b_{iL}^y)\big)\eta_i+b_{iL}^y \end{aligned} \tag{7.4}$$

where K_1 and K_2 are constant parameters normalizing and attenuating the shape of a function that filters η_i. Thus, the value of $\psi_i(\eta_i)$ updates the bias as new sensory information arrives.

Remark 2 (Convergence). The convergence rate of a system based on the Laplacian consensus feedback can be calculated depending on the eigenvalues of the Laplacian matrix $\mathcal{L}$ [89]:

$$\parallel p(t)-\mathbf{1}\parallel\leq\parallel p(0)-\mathbf{1}\parallel e^{-\lambda_2 t} \tag{7.5}$$

where $\mathbf{1}=[1,...,1]^T$ and λ_2 is the lowest non-zero eigenvalue of $\mathcal{L}$ that determines the speed of convergence. Under the assumption of the time-scale separation principle [95] and the constraint that the function ψ is chosen so that $\dot{\psi}\leq\tau e^{-\lambda_2}$, where $\tau\geq 1$, the system described in Equations 5.5 and 5.6 is assumed to achieve consensus. ■

The functions in Equation 7.4 satisfy the constraints discussed in Remark 1 and Remark 2. They have been tested in simulation with up to 10 robots, but their generalization to a larger number of robots is still to be proved.

It is important to note that it is not sufficient that the follower F_i changes its bias, but the modification must be propagated through the graph, otherwise the neighbors of F_i would accommodate the change by further distorting the formation. To address this issue, the follower F_i that initializes the local formation change communicates this fact and the degree of modification ψ_i to all its neighbors, which adjust the bias to F_i accordingly. In our case, this requires the communication of two variables, $\tilde{b}_{ij}^x$ and $\tilde{b}_{ij}^y$. This step could be omitted in sparse, minimum degree graphs were each follower is only connected to its local leader (for an example of such a formation refer to [25]). A solution not requiring information propagation through communication and yet retaining a connected graph could project a virtual robot at the default bias once a robot detects that the formation is being modified by one of its neighbors.

LFT complements obstacle avoidance in the sense that it reacts at farther distances through its multi-robot perception capabilities and has an implicit predictive component that estimates a collision course by monitoring the area to be visited shortly in time. While obstacle avoidance prevents collisions in a reactive manner, LFT allows navigation around structured obstacles such as corners or through narrow passages. It implicitly promotes a column formation, as superior for the teams of robots going through tightly confined areas (the superiority of such formations is shown in [37]).

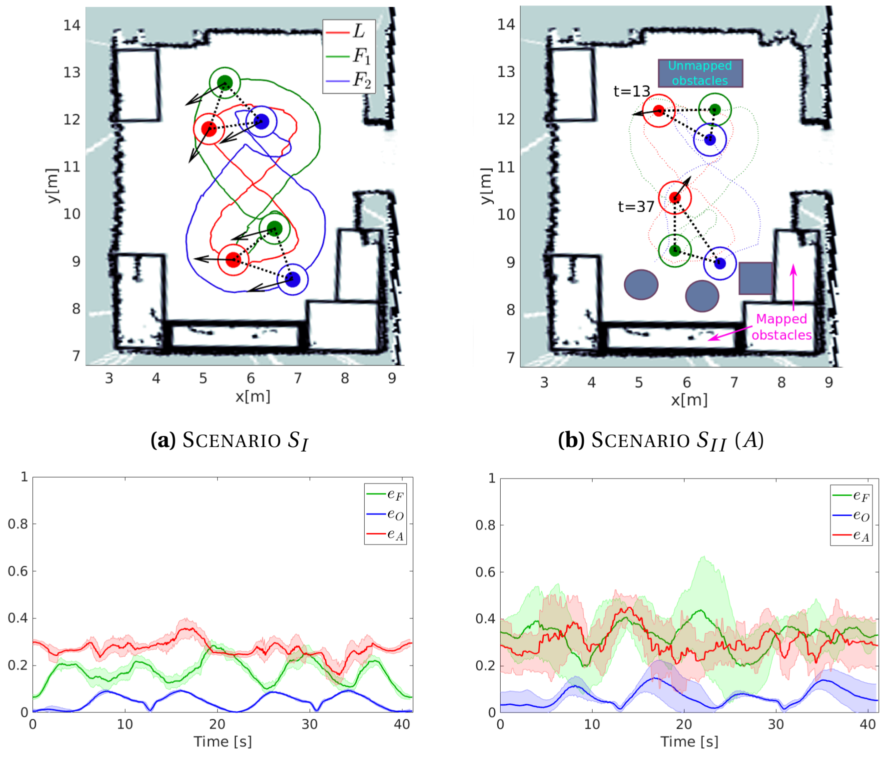

(a) SCENARIO S_I **(b)** SCENARIO S_{II} (A)

(c) SCENARIO S_I **(d)** SCENARIO S_{II} (A)

Figure 7.3 – Robot trajectories and performance evaluation in scenarios S_I (a) and $S_{II}(A)$ (b). Rectangles and circles in (a) and (b) represent static obstacles, including mapped and unmapped obstacles labeled in (b). Size of the outer robot circles indicates the true size of the robot. Plots in (c) and (d) show the formation error [m], orientation error [rad] and LFT reactivity (unitless). The shaded colors in (c) and (d) indicate the standard deviation averaged over 10 runs and the two followers. This convention remains valid for the performance evaluation plots throughout the chapter.

7.4 Experiments

Experiments are performed in the GR building of the EPFL campus described in Section 2.6.2, comprised of two rooms connected by a corridor. The objective is to assess the impact of environmental settings on the adaptive formation algorithm. Additionally to the metrics proposed in Section 5.4, we choose $e_A \in [0,1]$ as an indirect measure of the LFT reactivity in avoiding obstacles:

$$\text{LFT REACTIVITY} \qquad e_A = (N-1)^{-1} \sum_{i=1}^{N-1} \eta_i \tag{7.6}$$

where only the followers are taken into account as the leader is not contributing to the formation adaptation. The rationale behind analysis with e_A is that as soon as η starts to increase, the robot modifies locally its bias to lessen the chance of driving close to obstacles, leading

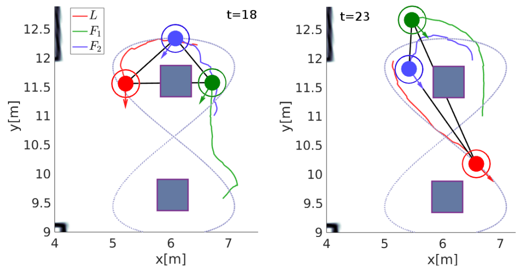

(a) SCENARIO S_{II} (B)

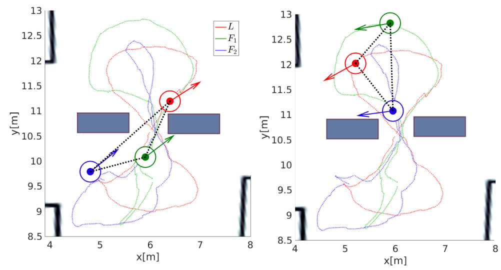

(b) SCENARIO S_{II} (C)

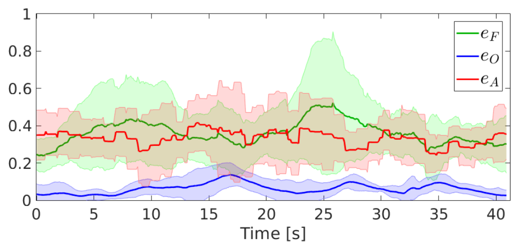

(c) SCENARIO S_{II} (B)

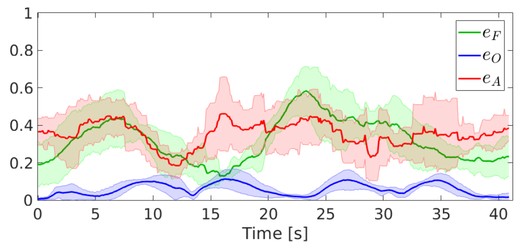

(d) SCENARIO S_{II} (C)

Figure 7.4 – Robot trajectories and performance evaluation in scenarios $S_{II}(B)$ (a) and (c) and $S_{II}(C)$ (b) and (d) respectively.

to a reduction of η. The value of e_A close to one means that all the robots undergo maximal degree of formation modification.

The parameters used in the experiments are: $K_1 = 10$, $K_2 = 1$ (in Equation 7.4) and $K_u = 0.8$ and $K_\phi = 0.8$. Weights have been tuned according to the distance between the follower R_i and the leader L using functions proposed in [45]. If the distance is large, the weight of the follower-follower edge w_{ff} is smaller than on the leader-follower edge w_{fl} and the followers have a higher potential to reach the leader first. If the opposite is true, the followers are forced to converge to the formation first.

7.5 Scenarios

We investigate the performance of the LFT algorithm in four scenarios. For each experiment we perform 10 runs with initial positions of the robots at the same locations. We consider a triangular formation with two followers F_1 and F_2 and one leader L forming an equilateral triangle with 1 m sides.

S_I – Trajectory Twist Scenario

We test the ability of the formation to move on a trajectory with sharp turns. This scenario serves as a baseline for performance comparison. The leader moves at a speed 0.5 m/s in an eight-figure trajectory shown in Figure 7.3a (the same trajectory is retained for scenarios S_I to

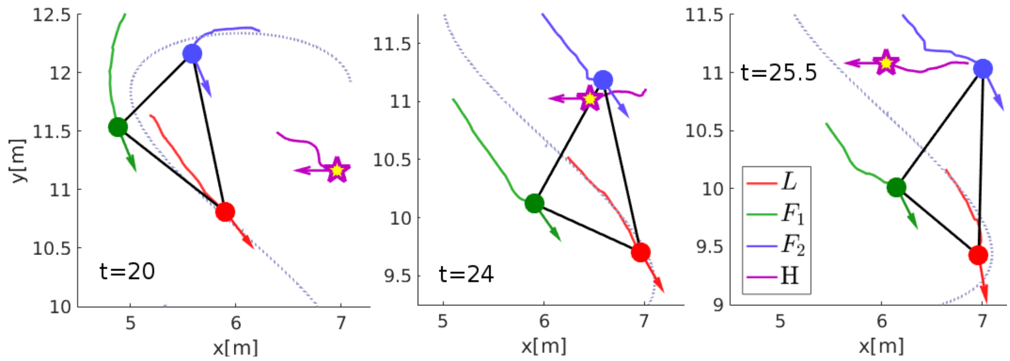

Figure 7.5 – Snapshots of the experiment S_{III} with the formation being interrupted by a person walking around the arena.

S_{III}).

S_{II} – Static Obstacles Scenarios

The second scenario evaluates the LFT algorithm in three static obstacle settings, where:

$S_{II}(A)$ – obstacles are scattered at the outer edges of the leader's trajectory (Figure 7.3b), at positions that require formation alteration by the followers as they cope with sharp turns.

$S_{II}(B)$ – obstacle setting (Figure 7.4a) requires the followers to deviate from their original formation as they pass over the inner circle of the eight-figure trajectory.

$S_{II}(C)$ – formation is to pass through a narrow passage of 1.5 times the diameter of the robot (Figure 7.4b).

S_{III} – Dynamic Obstacles Scenario

In S_{III} the formation encounters a dynamic obstacle - a human volunteer H moving on a collision course with the robots and tracked by an overhead camera. For repeatability of the experiment, the human follows a path delineated on the floor and timed with a stopwatch.

S_{IV} – Complex Environment Scenario

The final demonstration of the algorithm is in a realistic indoor environment. We consider a waypoint patrolling task designed so that it is necessary for the robots to change the formation during the run as well as rotate in a highly constrained space. The storyline is the following: the leader is to go out of the 1st and the 2nd door, pass through a narrow passage and then return to the initial position (see Fig. 7.6).

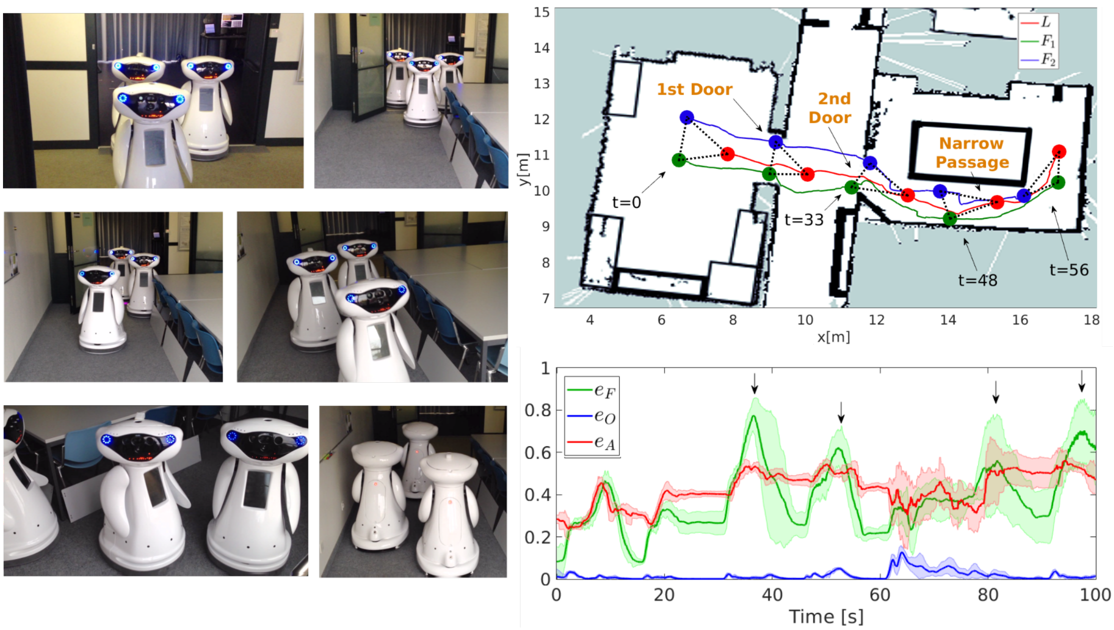

Figure 7.6 – Trajectories of the robots and pictures taken during one experimental run, and performance evaluation in Scenario S_{IV}.

7.6 Results

The performance metrics discussed in this section are averaged over the runs and over the two followers. The formation error e_F [m] indicates the average distances the followers are from the desired places in the formation, the orientation error e_O [rad] indicates the average difference between the follower's desired and actual orientation, and the LFT reactivity e_A [unitless], indicates the average degree of formation modification estimated locally by each follower. Additionally, for each scenario we provide trajectories of the robots carried out during one of the runs. Videos of the experiments are available at the link provided in the footnote[2].

S_I – Trajectory Twist Scenario

The trajectories presented in Figure 7.3a show that the robots converge to the desired formation and maintain it. The performance metrics presented in Figure 7.3c lead to the conclusion that the intricacy of the leader's trajectory has no significant impact of on the formation errors. Because of the size of the MBot robot, small rises of e_F during sharp turns are not visible in reality (the average formation error e_F = 0.27 m is less than half the robot radius). A small variation of performance over the runs suggests a strong repeatability of the results.

2 https://www.epfl.ch/labs/disal/research/InstitutionalRoboticsFormations

S_{II} **– Static Obstacles Scenarios**

$S_{II}(A)$ – As shown in Figure 7.3b, the followers visibly locally modify the formation during negotiation of the obstacles. The performance presented in Figure 7.3d only marginally decreases compared to Figure 7.3c, but has a larger variance.

$S_{II}(B)$ – Robots cope with the obstacles either by going around them or taking them in-between, as shown in Figure 7.4a. Scenario $S_{II}(B)$ is more challenging than $S_{II}(A)$ because of the obstacle being combined with a sharp turn. This is reflected by a decrease of performance, especially of the e_F component, which attains a peaky performance drops, as illustrated in Figure 7.4c. Additionally, the existence of two possible solutions the robots can adopt drives the large variance of e_F.

$S_{II}(C)$ – The followers detect the wall-like obstacles on their path and change to a column formation using the LFT algorithm as shown in Figure 7.4b. The two peaks of e_F in Figure 7.4d correspond to the two times the robots navigate through the narrow passage, but even then the mean e_F error is smaller than the robot diameter.

S_{III} **– Dynamic Obstacles Scenario**

S_{III} – The snapshots presented in Figure 7.5 show the situation before, during, and after the formation encounters a dynamic obstacle - a human. Notice that collisions with swiftly moving obstacles are primarily handled by the virtue of DWA, which is more reactive than LFT. Even though the DWA does not guarantee finding an optimal velocity, it prevents collisions by stopping the robot before hitting the obstacle. The performance is little affected as compared to S_I because of the instant nature of obstacle negotiation. The mean error values are $e_F = 0.25$ m, $e_O = 0.05$ rad and $e_A = 0.31$ with the maximum value of e_F never exceeding 0.45 m and the maximum value of e_A never exceeding 0.55.

S_{IV} **– Complex Environment Scenario**

S_{IV} – Figure 7.6 shows the trajectories of the robots and pictures taken during one experimental run. The team starts in a triangular formation which by the means of the LFT becomes more elongated and narrower as the robots navigate through the environment (in particular at the location indicated as a *Narrow Passage* at t=48 s). At the furthest waypoint (t=56 s), the formation rotates in a confined area with the diameter of the free space smaller than the diameter of the desired formation. The formation error e_F shown in Figure 7.6 attains its peaks at t=[38 s, 56 s, 81 s, 98 s], corresponding to the robots passing the 2^{nd} *Door*, *Narrow Passage*, the 2^{nd} and the 1^{st} *Door*. Note that without the LFT method, the followers moving in a triangle formation are not capable of negotiating an abrupt transition characterized by passage with a width smaller than that of the formation (such as door leading out of a room). Indeed, with a canonical implementation of a graph-based control law, the followers would head towards the wall surrounding the narrow passage and become trapped in the local minima. Only for smaller, convex obstacles, DWA allows the followers to safely navigate around the obstacles. The results indicate a high correlation between the density of the obstacles (reflected by e_A) and the formation shape error e_F. The orientation error remains largely unaffected by the experimental settings, mostly due to the fact that the heading control of the followers is decoupled from the position control.

Summary

In this chapter we presented an approach for realizing adaptive multi-robot formations for structured environments that yields local and gradual changes of the formation shape with the level of alteration proportional to the density of obstacles ahead. We concluded that the LFT enables the formation to navigate as a unit through demanding environments, such as narrow passages with abrupt entry points and the width smaller than the width of the formation. Reactive formation reconfiguration resulted in the ability of the team to cope with complex building features such as doors or other confined spaces in an environment with numerous uncertainties arising from the presence of static and dynamic obstacles as well as sensor and actuator errors. Motivated by human-populated environments, the LFT algorithm achieves desirable properties, including smoothness of motion and aesthetic negotiation of obstacles. While the LFT method is designed for enabling the passage of formations of large-sized robots through indoor structured spaces (with doors and corridors), the ability of adjusting the formation shape plays also a crucial role for the realization of social norms. In later chapters, we will illustrate how by modifying the formation bias we realize mixed human-robot formations, and how, through an augmentation of the graph with repulsive edges, we can achieve a formation reshaping to accommodate human-induced social space constraints.

8 Indoor Relative Localization for Multi-Robot Formations

THE settings in which we deploy our multi-robot systems are by far not standard for the majority of the spatial coordination methods, namely GNSS-denied, complex indoor environments populated with humans and obstacles. In this chapter we present a strategy for providing reliable state estimates that allow a group of robots to spatially coordinate even when communication fails and a relative localization system based on direct inter-robot measurements suffers from long-term occlusions and false detections. While the method proposed in this chapter can be generalized to other types of collective movements, our development has been targeting robot formations. We will therefore focus the remainder of the chapter on this specific way to spatially organize a group of robots.

We work with ID-based robot formations, where each robot is assigned a *role* (target position) in the formation, but the tracking data does not provide the identity of the robot. Roles are important to determine the range and the bearing that the robot has to keep with respect to the other robots in the formation, but since the robots are homogeneous, they can assume any role. However, the lack of ID information associated with the measurements does not allow for simple fusion of communication and tracking, especially when tracking is expected to sustain the formation dependably for some time, even when communication fails. Multi-target tracking methods without explicit data association such as PHD filters, require on-line role assignment, where the robot dynamically estimates the optimal matching of the estimates with the roles (including its own role that might change over time). When the assignment is not shared, the robots might diverge to different guesses of the assignments, making the formation ill-defined, which can lead to breaking of the formation. Wrong data association can be caused by as little as one track missing, imperfect detection, or long term occlusions, which are common occurrence in formations. These challenges make the tracking-only methods unsuitable for maintaining the entire formation for long periods of time.

To overcome such difficulties, we have developed a method that incorporates communication data (when available), tracking information, and knowledge about the desired formation geometry in GM-PHD filter [60]. The use of such filter allows us to combine data from multiple

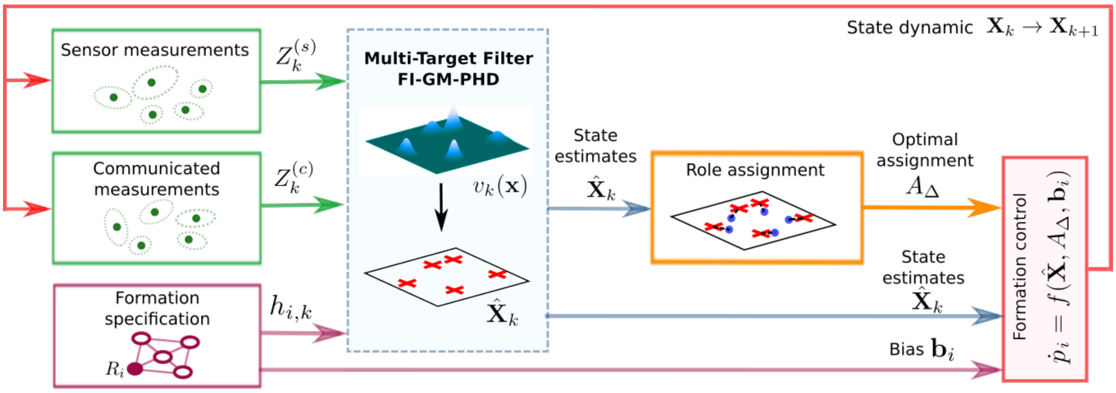

Figure 8.1 – Overview of the approach. We combine communication data, tracking information, and knowledge about the desired formation geometry in the FI-GM-PHD filter. The estimates are then assigned to the roles (target positions) in the formations to determine the desired bias to keep from the estimated robot location.

information sources without the need to use heuristic methods for data association. Moreover, a GM-PHD filter does not set the number of tracks a priori, therefore additional data regarding a target can be incorporated seamlessly [96]. To the best of our knowledge, our method called Formation Information GM-PHD (FI-GM-PHD) filter is the first approach attempting to improve tracking estimates of the robot poses based on specification of the desired formation geometry. Our method consists of two main components added to the GM-PHD filter: i) the *inception step*, which incorporates poses of the robots exchanged via communication, when such information is available, ii) the *coalition step*, which integrates the projection of the formation state based on the desired formation geometry. The projected formation state is either improving the current estimate or generating a new one, depending on the dissimilarity between the estimated formation state and the projected formation state. Once the estimates are obtained, we assign them to the roles in the formations to determine what is the desired bias from that estimate. A high-level overview of the method is shown in Figure 8.1.

8.1 Problem Statement

We address the problem of multi-robot tracking to provide absolute position estimates necessary for a team of robots to control and keep a desired formation geometry. The multi-robot system consists of N robots $R_1, ..., R_N$. The formation includes one leader robot, which moves on a pre-defined trajectory, and $N-1$ follower robots, all of which know the desired formation geometry, i.e. the desired range and bearing they should keep with respect to their neighbors, including the leader.

Each robot R_i independently estimates its own position $p_i = [x_i, y_i]$ and orientation α_i in a two-dimensional, GNSS-denied environment, based on a known map and onboard sensing. Since all the robots share the same map, they all share a common global coordinate frame $\mathbb{I}_W$. Each robot R_i is equipped with sensors that provide range and bearing to the other robots in its local reference frame $\mathbb{I}_{R_i}$, but the measurements do not include IDs of the detected robots. Since the relation of $\mathbb{I}_{R_i}$ w.r.t. $\mathbb{I}_W$ is known to R_i, the range and bearing sensor measurement j can be expressed in the global frame as $\mathbf{z}_j$ to constitute the measurements set Z_k. Because of the lack of IDs, the followers do neither have the means to distinguish among the neighbors, nor can they tell apart the leader from a follower based on the measurements alone. Therefore, the formation orientation is defined in $\mathbb{I}_W$, which is the only frame known to all the robots. The frame of the leader, $\mathbb{I}_L$, cannot be used for that purpose, as it is not known to the robots at all times. Robots are capable of communicating to each other their global self-localization positions in $\mathbb{I}_W$, but the communication is not reliable enough to be used as the only means to maintain the formation.

The state $\mathbf{x}_j = [x_j, y_j, \dot{x}_j, \dot{y}_j]$ of each target robot in the global reference frame $\mathbb{I}_W$ consists of its position and velocity. Each target follows the linear Gaussian dynamical model as proposed in [60]:

$$\mathbf{x}_{k|k-1} = F\mathbf{x}_{k-1} + n_k \tag{8.1}$$

with the process noise $n_k \sim \mathcal{N}(0, Q)$ and the matrices are:

$$F = \begin{bmatrix} I_2 & \delta I_2 \\ 0_2 & I_2 \end{bmatrix}, \quad Q = \sigma_f^2 \begin{bmatrix} \frac{\delta^4}{4} I_2 & \frac{\delta^3}{2} I_2 \\ \frac{\delta^3}{2} I_2 & \delta^2 I_2 \end{bmatrix} \tag{8.2}$$

where I_n and 0_n denote, respectively, the $n \times n$ identity and zero matrices, δ is the time step, and σ_f^2 is the standard deviation of the process noise.

The sensor measurement $\mathbf{z}_j = [z_j^x, z_j^y]^T$, expressed in the global frame $\mathbb{I}_W$ is a noisy version of the position of a target robot R_j and follows a linear Gaussian observation model:

$$\mathbf{z}_k = H\mathbf{x}_k + \nu_k \tag{8.3}$$

with the measurement noise $\nu_k \sim \mathcal{N}(0, U)$, where $H = [I_2, 0_2]$, $U = \sigma_\epsilon^2 I_2$, and where σ_ϵ^2 is the standard deviation of the measurement noise. The communicated information sent by R_j to the receiving robot R_i includes a position p_j in $\mathbb{I}_W$. Multi-target tracking is performed in the global frame $\mathbb{I}_W$.

8.2 Multi-Target Tracking

The objective of multi-target tracking is to jointly estimate the states of multiple targets, and sometimes the numbers of targets themselves, from a sequence of noisy observation sets. The existence of multiple targets and multiple measurements necessitates data association, a computationally expensive procedure dealt with either explicitly or implicitly. Therefore single-target tracking approaches are *not* readily extensible to the multi-target problems. From among many existing multi-target tracking algorithms that emerged in recent years, [97] classified the existing approaches in the following categories: *a) Non-Bayesian Approaches*, such as Nearest Neighbor (NN) data association, *b) Maximum A Posteriori approaches*, such as

Multiple Hypothesis Tracking (MHT), *c) Bayesian estimators*, such as Joint Probabilistic Data Association (JPDA) filter or Markov Chain Monte Carlo Data Association (MCMCDA), and *d) Finite Set Statistics (FISST)-based approaches*, such as Random Finite Sets (RFSs).

The NN methods deal with data association by assigning each measurement to the closest target based on a distance measure [98]. The NN filters assume one-to-one mapping between the measurements and the targets, therefore cannot deal with multiple observations of a single object, clutter, or occluded objects. MHT evaluates the likelihood that there is a target given a sequence of measurements. To restrict the exponential growth of the number of hypotheses, MHT requires pruning out spurious hypotheses for each track independently and discarding the deleted items, which makes it impossible for MHT to recover from errors [99]. The Probabilistic MHT (PMHT) uses soft association methods, but assumes that the number of targets is known, and that it is possible to initialize the states [100]. The JPDA filter is a sub-optimal Bayesian algorithm that calculates a marginalized probability on the joint data association space. To mitigate the computational burden, many of the heuristic techniques to an approximate JPDA sacrifice the tracking accuracy to make the algorithm computationally tractable and, as a result, the application domain is restricted to scenarios with few, well separated targets [101]. MCMCDA filter is a true approximation for the optimal Bayesian filter. However, the algorithm requires specification of numerous parameters, the target creation is accomplished heuristically and many particles are required for the method to perform well [102].

The Random Finite Set (RFS) approach to multi-target tracking is a novel and promising alternative to the traditional association-based methods. RFS is a theoretically optimal approach to multi-target tracking and a direct generalization of the single-target Bayes filter. Its main advantage is that it treats the problem of clutter and association uncertainty under a rigorous unified Bayesian filtering framework [103]. Moreover, it incorporates track initiation, a procedure that has mostly been performed separately in traditional tracking algorithms. Based on the RFS theory, the Probability Hypothesis Density (PHD) filter and its variations deal with the measurement-to-track association implicitly, resulting in higher robustness and accuracy in scenarios where the number of targets is not known in advance or changes over time.

Flexibility and ability to deal with challenging scenarios makes the PHD filters and the variations thereof increasingly popular in the robotics research. Among many other notable examples, the PHD filter is used for Simultaneous Localization and Mapping (SLAM) in [104], where it estimates vehicles trajectories and the encountered environment features, also in the multi-robot case in [105]. In [96], cooperative tracking of multiple vehicles is achieved by fusion of PHD hypotheses among the cooperative agents. Cooperative multi-target tracking with the PHD filter is also exploited in [106]. The output of the PHD filter in [107] is used not only for tracking of targets, but it also provides importance weighting so that the robots are drawn towards areas that are more likely to contain targets. In [108] visual detection and tracking with GM-PHD filter is performed to estimate the aircraft's position and velocity in 3D.

Multiple autonomous underwater vehicles in [109] are tracked using acoustic signals and a Monte Carlo PHD filter. An example of application-driven modification of a PHD filter in [56] reconstructs the identities of the unmanned ground vehicles by incorporating odometric data collected by an aircraft in the PHD filter.

8.2.1 Multi-Target Tracking using Random Finite Sets

The Random Finite Sets (RFSs) are natural representations of multi-target states and observations that provide a way to directly extend single-sensor, single-target Bayes statistics to multi-sensor, multi-target problems [103]. The RFS formulation treats the collection of individual targets as a *set-valued state*, and the collection of individual observations as a *set-valued observation*.

Let $M(k)$ be the number of the targets at time k with the states $\mathbf{x}_{k,1},\ldots,\mathbf{x}_{k,M(k)} \in \mathcal{X}$. At the next time step some targets may die, surviving targets evolve to the updated states and new targets may appear. Let $N(k)$ be the number of measurements $\mathbf{z}_{k,1},\ldots,\mathbf{z}_{k,N(k)} \in \mathcal{Z}$. Measurements can be generated by the targets or stem from the clutter (false positives). Missed detections can occur due to sensing imperfections (false negatives).

At time k, the collections of target states and the collections of measurements can be represented as finite sets:

$$\begin{aligned} X_k &= \{\mathbf{x}_{k,1},\ldots,\mathbf{x}_{k,M(k)}\} \in \mathcal{F}(\mathcal{X}) \\ Z_k &= \{\mathbf{z}_{k,1},\ldots,\mathbf{z}_{k,N(k)}\} \in \mathcal{F}(\mathcal{Z}) \end{aligned} \tag{8.4}$$

where $\mathcal{F}(\mathcal{X})$ and $\mathcal{F}(\mathcal{Z})$ are the collections of all finite subsets of targets $\mathcal{X}$ and measurements $\mathcal{Z}$ respectively[1].

The RFS State Evolution Model

An RFS model for the time evolution of the multi-target state can incorporate target motion, birth and death of a target [60]. The targets may be temporarily occluded or venture out of the field of view. Thus, given a multi-target state X_{k-1} at time $k-1$, the multi-target state X_k at time k is given by the union of the surviving targets and the spontaneous births:

$$X_k = \left[\bigcup_{\zeta \in X_{k-1}} S_{k|k-1}(\zeta) \right] \cup \Gamma_k \tag{8.5}$$

$S_{k|k-1}(\zeta)$ is the model of the behavior of the state ζ at the next time step and can take on either $\{\mathbf{x}_k\}$ when the target survives or $\emptyset$ when the target dies, corresponding to the cases where the target $\mathbf{x}_{k-1} \in X_{k-1}$ continues to exist at time k with probability $p_{S,k}(\mathbf{x}_{k-1})$ or dies with probability $1 - p_{S,k}(\mathbf{x}_{k-1})$ respectively. Targets that continue to exist transition from a state

1 Analogously to the single-target system, where uncertainty is expressed by modeling the state $\mathbf{x}_k$ and the measurement $\mathbf{z}_k$ as random vectors, the uncertainty of the multi-target system is expressed by modelling the multi-target state X_k and the multi-target measurement Z_k as RFS.

$\mathbf{x}_{k-1}$ to $\mathbf{x}_k$ with the *multi-target transition density* $f_{k|k-1}(\mathbf{x}_k|\mathbf{x}_{k-1})$. Γ_k is the RFS of spontaneous birth at time k.

The RFS Measurement Model

The RFS measurement model accounts for the uncertainty of detection and the clutter. Given a multi-target state X_k at time k, the multi-target measurement Z_k is given by the union of the measurements generated by the target and the clutter:

$$Z_k = \left[\bigcup_{\mathbf{x} \in X_k} \Theta_k(\mathbf{x}_k) \right] \cup K_k. \tag{8.6}$$

where $\Theta_k(\mathbf{x}_k)$ is an RFS generated by the target with a state $\mathbf{x}_k \in X_k$ at time k and can take on either $\{\mathbf{z}_k\}$ when the target is detected with probability $p_{D,k}(\mathbf{x}_{k-1})$ or $\emptyset$ when the target is not detected with probability $1 - p_{D,k}(\mathbf{x}_{k-1})$. For the detected targets, the probability density of obtaining an observation $\mathbf{z}_k$ from $\mathbf{x}_k$ is given by the *multi-target likelihood* $g_k(\mathbf{z}_k|\mathbf{x}_k)$. K_k denotes the clutter.

The Optimal Multi-Target Bayes Filter

The optimal multi-target Bayes filter propagates the multi-target posterior density $p_k(\cdot|Z_{1:k})$ in time using recursion:

$$p_{k|k-1}(X_k|Z_{1:k-1}) = \int f_{k|k-1}(X_k|X) p_{k-1}(X|Z_{1:k-1}) \mu_s(dX) \tag{8.7}$$

$$p_k(X_k|Z_{1:k}) = \frac{g_k(Z_k|X_k) p_{k|k-1}(X_k|Z_{1:k-1})}{\int g_k(Z_k|X) p_{k|k-1}(X_k|Z_{1:k-1}) \mu_s(dX)} \tag{8.8}$$

where $\mu_s(dX)$ is a reference measure on $\mathcal{F}(\mathcal{X})$. Computing (8.7)-(8.8) involves multiple integrals on the space $\mathcal{F}(X)$, which is intractable.

8.2.2 Probability Hypothesis Filter

To deal with intractability, the PHD filter, instead of propagating the multi-target posterior density in time, propagates the *posterior intensity*, a first order statistical moment of the multi-target posterior [103].

For RFS X on the collection of all finite subsets $\mathcal{X}$ with probability distribution P, the first order moment, called the *intensity*, is a non-negative function v on $\mathcal{X}$. In essence, for each region $S \subseteq \mathcal{X}$, the integral of v over that region, $\int_S v(\mathbf{x}) d\mathbf{x}$, gives the expected number of elements of X that are in S[2]. The posterior intensity approximation of the multi-target posterior, where v_k

2 The local maxima of the intensity v are the local concentrations of the expected number of elements and can be used to estimate the elements of X.

and $v_{k|k-1}$ are the approximations of p_k and $p_{k|k-1}$ respectively, can be propagated in time via the PHD recursion:

$$\begin{aligned} v_{k|k-1}(\mathbf{x}) &= \int p_{S,k}(\zeta) f_{k|k-1}(\mathbf{x}|\zeta) v_{k-1}(\zeta) d\zeta + \gamma_k(\mathbf{x}) \\ v_k(\mathbf{x}) &= [1 - p_{D,k}(\mathbf{x})] v_{k|k-1}(\mathbf{x}) + p_{D,k}(\mathbf{x}) v^D_{k|k-1}(\mathbf{x}) \end{aligned} \tag{8.9}$$

with

$$v^D_{k|k-1}(\mathbf{x}) = \sum_{\mathbf{z}\in Z_k} \frac{g_k(\mathbf{z}|\mathbf{x}) v_{k|k-1}(\mathbf{x})}{\kappa_k(\mathbf{z}) + \int p_{D,k}(\xi) g_k(\mathbf{z}|\xi) v_{k|k-1}(\xi) d\xi} \tag{8.10}$$

where $\gamma_k(\mathbf{x})$ is the intensity of the birth RFS Γ_k and $\kappa_k(\mathbf{z})$ is the intensity of the clutter RFS K_k. The PHD filter does not involve combinatorial computations, nevertheless it does not admit closed form solutions.

8.2.3 Gaussian Mixture Probability Hypothesis Filter

The GM-PHD [60] filter admits a closed form solution to the PHD recursion. Under linear, Gaussian assumptions on the target dynamics and birth processes, the posterior intensity is a *Gaussian mixture* of the form:

$$v_k(\mathbf{x}) = \sum_{i=1}^{J_k} w_k^{(i)} \mathcal{N}(\mathbf{x}; m_k^{(i)}, P_k^{(i)}) \tag{8.11}$$

where each Gaussian component i is associated with a weight $w_k^{(i)}$, J_k is the number of Gaussian components representing the intensity and $\mathcal{N}(\cdot; m, P)$ denotes a Gaussian density with mean m and covariance P.

The GM-PHD filter involves four steps: 1) prediction, where the previous intensity evolves according to the motion model and where new targets can appear 2) update, where the intensity is updated with the acquired measurements 3) selection, including merging and pruning, to reduce the number of Gaussian components and 4) state extraction from the posterior intensity.

Prediction

The predicted intensity at the time k is a Gaussian mixture of the form:

$$v_{k|k-1}(\mathbf{x}) = v_{S,k|k-1}(\mathbf{x}) + \gamma_k(\mathbf{x}) \tag{8.12}$$

where $v_{S,k|k-1}(\mathbf{x})$ is the survival intensity:

$$v_{S,k|k-1}(\mathbf{x}) = p_{S,k} \sum_{i=1}^{J_{k-1}} w_{k-1}^{(i)} \mathcal{N}(\mathbf{x}; m_{k|k-1}^{(i)}, P_{k|k-1}^{(i)}) \tag{8.13}$$

with $p_{S,k}$ being the probability of survival, and $\gamma_k(\mathbf{x})$ is the birth intensity with $J_{\gamma,k}$ components:

$$\gamma_k(\mathbf{x}) = \sum_{i=1}^{J_{\gamma,k}} w_{\gamma,k}^{(i)} \mathcal{N}(x; m_{\gamma,k}^{(i)}, P_{\gamma,k}^{(i)}) \tag{8.14}$$

The components of the survival intensity are computed from the previous intensity components according to a linear Gaussian motion model:

$$m_{k|k-1}^{(i)} = F_{k-1} m_{k-1}^{(i)} \tag{8.15}$$

$$P_{k|k-1}^{(i)} = Q_{k-1} + F_{k-1} P_{k-1}^{(i)} F_{k-1}^T \tag{8.16}$$

and $w_{k|k-1}^{(i)} = w_{k-1}^{(i)}$, where F_{k-1} is the state transition matrix and Q_{k-1} is the process noise covariance. The mean values of the birth intensity components, $m_{\gamma,k}^{(i)}$, represent places, where new targets are likely to appear.

Update

Given a set of measurements Z_k, the posterior intensity is updated as follows:

$$v_k(\mathbf{x}) = v_{T,k}(\mathbf{x}) + \sum_{\mathbf{z} \in Z_k} v_{D,k}(\mathbf{x}; \mathbf{z}) \tag{8.17}$$

where

$$v_{T,k}(\mathbf{x}) = \sum_{i=1}^{J_{k|k-1}} (1 - p_{D,k}) w_{k|k-1}^{(i)} \mathcal{N}(\mathbf{x}; m_{k|k-1}^{(i)}, P_{k|k-1}^{(i)}) \tag{8.18}$$

$$v_{D,k}(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^{J_{k|k-1}} w_k^{(i)}(\mathbf{z}) \mathcal{N}(\mathbf{x}; m_{k|k}^{(i)}(\mathbf{z}), P_{k|k}^{(i)}) \tag{8.19}$$

where $p_{D,k}(m_{k|k-1}^{(i)})$ is the state-dependent probability of detection. Intuitively, $v_{T,k}(\mathbf{x})$ is the missed-detection term, where the weight of each Gaussian component of the predicted intensity is discounted according to $p_{D,k}$ (see Equation 8.18). The $v_{D,k}(\mathbf{x}; \mathbf{z})$ term, one for each measurement $\mathbf{z} \in Z_k$, is the detection term, which provides closed form expressions for computing the means, covariances and weights of v_k from those of $v_{k|k-1}$ when a new set of measurements arrives. The complete expressions for $w_k^{(i)}$, $m_{k|k}^{(i)}$ and $P_{k|k}^{(i)}$ in Equation 8.19 are:

$$w_k^{(i)}(\mathbf{z}) = \frac{p_{D,k} w_{k|k-1}^{(i)} q_k^{(i)}(\mathbf{z})}{\kappa_k(\mathbf{z}) + p_{D,k} \sum_{l=1}^{J_{k|k-1}} w_{k|k-1}^{(l)} q_k^{(l)}(\mathbf{z})} \tag{8.20}$$

$$m_{k|k}^{(i)}(\mathbf{z}) = m_{k|k-1}^{(i)} + K_k^{(i)}(z - H_k m_{k|k-1}^{(i)}) \tag{8.21}$$

$$P_{k|k}^{(i)} = [I - K_k^{(i)} H_k] P_{k|k-1}^{(j)} \tag{8.22}$$

with

$$q_k^{(i)}(\mathbf{z}) = \mathcal{N}(\mathbf{z}; H_k m_{k|k-1}^{(i)}, U_k + H_k P_{k|k-1}^{(i)} H_k^T) \tag{8.23}$$

$$K_k^{(i)} = P_{k|k-1}^{(i)} H_k^T (H_k P_{k|k-1}^{(i)} H_k^T + U_k)^{-1} \tag{8.24}$$

where H_k is the observation matrix, U_k is the observation noise covariance and $\kappa_k(\mathbf{z})$ is the expected clutter level.

Selection

The update step yields a quadratic increase in the number of Gaussian components the posterior intensity is composed of. To keep the problem tractable, components with weak weights are pruned: $I = \{i = 1, ..., J_k | w_k^{(i)} > T_S\}$.

Furthermore, all Gaussian components close to each other are merged into a single Gaussian as follows. At first, a Gaussian component with the highest weight is selected with $j = \text{argmax}_{i \in I} w_k^{(i)}$. Then all Gaussian components within the Mahalonobis distance U_S from j are forming a set of Gaussian components:

$$L = \{i \in I | (m_k^{(i)} - m_k^{(j)})^T (P_k^{(i)})^{-1} (m_k^{(i)} - m_k^{(j)}) \leq U_S\} \tag{8.25}$$

that are merged into a single component:

$$\tilde{w}_k^{(l)} = \sum_{i \in L} w_k^{(i)}, \qquad \tilde{m}_k^{(l)} = \frac{1}{\tilde{w}_k^{(l)}} \sum_{i \in L} w_k^{(i)} m_k^{(i)}, \tag{8.26}$$

$$\tilde{P}_k^{(l)} = \frac{1}{\tilde{w}_k^{(l)}} \sum_{i \in L} w_k^{(i)} (P_k^{(i)} + (\tilde{m}_k^{(l)} - m_k^{(i)})(\tilde{m}_k^{(l)} - m_k^{(i)})^T)$$

Finally, the number of Gaussian components is truncated to J_{max} components with the highest weights, resulting in a posterior intensity shaped as a Gaussian mixture:

$$v_k(\mathbf{x}) = \sum_{i=1}^{J_k} w_k^{(i)} \mathcal{N}(x; m_k^{(i)}, P_k^{(i)}) \tag{8.27}$$

State extraction

The means of the Gaussian components are the local maxima of the posterior intensity v_k. Extraction of multi-target state estimates comes down to selection of the Gaussian means that have weights greater than a threshold T_{SE}.

8.3 GM-PHD Filter with Formation Information

In this section we introduce the FI-GM-PHD filter, overview of which is shown in Figure 8.2. The FI-GM-PHD filter consists of two steps for supplementing additional Gaussian compo-

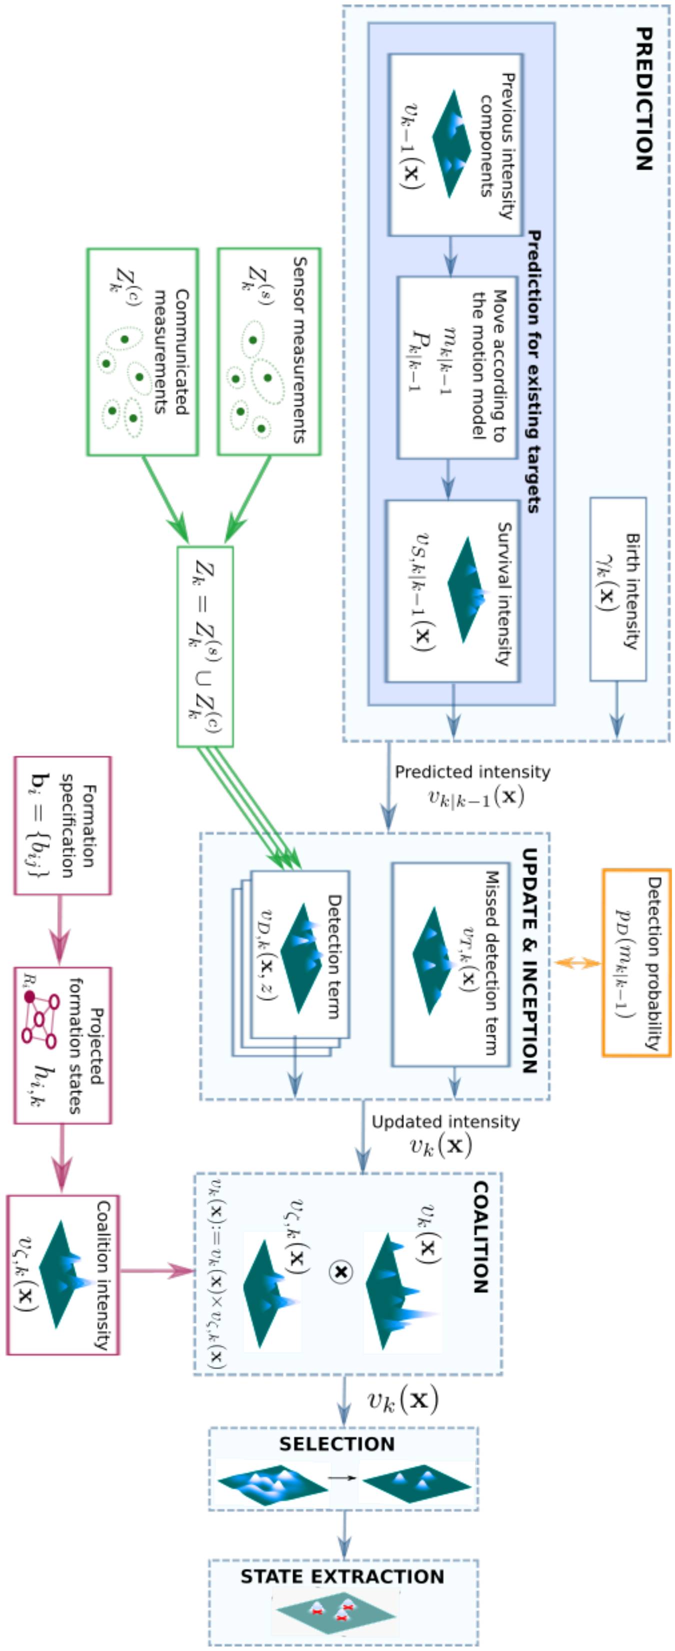

Figure 8.2 – An overview of the FI-GM-PHD filter.

nents to the filter intensity:

- ***The inception step*** - incorporation of communication data;
- ***The coalition step*** - using spatial configuration of the formation as a prior for the PHD filter.

To provide the robots with the necessary information to maintain the formation, each follower robot R_i runs its own FI-GM-PHD filter. The leader robot moves independently of the other robots and therefore does not rely on other robot estimates.

The filters are decentralized from the computational viewpoint, as they are run locally and independently by the different follower robots. Role assignment and formation control procedures are also run individually by the follower robots. However, because of the global networking and the implemented information sharing strategy, each follower runs a FI-GM-PHD filter instance with information about the locations of all its teammates.

8.3.1 Prediction in the FI-GM-PHD Filter

The components of the survival intensity $v_{S,k|k-1}(\mathbf{x})$ constituting the predicted intensity $v_{k|k-1}(\mathbf{x})$ (Equation 8.12) at the time k are computed from the previous intensity components according to a linear Gaussian motion model: $m^{(i)}_{k|k-1} = F_{k-1} m^{(i)}_{k-1}$ and $P^{(i)}_{k|k-1} = Q_{k-1} + F_{k-1} P^{(i)}_{k-1} F^T_{k-1}$ (corresponding to Equation 8.15 and Equation 8.16 respectively). The motion model of a target with the state $\mathbf{x}_j$ composed of position and velocity is a linear Gaussian dynamical model with the state transition matrix F and the process noise covariance Q, described in detail in Section 8.1. At initialization, the filter run by robot R_i is supplied with a birth RFS $\gamma_0(\mathbf{x})$ at the initial detections. At time $k \neq 0$, birth intensity in Equation 8.12 is $\gamma_k(\mathbf{x}) = \emptyset$.

8.3.2 Update in the FI-GM-PHD Filter

The update step in the FI-GM-PHD filter differs from the procedure of the original filter described in Sec. 8.2.3 with regard to the source of the measurements. While in the GM-PHD filter the set of measurements Z_k consists of direct measurements obtained from the sensors of the tracking robot, the detection term $v_{D,k}$ in the FI-GM-PHD filter also includes communicated state information, added in what we refer to as the inception step.

Measurement Update

The missed-detection term $v_{T,k}(\mathbf{x})$ constituting the update intensity in Equation 8.17 is calculated using the predicted intensity $v_{k|k-1}(\mathbf{x})$, discounted according to the probability of detection $p_{D,k}$ as shown in Equation 8.18. The state-dependent probability of detection used in this work is described in detail in Section 8.3.2; it suggests how likely it is that the tracked robot is detected given the occlusions and its position with respect to the field of view of the tracking robot. The detection term $v_{D,k}(\mathbf{x};\mathbf{z})$ (Equation 8.19), one for each measurement of the position $\mathbf{z}_j$ of the target R_j, follows a linear Gaussian observation model with the observation matrix H and the observation noise covariance U, as detailed in Section 8.1. The sensor measurements in Z_k are obtained from onboard range and bearing sensors described in Section 2.4.2, where the range and bearing data is used to determine the position of target R_j in the global frame $\mathbb{I}_W$ as explained in Section 8.1.

Inception of the Communicated Data

Even when communication between the robots is possible, it may suffer of message losses, be of low rate or break occasionally. For the receiving robot R_i to perform inception in the update step at time k, the communicated position information $p_{j,k}$, one per each communicating neighbor robot $j = 1, ..., N_k$, forms a measurement $\mathbf{z}_k^{(j)}$. The measurement is added to Z_k to form a new measurement set:

$$Z_k := Z_k \cup \sum_{j=1}^{N_k} \mathbf{z}_k^{(j)} \tag{8.28}$$

There is no need to associate the position messages with the existing Gaussian components from the prediction step, as the PHD filter does not require data association, but each measurement generates a new set of components updated from the predicted intensity (see Equation 8.19). In a general case, including communicated measurements is not suitable for the update step, because if only a fraction of the robots is exchanging the information, then performing the update would delete from the output the tracks of the robots that did not communicate. This can be prevented by ensuring that when merging the Gaussian components generated by the transmitted data of the robots that *do* communicate and the Gaussian components from the prediction step that might correspond to those robots, the weight of the merged component always remaining between $(0, 1)$.

By curbing the merged weight, the predicted components that are not strengthened by additional communicated data are prevented from pruning. The selection step thus includes an additional merging procedure. Now, after performing the standard merging step (Equation 8.2.3 to Equation 8.26), we again select a Gaussian component with the highest weight $j = \mathrm{argmax}_{i\in I} w_k^{(i)}$. Then all the components within an Euclidean distance $2r_r$ between the position part of the Gaussian mean, $m_{0:1}$, and the position part of the mean of j form the set $L' = \{i \in I | \| m_{0:1,k}^{(i)} - m_{0:1,k}^{(j)} \| \leq 2r_r\}$, where $2r_r$ is the robot diameter. The set L' is merged into a single component with the weight $\tilde{w}_k^{'(l)} = \min\left(\sum_{i\in L'} w_k^{(i)}, 1\right)$. The calculation of $\tilde{m}_k^{'(l)}$ and $\tilde{P}_k^{'(l)}$ proceeds as in Equation 8.26.

Note that if the model for the communicated measurements differed from the model of the sensor-based measurements, Equations. 8.20-8.22 would have to be applied with each measurement updated according to its respective model.

Probability of Detection

In order to reduce the risk of losing track of a robot when it enters an occluded area or escapes the field of view of the detecting robot, the probability of detection $p_{D,k}(m^{(i)}_{k|k-1})$ is state dependent. Constant probability of detection would result in quick decrease of the weights of the components without associated measurements.

The method for determining $p_{D,k}(m^{(i)}_{k|k-1})$ is summarized in Algorithm 1 and illustrated in Figure 8.3. If the estimated position $\hat{\mathbf{p}}_j$ of the tracked robot R_j is outside the sensing range r_s of the detecting robot R_i, a minimum detection probability p_D^{min} is assigned (line 4 in Algorithm 1). Additionally, the object of interest is considered occluded if any other object k is located between the detecting robot and the object of interest. In this case, the decrease in the probability of detection is proportional to the weight of the occluding object and proportional to the estimated degree of occlusion.

In the Inception step (Sec. 8.3.2), the probability of detection associated with the components stemming from the communicated measurements p_D^{Inc} is state-independent and constant.

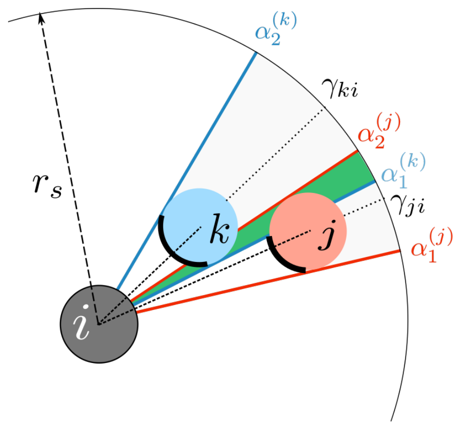

Figure 8.3 – Illustration of the occlusions and field of view model. The detecting robot R_i determines the degree of occlusion of the robot R_j. Angles $[\alpha_1^{(o)}, \alpha_2^{(o)}]$ are the boundaries of the region occluded by object o. The green region is the intersection of the occlusion regions of R_j and R_k.

8.3.3 The Projected Formation State

Given the absolute position $p_{i,k} = [x_{i,k},\ y_{i,k}]$ of the robot R_i at time k, the position of the robot R_j is projected in the global frame $\mathbb{I}_W$ based on the desired formation geometry:

$$\begin{bmatrix} h^x_{ij,k} \\ h^y_{ij,k} \end{bmatrix} = \begin{bmatrix} b^x_{ij} \\ b^y_{ij} \end{bmatrix} + \begin{bmatrix} x_{i,k} \\ y_{i,k} \end{bmatrix} \tag{8.29}$$

where b_{ij} is the bias between the robot R_i and R_j. The collection of the projected positions with respect to the robot R_i of all the other robots in the formation is denoted by:

$$h_{i,k} := \{\{h^x_{ij,k}, h^y_{ij,k}\}, \mid j = 1, ..., N;\ j \neq i\} \tag{8.30}$$

where N is the number of robots in the formation and so, $|h_{i,k}| = N - 1$.

Algorithm 1: PROBABILITY OF DETECTION

Given the position of the detecting robot $p_i = [x_i,\ y_i]$, the estimated position of the object of interest $\hat{p}_j = [\hat{x}_j, \hat{y}_j]$ and the set of all object estimates $\{\hat{\mathbf{x}}_k\}_{k=1}^{N}$ with weights $w^{(k)}$
Denote range $r_{ab} = \| p_a - p_b \|$
Denote bearing $\gamma_{ab} = tan^{-1}((y_a - y_b)/(x_a - x_b))$
  $p_D^{(j)} = p_D^0$
  $\beta_{TOT} = \emptyset$ *(total occluded region)*
  1) Check if object within the sensing range
  if $r_{ij} > r_s$ *(outside sensing range)*
    $p_D^{(j)} = p_D^{min}$
  else
    $\alpha^{(j)} = tan^{-1}(r_r / r_{ij})$
    $[\alpha_1^{(j)}, \alpha_2^{(j)}] = \gamma^{ji} \pm \alpha^{(j)}$ *(region where R_j can be occluded)*
    for $k = 1, \ldots, N$
      if $r_{ik} < r_{ij}$
        $[\alpha_1^{(k)}, \alpha_2^{(k)}] = \gamma^{ki} \pm \alpha$ *(region which R_k can occlude)*
        2) Check if object k occludes the object of interest
        $\left[\beta_1^{(jk)}, \beta_2^{(jk)}\right] = \left[\alpha_1^{(j)}, \alpha_2^{(j)}\right] \cap [\alpha_1^{(k)}, \alpha_2^{(k)}]$
        if $\left[\beta_1^{(jk)}, \beta_2^{(jk)}\right] \neq \emptyset$ *(R_j is occluded by R_k)*
          3) Check if not double counting occlusion
          $\left[\beta_1^{(jk)}, \beta_2^{(jk)}\right] = \left[\beta_1^{(jk)}, \beta_2^{(jk)}\right] - \beta_{TOT}$
          $\beta_{TOT} = \beta_{TOT} \cup \left[\beta_1^{(jk)}, \beta_2^{(jk)}\right]$
          4) Decrease the probability of detection
          $p_D^{(j)} = p_D^{(j)}\left[1 - w^{(k)}\left(1 - \frac{\left|\left[\beta_1^{(jk)}, \beta_2^{(jk)}\right]\right|}{\left|\left[\alpha_1^{(j)}, \alpha_2^{(j)}\right]\right|}\right)\right]$
        end if
      end if
    end for
  end if
return $p_D^{(j)}$

8.3.4 Coalition of the Projected Formation States

The *coalition* step extends the GM-PHD filter with an additional block, added after the update step. It combines the intensities obtained during the update step with the coalition intensity derived from the projected formation states. Thus, the Gaussian components constituting the coalition intensity serve as an outline of where the tracked robots are to be expected.

The coalition step is detailed in Algorithm 2. The set of states $h_{i,k}$ projected by the robot R_i is used to approximate the means of the components of the coalition intensity $v_{\zeta,k}$ as follows. For each tracked robot R_j the mean is placed at the projected position of that robot $m_{\zeta,k}^{(j)} = [h_{ij,k}^{x},\ h_{ij,k}^{y},\ 0,\ 0]^T$, where $j = 1, \ldots, N;\ j \neq i$ and so $J_{\zeta,k} = N - 1$. Each component

$j = 1, \dots, J_{\zeta,k}$ of the coalition intensity $v_{\zeta,k}$ at time k is given an initial budget $\Phi^{(j)}_{\zeta,k} = \Phi_{\zeta,0}$. Then, all the components forming the posterior intensity $v_k(\mathbf{x})$, $l = 1, \dots, J_k$, are compared against j to find the matching that maximizes some criteria (line 7 of Algorithm 2). We choose to put emphasis on minimizing the distance between the position part of the component means ($m_{0:1}$ corresponds to the position part of the state, while $m_{2:3}$ corresponds to the velocity part of the state), while choosing components of the posterior intensity with significant weights. By sorting the posterior intensity components according to the $o^{(j,l)}_k$ measure, we first evaluate the best candidates for good matching.

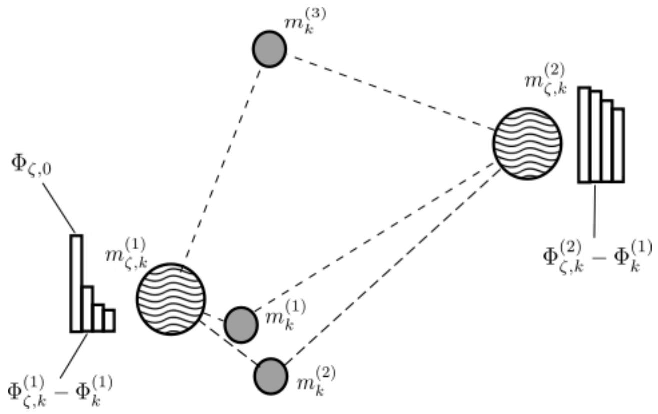

Figure 8.4 – Illustration of the coalition step. Three posterior intensity components, $m^{(1)}_k, m^{(2)}_k, m^{(3)}_k$ are compared against two coalition components, $m^{(1)}_{\zeta,k}, m^{(2)}_{\zeta,k}$ with the corresponding budgets $\Phi^{(1)}_{\zeta,k}, \Phi^{(2)}_{\zeta,k}$. Closeness of the $m^{(1)}_{\zeta,k}$ to the $m^{(1)}_k$ decreases significantly its budget, while for $m^{(2)}_{\zeta,k}$, the budget does not become depleted and a novelty is created.

The components are coalesced as follows. Lines 16-18 calculate a temporary budget for the posterior component ($\Phi^{(l)}_k$) using a sigmoid function $f \in \langle 0, 1 \rangle$ of the distance between the means of the components (line 15). If the distance is small, the possibility that both components correspond to the same target is high and the components j and l are coalesced to form a new Gaussian component, with a mean of intensity being a combination of the means of the two, the covariance and the weight being the covariance and the weight of l modified as a function of the divergence. If the distance between j and l is large, the likeliness that l is associated with j is small and l is propagated further with little change. This diverse behavior is assured by using a sigmoid-shaped function of the distance when comparing the components.

The budget of the coalition component is decreased with every posterior component that has been found close to it, and has two major advantages. Firstly, it limits the number of new components that can originate around it. Secondly, a budget left at the end of iteration indicates that one of the coalition components did not have a corresponding component in the posterior, whether because of the missed detection, occlusion, or field of view limits. In that case, a new component, called the *novelty* is created with the mean at the area where a robot is expected to be. Only the coalition components are allocated the initial budget; the components of the posterior intensity come with an associated weight and their budget is computed during the coalition process (line 15).

While in our previous work [4] we assumed that the formation always stays close-to-desired, here we relax this assumption by moderating the importance of the novelty components

Algorithm 2: THE COALITION STEP

Given the components means approximating the projected formation states $\{m^{(j)}_{\zeta,k}\}_{j=1}^{J_{\zeta,k}}$ and the components of the posterior intensity $\{w^{(l)}_k, m^{(l)}_k, P^{(l)}_k\}_{l=1}^{J_k}$

$n = 0$
for $j = 1, \dots, J_{\zeta,k}$
  $\Phi^{(j)}_{\zeta,k} = \Phi_{\zeta,0}$
  for $l = 1, \dots, J_k$
    1) (Evaluate best-matching criteria)
    $o^{(j,l)}_k = \exp(\| m^{(l)}_k - m^{(j)}_{\zeta,k} \|) + (w^{(l)}_k + \epsilon)^{-1}$
    $d^{(j,l)}_k = \| m^{(l)}_{0:1,k} - m^{(j)}_{0:1,\zeta,k} \|$
  end for
  2) (Sort posterior components)
  $\mathrm{sort}(\{w^{(l)}_k, m^{(l)}_k, P^{(l)}_k\}_{l=1}^{J_k}, o^{(j,l)}_k)$
  for $l = 1, \dots, J_k$
    if $\Phi^{(j)}_{\zeta,k} > \Phi_{\zeta,min}$
      3) (Coalesce components)
      $\Phi^{(l)}_k = f((d^{(j,l)}_k)^{-1})$
      $\bar{m}^{(n)}_k := \Phi^{(l)}_k m^{(l)}_k + (1 - \Phi^{(l)}_k) m^{(i)}_{\zeta,k}$
      $\bar{P}^{(n)}_k := (\Phi^{(l)}_k + \epsilon)^{-1} P^{(l)}_k$
      $\bar{w}^{(n)}_k := \Phi^{(l)}_k w^{(l)}_k$
      $n := n + 1$
      4) (Update budget)
      $\Phi^{(j)}_{\zeta,k} = \Phi^{(j)}_{\zeta,k} - \Phi^{(l)}_k$
      5) (Estimate association error)
      $\hat{e}_{S,k} = (1 - \Phi^{(l)}_k)(w^{(l)}_k + \epsilon)^{-1}$
    end if
  end for
end for
$\hat{e}_S = \mathrm{median}(\hat{e}_{S,k})$
6) (Novelty)
for $j = 1, \dots, J_{\zeta,k}$
  if $\Phi^{(j)}_{\zeta,k} > 0$:
    $\bar{m}^{(n)}_k := m^{(j)}_{\zeta,k}$
    $\bar{P}^{(n)}_k := \hat{e}_S P^{(j)}_{\zeta,k}$
    $\bar{w}^{(n)}_k := (\hat{e}_S + \epsilon)^{-1}$
    $n := n + 1$
  end if
end for
$\bar{J}_k = n$
return $\{\bar{w}^{(n)}_k, \bar{m}^{(n)}_k, \bar{P}^{(n)}_k\}_{n=1}^{\bar{J}_k}$

according to how well the matching between the posterior and the coalition components has been performed on the whole. The matching is estimated using the association error $\hat{e}_S$, calculated in line 23 of Algorithm 2, where a good matching means that a posterior component with a substantial weight is found close to a coalition component. The quality of the overall matching is used to determine the weight of the novelty components in line 33. The larger the error, the lower the weight of the novelty component. In case the overall matching is poor, the weights of the coalition components are negligible and become discarded in the prediction step of the next iteration. This property modulates the impact of the coalition step on the role assignment procedure, making the latter rely on a) the Gaussian components of the full FI-GM-PHD filter when the matching is good or b) the Gaussian components created based on communicated and sensed data otherwise, for example when formation shape is not perfect.

As an example, depicted in Figure 8.4, consider three posterior intensity components, $m_k^{(1)}$, $m_k^{(2)}$, $m_k^{(3)}$ compared against two coalition components, $m_{\zeta,k}^{(1)}, m_{\zeta,k}^{(2)}$ with the corresponding budgets $\Phi_{\zeta,k}^{(1)}, \Phi_{\zeta,k}^{(2)}$ initially set to $\Phi_{\zeta,0}$. Closeness of $m_k^{(1)}$ to $m_{\zeta,k}^{(1)}$ decreases significantly its budget $\Phi_{\zeta,k}^{(1)}$, and a new component is created as a combination of the two. For $m_{\zeta,k}^{(2)}$, the budget does not become depleted and a novelty is created. Finally, component $m_k^{(3)}$ has no close correspondence in v_{ζ_k} and it is propagated with no modification.

The coalition step and the inception step supplement the intensity with additional Gaussian components. In case the robot has not been tracked, adding a new Gaussian component creates a new target based on the high probability that the target is there. This is analogous to the birth process of the GM-PHD filter in [60]. In case the robot has already been tracked, adding a new Gaussian component that corresponds to that robot increases the likelihood of the robot being present at that position, provided that the state extraction takes into account the fact that multiple targets cannot occupy the same physical position, i.e. the components with weight above one are extracted as single targets.

8.4 Role Assignment

The multi-target state estimates of the FI-GM-PHD filter provide the range and bearing information, but in order to be used in the ID-based formation algorithm they are dynamically assigned with the target positions (roles) in the formation.

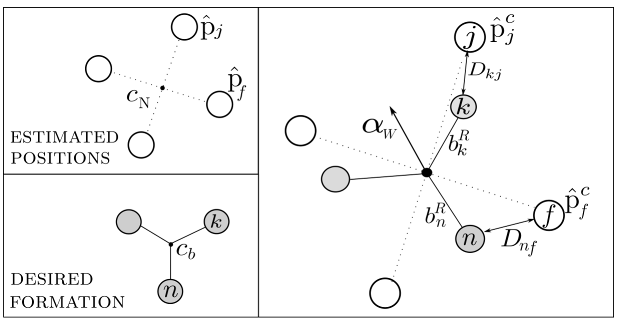

Figure 8.5 – Illustration of the role assignment step. The center of the estimates $\mathbf{c}_N$ and the center of the desired formation $\mathbf{c}_b$ are brought to a common reference frame at the origin, where the translated position estimates are $\hat{p}_N^c$, and the bias, translated and rotated according to the formation orientation α_W, is b^R. D_{kj} is the cost of associating bias b_k^R with estimate $\hat{p}_j^c$.

Algorithm 3: ROLE ASSIGNMENT

Given formation orientation α_W, position of detecting robot R_i, $p_i = [x_i, y_i]$, estimated positions $\hat{p}_{N-i} = \{[x_j, y_j]\}_{j=1}^{\hat{N}-1}$ and bias $b_i = \{[b^x_{ij}, b^y_{ij}]\}_{j=1}^{N}$
1) Find common coordinate frame
    $\hat{p}_N = [p_i, \hat{p}_{N-i}]$
    $\mathbf{c}_N = \text{mean}(\hat{p}_N)$ *(find center of the robots)*
    $\hat{p}^c_N = \hat{p}_N - \mathbf{c}_N$ *(robot positions w.r.t the origin)*
    $\mathbf{c}_b = \text{mean}(b_i)$ *(find center of the formation)*
    $b^R = \text{rotate}(\alpha_W)(b_i - \mathbf{c}_b)$ *(rotated global bias)*
2) Find best assignment
    $D = \text{cost}(\hat{p}^c_N, b^R)$ *(find assignment cost)*
    $A_\Delta = \text{hungarian}(D)$ *(choose best assignment)*
return $A_\Delta = \{\text{link } \epsilon_L = (\hat{p}_j, s)\}_{j=1,s=1}^{N}$

The role assignment procedure finds a permutation that assigns the estimates and the position of the detecting robot to the target positions in the formation. Although we use the graph-based framework notation, the method can be generalized to other formation algorithms. The procedure is sketched in Algorithm 2. Consider a formation with its geometry uniquely determined by a bias matrix b, a set of desired inter-robot distances. The heading of the formation is defined in a global reference frame and is assumed to be known by all the robots. In our case the formation orientation α_W is along the y-axis of the world frame. The bias[3] and the estimates are brought to a common reference frame by matching the center of the rotated global bias b^R with the center of the estimated positions $\hat{p}^c_N$ (see Figure 8.5). The combination of the estimates and the roles resulting in the smallest cost is computed using the Hungarian algorithm [110]. The role assignment procedure, run by the robot R_i, provides a set of links, $A_\Delta = \{\epsilon_L = (\hat{p}_i, s),\}_{s=1}^{N}$. The link $\epsilon_L = (\hat{p}_j, s)$ is used for correct matching of the bias $b_{is}(t)$ with $d_{is}(t)$ and $\gamma_{ij}(t)$ in Equation 5.5, while its cost gives the confidence about the assignment. More precisely, estimate $\hat{\mathbf{x}}_j$ (akin to position $\hat{p}_j$) is associated with role "s" and its subscript is changed to $\hat{\mathbf{x}}_s$, with the range d_{is} and the bearing γ_{is} with respect to the robot R_i. The estimate $\hat{\mathbf{x}}_s$ is coupled with bias b_{is} that corresponds to the "s^{th}" place in the formation.

The role assignment procedure provides the tracking robot R_i with its own role. In the next iteration of the FI-GM-PHD filter this role is used to determine the projected formation state that affects the coalition step, which in turn after state extraction is used for the next role assignment. Since the coalition step introduces novelty components for robots that can neither be communicated with nor tracked, it is important to note that without the fundamental assumption of our work, namely that only a subset of robots experiences communications outage at the same time, errors are likely to occur and have a tendency to escalate. With relatively small teams of robots, the formation recovers from sporadic inconsistencies.

3 Note that the bias b_i specifies the relative distances between the robot R_i and all the other robots in the formation. It is possible to calculate the global formation shape based on that information.

8.5 Experimental Setup

Thorough evaluation of the FI-GM-PHD filter is first performed in the Webots simulator, where the two LIDAR sensors are accurately simulated and calibrated using real data. The second set of experiments with real robots is carried out at the Jordils experimental facility. In both cases the implementation details, including parametrization of the algorithms, remain the same with the exception of the probability of detection model, which in real experiments is updated based on empirical evaluation to take into account imperfection of the detection system caused by overlapping laser scans.

8.5.1 Implementation

For robot detection, we use the on-board relative localization system described in Section 2.4.2, which returns a measurement of a robot position in the form $\mathbf{z}_i = [z_i^x, z_i^y]^T$. The position of the leader is not known globally to all the robots, but the robots must track the leader in order to follow it in the formation. Since the pose of the leader is not globally shared, the followers align their orientation to the positive y-axis of the map frame $\mathbb{I}_W$[4].

The tracking filter is run in a global frame $\mathbb{I}_W$. The tracking parameters of the filter are as follows. The state of the target is $\mathbf{x}_j = [x_j, y_j, \dot{x}_j, \dot{y}_j]$. Each target has the survival probability $p_{S,k} = 0.95$, and follows the linear Gaussian dynamical model (see Sec. 8.1) with $\sigma_f^2 = 1.0$, which accounts for motion uncertainty associated with holonomic motion of robots, and $\delta = 0.5s$. The measurement $\mathbf{z}_i = [z_i^x, z_i^y]^T$ is a noisy version of the position and follows a linear Gaussian observation model with the standard deviation of the position $\sigma_\epsilon^2 = 0.2$, determined empirically in simulation.

At initialization, the filter run by robot R_i is supplied with a birth RFS at the initial detections with the covariance $P_{\gamma,0} = diag([\sigma_\epsilon^2, \sigma_\epsilon^2, \sigma_\gamma^2, \sigma_\gamma^2]^T)$, where $\sigma_\gamma^2 = 1.0$, and the weight $w_{\gamma,0} = 1.0$. We choose σ_ϵ^2 as it is the expected measurement error and σ_γ^2 as the measurements do not include the velocity. The parameters for merging and pruning are: $J_{max} = 10$, $U_S = 0.3$, $T_S = 10^{-4}$ and we select Gaussian components above $T_{SE} = 0.5$. Furthermore, to take into account physical dimensions of the robots, the selection step (see Equation 8.2.3) additionally merges components if their means are closer than twice the radius of the robot, $2r_r = 0.65\ m$, $L = \{i \in I \| \| m_k^{(i)} - m_k^{(j)} \| \leq 2r_r\}$. The poisson distributed clutter level is $\kappa_k(z) = 0.015$ within the sensing range of the robot. The parameters for the detection probability are $p_D^0 = 0.9$, $p_D^{min} = 0.02$. The sensing radius $r_s = 4.0\ m$ is the range of the laser scan.

Inception Given that the sending robot R_i at time k communicates its estimated position $p_{i,k} = [x_{i,k},\ y_{i,k}]$ and the orientation, this information is used in the inception step in the form of a measurement, where $\mathbf{z}_k^{(i)} = [x_{i,k},\ y_{i,k}]$. The robot orientation is not included in the

4 The robots self-localize on a known map, therefore the orientation of the body with respect to the map frame is estimated online.

state estimation as our current laser detection cannot provide orientation information. The constant probability of detection used in the model is $p_D^{Inc} = 0.9$.

Coalition The coalition mean intensity is based on the projected formation state given by $m_{\zeta,k}^{(j)} = [h_{ij,k}^x,\ h_{ij,k}^y,\ 0,\ 0]^T$, $j = 1, \ldots, J_{\zeta,k}$. The initial budget is $\Phi_{\zeta,0} = 0.5$ and the minimal budget is $\Phi_{\zeta,min} = 0.1$. The function used for coalescing the components is $f(x) = \left(1 + \exp(-10(x-0.5))\right)^{-1}$.

8.5.2 Performance Evaluation

We study the tracking performance using the Optimal SubPattern Assignment (OSPA) metric [111]. OSPA is comprised of two components, one accounting for localization accuracy and the other for the cardinality error:

$$\bar{o}_p^{(c)}(X,Y) = \Big(\frac{1}{\delta}\big[\min_{\pi\in\Pi}\sum_{(i,j)\in\pi}\min(\|x_i, y_j\|, c)^p + c^p(|n-m|)\big]\Big)^{1/p} \tag{8.31}$$

where $n = |X|$, $m = |Y|$ and $\delta = \max(m, n)$. Π is the best permutation between X and Y found using the Hungarian algorithm, p penalizes the estimated position error and c is the cut-off parameter for penalizing cardinality errors. In our experiments, OSPA is computed between the ground truth positions and the estimated positions, unless stated otherwise. Each robot calculates the OSPA metric individually based on the outcome of its own instance of the filter. We present the mean results of the robots, averaged over the runs.

8.6 Simulation Experiments

With four scenarios, we conduct a thorough evaluation of the FI-GM-PHD filter, and present a comparison with respect to the full-communication-no-tracking situation and with respect to the standard GM-PHD filter.

8.6.1 Scenarios

We study the performance of the methods under the following four scenarios: I) for tracking purposes only, i.e. the robots do not use the tracking data for control; II) for formation initialization and convergence; III) in challenging situations, where the robots navigate around obstacles scattered in the environment or the measurement error is large; IV) in realistic scenario where communications suffers from the periods of outage.

Scenario I: Multi-Robot Tracking

The dataset is collected when five robots, one leader and four followers, move in a cross-shaped formation on a circular trajectory. Robots maintain the formation using communicated self-localization information, while simultaneously collecting sensory data. With the collected data we perform multi-robot tracking with i) the standard GM-PHD filter, and with the variations of the FI-GM-PHD method: ii) with inception step only, with iii) the coalition step only, and with iv) the full FI-GM-PHD system (both inception and coalition steps). We test the methods that combine communications (ii and iv) with message drop probabilities $p_{md} = 0.0$, $p_{md} = 0.5$ and $p_{md} = 0.9$ and in a situation when two robots are not communicating and the other robots communicate with $p_{md} = 0.0$, $p_{md} = 0.5$ and $p_{md} = 0.9$. For each experiment we perform 10 sequential runs, each lasting approximately 180 s.

Scenario II: Initialization

We generate ten worlds (corresponding to ten experimental runs) with four robots placed at random initial positions and with random orientations. The sensing network is initially connected. In the experiments, both the tracking and the control steps of our method are applied, which means that the tracking data is used for initialization (bringing the robots from random initial positions to positions in the formation) and maintenance of a diamond formation. As a baseline for comparison, we use an experiment where robots rely on i) only communication, but where robot roles are given a priori, ii) only communication with dynamic role assignment, so that the robots take optimal paths to the formation. We study the performance of iii) the standard GM-PHD filter, iv) the inception step only, v) the coalition step only, and vi) the full FI-GM-PHD system. The methods (i-ii, iv and vi) are tested with message drop probabilities $p_{md} = 0.0$, $p_{md} = 0.5$ and $p_{md} = 0.9$ and in a situation when two robots are not communicating (and where for the other robots $p_{md} = 0.0$). For each of the ten worlds we allow maximally three trials for testing whether one of the variants above can complete successfully.

Scenario III: Limitations

Limitations of our methods are studied with regard to two aspects, A) challenging environments with obstacles and B) precision of robot detection.

Scenario III-A: Challenging Environments The leader robot guides a diamond-shaped formation of four robots in an environment populated with obstacles of various shapes and sizes. We use an experiment with i) perfect communication with dynamic role assignment as a baseline for comparison. We study the performance of ii) the standard GM-PHD filter and iii) the full FI-GM-PHD system. In iii), only three robots communicate. For each experiment, 5 sequential runs of approximately 400 s are performed.

Scenario III-B: Measurement Error Three robots move on a circular trajectory in a triangle-shaped formation. We compare i) perfect communication with dynamic role assignment, ii) the standard GM-PHD filter and iii) the full FI-GM-PHD system. For ii-iii) we add a random uniform error to the original measurement, with the magnitude of a) $e_M = 0.0\ m$, b) $e_M = 0.3\ m$, c) $e_M = 0.6\ m$, d) $e_M = 0.75\ m$ and e) $e_M = 1.0\ m$. For iii) only the leader robot communicates. We perform 10 sequential runs of around 120 s.

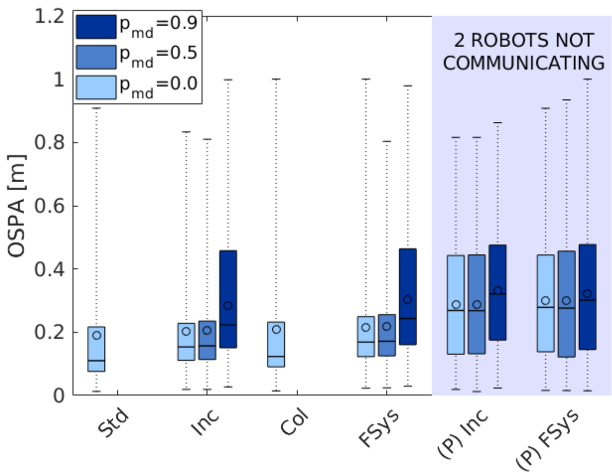

Figure 8.6 – Scenario I. The OSPA metric. Note that the *Std* and *Col* do not use the communicated data. *Std* stands for the standard GM-PHD filter; *Inc* stands for FI-GM-PHD filter with inception step only; *Col* stands for FI-GM-PHD filter with the coalition step only; *FSys* stands for the full FI-GM-PHD system; *(P)* stands for a situation, when a subset of robots is not communicating.

Scenario IV: Final Demonstration

The goal of the final experiments is to showcase the nominal situation for which our method is targeted, namely when communication has temporary problems that can be overcome by taking advantage of the tracking. Four robots move on a circular trajectory. At $t_1 = 10\ s$, one of the robots loses communication for 20 s. At $t_2 = 50\ s$, there is no communication between any pair of robots, for a period of 10 s. Finally, at $t_3 = 110\ s$, one robot loses communication for 10 s. The total length of the run is 120 s, and we perform 10 sequential runs. A i) perfect communication with dynamic role assignment is the baseline for comparison. We study the performance of ii) the standard GM-PHD filter and iii) the full FI-GM-PHD system.

8.6.2 Results

In the following section we use the following acronyms for labeling the methods. *Std* stands for the standard GM-PHD filter; *Inc* stands for FI-GM-PHD filter with inception step only; *Col* stands for FI-GM-PHD filter with the coalition step only; *FSys* stands for the full FI-GM-PHD system; *NT* stands for experiments where robots do not perform tracking, but use only communicated data; *RA* stands for experiment where dynamic role assignment is used (indicated only for *NT*); *(P)* stands for a situation when a subset of robots is not communicating. Videos of the experiments are available at the link provided in the footnote[5].

5 https://www.epfl.ch/labs/disal/research/InstitutionalRoboticsFormations

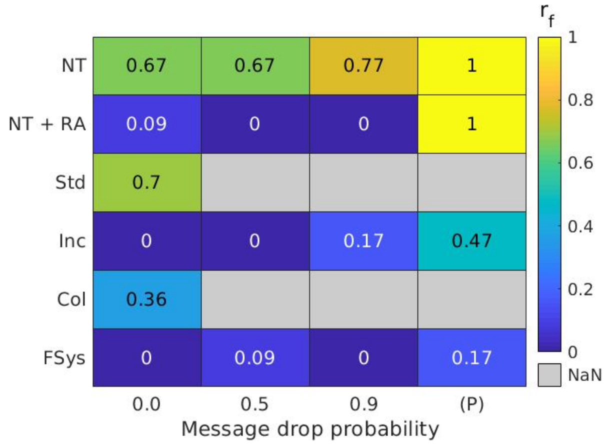

Figure 8.7 – Scenario II. Failure rate r_f. *NT* stands for experiment where robots do not perform tracking, but use only communicated data; *RA* stands for experiment where dynamic role assignment is used (indicated only for *NT*).

Scenario I: Multi-Robot Tracking

Figure 8.6 shows the average OSPA results calculated between the ground truth and the estimates of the five robots. The results suggest that the standard GM-PHD filter and the *Col* method achieve the best performance, irrespectively of whether the other methods use additional data. This is counterintuitive, as one might expect that when communicated self-localization information is added to the tracking data, this additional source of information should improve the estimates. Since the poses shared by the robots and used in the methods *Inc* and *FSys* are the self-localization poses, this data inherently incorporates the self-localization error of the tracked robots (in addition to the self-localization error of the tracking robot), so the above methods naturally perform worse than the methods that do not incorporate the error. However, even though according to the ground truth the tracking data in methods *Inc* and *FSys* appears to be inferior when communication is sparse (with high message drop probability and when not all the robots communicate), in reality the robots perform formation control based on their own self-localization data and have no access to ground truth, neither of their own nor of the neighboring robots. Therefore, formation control based on estimates that include the self-localization data in practice does not perform worse than if it relied purely on tracking. Moreover, in Scenario I the robots do not rely on the tracking data for navigation, but only passively collect it. The FI-GM-PHD filter has been designed to be used for active use in formation control, where collective motion control and tracking have mutual influence on each other. In the ensuing experiments, it will be shown that the proposed approach exceeds the performance of the original method in its designated applications.

Table 8.1 – Scenario II: OSPA metrics

	Std	**Col**		**FSys**		
p_{md}			0.0	0.5	0.9	(P)
OSPA mean	0.4	0.4	0.20	0.22	0.26	0.32
OSPA std	0.07	0.3	0.05	0.05	0.10	0.03

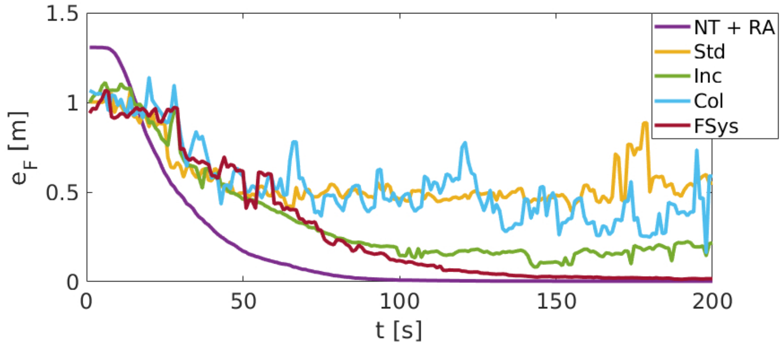

Figure 8.8 – Scenario II. Formation error.

Scenario II: Initialization

Figure 8.7 shows the failure rates r_f for each of the tested methods. An experimental run is labeled as *failed* if the robots are not capable of converging to a stable formation state, caused by the lack of data or collisions. We consider a formation to be stable when the formation error falls below a pre-established threshold and attains a steady state. If the robots are not capable of converging to a formation after initialization, or if the formation breaks after convergence without external disturbance, the formation is considered unstable. It is clear that when robots do not use tracking (*NT*, *NT+RA*) and some of the robots are not communicating *(P)*, the robots always fail, hence $r_f = 1.0$. The lack of dynamic role assignment in *NT* leads to frequent collisions, as the robots must shuffle when moving to designated places in the formation. In general, methods *Inc* and *FSys* perform at least as well as the baseline *NT*, but they also allow for formation initialization even when not all the robots communicate *(P)*. Additionally, using the formation prior in *FSys* leads to higher chance of success than not using it in *Inc*. Using the formation prior but no communication in *Col* leads to slightly worse performance than *FSys*, with $r_f = 0.36$, however *Col* still significantly outperforms the standard GM-PHD filter (*Std*) in terms of its success rate.

The tracking performance, summarized in Table 8.1 shows that communication always improves tracking in the FI-GM-PHD method (*FSys*) as compared to when communication is not included (*Std* and *Col*). The formation convergence as well as the precision of the formation can be deduced from Figure 8.8, which shows the average time-wise evolution of the formation error e_F for $p_{md} = 0$. The convergence rate is slower for the tracking methods (*Std*, *Inc*, *Col* and *FSys*) than for the baseline *NT* methods. Among the tracking methods, the steady state with $e_F \sim 0$ is only achieved by the full FI-GM-PHD (*FSys*), while lack of communication (*Std* and *Col*) leads to slight deformation of the final formation shape.

Scenario III-A: Challenging Environments The trajectories of the robots using the FI-GM-PHD tracker in a challenging environment can be seen in Figure 8.9. Even though a comparison between the FI-GM-PHD filter and the GM-PHD filter was not the aim of the experiment, the tests performed with the standard GM-PHD filter give a better insight as of why a tracking-

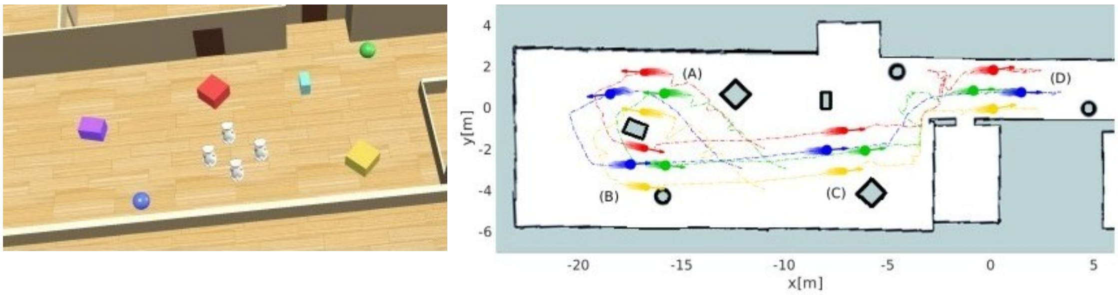

Figure 8.9 – Scenario III-A. Simulation screenshot (a). Trajectories of the robots at $t = 62$ s (A), $t = 114$ s (B), $t = 190$ s (C) and $t = 270$ s (D).

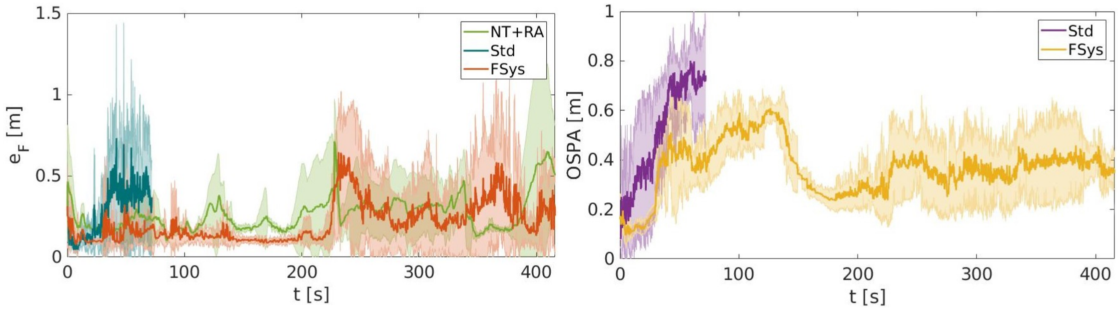

Figure 8.10 – Scenario III-A. Formation error (left) and OSPA (right) performance measures.

only method (*Std*) is not sufficient for robust navigation of multi-robot formation in realistic environments. Figure 8.10 shows that for the GM-PHD filter (*Std*) the e_F and OSPA metrics on average keep rising until the point when the formation breaks apart at around $t = 80$ s (caused by the robots losing track of the leader or drifting apart form each other until they move out of field of view). This is not the case for the FI-GM-PHD (*FSys*) method, which enables periodic correction of the tracker with the communication data, even when only a subset of robots communicates. Note that the presence of obstacles that are close to the path of the robots forces a deformation of the formation shape, which in turn leads to lowering the weights of the filter components that stem from the formation prior, so the coalition step of the FI-GM-PHD has a small effect on improving the tracking performance. Therefore in challenging environments the FI-GM-PHD filter with no communications is expected to perform very similarly to the standard GM-PHD filter.

Scenario III-B: Measurement Error Results summarized in Table 8.2 indicate, that even though larger measurement error has a negative impact on the performance of the FI-GM-PHD filter, the method is still capable of providing good enough estimates to sustain the robot formation even with the measurement error of up to 1 meter. Precision of the formation, expressed by the formation error e_F, deteriorates steadily with the increase of the measurement error e_M, however it remains within reasonable bounds of less than half of robot radius. There is a similar trend in the tracking performance. For $e_M = 1\ m$ the robots, however, begin to have a large offset from the desired places in the formation, so the likelihood of collision increases, up to the point where for $e_M > 1\ m$ the formation is no longer guaranteed. One can also notice that even though the values of OSPA in Table 8.2, in the range $\sim (0.7-0.8)$ are significantly larger than in Scenarios I to III-A, the robots are still capable of maintaining the formation based on the estimates.

Table 8.2 – Scenario III-B: Metrics

	NT+RA	**FSys**				
$e_M[m]$		0.0	0.3	0.6	0.75	1.0
e_F mean	0.2	0.27	0.32	0.36	0.32	0.34
e_F std	0.03	0.1	0.1	0.1	0.16	0.1
OSPA mean	-	0.7	0.67	0.76	0.77	0.83
OSPA std	-	0.04	0.04	0.05	0.07	0.05

Scenario IV: Final Demonstration

Similarly as in Scenario III-B, the GM-PHD tracker alone is not capable of providing robots state information reliable enough for maintaining a formation for extended amount of time. Even in an environment free of obstacles, rotation of the formation and the associated temporary loss of neighbor tracks causes divergence in the local robot perception of the optimal assignment, which when uncorrected, leads to instability of the formation (see Figure 8.11, $t = 42\ s$). The performance of the FI-GM-PHD method (*FSys*) remains stable even when one robot is disconnected and when none of the robots communicates. The tracking error, shown in Figure 8.11, stays around a constant value, irrespectively of the communications status. The formation error in Figure 8.11 is comparable to the baseline error of formations that rely on perfect communication (*NT+RA*). Slightly divergent behavior between *NT+RA* and *FSys* at times $t = 80-120\ s$ occurs due to the fact that during this part of the experiment the robots are navigating closely to a wall. The *NT+RA* method leads to higher temporary formation distortion, while the *FSys* method results in the formation staying closer to the desired shape, but taking longer to recover.

8.6.3 Discussion

Highly dynamic scenarios with multiple robots moving nearby in a coordinated fashion are decidedly challenging for tracking. While the multi-target tracking methods work well with well-separated targets and reliable measurements [60], long-term occlusions, convoluted tracks, and sensor-induced clutter hardly distinguishable from the real tracks cause deterioration of tracking performance up to a point, where the reliability of estimates is insufficient for maintaining a formation. Structured indoor environments, where the robots must navigate along complex paths and around obstacles, increase the difficulty of the problem even further.

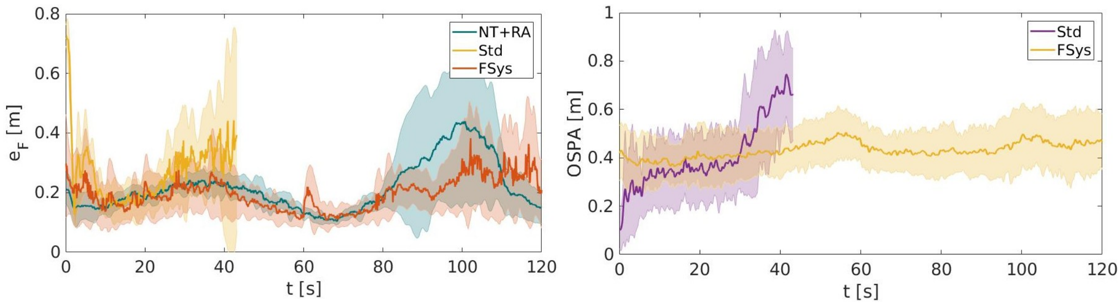

Figure 8.11 – Scenario IV. Formation error (left) and OSPA (right).

Furthermore, in case the estimates must be matched with the formation roles, and when the role assignment is necessary, any tracking errors are escalated. If the assignment is inconsistent among the robots, the different perception of the overall formation state may not allow consensus methods to converge. For this reason only a subset of robots can maintain formation with tracking-only data for longer periods of time. With a small numbers of robots, the complexity of association and probability of wrong assignment is low, therefore a larger proportion of non-communicating robots is allowed. With an increasing numbers of robots, this proportion becomes smaller. Therefore for groups of robots larger than the size demonstrated in this chapter, one should consider anonymous (ID-less) formations (e.g., potential fields), where each robot keeps a constant distance from each estimated neighbor irrespective of the identity of that robot. Such methods are much less prone to tracking errors, at the expense of the precision of the formation shape.

Incorporating communicated data in the tracking filter can provide reliable data used for reinforcing the existing targets or for adding targets that cannot be tracked using other means. In our work [4], the communicated data is included as intensity in the prediction step (analogously to the birth intensity), while in this work the communicated data is integrated as measurements in the update step. We compared the two approaches in a series of tests (not included in this manuscript), the results of which show that the latter systematically yields moderately better performance. More importantly, the latter method is simpler and more intuitive. It does not require complex parametrization that depends on the experimental setup, as it is the case with the first method. Irrespective of the method, when adding communicated data in the filter it is important to adapt the selection step of the PHD filter to discourage double-counting of the targets and prevent inconsistencies in the targets weights.

Finally, incorporating formation geometry in the tracker might lead to worse performance than if such information was not used. While the posterior enforcement stage of the coalition step (step 1-4 in Algorithm 2) is robust to challenging scenarios, and in the worst case it does not improve performance of tracking, the novelty stage (step 6 in Algorithm 2) can create a virtual robot in place of a robot that singles out from the formation. This can happen when all the robots except one are close to desired places in the formation, and is common with leader-follower formations, when the leader moves too fast for the formation to converge. One

Figure 8.12 – Representative pictures of experiments in Scenario II-A with $p_{md} = 1.0$ (left) and Scenario III (right).

possible solution to this problem is to reinforce the follower-leader edges with higher weights. However, in some environments characterized by high clutter, where the followers might get stuck behind an obstacle, one might have to consider disabling the novelty step. Using formation geometry is particularly advantageous in situations where a subset of robots cannot communicate and measurements are unreliable, but the formation is capable of physically maintaining a close-to-desired shape. In such instances, the coalition step can determine whether the formation continues to function or if it fails.

8.7 Real Robot Experiments

As our objective is to evaluate the robustness of the methods that were first tested in high-fidelity simulation, we keep the parametrization used in the simulated experiments of Section 8.6, with the notable exception of calibrating the sensor model according to the empirical data (the details are given in the sections that follow).

8.7.1 Self-Localization and Measurement Errors

The performance of our methods is affected by two sources of stochasticity. First, the self-localization error e_L, which is included in the formation projections and in the positioning information communicated by the robots. Second, the measurement error e_M, which is independent of e_L and affects the sensory data. Before evaluating our FI-GM-PHD method in reality, we carried out a series of tests to understand to what extent these two errors may affect the performance of our system.

The self-localization error e_L is the difference between the self-localization positions and ground truth data from the MCS. To calculate the e_L in test T_I we move a robot around the arena for 960 s and average the results.

The measurement error e_M is the difference between the estimated position of detected robot and its actual position. In our system the error is higher in dynamic situations, where both the detecting and the detected robots are moving [5], therefore to determine e_M in test T_{II} we moved two robots independently, keeping them within sensing radius, with the range and the bearing between the two varying throughout an experiment that lasted for 960 s.

The results are summarized in Table 8.3. The self-localization error can be seen as insignificant given the size of the robot with the diameter of 0.65 m. The measurement error is on average close to the size of the robot radius, posing an additional challenge of substandard robot sensing that has a direct effect on the tracking performance.

e_L		e_M	
mean	std	mean	std
0.18	0.051	0.33	0.24

Table 8.3 – The self-localization error e_L and the measurement error e_M of our setup, determined empirically.

8.7.2 Model-Based Probability of Detection

In Section 8.3.2, we explain how the probability of detection p_D reduces the risk of loosing a track. For this purpose, in our model we integrate the sensor FOV and occlusions. For real robot experiments, we additionally model sensor-dependent probability of missing a detection and incorporate it in p_D.

Our model is based on empirical data collected in test T_{II}. The model characterizes specific sensors, therefore it is different for each robot. Recall that the MBot robots are equipped with two LRFs. The sensors, each of them providing 240° field of view, are located at the front and at the back of the robot, while on the sides their ranges overlap. The overlapping however is skewed, which results in higher probability of detection loss around the angles $-\pi/2$ and $\pi/2$. To the resulting distributions, we fit Gaussian models using the curve_fit method from SciPy optimize[6]. The data and the models fitted to it are shown in Figure 8.13. The spikes indicate the portion of the lost detections $p_{D,s}$ for a given angle. The resulting probability model is added directly to the probability of detection p_D of each robot.

8.7.3 Scenarios

The FI-GM-PHD filter is evaluated in three scenarios: (I) tracking decoupled from formation control, where robots do not use the tracking data for control, (II) tracking for formation control, where we alter the quality of communication and simulate augmented detection error, and (III) in a realistic scenario, where robots navigate among obstacles scattered in the environment. Our methods are compared with the standard GM-PHD filter and with respect to the baseline formation control with full communication an no tracking.

6 SciPy library, https://www.scipy.org

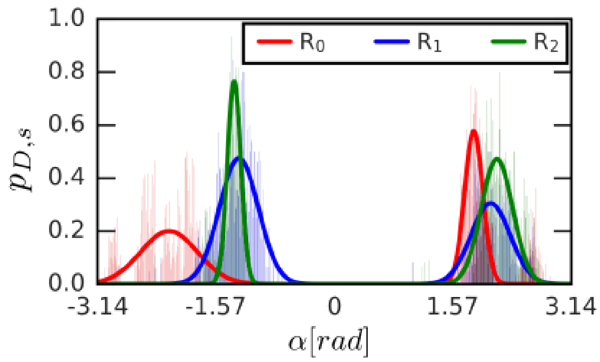

Figure 8.13 – The model of sensor-dependent missed detection probability determined empirically for each robot.

Scenario I: Multi-Robot Tracking

We collected a dataset (raw sensor data, positioning information and the formation state) in baseline experiment with formation relying on regular WiFi communication. The formation followed an eight-shape trajectory and involved three robots forming a triangle shape with the inter-robot spacing of 1.75 m. We performed multi-robot tracking with the collected data *offline*, with the standard GM-PHD filter, and with the FI-GM-PHD method with simulated message drop probabilities of $p_{md} = 0.0$ (i.e. regular communication), $p_{md} = 0.5$, $p_{md} = 0.9$ and $p_{md} = 1.0$ (i.e. no communication). For each experiment, we performed 11 sequential runs, each lasting 120 s.

Scenario II: Tracking for Formation Control

In contrast to Scenario I, in the following experiments tracking is run *online*, and used for formation control directly. In other words, the performance of the tracking system affects the formation control efficiency, which in turn has an effect on tracking through the coalition step. We distinguish two sub-scenarios:

Scenario II-A: Message Drop Probability where we simulated message drop probabilities by varying $p_{md} \in \{0.0, 0.5, 0.8, 1.0\}$ and compared to a regular communication baseline case.

Scenario II-B: Measurement Error where we altered the precision of the robot detection by adding a random uniform error to the original measurement, with the magnitude of $e_M = \{0.0, 0.5, 0.9, 1.0\}$ m. The probability of message drop was $p_{md} = 0.0$ so as to decouple the effects of communication quality and the sensing factors. The experimental settings, including the number of robots, the desired formation shape and the prescribed trajectory were identical to those of Scenario I.

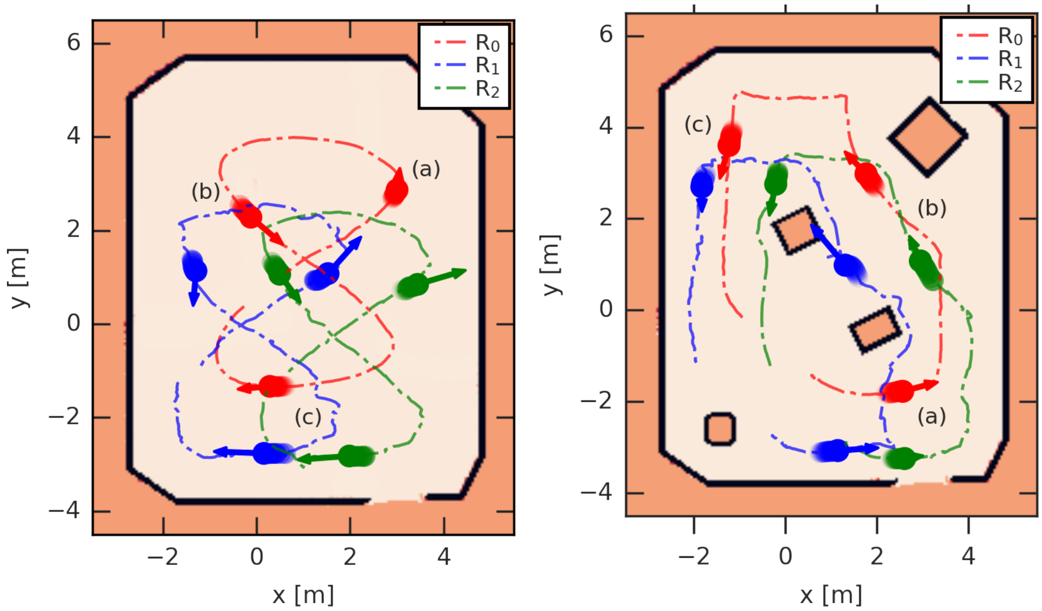

Scenario II-A. Scenario III.

Figure 8.14 – Trajectories of the robots using the FI-GM-PHD filter with $p_{md} = 1.0$, i.e. with no communication. For Scenario II-A trajectories of the robots are given at (a) $t = 15$ s, (b) $t = 55$ s, (c) $t = 95$ s. For Scenario III, the trajectories are given at (a) $t = 15$ s,(b) $t = 45$ s,(c) $t = 75$ s.

Scenario III: Challenging Environment

In the final set of experiments, the robots move in a triangular formation with the inter-robot spacing of 1.6 m in the arena scattered with obstacles. The leader robot planned the trajectory using a FMM [112]. For each experiment, we performed 11 sequential runs of approximately 100 s. We perform a perfect-communication baseline experiment and successive runs with varying communication quality characterized by $p_{md} \in \{0.0, 0.5, 0.8, 1.0\}$.

8.7.4 Results

We use the following acronyms for labeling the methods. *NT* stands for the baseline experiments with the formation relying on regular communication and no tracking, *Std* stands for the standard GM-PHD filter and *FSys* stands for the full FI-GM-PHD system. Videos of the experiments are available at the link provided in the footnote[7].

Scenario I: Multi-Robot Tracking

The OSPA performance is summarized in Table 8.4, from which we draw two conclusions. First, the tracking performance of the FI-GM-PHD filter degrades gracefully with the drop of

7 http://disalw3.epfl.ch/research/alicja/Chapter_8-7.mp4

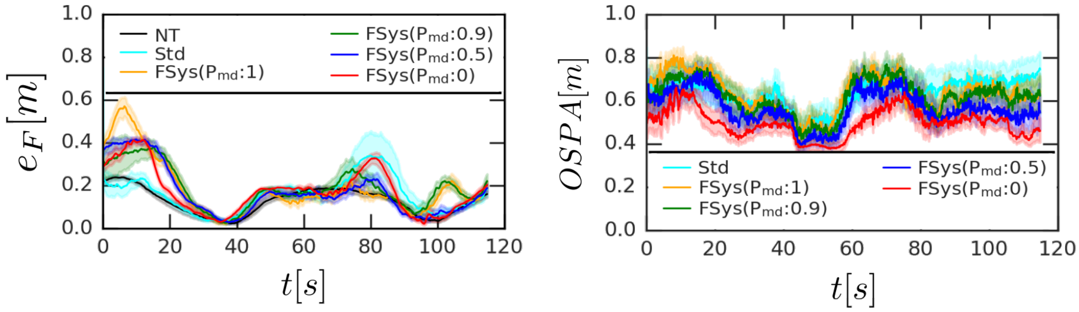

Figure 8.15 – Scenario II-A: OSPA and formation error e_F.

the communications quality. Compared to when the positioning data is received at 10 Hz, in the case of no communication the performance of FI-GM-PHD method is only 37% worse. Second, in the case of $p_{md} = 1.0$, i.e. where no communication occurs, the FI-GM-PHD filter outperforms the standard GM-PHD filter. This is a fair comparison, as both methods rely on the same data, but the FI-GM-PHD filter performs an additional coalition step.

Scenario II-A: Message Drop Probability

The formation error, shown in Figure 8.15, remains bounded for all the tested cases. It oscillates between as low as 0 m and up to 0.4 m, with a short-term peak in the *FSys* and $p_{md} = 1$ case. Higher values of e_F are resulting from the fact that during part of the experiment the leader robot is situated behind the followers, and the "pushing" forces it exerts have a smaller effect than the "pulling" ones (they act against the follower-to-follower forces, not with them). For the majority of the run duration the formation error of all the methods follows that of the *NT* baseline.

Shown in Figure 8.15, on average the OSPA error is the lowest for the *FSys* method and it gracefully degrades with the reduction of the communications throughput. As summarized in Table 8.4, the rise of the OSPA error with respect to the p_{md} is moderate, with the difference between the $p_{md} = 0.0$ and $p_{md} = 1.0$ amounting to 27%. This confirms the results we obtain in Scenario I, but in the case where the tracking is performed online. On average, the OSPA error of the FI-GM-PHD with no communication ($p_{md} = 0.0$) is almost identical to that of the standard GM-PHD filter. However, during our experiments, the *Std* method resulted in *3 formation failures* out of the total of 11 runs. A run is labeled as failed when at least one of the robots stops keeping the formation with the other robots and falls behind. This phenomenon is typically caused by a lost estimate, an estimate fixed to an object in the area, mistaken role association or a combination of the above. No failures occur in the *FSys* case, even when no communication is allowed. An example of a trajectory of the FI-GM-PHD filter in the $p_{md} = 1.0$ case is shown in Figure 8.14.

Table 8.4 – OSPA metrics for Scenarios I-III.

Scenario I					
	Std	**FSys**			
p_{md}		0.0	0.5	0.9	1.0
OSPA mean	0.63	0.41	0.49	0.56	0.57
OSPA std	0.24	0.11	0.16	0.18	0.19
Scenario II-A					
	Std	**FSys**			
p_{md}		0.0	0.5	0.9	1.0
OSPA mean	0.64	0.50	0.57	0.62	0.63
OSPA std	0.21	0.14	0.18	0.19	0.19
Scenario II-B					
	FSys				
e_M	0.0	0.1	0.3	0.6	1.0
OSPA mean	0.50	0.48	0.49	0.52	0.55
OSPA std	0.14	0.14	0.14	0.14	0.14
Scenario III					
	Std	**FSys**			
p_{md}		0.0	0.5	0.9	1.0
OSPA mean	0.74	0.53	0.63	0.67	0.67
OSPA std	0.23	0.16	0.20	0.20	0.20

Scenario II-B: Measurement Error

Based on the results summarized in Table 8.4 we can deduce that once the communication quality is marginalized, the measurement error has little effect on the performance of our method. Recall that our preliminary evaluation determined the baseline detection error of our setup with two LRF to be around 34 cm. The addition of a random uniform error of less than that value (as in the $e_M = 0.1$ and $e_M = 0.3$ cases) has no effect on the tracking performance, while injection of the error as high as 1 m (one and a half times the robot diameter) results in 14% decrease of OSPA compared to the $e_M = 0.0$ case, confirming robustness of our methods to sensory imperfections.

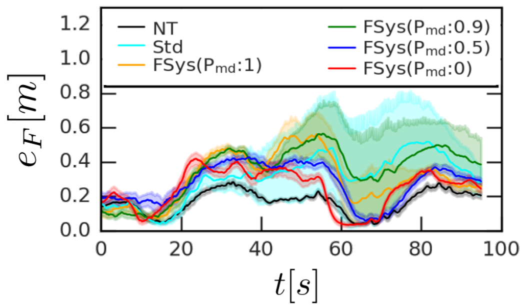

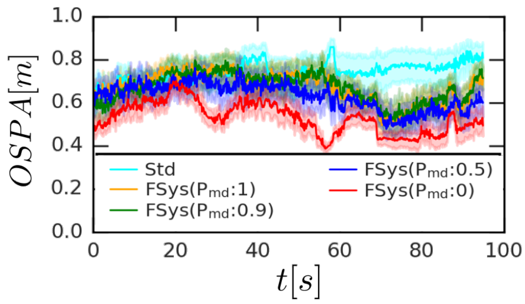

Figure 8.16 – Scenario III: OSPA and formation error e_F.

Scenario III: Challenging Environment

The experimental setup with the obstacles scattered around the arena and the robot trajectories recorded during one run of the FI-GM-PHD filter with $p_{md} = 1.0$ is shown in Figure 8.14. The OSPA metrics, plotted in Figure 8.16 and summarized in Table 8.4 once more confirm stability of the tracking performance of our methods, even when the formation experiences deformities resulting from the environmental factors. Once more we observe the trends recognized in Scenario I and Scenario II-A, namely that the increase of p_{md} has a bounded effect on the quality of tracking (with the OSPA in the $p_{md} = 1.0$ case being 27% worse than in the $p_{md} = 0.0$ case) and that the FI-GM-PHD filter outperforms the standard filter even in the case when communication is not used (with OSPA of *FSys*, $p_{md} = 1.0$ being 10% lower than *Std*).

The formation error shown in Figure 8.16 remains close to the *NT* baseline, with the exception of the *Std* and the *FSys* with $p_{md} = 1.0$ methods, the variance of which rapidly increases around $t = 50$ s, at the time where both methods experience formation failures. Out of all the tested cases, the *Std* and the *FSys* with $p_{md} = 0.9$ and with $p_{md} = 1.0$ each resulted in failure to maintain the formation in one out of 12 runs. Each of these methods incorporates very little (one message per second) to no communication. Once a robot falls slightly behind during a maneuver of obstacle avoidance such robot has no means to recover if the obstacle occludes the other formation members, while the impact of including the formation geometry is reduced because of the actual formation is drifting from the desired set point.

8.7.5 Discussion

The presented results consistently lead us to two conclusions. Firstly, our methods are robust to the deterioration of communications quality (Scenario I and II-A), sensory imperfections (Scenario II-B) and the environment complexity (Scenario III), with the tracking performance degrading gracefully with the increasing levels of experimental difficulty. Second, our FI-GM-PHD method outperforms marginally (Scenario II-A) or significantly (Scenario I and III) the standard GM-PHD filter, even in the cases when no communication occurs. One should note that although it may seem that the FI-GM-PHD filter has obvious advantages over the standard filter, as it combines data from multiple information sources as opposed to the single one used in the canonical GM-PHD filter, achieving effective fusion is nontrivial because of the inconsistencies introduced by the self-localization (incorporated in the communicated positioning information) and the detection errors. Fusion, if done inappropriately, can result in track splitting and ambiguity of estimates, which in turn can lead to erroneous role assignment, ill-defined formation, and finally, breaking of the formation. The GM-PHD filter facilitates fusion of data from multiple heterogeneous sources, but care must be taken so that the advantageous properties of the original method are not sacrificed.

The primary objective of the experiments above has been to validate the FI-GM-PHD filter with real robots and within a realistic environment. We conclude that our methods prove to be highly robust, and do not require fine-tuning when moved from simulation to reality (recall that we do not perform re-parametrization except for updating the sensor-dependent probability of missed detection).

Through the experimental validation not only did we prove the robustness of the FI-GM-PHD method, but we tested it in settings more challenging than what the filter has been originally designed for – situations where communication suffers from short term outages. The FI-GM-PHD is shown to be able to successfully sustain the formation even in situations of total communication loss, keeping the probability of formation failure marginal even in environments cluttered with obstacles.

Summary

In this chapter we have presented a strategy for providing reliable and robust robot state estimates to be used for formation control when the communications throughput is low or even when communication fails. For safety and acceptance reasons, such backup system is necessary for establishing cooperative multi-robot navigation in human-populated environments.

Our method combines absolute positions exchanged by the robots, information about the formation geometry and sensory detections in an extension of the GM-PHD filter. The experiments performed in a high-fidelity simulator and with real robots demonstrated that our approach, the FI-GM-PHD method, is capable of maintaining the state estimates even when long-duration occlusions occur, and improves awareness of the situation when communication is sporadic or suffers from short-term outage. Moreover, the results have confirmed that the proposed tracking strategy allows for sustaining formations in cluttered environments, with high measurement uncertainty and low quality of communication. We have studied the limitations of the method when the spatial configuration of the robots is far off from the desired formation geometry, including initialization and cluttered environments. The proposed method not only outperformed the standard tracking, but also proved comparable to methods relying on perfect communication.

Our final validation of the FI-GM-PHD filter is presented in Chapter 15. In experiments governed by institutions a group of robots shares an environment with humans and modifies the formation shape in order to respect social norms. Although the high dynamics of the experiments caused by the presence of a human and a purposeful alteration of the formation geometry to meet social norms present additional challenges to our method, we show that complementing faulty communication with tracking through a FI-GM-PHD filter significantly reduces the chance of formation breaking and the probability of mission failure.

Institutions and Norms Part III

When things are simple, fewer mistakes are made.
The most expensive part of a building is the mistakes.

Ken Follett, The Pillars of the Earth

9 Introduction

All human interactions, in both social and economic life, depend on some sort of trust in the behavior of the others. Trust is based on a social order, which inspires reliance and confidence, and where existing rules of conduct ban unpredictable, erratic behaviors and suppress opportunism [113]. Social order – a predictable and comprehensible pattern of human actions and reactions – is achieved through *institutions*, while the core of each institution relies on *social norms*. To become integrated into our society and effectively engage in human-robot interactions, robots need to be aware of the norms and able to act according to the expectations that such norms create. It is believed that better understanding of the social norms can lead to higher acceptance of robots in our everyday lives [114], while conformance to human expectations through normative behaviors can enhance human-robot interactions and human perception of cognitive capabilities of the robots.

The norms of human societies are formulated in human language, and it is the language that every human comprehends, irrespective of the age, occupation or culture. However, norms stated in a form of the human language are inherently abstract and open to interpretation. The abstraction implies generality that allows for applying norms in a variety of situations, but if the robots are to adhere to social norms, they must be capable of translating abstract norms to the robotic language.

In this chapter we introduce the institutional and norm-related concepts and discuss the challenge of dealing with generic norms.

9.1 Institutions and Norms as Essential Elements of a Society

Institutions sit at the the very foundations of our living standards and our sense of security and community. There are many different forms and types of institutions. The general simple understanding is that institutions are formal, legally-bound structures, but the notion of institution studied across the disciplines including philosophy, sociology, economics, politics, and legal systems encompasses a much broader concept. Language, money, law, table manners, religion and family – those are all systems of social order, and so, also institutions. Such institutions create prescriptions – *norms* for social behaviors, facilitating decision-making in unknown situations and regulating how humans communicate, act and cooperate during social engagements. They reduce the cost of coordinating human actions by imposing form and consistency on human activities and creating stable expectations of the behavior of others [115].

In view of the above description, every social interaction in human societies is steered by some sort of social norm embedded in our minds since the early childhood. Similarly, social norms and institutions can be identified in the behavior of gregarious animals. Social insects such as ants and bees prescribe each individual an institutional role with its associated expectations. And so, every individual knows its mission – the worker ant is to collect seeds, the bee queen is to lay eggs, while the social norms embedded in the animals allow the other individuals to believe that the queen will not leave the apiary and that the worker will bring the food to feed the minors. This mutual understanding of the institutional environment assures social order and stability, and facilitates cooperation at the most fundamental level. One can conclude that norms (and institutions that are build around them) are essential for all stable societies, irrespective of the cognitive capabilities of the individuals that form them.

9.2 What are Norms? What are Institutions?

Social norms represent the common understanding that govern the behavior of the members of a society. As such, they regulate communication, cooperation and other social interactions. The definition of norm varies among the different areas of study such as sociology, game theory, psychology, and legal theory. In social sciences, norms are informal rules and behavior standards that are shared among the members of the society. They constrain social behavior without the force of laws and are considered valid by the majority of a social group [116]. Norms represent desirable behaviors for a population and indicate actions that are expected to be pursued that are either obligatory, prohibitive, or permissive given the situation [117].

Institutions are mechanisms for reducing uncertainty, simplifying decision-making and promoting cooperation [118]. To be effective, they should be simple, certain, abstract and reasonably stable [113]. Institutions are a product of deliberate design (law or political systems), arise spontaneously as a set of informal norms on the basis of self-interest of individuals (table manners) or are a combination of the above. They preside over the individual and collective behaviors by obliging everyone to act according to the norms.

9.3 The Issue of Norm Interpretation

In sociology, norms define the behavioral expectations within a society. A norm is a general rule of conduct forming a link between the abstract values of the group and the concrete behaviors that are to respond to such values [119]. As such, norms are the explicit or implicit rules that a group uses to determine appropriate or inappropriate beliefs, attitudes, and behaviors [120]. As stressed by [121], the concept of social norm means different things to different scholars, and since there is no common definition of social norms, there can be little agreement about how to formalize them. For this reason we focus on the most general form of norms – norms formulated in *human language*. Nonetheless, norms of human societies formulated in human language allow for a certain maneuvering margin in terms of interpretation, or even, as stated in [118], they "share problems of lack of clarity, misunderstanding, and change that typify any language-based phenomenon". Even for humans, interpreting norms is more challenging than writing them down. For robots that operate on robot-understandable commands, interpretation of norms defined in human language is close to impossible.

Summary

Social order – a predictable and comprehensible pattern of human actions and reactions – is achieved through *institutions*, while the core of each institution relies on *social norms*. It is believed that robots capable of reasoning about social norms are more likely to be recognized as an element of the human society. Nonetheless, behaviors of such robots and the norms they follow must be understandable. Norms in human societies are defined in human language, and are therefore inherently abstract and interpretable. If robots were to adhere to such human-understandable norms, it is necessary that they have means to interpret them. However, for robots that operate on robot-understandable commands, interpretation of norms defined in human language is close to impossible. In the next chapter, we review the literature addressing the issue of norm interpretation across the disciplines, and discuss how our institutional formalism draws inspiration from these approaches to provide the robots with the means to interpret abstract, language-defined norms.

10 Related Work

THE concepts of norms, institutions, and methods for achieving social order are interpreted, analyzed, and modeled across the domains. While most of the methods are disconnected from physical execution and, as a consequence, cannot be easily applied to robotic scenarios, they provide valuable insights into the design of normative frameworks. In this chapter, we describe how the institutional and norm-related concepts are viewed in different areas of study, ranging from social sciences, through economics, to computational multi-agent frameworks and robotics. How our work is situated with respect to these studies is illustrated in Figure 10.1.

10.1 The Economics View

The definition of institution is an object of discussion among the social scientists, where numerous interpretations and alternative meanings do not seem to converge to a generally valid definition, other than the institution having a loose association with regularity of behavior [113]. An ongoing debate within the new institutional economics discusses whether institutions should be regarded as equilibria, norms, or rules, while the different terminology appearing across the disciplines is compared to the Tower of Babel [118].

10.1.1 Institutional Approaches

One interpretation declares institutions as rules with sanctions, that have normative influence on human behavior, constrain arbitrary and opportunistic actions and structure interactions [122][113]. An alternative, and more general view, defines institutions as systems of established and prevalent social rules, rather than rules as such [115]. In short, institutions are systems of social rules, not simply rules. As systems, institutions include enforcement arrangements [116], specify the roles, information shared by the participants, allowable actions and their outcomes as well as costs and benefits of participation [118].

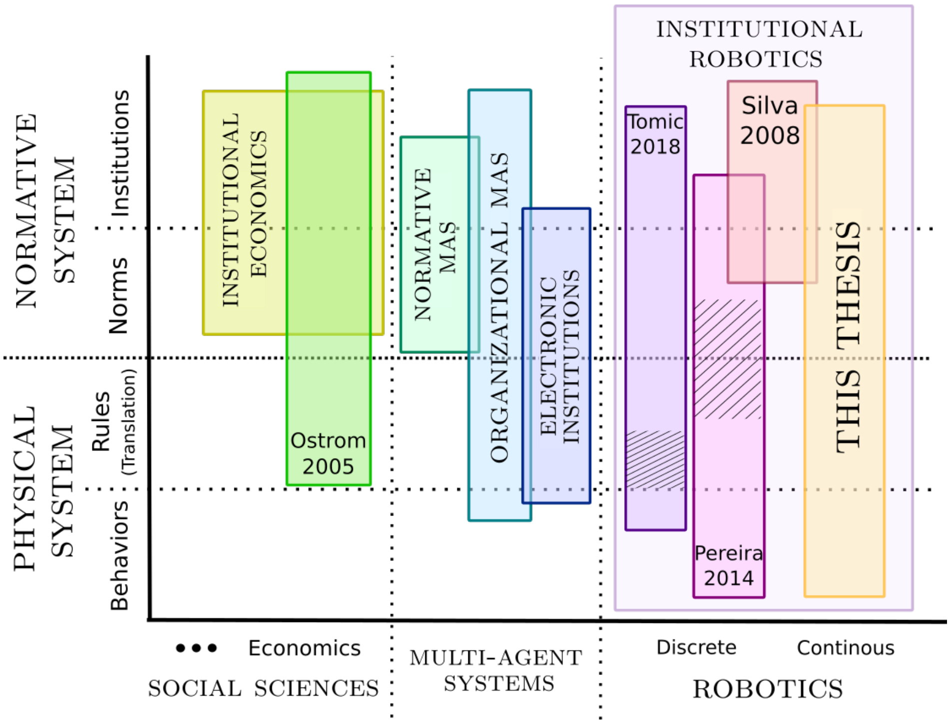

Figure 10.1 – The situatedness of our work with respect to the studies ranging from social sciences, to computational multi-agent frameworks and robotics. The dashed areas indicate lack of representation of the given level.

The common view is that modeling and understanding of a social framework with all its relevant variables and the their immense number of combinations, all existing at different levels of abstraction, is a complex endeavor [118].

Attempts at developing a general modeling framework can prove a valuable inspiration for designing socially aware robots. Social frameworks proposed by economic studies identify the universal elements and the relationships that one needs to consider for institutional analysis and provide meta-theoretic languages to support interpretation of social interactions. Conventional approaches, however, tend to focus on game-theoretic models or utility-driven decision-making in an isolated context [119]. A notable exception is the Institutional Analysis and Development (IAD) framework proposed by Elinor Ostrom [118], which embraces the complexity of real-world situations with its countless variables and contexts within contexts and makes precise assumptions about a limited set of parameters and variables.

10.1.2 Norm Abstraction and Interpretation

In economics, the issue of norm interpretation is identified at different levels of their operation. In [113] it is stated that rule systems work better in ordering human actions if they form a hierarchy running from general, universal rules, which are often abstract, to specific rules. A definition of rule that hinges solely on behavioral regularities proposed in [123] neglects the ontology of rules and suggest that a rule does not have to be known to the individuals in any other sense than they normally act in accordance with it. In [124] it is pointed out that rules must be codifiable, but part of knowledge can never be fully articulated. A discussion in [116] takes a step forward by implying that gaps in formal constraints are covered to some degree by informal rules and that institutional incompleteness is inevitable, causing problems for the design of behavioral rules for a given institution.

In an attempt on defining norms formulated in human language, the IAD framework in [118] defines the term *institutional statement* – a broad concept encompassing three normative media – rules, norms, and shared strategies. These statements describe a broad set of shared linguistic opportunities and constraints that create expectations about other actors' behaviors and prescribe, permit, or advise actions or outcomes for the group members. A general syntax used for analysis of similar statements includes: ATTRIBUTES – to whom the statement applies, DEONTIC – modal verb, such as must, may or must not, AIM – is the action or action outcome, CONDITIONS – are variables defining when and where an action or outcome is permitted, obligatory or forbidden, and OR-ELSE – are consequences of not following the rules. Elements of syntax are used for distinguishing rules, norms (rules without sanctions), and strategies (norms without deontic operators). Typical institutional statements are: a) "If you use a microwave, you must clean up your own mess!" or b) "The person who places a phone call, calls back when the call gets disconnected." where a) is a norm (it missed the OR-ELSE element), while b) is a shared strategy (it misses the modal verb).

There are no uncontrollable interactions in robotics, as there are no problems with opportunistic behaviors or conflicts over limited resources. In this sense, the purpose of institutions in robot and human societies differ. However, as we will describe in later chapters, the powerful formalization of the IAD framework became a strong inspiration for our definitions of norms and institutions.

10.2 The Multi-Agent Systems View

The field of Multi-Agent Systems (MAS) started showing increasing interest in social theories as the focus of research expanded from the individual agent models to models of socially situated agents.

MAS research has addressed the problem of structuring agent interactions by attempting different approaches: agent communication languages, communication and interaction protocols, teams and coalitions, modeling of negotiations, institutions, organizations and norms [125]. Similarly as the concepts of institutions, organizations, and norms are innately bound, the MAS studies on organizational approaches, electronic institutions, and normative MAS are closely related and overlapping in some aspects. Nevertheless, they provide invaluable tools for analysis of interactions in artificial systems.

10.2.1 Institutional Approaches

In this section, we take a closer look at the the normative MAS and the models of institutions and organizations, where the concepts of organization and institution are used as abstractions for analysis and modeling of cooperation and coordination in agent societies, while the normative systems study the relation between norm-governed agent behavior and macro-level system effects, and how the heterogeneous micro-world of individual behaviors generate the global macroscopic regularities of the society.

Normative Multi-Agent Systems

Normative MAS provide the means to integrate agent mechanisms at both social and individual level, resulting in increased fidelity with respect to modeling complex social phenomena such as cooperation, coordination, group decision-making, and organization, in both human and artificial systems [114].

The key idea behind the normative systems is that individual and collective behavior is affected by norms, which serve to guide, control and regulate proper and acceptable behavior. Normative MAS offers mechanisms to represent, communicate, distribute, detect, create, modify and enforce norms, and mechanisms to deliberate about norms and detect norm violation and fulfillment [126].

Norms can be hard or soft constraints on actions that an agent can perform. In the latter case, norms allow for the possibility that actual behavior may at times deviate from the ideal, i.e., that violations of obligations, or of agents' rights, may occur [127]. This is particularly interesting in robotics, as physical situatedness in non-deterministic worlds does not admit perfect norm compliance.

Normative MAS represent norms with simple data types, such as deontic logic, binary strings, conditions-action pairs (in rule-based systems) and game theory [128]. Simple representation schemes reduce the complexity of the models and their computation requirements and allow for exploration of certain system properties which may not be easily understood otherwise [129] at the expense of realism and power of representation. Because of these properties normative MAS approaches gain popularity in the planning aspect of robotics, where complexity and simplicity of representation are essential to real-time planning, but they lack elaboration necessary for representing continuity and complexity of multi-robot behaviors. As we will discuss in the later chapters, abstract norm formulation and lack of precise connection between the abstract normative statements and any computational model makes it difficult or even impossible to directly connect this kind of norms with the practice [130].

Organizations

Organizational MAS approaches incorporate organizational abstractions into the computing systems. Organizations are viewed either as processes of organizing a set of individuals, or as separate entities, with their own requirements and objectives [130], but the primary goal of organizations is to provide scope of interactions among a set of agents and coordinate their behaviors to achieve some collective goal.

Organizations are conceptualized in terms of their structures, whereby a structure is a set of roles, groups, and links specifying the structure of the system independently from the individual agents. Structure can be explicitly implemented in the form of a social artifact existing independently of the implementation of the agents, may be included in the implementation of the individual agents, or may exist only intangibly, in the form of the behavior patterns exhibited by the collective of agents during interaction [131]. A concrete organization is an instantiation of an organizational structure, where roles are filled in by specific individuals [119]. That is, organizations describe objectives of the society, the roles within, rules of interactions and the coordination protocols [130] without considering the particular characteristics of the individuals involved.

Electronic Institutions

The notions of institutions and organizations in MAS are closely related. At the very basic level of understanding institutions address the question of *what can be done?*, while the organizations address the question of *who does it?*.

Electronic Institutions (EI) [132][133] are the electronic counterpart of human institutions – they establish the expected behavior of agent societies [134]. Institutions are viewed as coordination artifacts external to the agents meant to facilitate agent interactions by establishing an interface between the internal, rational decision-making capabilities of agents and the social effects of their interactions [119].

EI define a controlled environment where heterogeneous agents, humans and software, can interact by means of speech acts [134]. EI define constraints on agent interactions, regulating them by the means of multi-agent protocols, while the individual agents are abstracted by their roles, allowing participants with the same role to be treated collectively.

10.2.2 Norm Abstraction and Interpretation

Studies on normative MAS understand the importance of decoupling norm abstraction from concrete system representation, and study general and domain-independent properties of norms [114]. Existing MAS approaches distinguish several levels of abstraction, starting from an individual agent, its role and the group it belongs to, and concluding with the overall organization level. The HarmonIA framework [135] models organizations from the most abstract level, where the norms are defined, down to the concrete protocols and procedures that implement these norms. OperA [136] distinguishes between the mechanisms that describe and coordinate the global behavior of the model, and the goals and the behaviors of the agents that populate the model. The Organisational Model for Normative Institutions (OMNI) [137] separates three levels of abstraction, with generic organization definition and model ontology at the abstract level, definitions of norms, rules and roles at the concrete level and translation of abstract norms into actions and agent low-level protocols at the implementation level. The ISLANDER framework [138] proposes architecturally-neutral e-institutions, where no particular agent architecture or language is assumed. Computational Organization Theory [139] uses mathematical and computational methods to study both human and automated organizations, with the aim to build concepts and theories about organizations at an abstract level and to develop tools and procedures for the validation and analysis of the models. Other normative MAS frameworks, including GAIA [140], MOISE [141] and MOCA [142] explore the organizational metaphor that promotes both micro-level (agents level) and macro-level (system level) control over the design and understanding of the overall system behavior.

Normative systems and MAS are tightly related with deontic logic. Authors in [114] recognize issues related to representation of norms as domain-dependent constraints, where the distinction between a normative behavior and an actual behavior is being disregarded and where it is not possible to specify that some behaviors are illegal but nevertheless possible except by ruling out these illegal behaviors by specification. The same study emphasizes that norms should be represented using a domain-independent theory, such a deontic logic. Deontic logic provides a means to specify what should happen if illegal but possible behaviors occur by using special modal operators that indicate the status of behavior, namely whether it is legal (normative) or not [143].

Insights from deontic logic can be used to represent and reason with norms and represent norms as rules or conditionals. However, there are several aspects of norms which are not covered by constraints nor by deontic logic, such as the relation between the cognitive abilities of the agents and the global properties of norms. It is hard to directly connect the norms that are formalized in the form of deontic logic with the practice [130]. The reasons listed in [144] are:

- Norms in law are formulated in a very abstract way (vague and ambiguous).
- Norms are declarative and have no operational semantics (they express what is acceptable, but not how to achieve it).
- There is no precise connection between the abstract normative statement and any computational model.

Furthermore, formalization based on semantics and consequences of norms does not indicate how the norm should be interpreted within a certain institution. As pointed out in [145], an action mentioned in the norm statement is far more abstract than the level on which an agent operates on, and that it is very unlikely that the agents functioning within the institution will explicitly have such an action available. In other words, the level on which the norms are specified is more abstract and general than the level on which the processes and structure of the institution are specified. Therefore, according to [145], we need to translate the norms. Moreover, this translation is dependent on the domain (i.e. concrete system) and therefore the translation rules depend on the ontology for that domain. And because the concrete translations of norms do not have a direct counterpart in the institution, another translation is needed on the level of institution to indicate how the norm is implemented. An approach to norm translation proposed in [144] is inspired by how the gap is bridged in human rule-based legal systems, where human laws are expressed in a very abstract way and are hard to use in practice, but where regulations provide interpretation of the law and offer operational constraints to be met in practice. Similarly, in [144] an intermediate level between institutional norm specifications and institutional protocols if formed by the landmarks – sets of states partially ordered in directed graphs, which link the institutional abstraction with the concrete system.

Perspective

The two dominant views of institutions in economics are that 1) institutions are norms and 2) institutions are systems in which the norms are embedded. Our approach follows the second definition - in robotics it is not only necessary to know what are the rules, but also who has to follow them, when and how.

The institutional formalism we present in this thesis is founded on the solutions discussed in this chapter.

First, attempts at modeling and analysis of real world social phenomena undertaken by the economists provide valuable insights into building a framework for normative robots. Models such as IAD proposed by Ostrom [118] encompass a large variety of interactions and social contexts, while admitting that not all system components can be represented or classified. Moreover, by studying actual norms existing in human societies, which are abstract in their nature and defined in human language, Ostrom acknowledges the need of translating abstract norms into concrete system representations. It is for those two reasons our approach, especially the definitions of institutions and norms, is heavily inspired by the work of Ostrom [118].

Second, normative systems propose that norms can be soft constraints and acknowledge that the actual behavior of an artificial agent may at times deviate from the ideal – which is granted in highly stochastic robot environments. In the organizational approaches, abstraction of social structure from the particular implementation of agents and instantiation of the organization by filling the structure with the elements of a concrete system is akin to what we refer in our work as the grounding.

The MAS frameworks however are disconnected from physical execution and cannot easily be applied to robots. They do not provide means for representation of physical world, its physical actions and perceptual capabilities; furthermore, the models cannot encompass the complexity required for social robot behaviors. Consequently, it is impossible to use any of the aforementioned MAS methods for real robot applications as the ones targeted in this thesis.

10.3 The Robotics View

The need for a normative framework in robotics is excellently summarized in [31]:

> "(...) one main challenge in the [human-aware navigation] research area is to unify different methods and solutions technically and semantically. This unification requires a grounding of methods in a semantic context of behavior. Such a context needs to define each spatial human-robot encounter as an interaction following social rules. To find and evaluate such rules, and to map them to suitable software processes remain the main challenge of human-aware navigation planning.".

Furthermore [31] argues that the vast majority of surveyed publications deal with individual domains and challenges, but there are no attempts at developing a holistic theory of human-aware navigation. In this section, we review the related work on Institutional Robotics (IR), in particular research in two distinct robotic domains: in the context of swarm robotics and planning using Artificial Intelligence (AI) methods. Then we broaden our scope to review the state-of-the-art approaches to norm-following robots.

10.3.1 Institutional Approaches

The emergent field of IR [6] is a union of disciplines studying human societies and formalization requirements of robotics. Institutions are introduced as coordination artifacts in multi-robot systems for specifying social interactions among robots and humans and intended to facilitate the integration of robots in human societies. Since the robots controlled using the IR approach abide by the norms of institutional environments created by the humans, the collective performance during human-robot interaction is expected to surpass the existing methods. The human relationship with a robotic system may not be necessarily focused on verbal and gestural communication skills, but rather focused on understanding of the intentions of artificial agents or robots, so that humans no longer need any specific training to interact with them, because when dealing with robots they adopt the same attitudes as when dealing with other humans [146].

The principles of IR address the aspects traditionally neglected in AI, namely relevance of time and chronology, the agent body, the world where the agent is situated, and the other agents.

While distributed robotics approaches make progress toward embodying and situating the agents, the sophistication of their social environment models did not reach their full potential. The IR approach aims to meet this potential by adding the concepts of physically and socially bounded autonomy of cognitive agents, uncoupled interaction among them and deliberately set up coordination artifacts – institutions. Importantly, IR acknowledges the active character of the environment and the unpredictable aggregate effects of multiple simultaneous actions carried out by multiple agents in both physical and social (institutional) environments. Institutions are defined as artificial modifications to the environment that influence the collective order. Humans and artificial agents are situated not only in a physical but also in an institutional environment, where their interactions are being guided by a network of institutions. An institution can be a norm, a role, a behavioral routine, physical device or any other type of coordination artifact implemented as a material object or a mental construct. The existence of coordination artifacts implies that an institutional setup is neither purely centralized, nor fully decentralized, nor purely distributed, but rather a mixture of centralization and distributiveness. By principle, institutions are *generic*: they are not designed to any specific set of robots. But it is also acknowledged that the boundaries between institutional and purely physical aspects of the world are not sharp, and subject to interpretation.

The original work on IR [6] is a collection of insights on introducing institutions in multi-robot systems, but it is still far away from a formal framework. A first abstract definition of institution is presented in [7], where institution is as a tuple *(ID, Rationale, Modifiers, Network, Institutional Building, History)*, with each component capturing the main constitutive elements of the social order dynamics. As noted in [147], the generality of this definition makes it insufficient for realization in control of robotic systems.

Institutional Robotics for Robot Swarms

In the context of multi-robot systems, institutions have been introduced [6], formalized [148], and used for modelling and implementation of simple robotic behaviors [30]. For formal representation of robot institutions, authors in [148] use Petri Nets (PNs), which encapsulate the behavioral rules and allow for concurrent, regulated execution. An institution is defined as a tuple consisting of conditions for institution activation and deactivation, the associated deontic operators (stating whether the institution can be composed with other behaviors) and a PN representation used for such composition. Composition of individual robot behaviors and institutions organized in a multi-layer methodology results in an Institutional Agent Controller (IAC) - a blueprint that each robot in the collective is equipped with.

Institutions are kept rather simple – an example of an institution is a 180° turn. An experience we gained during our preliminary study in [10] made us understand that the use of similar compact, discrete representation with PNs lacks the refinement needed for representation of continuous multi-robot behaviors, and would quickly grow in complexity of design when used for such systems. Furthermore, institutions are introduced to facilitate robot-robot interaction, but not human-robot interaction, and so, there is no incentive to define norm in human language. Nonetheless, authors in [147] recognize the need for encapsulation, modularization and abstraction – the key properties we aim to achieve in our institutional framework. Encapsulation allows for heterogeneity of the actors, and so, also for mixed societies of humans and robots. Modularity allows for hierarchical representation, where rules (institutions) are triggered based on the events in the environment. This is equivalent to the norm activation in our framework, where norms are activated due to conditions. Finally, institutions as presented in [147] can be thought of as akin to the norms in our framework, whereas the IAC, which composes institutions and individual behaviors by the means of deontic logic, resembles an equivalence of one institution in our formalism.

Institutional Robotics for Planning

The framework presented in [13] is the first work based on IR concepts to address the need of separation between the social structures and the concrete system. The notion of institution, defined as a set of artifacts, roles, actions and norms, is used for norm encapsulation, whereas norms are predications over the institutional statements, e.g. a norm "the customer pays money before the waiter serves food" is a relation between two statements: $\mathsf{before}((\mathsf{customer}, \mathsf{pays}, \mathsf{money}), (\mathsf{waiter}, \mathsf{serves}, \mathsf{food}))$. The framework distinguishes between abstract norms and their instantiation into a concrete system, hence allowing for the use of the same institution across different physical domains. Norms are used for restricting the system state evolution in the planning process formulated as a constraint satisfaction problem. Continuing with the above example, all possibilities of state evolution, where customer pays before receiving food (at the domain level), are compliant with the norm.

Although in our methodology we adapt the separation of the institutional abstraction from a concrete system using grounding as in [13] [149], our focus lies on coordinated multi-robot behaviors in continuous state space, and so, the norm semantics used in [13] does not provide sufficient sophistication and representation power to encompass the complexity of the behaviors we deal with. The institutional approach presented in this thesis diverges from [13] at the point where institutions are applied in physical system and the differences mostly emerge due to different objectives – in our case for continuous collective behaviors in human environments, in [13] for planning. However, the definition of institution, domain and their link through grounding are similar.

10.3.2 Norm Abstraction and Interpretation

Only relatively few examples of using social norms in robotics can be found in the literature. In [150], social rules are specified at the concrete system level using domain-dependent and language-dependent formalism. In [151] and [152] norms are used for designing what the robots should do, but not for *how to do it*, meaning that there is disconnection between the abstract normative layer and the concrete system. To address the question of *how* to shape robot behaviors according to social norms, the majority of social robotics research refers to the models of Proxemics [63] and Social Forces [18], based on which norms are implicitly used for social path planning [65][153][154][155], human guidance [67] and behavior selection [79]. Lists of descriptive norms written in a natural language are employed in [156], [157] and [75], but the translations of norms are method-specific, and targeting primarily the restricted context of single-robot navigation with little emphasis on robot-robot or robot-human co-operation. The above studies share the problem of being domain specific and do not have abstract, reusable models. A model aiming at encoding guidelines for culturally competent behaviors proposed in [78] uses specific ontology which operates at an abstract and concrete level, but norms have no explicit representations.

Two closely related areas of research on *joint actions* [158] [159] and *social practices* [160] draw inspiration from psychology to reason about the mechanisms that allow multiple agents to coordinate their actions in space and time. Studies on joint actions of humans and robots focus on the practical aspect of *how* to achieve successful collaboration by providing the agents with the means for joint planning, shared task representation, intention inference and representation of the mental model of the collaborator [161]. Social practices, on the other hand, describe physical and social patterns of joint actions using a high-level representation, standardized for a given context, by combining aspects such as roles, plans, norms and resources, which can be used to construct interactions. These abstract representations are then to be filled when they are needed to determine a course of action. A consolidated framework bringing together solutions inspired by the work on joint actions and social practices has the potential for achieving a powerful tool for human-robot cooperation, but up until now solutions have been proposed only at a conceptual level [162].

Efforts to consider higher levels of abstraction are undertaken in the broad field of human-aware task planning, focused on scheduling robot tasks and actions so as to accommodate human presence or to collaborate with a human. Authors in [163] identify and address the key challenge of understanding and interpretation of a broad variety of situations with rich semantics. To this end, they propose a modular and extendable framework, where new types of constraints can be added and solvers can be exchanged and re-arranged. In [164] planning abilities of a social robot are designed and implemented in a task-independent manner, and provide high levels of parametrization, so that a robot can adapt to various environments, different tasks and variable levels of engagement. A system presented in [165] allows the robot to elaborate and execute shared plans that are flexible enough to be achieved in collaboration with a human in a smooth and non-intrusive manner, while building upon algorithms generic enough to be used for other tasks and contexts. Studies on human-aware task planning focus on *what* to do and *when*, but not on *how* to do it. For this reason, the key research questions we have identified, such as interpretation of norms defined in human language and their translation into behaviors relying on potentially complex, continuous control laws, are not considered.

In summary, although the main purpose of norms in robotics is to guide social interaction, no systematic framework exists to this day that brings all the elements necessary for introducing the normative aspect to multi-human, multi-robot behaviors. Such key elements include general, human-understandable representation of norms, distinction between abstract, normative layer and the physical system and the connection between the two, which should provide an answer to the question of *how* to translate abstract norms to concrete robot actions. In Chapter 12, we will explore the state-of-the-art approaches to introducing language-defined norms in robotics, particularly in navigation methods, which deal with low-level, continuous behaviors.

Summary

Translation of generic norms defined at an abstract, institutional level into a terminology of the concrete system is a challenge identified across the disciplines. In this chapter we described how we draw inspiration from these findings to propose our solution to norm representation and realization in robotic systems. In particular, the work of Ostrom in [118] strongly influences our deliberation of the institutional components that form a normative statement, whereas the study of Dignum on norm translation in [145] motivates our work on norm realization.

We took this multi-disciplinary perspective to survey the recent developments on normative robotics and draw two conclusions. First, the need for abstraction of the normative layer has been identified in the robot planning approaches, but their link to a physical system is ignored or oversimplified. An emphasis is placed on clear semantics needed for plan representation, but the same requirement implies need for low complexity and thus, poor power of portrayal of complex behaviors. Indeed, in planning, most of social norms do not answer the question of "how?", but only "what?" and "when?" and are, in spite of claiming to be abstract, defined at the level of a concrete system. Second, methods and frameworks for social navigation focus on answering the question "how?" to apply social norms to robot behaviors, but the operation of norms at a concrete system level leads to poor reusability, modularity and scalability of the solutions. Furthermore, although proposed with the objective of being general, it is clear that the methods are restrictive to the context of single-robot navigation.

From this perspective, our work brings the two worlds and their advantages together in an integrated framework. In the next chapters we describe our formal framework inspired by the planning approaches, where we distinguish the institutional layer from the concrete system. Furthermore, we introduce the concept of *norm realization*, where we answer the question of "how?" missing in the planning approaches, but addressed by the social navigation community. We showcase our framework in a number of case studies, where we apply norms to achieve social multi-robot behaviors in presence of humans, and in mixed human-robot teams, in both navigational and emotive contexts.

11 Institution Formalization

In view of the literature review presented in Chapter 10 we conclude that although multiple models for normative behaviors are proposed in the field of multi-agent systems and robot planning, because of their disconnection from physical execution they cannot easily be applied to robot behaviors. Conversely, in the field of social robotics, methods exploring norms are domain specific and do not have abstract, reusable models. In this thesis, we propose a framework that bridges the above solutions by integrating abstract normative models and physical execution in a model-based approach.

In this chapter, we identify the concepts that lay down the foundations of our formalism of norm-following robots situated in institutional environments, the overview of which is presented in Figure 11.1. First, we provide the definition of institution and discuss the choice for the components that together with norms form an institution. In a second step, we present a formal definition of a concrete system – a *domain*, and discuss the link between the abstract institutional layer and the domain achieved through *grounding*. So defined institution can be reused on several physical systems. While the grounding defines the relationship between the institutional components and the domain, it does not allow for translation of language-defined norms onto the terminology of concrete robotic systems. Such translation is performed by means of *norm realization*, which we will discuss in the next chapter.

11.1 Definition of Institution

The view we choose to adopt defines an institution as a *systems of norms*, as opposed to a collection of norms only. In other words, institutions are norms and components necessary for translation of such norms into practice. Our choice stems from the fact that, in order to realize social norms, it is necessary to define to whom they refer to, in what situations (or under which circumstances), and, at times, how to apply certain norm to the behavior. The questions of "*who?*", "*when?*" and "*how?*" concern both natural or artificial systems – humans know that dressing code applies in formal situations, but not at home ("*when?*"), that only children above certain age are allowed in the bouncing castles ("*who?*"), and that before being allowed to study at a university, one must fill in an application form and be accepted ("*how?*"). In order to answer such questions in an artificial system, agents (or robots) must have access to structured information necessary to enforce social norms. In our formalism, institutional components hold such information.

Definition 11.1. Formally, an institution is a collection:

$$\mathcal{I} = \Big\langle \textit{Norms}, \textit{Roles}, \textit{Actions}, \textit{Conditions}, \textit{Knowledge} \Big\rangle$$

where a role in the *Roles* set is a function assigned to an agent, the set *Actions* comprises the actions that have to be taken when playing a given role, conditions in the *Conditions* set regulate whether norms are active and evaluate whether the same norms are satisfied, and the *Knowledge* set gives the agents the means to comply with the norms by providing basis on which they can act.

The above definition of institution has emerged as a result of unifying two views – a) our preliminary work on institutions in [10], where an institution is defined as a collection of rules, actors, actions, knowledge, memory and payoffs, and b) definition in [13], where an institution is a tuple of norms, roles, actions and artefacts (objects in the environment).

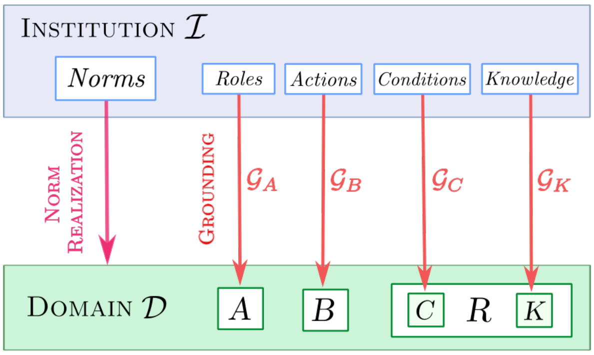

Figure 11.1 – Institutions are defined at the highest level of abstraction. A concrete system is represented as a domain. Grounding is a link connecting the abstract layer to the concrete one, however, it is insufficient for translating abstract norms in terms of concrete system representation. The translation is achieved through norm realization, which we describe in Chapter 13.

Roles. The purpose of assigning roles is twofold. First, roles serve to specify the mission of an agent: for instance, a worker ant is to collect seeds, or in our formations Leaders are to guide the Followers). Second, roles are important for understanding the classes of interactions agents engage in: for instance, a robots should behave differently around hospital Staff than hospital Patients, where for the first group they should minimize disturbance (also in terms of interactive features, lights, sounds etc.), while for the latter one more lenience is allowed and the primary objective of the robots is to provide entertainment. One should also note that preferences can be defined for concrete agents, or the attributes thereof. For example, we can specify that all children prefer when the robot moves slowly next to them and so does an individual named Steven. Personalization for a given role or a given agent is further achieved through institutional knowledge.

Actions. An action is a placeholder abstracting specific agent behavior. It allows us to reason on a high level without the need to consider implementation details. For example, the two collective behaviors – formation control and flocking introduced in Chapter 5 are abstracted away to be the same action – Follow. The differences between algorithms are taken care of at the translation layer, but the same institution with the same norms applies to both.

Conditions. Defining conditions over the world state allows the agents to understand and interpret the social context. For example, a condition evaluating whether there is an exam or a class break in school situates a robot in a different social context and so, different sets of norms are to be employed.

Knowledge. Institutional knowledge is a structure for providing the agents with information necessary to address the question "*how?*". When concretized in an actual system, knowledge incorporates symbols, relations, facts and beliefs imperative for cooperation and provides recipes for socially adequate interactions with humans [113]. For example, any robot moving in a human-populated environment should know that there is a maximal allowable speed, and what the value of that speed is. Knowledge therefore is used for encoding information necessary for interpretation of abstract norms and should be available to all agents so that every participant knows how to act, and knows that the others know how to act as well [166]. Knowledge also allows for personalization of interactions by specifying the preferences of concrete agents (such as the values of the maximal speed to keep near the concrete human), and for customization of actions with regard to the social context (for example, a robot can make cheerful sounds to children playing a game, but not to a teacher giving a class).

11.1.1 Norms

Norms are inherently abstract statements [130], general enough to be interpreted concretely in diverse situations. They encompass a broad set of shared linguistic constraints and opportunities that prescribe, permit or advise actions for participants. Institutional *Norms* denoted with $\mathcal{N} = \{n_1, n_2, \ldots\}$ take the form of a human readable sentence with a specific syntax.

Definition 11.2. Formally, a *norm* is a statement, where a deontic expression forms a relation:

$$\mathcal{N} : \mathit{Conditions} \rightarrow \mathsf{deontic}\left(\mathit{Roles} \times \mathit{Actions} \times \mathit{Knowledge}\right)$$

The deontic includes obligations, permissions, and related concepts, e.g., *must*, *should*, *must not*, etc. Not every norm sentence must explicitly include all the above elements. Statement "no smoking" has implicit conditions (in this place, at all times) and implicit roles (everyone), providing only the deontic operator and an action of smoking. A speed limit road sign showing "30" is a type of norm providing only the knowledge element (the limit value), but it is generally understood to whom it applies to and under which circumstances.

The above relation is inspired by the syntax proposed in [118] by Elinor Ostrom, winner of the Nobel Prize in Economics in 2009, where norms are referred to as regulatory rules and take the following syntax:

> "ATTRIBUTES of participants who are OBLIGED, FORBIDDEN, OR PERMITTED to ACT under specified CONDITIONS, OR ELSE".

The similarities of our proposal include the use of deontic elements, the relation between ATTRIBUTES and *Roles*, the ability of participants to ACT or perform *Actions*, under certain CONDITIONS or institutional *Conditions*. The OR ELSE part of Ostrom's regulatory rule encompasses the penalty for rule infractions – an institutional aspect that we do not address in our work.

As norms operate over actions that are linked to the roles, we find it useful to define a *feasibility relation*, i.e. a relation between a role and the action assigned to the participant assuming that role (subsuming the actions that are possible for such agents). For example, a role of a teacher is linked to an action of giving a class.

Definition 11.3. The feasibility relation takes the form:

$$\mathcal{F} \in \left(Roles \times Actions \right)$$

The feasibility relation is useful when multiple heterogeneous agents are to deliberate on the role assignment within the team. The link between an agent and the role it can assume given its capabilities is provided by grounding, which we discuss in the sections that follow.

11.1.2 Domain

An institution is an abstraction that can be instantiated in concrete systems that are physically different but can be described by the same structure. Such a concrete system is called a *domain*.

Definition 11.4. A *domain* is a tuple

$$\mathcal{D} = \langle A, B, R \rangle$$

where A is a set of *agents*, B is a set of *behaviors* and R is a finite set of *state variables*.

The set A can include humans and robots. B is the collection of all behaviors that agents can perform. The state variables R define properties pertaining to the agents or objects in the domain. They may indicate the position of an agent or an object, the activation of a behavior, etc. The state variables also serve to evaluate whether an institutional condition is satisfied and ground institutional knowledge to its concrete representation in the system.

11.1.3 Grounding

Grounding provides the key to reusability of the same abstract institution for regulation of different physical systems [13].

Definition 11.5. *Grounding* of institution $\mathcal{I}$ into a domain $\mathcal{D}$ is a tuple:

$$\mathcal{G} = \left\langle \mathcal{G}_A, \mathcal{G}_B, \mathcal{G}_C, \mathcal{G}_K \right\rangle$$

where:

- $\mathcal{G}_A \subseteq Roles \times A$ is a *role grounding*,
- $\mathcal{G}_B \subseteq Actions \times B$ is an *action grounding*,
- $\mathcal{G}_C \subseteq Conditions \times 2^R$ is a *condition grounding*,
- $\mathcal{G}_K \subseteq Knowledge \times 2^R$ is a *knowledge grounding*.

Grounding establishes the relation between an abstract institution and a specific domain. It relates roles to agents, generic actions to behaviors of agents, and conditions and knowledge to the factual state of the environment.

The grounding $\mathcal{G}_B$ between the institutional *Action* and the behavior B at the domain level does not restrict to one-to-one mapping. It is entirely possible that agents within the same team ground the same action to different behaviors, for instance, part of the robots moving as a group might ground an action Follow to a MoveInFormation behavior, while others to a MoveByFlocking behavior. Similarly, the above definitions do not prevent dynamic re-grounding, where for a given robot, the action grounding changes depending on the situation.

For convenience, we will distinguish *grounded conditions* as C and *grounded knowledge* as K. Grounded conditions C are the results of evaluating boolean functions $f : 2^R \rightarrow \{ \, True, \; False \, \}$ over state variables R. For example, a condition LEADER_IN_KITCHEN evaluates whether position of the agent that plays the role of Leader is within the area designated as the kitchen in the given environment. Knowledge includes a-priori information encoded by the institution, and grounded to concrete values for a given domain. For example, institutional knowledge PersonalSpace can be grounded to a specific distance threshold once it is known that the individual has a specific spatial preference, such as disliking when large robots are moving too close by.

Grounding allows the institution to be universal and applicable to an unknown and indeterminable number of persons and circumstances. The importance of grounding is widely acknowledged in social sciences [118] [113] and put into practice by the MAS community [119], but the proposed solutions are not adequate for translating language-defined norms into robot terms. In our formalism norms are translated and applied to the domain by the virtue of a special type of grounding called *norm realization*, which we will describe in the next chapter.

11.2 Networks of Institutions

In our preliminary study [10], we investigated the methods for specifying relations between institutions. Although the study and this thesis bear significant differences, in this section we summarize the main concepts related to the networks of institutions, the implementation and validation of which can be found in [10].

First, we formally define the concept of the *social context*. Social context is the immediate physical and social setting of the environment within which participants (humans and robots) function and interact [167].

In our formalism, an institution operates in a specific social context it has been designed for. Similar patterns can be observed in human environments, where some institutions (especially the formal ones) function within clear boundaries. For instance, an educational institution has control over the behavior of children when they are physically at school, but its norms do not apply at home or during vacation time. The recognition of the social context is not trivial and is considered one of the challenges of IR [13]. The same context might involve different participants, occur in physically different place or involve other characteristics that make it difficult to classify it.

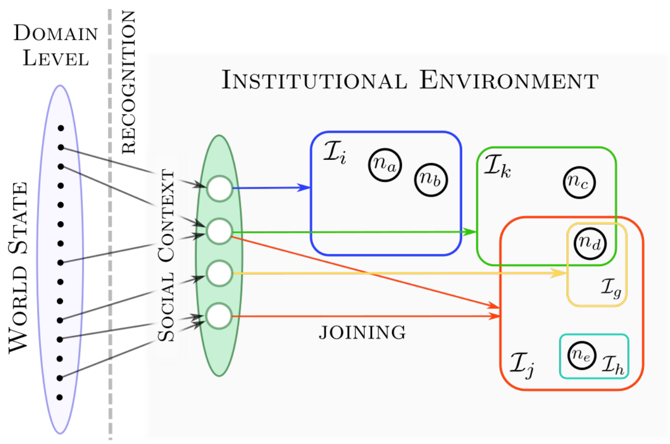

Figure 11.2 – Illustration of the institutional environment. Social context is an abstraction of the world states, classifying them in terms of socially-relevant aspects. Once a social context is recognized, an agent joins an institution with its norms $n_k \in \mathcal{N}$ that has been designed to govern the agent's behavior in such context.

Once a normative agent recognizes the social context, it joins the corresponding institution – it is assigned a role within that institution and it is obliged to act according to the institutional norms. We say that such an agent is immersed in an *institutional environment*, which encompasses the social contexts the agent is capable of recognizing and the network of institutions that correspond to these contexts. An agent joins an institution upon recognition of a social context and leaves an institution when the context is no longer perceived. For example, a robot entering a classroom joins an institution responsible for educational activities; once a class concludes, it switches to an institution for entertaining children. It is possible to belong to multiple institutions: for example, an educational institution can be run in parallel with another institution responsible for monitoring children's safety. One should, however, notice that the granularity of institutions is a design choice, and simultaneous execution of multiple institutions necessitates methods for resolving conflicts between them. The summary of the above definitions is shown in Figure 11.2.

A similar view of institutional environments, where execution of institutions is triggered by the environmental conditions, is adopted in [147], where the composition of institutions preventing their inadequate concurrent execution is controlled by a set of deontic operators and implemented in the Institutional Agent Controller (IAC). In our study in [10] we introduce the concept of a Relational Graph (RG) - a directed graph defining the relationship between institutions, regulating which of them can be active at a given time and assuring that the actions of conflicting institutions do not run concurrently. Both IAC and the RG are formulated as mathematical models used for forming dependencies between complex behaviors and can be readily proposed as the robot controllers.

The objective of the above discussion is to provide preliminary considerations about how multiple institutions can coexist and form a balanced network, which, once embedded in the robot's code, would allow for building a robot's controller as in [147] and [10]. However, it is out of scope of this thesis to deal with institutional concurrency and conflicting institutions. Nevertheless, we show a simple example of two linked institutions in our case study $\mathcal{C}_{II}$ in Chapter 16, where a condition for termination of one institution leads to the activation of another.

11.3 Granularity of Institutions

There is not a unique approach for defining the scope of institutions. They can be either very general, with a large number of norms that are filtered by conditions, or very specialized, which necessitate presence of protocols defining relations between them.

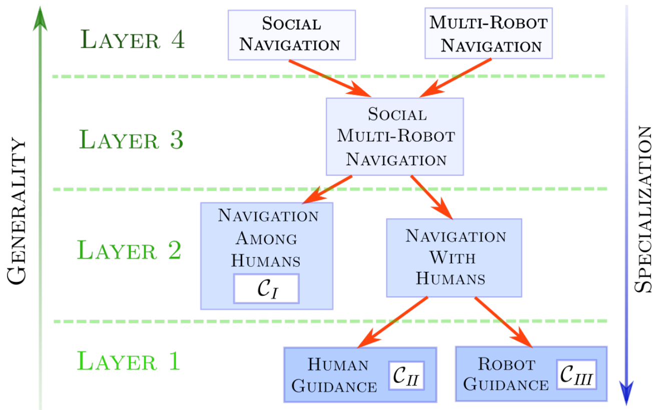

Figure 11.3 – Inheritance relations between the institutions deployed in this thesis and more general institutions that can be though of as meta-institutions.

In our implementation the institutions are neither general, nor specific, and defined at several levels of generality, as shown in Figure 11.3. In our particular application of deploying multi-robot systems in human-populated environments we distinguish two sub-cases (Layer 2 in Figure 11.3) that differ enough to embrace them in separate institutions, namely *navigation among humans*, used when humans and robots co-exist but do not cooperate, and *navigation with humans*, used for mixed groups of humans and robots. An institution for navigation among humans is described in the Case Study $\mathcal{C}_I$ in Chapter 15.

We further decide to separate the case of mixed groups where the humans are guided by the robots from the case where the humans are the guides, as most of the norms in the two social contexts are dissimilar. An institution for the first case is deployed in Case Study $\mathcal{C}_{II}$ in Chapter 16, while an institution for the latter context in Case Study $\mathcal{C}_{III}$ in Chapter 17. The two institutions for mixed groups form the most specialized bottom layer in Figure 11.3.

The aforementioned institutions can be thought of as specializations of one general institution for social multi-robot navigation (Layer 3 in Figure 11.3), which in turn inherits all the norms related to human comfort from an institution for social navigation and all the norms related to robot-robot interactions and collective behaviors from an institution for multi-robot navigation (Layer 4 in Figure 11.3). It is clear that the proposed fragmentation is not the only possible option, and is a part of a much larger inheritance relation.

Summary

In this chapter we introduced a model-based approach for abstraction, encapsulation and formalization of social norms for robotic systems, where we systematize generic norms into reusable structures, called institutions, and then ground these institutions into concrete systems – domains.

We focus on transparency and generality of the formalism, where its abstract representation allows for the use of miscellaneous robot behaviors and integration of social constraints of diverse nature. The properties of institutions allow us to seamlessly reuse the same institution across the domains with a variety of agents capable of performing different behaviors, as long as the social context is the same. The proposed definitions meet our objective of providing designers of social robot behaviors with a simple tool that allows for encoding behavior specifications in a plug-and-play principle instead of programming hard-wired social compliance in ad-hoc behaviors. We will demonstrate how our approach can be used to introduce normative aspects into robot planning and control for participating in mixed human-robot societies in the chapters that follow.

12 Norm Formalization

WHEN deploying robots in social environments with the aim of interacting with humans, it is only natural that the designer of the robot architecture combines a number of procedures in a unified software stack to achieve the desired behavior. For example, a robot that has to follow the norms of socially avoiding people, slowing next to them, and planning its path so as to remain visible, typically executes an overall behavior that integrates all the three norms or arbitrates among them. When a norm is to be changed, or another norm is to be added, the algorithm must be modified accordingly. From that perspective, social compliance is hard-wired in ad-hoc behaviors, which are difficult to reuse or generalize. In this chapter, we revisit a body of work on social robots that explicitly deals with norms and their representation and explain the advantages that can be gained when instead of the behavior design is being driven by norms, norms are imposed as constraints operating over the parametrization of already existing behaviors.

12.1 Formalization of Norms in Robotics

Current research on norms for social navigation treats them as prescriptions for behavior designers on how to select or modify robot behaviors, and the robots have no choice but to follow those prescriptions. The bodies of work that endeavor to bring together different aspects of social navigation under a unified framework include the Human Aware Motion Planner [154], the COMPANION method [79], the human-centered sensitive navigation approach [156], further pursued by [157] and the socially aware robot motion framework [75]. In these studies, norms are described in human language, e.g. “(...) robot should not enter the personal space of a human”, or “the robot remains as visible as possible along the path”, but not formally represented - they are directly encoded by the designer in the robot behaviors through a motion planner. In [153] norms are specified formally as constraints on robot trajectory, limiting them to a narrow context of single robot navigation. Studies on human-aware robot navigation in crowded places characterized by high complexity of interactions with large number of participants propose more general schemes for motion planning under social constraints [168] [169] [65], however abstraction from concrete implementation is not considered.

None of the aforementioned approaches provides means to represent norms independently of the domain, and usually any abstraction of normative rules is very simple or non-existent. However, they also largely identify, implicitly or explicitly, elements necessary for concretization of abstract norms, namely a) *conditions* or events that determine when a given norm should be satisfied, e.g. robot should slow down when close to human, b) *knowledge-base* – a type of lookup table encoding behavioral specifications for satisfying a given norm, e.g. the preferable distance from the robot (i.e. proxemics), and c) the *behaviors* the norm has an influence on or must be undertaken in order to satisfy the norm.

In our formalism, conditions, knowledge and actions are integral part of the normative system layer, namely of an institution. Therefore, using the terminology of our formalism, the integration of norms and the institutional components within the normative behavior design of the above approaches could be represented as follows:

$$\text{BEHAVIOR} = \text{FUNCTION OF } (\text{NORMS} \times \text{CONDITIONS} \times \text{KNOWLEDGE})$$

The problems associated with similar design principle, where constraints drive the design of a behavior include:

- Poor reusability – since the norms are often designed for a specific behavior, and not for general purpose; therefore, their use is limited to the current application.
- Lack of modularity – since behavior design is driven by norms, every change of a norm or adding a new norm might require a re-design of the behavior.
- Poor scalability – complexity of the design rapidly grows with the problem space, as introducing a new norm might require heuristics on how to merge it with the solution established by the existing norms.

As we will describe in details in the chapters that follow, our formalization takes the following perspective:

$$\text{NORMS}: \text{CONDITIONS} \rightarrow \text{DEONTIC}(\text{ROLES} \times \text{ACTIONS} \times \text{KNOWLEDGE})$$

where norms operate on the existing system components. Actions, defined at an abstract, institutional level represent an infinite set of possible behavioral modalities, while norms allow only a subset that satisfies the constraints. It means that we have reversed the process – instead of the design being driven by norms, norms are imposed on already existing baseline in a plug-and-play manner. The difference in the approaches is shown in Figure 12.1. Although

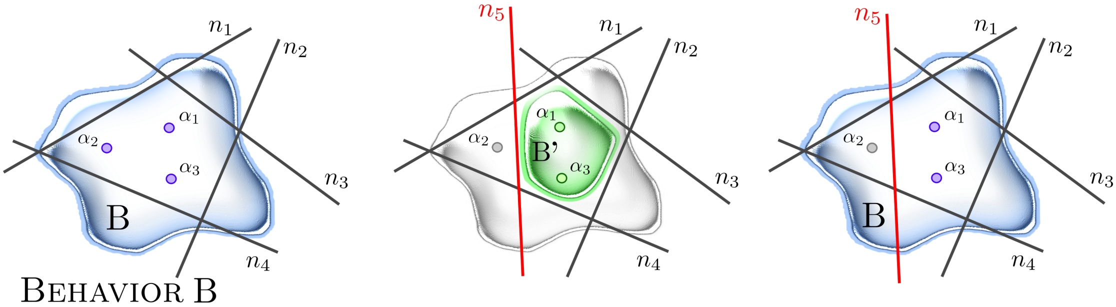

Figure 12.1 – Illustration of the difference between the design approaches typically adopted in the literature and the method we propose. Given the original behavior B with the norm-induced constraints $n_1 - n_4$ (left), introducing a new norm n_5 in the earlier case might necessitate re-design of the behavior (middle). In our formalism we follow a different strategy, where the norms restrict the parameter space of the behavior (right).

the final result in terms of achieved behavior complexity might be the same in both cases, our method, as we will show through several case studies, addresses the above problems of reusability, modularity and scalability.

In the above discussion we consistently use the term behaviors when describing the approaches proposed in the literature and actions in the case of our normative model, as in our formalism we differ between these two elements. In particular, actions are defined at an abstract, institutional level and behaviors are concrete realizations of such actions in the domains. In the state-of-the-art methods the norms typically operate directly on concrete behaviors, while in our case they do so indirectly via a translation layer. Eventually, however, in both cases norms result in constraints on the concrete system level as presented in Figure 12.1.

To place our work with respect to state-of-the-art literature, in Table 12.1 we outline the approaches discussed in this chapter, highlighting the aspects relevant to this thesis, namely whether the normative layer is separate from a concrete system, whether a link between the abstraction and the concrete has been made, and the choice of norm representation and validation. The second part of the table outlines the normative frameworks in robotics, including the studies on institutions we took inspiration from [147] [13]. We also mention the IAD model of Ostrom [118], for its close link to our objectives on the power of norm representation.

12.2 Formalization of Norms in this Thesis

Approaches on human comfort or naturalness can readily be framed in the form of a social norm. For instance, methods for keeping a distance from a human address a norm "*A robot must not enter human personal space*", while imitation of human motion by a robot is equivalent to a norm "*A robot should replicate the way humans move*". Our choice to use human language for norm representation is dictated by the fact that such form is the most generalizable and understandable, also for the end users that are to share environments with the

			NORM REPRESENTATION			VALIDATION		
	ABSTRACT NORMATIVE FRAMEWORK	ADDRESSES QUESTION "HOW?"	IMPLICIT	PREDICATE	HUMAN LANGUAGE	CONTINUOUS BEHAVIORS	MULTI-ROBOT	WITH HUMANS
Sisbot 2007 [154]	✗	✗	○	○	●	●	○	●
Kirby 2009 [79]	✗	✗	○	○	●	●	○	●
Alili 2009 [150]	✗	✗	●	○	○	○	○	●
Pandey 2010 [75]	✗	✗	○	○	●	●	○	●
Lam 2011 [156]	✗	✗	○	○	●	●	○	●
Gomez 2013 [157]	✗	✗	○	○	●	●	○	○
Carlucci 2015 [151]	✓	✗	○	●	○	○	○	●
Triebel 2016 [155]	✗	✗	●	○	○	●	○	○
Okal 2016 [153]	✗	✗	○	●	○	●	○	●
Fererra 2017 [152]	✗	✗	●	○	○	○	●	○
Chen 2017 [65]	✗	✗	●	○	○	●	●	●
Ostrom 2005 [118]	✓	✓	○	○	●	○	○	○
Pereira 2014 [147]	✓	✗	○	●	○	○	●	○
Wasik 2017 [10]	✓	✗	○	○	●	●	●	●
Sgorbissa 2018 [78]	✓	✓	●	○	○	○	○	●
Tomic 2018 [13]	✓	✓	○	●	○	○	●	●
This thesis	✓	✓	○	○	●	●	●	●

Table 12.1 – Overview of the literature on norms in robotics, with the notable exception of the IAD framework shown for comparison. The first part of the table lists the approaches in navigation and planning, while the second part the normative frameworks.

norm-following robots. Such language-defined norms are generic, so stating them explicitly in such form does not bring benefits, unless they can be translated in robot-understandable terms. The problem of vague norms and norm translation is critical in this thesis and here we shortly describe the single-robot human-aware methods we adapt in our case studies to exemplify the scope of approaches to human-robot interaction we can achieve. A selection of norms defined in human language is outlined in Table 12.2.

We would like to emphasize that although social awareness is an essential part of our work, we use the existing, well-studied methods to demonstrate the power of representation of our framework. It is not our goal to advance the state-of-the-art on social navigation, although in some cases we are the first ones to apply existing single-robot approaches to a multi-robot context.

12.2.1 Methods for Achieving Human Comfort

To make the robot move in a way that feels safe and considerate to the human, most of the related work focuses on respecting human personal and activity spaces and velocity control near humans.

Proxemics Representation of Human Personal Spaces

To represent human comfort spaces we use the Proxemics Model (PM) [63], described in Chapter 5. The radii of human spaces serve to determine the degree of compliance the robot must apply to assure human comfort. We vary them depending on the attributes of a specific human (e.g., securing larger personal space near a child) or a human role (e.g., smaller zone for humans with previous acquaintance with the robots). We label the extent of the comfort zones for human H_h with $\Delta_{S,h}$ for social space, $\Delta_{P,h}$ for personal space, and $\Delta_{I,h}$ for intimate space. The norm for respecting personal space is stated in Table 12.2, while the examples of similar normative constraints formulated based on human social spaces can be found in [154] [75] [153] [156] [79].

Cost Map Representation of Human Activity Spaces

In our approach, activity-critical areas, also referred to as the affordance spaces (see Section 4.4 for details) are represented as cost functions and added to robot's occupancy grid map. Such socially-augmented cost map is further used in path planning by the formation leader. Examples of similar norms related to affordance spaces are found in [153][154].

Velocity Control Near Humans

The maximum possible speed of the MBot robot of 2 m/s can be dangerous in constrained environments populated with humans, therefore the overall speed in our experiments is limited to a physically safe (but not social[1]) bound. Instead of enforcing further speed limits in the vicinity of humans, we use reduction factors to lower the speed of the robots with respect to the speed resulting from the multi-robot algorithms. The respective norm is stated in Table 12.2.

1 The importance of speed modulation in the vicinity of humans has been described in Section 4.3.

NORMS RELATED TO HUMAN COMFORT			
PERSONAL SPACES	[75] [156]	*"Robots must respect human personal spaces"*	$\mathcal{C}_{I-III}$
AFFORDANCE SPACES	[153] [154]	*"Robots should not enter activity-critical areas"*	$\mathcal{C}_I$
VELOCITY CONTROL	[79]	*"When navigating among humans, robots should assume appropriate speed"*	$\mathcal{C}_I$
NORMS FOR MIXED FORMATIONS OF HUMANS AND ROBOTS			
WAITING	[170]	*"When humans fail to follow behind the leader during guidance, the leader should wait for them"*	$\mathcal{C}_{II}$
RE-ENGAGING	[93] [87]	*"When humans guided in a formation are not following, the followers should encourage them to return"*	$\mathcal{C}_{II}$
NORMS FOR HRI-AUGMENTED INTERACTION			
FACIAL EXPRESSIONS	[33]	*"Robots should indicate their intentions through facial expressions, gestures and sounds"*	$\mathcal{C}_I$
GAZE DIRECTION	[64]	*"When interacting with humans, robot should direct its gaze towards the interaction partner"*	$\mathcal{C}_I$
SOUNDS & SPEECH	[33]	*"Robots should offer informative announcements"*	$\mathcal{C}_I$

Table 12.2 – Examples of norms used in the case studies $\mathcal{C}_I$, $\mathcal{C}_{II}$ and $\mathcal{C}_{III}$ presented in Chapter 15, Chapter 16 and Chapter 17, respectively. The second column is a reference to work from which we took our inspiration, while the third column states the norm in human language – a representation that in our formalism we deem to be the most generalizable and understandable, also for the end users that are to share environments with the norm-following robots.

12.2.2 Norms Relevant for Mixed Groups of Humans and Robots

In our case studies, we employ a number of norms for mixed groups – for human guidance (Chapter 16) and for human following (Chapter 17). A more detailed review on related work can be found in Section 4.4. In the case study of Chapter 16, social norms exemplified in Table 12.2 are inspired by the methods presented in [93] and [170], where in [93] a robot re-engages a strayed person by exerting a repulsive social force that urges that human towards a goal, while in [170] the robot adapts its trajectory using a learning-based approach to slow down or stop when the human is not following. Methods for mixed groups on humans and robots have been discussed in Section 5.2.2.

12.2.3 Norms for HRI-Augmented Interaction

The approaches presented so far focus on normative robot navigation. To augment user experience, we employ HRI features, such as lights, sounds, gaze control and facial robot expressions.

Gaze Direction

We explore robot gaze in the case study in Chapter 15, where a group or three robots moves in human populated environment. Since gaze of multiple robots focusing on one human can be menacing, only one robot – the leader – directs its gaze towards passing a human and only for a short time.

Signaling Robot Intentions

While robot intentions can be conveyed through navigation [31], we choose to use sounds and robot facial expressions, including LEDs for eyes, mouth and cheeks as an attempt to provide feedback to the user, convey information about internal state of the robot, and facilitate believable human-robot interaction [33]. Our choice of robot expressions and sounds is based on what we believe conveys a clear message about robot's state.

Summary

The state-of-the-art research on human-aware robotics offers a number of powerful methods providing the means to achieve socially-compliant robot behaviors. However, most of the current research is focused on finding a solution to an isolated problem, while integration is less of a priority. Large collaborative projects such as MOnarCH provide an opportunity of integrating such disjoint bodies of work, and understanding the difficulties associated with such integration, but they are rather an exception than a rule. In this chapter, we have identified the need of abstraction of the social layer from that of its implementation and highlighted the advantages that can be gained not only by the designers of norm-following behaviors, but also the community.

Furthermore, we described how we take inspiration from research on human-aware robotics to motivate the selection of norms we choose to employ in our case studies. In particular, we focused on methods related to human comfort, approaches for establishing mixed human-robot groups and techniques for augmenting interactions with simple HRI. Furthermore, we gave a glimpse at how we formulate those methods as human-understandable norms. It is not our goal to encode every social robotics method in our normative framework. On the contrary, the selection presented in this chapter is intended to represent different aspects of human-robot interaction (from navigation and gaze control to expressions and speech) and different approaches to achieve it. How the norms are further applied in our methods is the subject of later chapters, where we deal with realization of abstract norms in real robotic systems.

13 Norm Realization

We believe that a general framework for using norms in mixed human-robot societies requires three key features: (1) the ability to abstract norms from the concrete systems; (2) the possibility to encapsulate norms into a model that is reusable across such systems; and (3) a means to interpret abstract norms in terms of robot-understandable language, readily implementable onto concrete restrictions of robot behaviors. In the previous chapters we specified the abstraction of institutions (and norms) from the concrete systems they regulate (1). Furthermore, the encapsulation of social norms and the context in which they operate in the notion of institution makes them reusable across situations (2). In this chapter, we describe the final missing element – *norm realization*. Similarly to the way the grounding allows us to reuse institutional abstraction across different domains, norm realization provides the means to do so with the norms. Translation of abstract norms in terms of robot-understandable language makes them readily implementable onto concrete restrictions of robot behaviors, and executable in real physical systems (3).

13.1 Norm Realization

Norm realization provides the steps needed for bridging the abstraction-to-implementation gap, inherent of the normative systems [130]. Abstract norms of institution are translated in terms of low-level parameters of complex behaviors to be realized to act on concrete behaviors of concrete robots in a real physical system.

The advantages of norm realization are threefold. First, they allow us to define norms at an abstract level, making them reusable across different domains. Second, norms can be formulated in an inherently generic human language, in a form easily understandable by humans, which is particularly important for forming mixed societies of humans and robots. Third, abstracting norms from the concrete behavior implementation promotes modularity, as norms can be added or removed in a plug-and-play principle, while the norm realization forms a cross-compiler that put the norms on top of the existing behavior.

The translation is achieved by the means of the *rules* – a tool chain linking the domain and the institutional abstraction. Once the elements of the normative sentence (i.e. *Roles*, *Actions*, *Conditions* and *Knowledge*) are classified in terms of institutional components according to the Definition 11.2, and the components are grounded to the domain by the means of grounding, the rules leverage these elements at the concrete domain level to formulate constraints of the existing behavior.

Before introducing the principles of norm realization it is important to characterize the terms that robots operate on, i.e. the low-level parameters of complex robot behaviors.

13.1.1 Parameters and Values

An institution $\mathcal{I}$ with norms $\mathcal{N} = \{n_1, n_2, ...\}$, regulates behaviors in $B = \{B_1, B_2, ...\}$ of agents in A. Behaviors are parametric, i.e. they have multiple *behavioral modalities* $\lambda = B_k(p \in P, v_p \in V_p)$ that depend on a number of parameters P and their values V_p. In general, P is an index set of parameters, where for each $p \in P$ we associate a set of values V_p from a family of sets $V(P) = \{V_p\}_{p \in P}$. The norm realization, by setting behavior parametrization to concrete values, constraints the choice of all possible behavior modalities to only one viable option. This is illustrated in Figure 13.1, where three parameters of behavior B_k, p_1, p_2 and p_3 take the values v_{p_1}, v_{p_2} and v_{p_3} respectively, to result in B_k assuming the behavioral modality $\lambda \in \Lambda_k$.

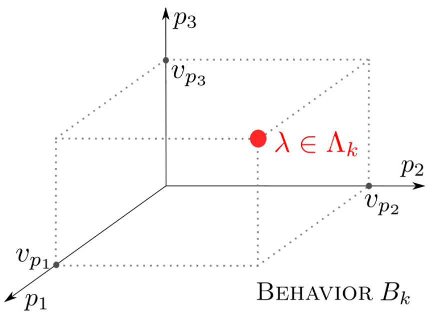

Figure 13.1 – Three parameters p_1, p_2 and p_3 of behavior B_k take the values v_{p_1}, v_{p_2} and v_{p_3} respectively, to result in the behavior B_k assuming the modality $\lambda \in \Lambda_k$.

13.1.2 Rules as Constraints over Behaviors

Rules of norm realization generate constraints on behavioral modalities, restricting Λ_k to a subset, the execution of which is norm-compliant. An illustration of the constraining nature of the rules is shown in Figure 13.2, where the horizontal axis represents the parameter space of a behavior and the vertical axis the value space of the parameters. In Figure 13.2 a) the rules are shown to be restrictive over the parameter value space, giving an allowable range of parameters, for which the rules (and so, the norm) will be satisfied. A behavioral modality is an arrangement of parameters and values for a given behavior, as shown in Figure 13.2 b), where three relations delineate three behavior modalities.

As a concrete example, consider Figure 13.2 c), where a behavior has four parameters, namely

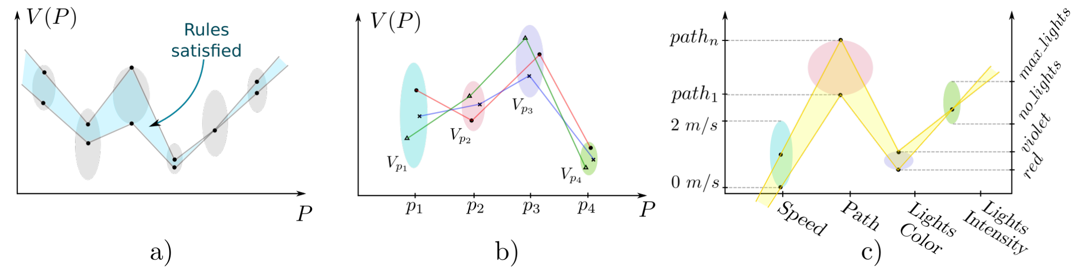

Figure 13.2 – Illustration of the constraining nature of the rules. The horizontal axis represents the parameter space of a behavior and the vertical axis the value space of the parameters. In a) the rules restrict the parameter value space, giving a range for which the rules are satisfied. In b) a behavioral modality is an arrangement of the parameters and the values for the given behavior. As a concrete example in c) shows how rules restrict the behavioral modality for satisfying a norm "During class time robot should move slowly and dim the lights".

Speed, *Path*, *LightsColor* and *LightsIntensity*. Let us assume that a robot is to apply a norm stating *"During class time robot should move slowly and dim the lights"*. To comply with the norm, the robot is constrained by the rules of norm realization that only allow a *Speed* of up to 1 m/s, and intensity of lights at around 30% of the maximal value, but the choice of the path and the color of the lights is free. Any modality that lies within the shaded yellow space in Figure 13.2 c) conforms to the norm.

Next, we provide the details on the normative power of the rules.

13.1.3 From Abstract Norms to Concrete Behaviors

The process of translating norms to concrete behaviors through norm realization involves steps operating at both institutional and the domain levels, the relation between which is shown in Figure 13.3.

Definition 13.1. Given that the norm $n_i \in \mathcal{N}$ is provided in the form of a human language sentence, the norm realization at the institutional level involves:

NORM REALIZATION: INSTITUTION LEVEL

- **Stage I.** Identification of the elements of the sentence that correspond to the institutional components, namely *Roles*, *Actions*, *Conditions* and *Knowledge*. The components constitute the semantics of norms given in Definition 11.2, i.e. $\mathcal{N}: \mathit{Conditions} \rightarrow \text{deontic}\left(\mathit{Roles} \times \mathit{Actions} \times \mathit{Knowledge}\right)$.
- **Stage II.** Once the components of the norm definition are identified, they particularize the sequence summarized in Figure 13.3 and detailed as follows:
  - An *agent* playing a *role* $\in$ *Roles* determines whether the norm applies to that role.
  - A set of activation conditions $C_i \in \mathit{Conditions}$ is verified to identify whether n_i is active.

- An active norm $n_i \in \mathcal{N}^A$ modifies an action $\in$ *Actions* of the aforementioned agent based on a subset of elements in *Knowledge*.
- The agent executes the normative action and evaluates its performance to establish whether the norm is satisfied and $n_i \in \mathcal{N}^T$.

With the exception of the first bullet point, the second stage of the institutional level is purely descriptive and does not take place in practice. Instead, once it has been determined whether the norm n_i applies to the given agent at the institution level, the sequence of the Stage II of that level is realized for a concrete system at the domain level as shown in Figure 13.3 and detailed as follows.

Definition 13.2. The norm realization of norm $n_i \in \mathcal{N}$ at the domain level involves five fundamental steps:

NORM REALIZATION: DOMAIN LEVEL

A) Norm n_i is active when the required *conditions* $\in C$ are satisfied

B) Each active norm $n_i \in \mathcal{N}^A$ can activate/modify/disable a subset of *parameters* $P(B_k)$ of agent's behavior B_k

C) Parameters in P attain *values* $V(P)$ that depend on the knowledge K and state variables R

D) The parameter values are directly applied the behavior to result in norm-complying behavior modality $\lambda \in \Lambda_k$

E) Norm n_i is satisfied if the outcome conditions in C are satisfied

---

13.1.4 Rules

Rules are the binding material connecting norms to behaviors. In a one-to-one relation with the steps of the domain level listed above, we distinguish the following rules:

A) REQUIREMENT RULES	$r^N : 2^C \to \mathcal{N}^A$
B) CHOICE RULES	$r^P : \mathcal{N}^A \times B \to 2^P$
C) VALUE RULES	$r^V : P \times 2^K \times 2^R \to V$
D) APPLICATION RULES	$r^B : 2^P \times 2^V \to \Lambda$
E) OUTCOME RULES	$r^O : \mathcal{N}^A \times 2^C \to \mathcal{N}^T$

where $\mathcal{N}^A$ is the set of active norms, $\mathcal{N}^T$ is the set of satisfied norms and $\Lambda = \{\lambda_1, \lambda_2, ...\}$ is the set of behavioral modalities of a behavior in the set B.

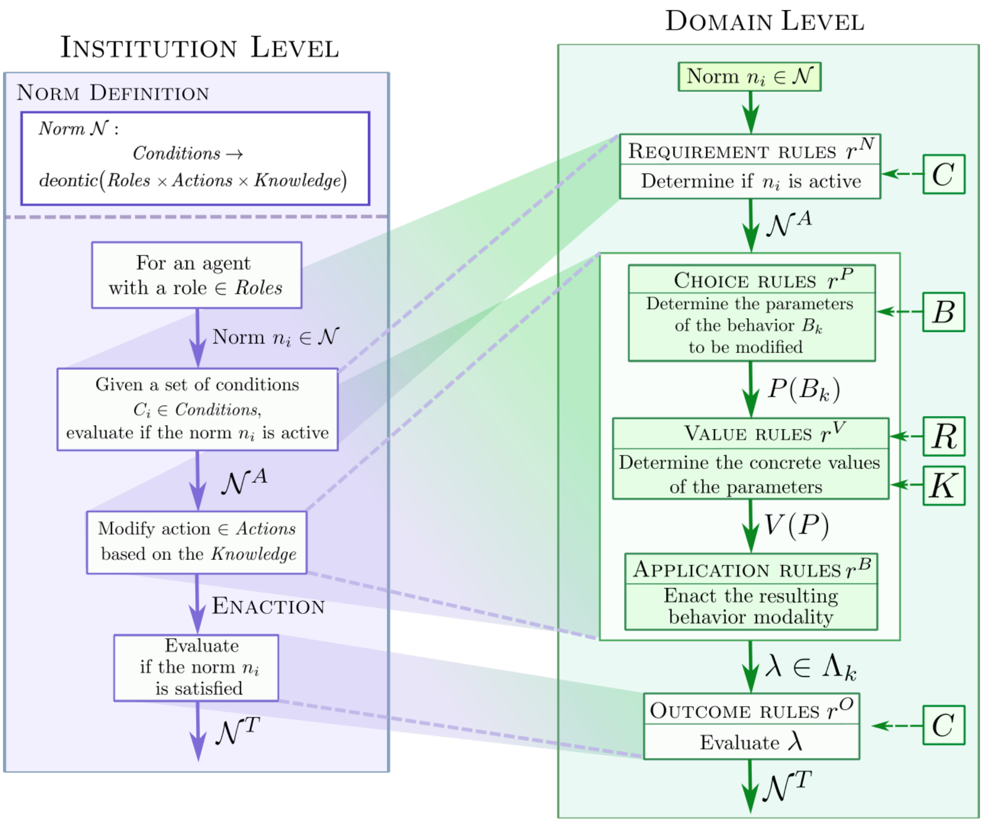

Figure 13.3 – Illustration of norm realization.

Rules r^N activate norms based on grounded conditions C, r^P give a choice of parameters that should be altered for the given behavior, r^V provide methods to set the parameter values based on knowledge K and state variables R, r^B enact behavior procedures for applying parameter values, establishing the behavior modality, and r^O verify norm compliance after the behavior changes the world state. For convenience, the conditions of the requirement rules will be referred to as the *activation conditions*, and the conditions of the outcome rules as the *outcome conditions*. The relations between the rules and the other elements of the domain are shown in Figure 13.3.

Definition 13.3. Formally, the norm realization at the domain level is a collection of the rules:

$$\mathcal{K} = \langle r^N, r^P, r^V, r^B, r^O \rangle$$

Next, we discover how norm realization is carried out in practice by realizing an example of norms in one of the collective behaviors we employ in this thesis.

13.1.5 Illustrative Example

Consider the following norm: *"all robots must respect human personal spaces"*. According to the Stage I of the institutional level, we can identify the institutional components forming the norm semantics of Definition 11.2:

$$\mathcal{N}: \ \textit{Conditions} \rightarrow \text{deontic}(\textit{Roles} \times \textit{Actions} \times \textit{Knowledge})$$
$$n: \ \text{(when close to human)} \rightarrow \text{must}(\textit{robots},\ \textit{respect},\ \textit{personal spaces})$$

Now consider a robot F_i with a role of a Follower moving in a formation using the MoveInFormation behavior described in Chapter 5. As the above norm n refers to all robots irrespective of their roles, we can conclude that n applies also to F_i as well.

The remaining steps of the norm realization are performed at the domain level.

A) A requirement rule r^N evaluates the condition CLOSE_TO_HUMAN directly on the state variables (for example, by calculating the distance between F_i and every human in the environment. If the condition is satisfied, norm n is active.

B) Norm n does not state explicitly the action it refers to. One could state explicitly that "*all the robots moving in a formation must respect...*", but this would defeat the purpose of maximal generalization of norms. For the MoveInFormation behavior a choice rule r^P decides that it is the parameter *RepulsionWeights* $\mathcal{W}_{\mathcal{R}}$ – a parameter that generates a repulsive force driving F_i away from the human H_h – that should be modified according to Equation (5.8).

C) For a human H_h, the robot F_i calculates the repulsion forces through the value rules r^V. They determine the value of the parameter *RepulsionWeights* based on the distance between H_h and F_i, which can be obtained from state variables R and based on the value of *personal space* of that particular person that is an element of the institutional knowledge grounded to the domain.

D) The chosen parameter value is applied directly to the MoveInFormation behavior with the application rule r^B to result in behavior modality that is expected to satisfy the norm n.

E) Norm is verified at the next time step with the outcome rules r^O by evaluating the outcome conditions, in this case RESPECTED_PERSONAL_SPACE, evaluated based on the distance state variable.

13.2 Norm Adaptation

The above considerations regard an institution as a static entity existing well before and well after the agents employ the institutional norms. However, as institutions of human environments are subject to constant alterations, similarly robot institutions can be modified by cooperative decision-making, or by an individual with a power to do so, while the experience gained by the robots participating in a social environment can lead to gradual modification of the existing rules.

In this section we propose adaptive norm realization as a mechanism by which the robots can modify their institutional interpretation through a self-regulated process based on the experience gathered in a given domain. This can be seen as equivalent to adjustment of how humans interpret norms, where after familiarization with an environment, we slowly learn to modify our behaviors to better fit the circumstances. For example, even if a speed limit in a given area is 50 km/h, we might decide to go slower once we learn that children frequently play nearby.

Institutions can dynamically accommodate the evolution of a given domain by the means of adaptation rules – rules allowing to leverage robots' experience and humans increased confidence in robots' actions, as well as adjust to continuous environmental changes. With the adaptive norm realization it is possible for the robot to track and adjust to these changes through a self-regulated process.

To accommodate norm adaptation, one has to introduce a set of measures quantifying applicability of the norm in the environment. The set $L: 2^R \rightarrow R$ referred to as the robot's *experience measure* is defined over state variables and includes measures used for norm evaluation. For example, one can monitor the acceptance[1] $\textsc{acceptance}_{h,i} = \{0,1\}$ of a robot R_i that is believed to be increasing with the amount of time human H_h interacts with it. Based on this measure one can modify the values of the behavior parameters, such as the strength of *RepulsionWeights*, accordingly.

The value rules encode information customized to the given domain by the grounding, but they are given a priori and do not reflect the response of the environment. Moreover, it is not known whether the chosen value is optimal in the given situation before testing it and evaluating the result. In order to dynamically modify the values of the behavior parameters we introduce a new set of rules – *adaptation rules* r^L. With the adaptation rules, parameter values that were chosen by the value rules are modified to accommodate the reaction of the environment.

More precisely, if at time t the value rules determine the value of the parameter p_k of a norm n_i to be $v_{p_k}(t)$, the adaptation rules change the value $v_{p_k}(t)$ as a function of the experience measure $l(\tau) \in L_i$ collected at time τ in the past, i.e. $v_{p_k}(t) \leftarrow g(l(\tau))v_{p_k}(t)$. The function g from now on will be referred to as the *adaptation function*.

Formally, we introduce a new step in the norm realization at the domain level given in Definition 13.2:

C') Parameter values can be adapted on the basis of the behavior's performance

with the corresponding adaptation rules:

1 Such measure can be obtained from video annotations, questionnaires, or other evaluation methods we will use in later chapters.

ADAPTATION RULES $r^L : L \times V \rightarrow V$

Note that the adaptation rules are applied right after the value rules.

Definition 13.4. Norm realization at the domain level with the adaptation rules takes the form:

$$\mathcal{K} = \langle r^N, r^P, r^V, r^B, r^O, r^L \rangle$$

With norm adaptation we effectively provide a feedback to the norm realization about the outcome of enforcing the norm, as the values are updated depending on the output of the system in the previous time step.

Coming back to the example of the norm for respecting human PersonalSpace, if we assume that the value of the weight of the repulsive force preventing the robot from invading that space is w_R, we can modify the strength of the force as follows. If the determined value $w_R(t-1)$ resulted in the robot failing to obey the norm, i.e. when IN_PERSONAL_SPACE was true $l(t-1) = 1$, we could take a simple function $w_R(t) \leftarrow (1 + l(t-1))\, w_R(t) = 2 w_R(t)$, which doubles the force strength every time the robot intrudes the PersonalSpace of the human. Note that even though the procedure for calculation of the parameter value is determined a priori through the value rules, the values themselves can vary with time. In the example above, the strength of the force depends on the distance between the robot and the human, while the adaptation rule further increases the weight to improve the safety margin.

Summary

With this chapter we concluded the description of our framework for abstraction, encapsulation and formalization of social norms in robotics systems. We have clarified that while the institutional abstraction of the normative concepts from the concrete domain and their association through grounding allows for reuse of institutions, it is also not enough to accomplish realization of abstract norms, formulated in human language, in robotic systems.

Norm realization – a mechanism for translation of general, language-defined norms has three main strengths. First, it provides the means for interpretation of a norm given in the form of a sentence in the institutional terms. Second, akin to the grounding, it establishes the association between the institutional level and the domain level. Finally, the rules of the norm realization at the domain level allow for analysis and interpretation of the norm elements at the level of robotic language, i.e. parameters, behaviors and algorithms. Instead of the design being driven by norms, norms are imposed as constraints on the parametrization of already existing behaviors in a plug-and-play manner.

The proposed norm realization not only makes it possible to impose language-defined norms onto the robot behaviors, but also to reuse norms across different domains and across different institutions. In the chapters that follow, we show that norm realization addresses the problems of reusability, modularity and scalability, the main challenges of the state-of-the-art approaches to normative robot behaviors we discussed in Chapter 10.

14 Overview of the Case Studies

OUR institutional framework is validated through three extensive case studies, each aimed at investigating different aspect of our methodology. This chapter provides an overview of the case studies $\mathcal{C}_I$, $\mathcal{C}_{II}$ and $\mathcal{C}_{III}$ presented in Chapter 15, Chapter 16 and Chapter 17, respectively. The overview is summarized in Table 14.1 and will be detailed in the sections that follow.

The case studies presented in this thesis are shaped so as to showcase different aspects of our formalism, including a) the ability to reuse the norms and institutions in a plug-and-play manner, b) the capacity of translating abstract norms into the robotic language, and c) modularity and scalability of the norms that can be composed together without the need of redesigning the behaviors.

The common scope of the case studies is the context of social, multi-robot navigation. In $\mathcal{C}_I$, we focus on navigation among humans in situations when humans and robots co-exist but do not cooperate. In $\mathcal{C}_{II}$ and $\mathcal{C}_{III}$, we deploy mixed groups, with robots guiding humans in $\mathcal{C}_{II}$, and with humans guiding robots in $\mathcal{C}_{III}$. The common objective of the case studies is that of validation and demonstration of the institutional framework.

14.1 Institutions, Norms, Domains

As summarized in Table 14.1, in each case study we deploy a combination of institutions and norms, and evaluate them in the respective domains. In $\mathcal{C}_I$, the NAVIGATING-AMONG-HUMANS institution is evaluated in three different domains. In $\mathcal{C}_{II}$, we deploy two different institutions, a STEERING institution for guiding humans to a destination, and TUTORING, where robots give a tour of the environment, providing information about objects of interest. In addition, the norms of the STEERING institution are tailored to operate over two possible cases, first for when humans are following, second for when at least one human fails to follow the robot. Finally, in $\mathcal{C}_{III}$ the ROBOT-GUIDANCE institution is evaluated with two different behaviors (and so, also two different domains) – formation and flocking. The domain of $\mathcal{C}_{III}$ is the same as one of the domains of $\mathcal{C}_I$ purposefully – that way we can demonstrate that it is not only possible to apply one institution to different domains, but also that different institutions can be plugged-and-played in the same domain.

14.2 Experimental Settings

As shown in Table 14.1, the first case study $\mathcal{C}_I$ focuses on navigation among humans, while the studies $\mathcal{C}_{II}$ and $\mathcal{C}_{III}$ are designed to operate over mixed groups of humans and robots. The relations between the institutions of the case studies have been described in Section 11.3 and illustrated in Figure 11.3. In particular, the NAVIGATING-AMONG-HUMANS institution is designed for situations, where humans and robots share the same environment, but do not cooperate. In contrast, in $\mathcal{C}_{II}$ and $\mathcal{C}_{III}$ humans and robots form mixed groups. In $\mathcal{C}_{II}$ robots are guiding a group of humans. In $\mathcal{C}_{III}$ humans are helping the robots to navigate through a complex area by serving as the robot leaders. In all three case studies, the robot behaviors are based on the formation or the flock behaviors, as shown in Table 14.1.

Methods in $\mathcal{C}_I$ and $\mathcal{C}_{II}$ are evaluated with volunteers who agreed to participate in our experiments and provide invaluable feedback regarding their perception of robots' sociality. Norm adaptation was employed in case $\mathcal{C}_{III}$ with the aim of improving robustness of navigating in formation with human leaders in complex spaces. In $\mathcal{C}_I$ we evaluated the multi-robot cooperative localization algorithm from Chapter 8 in a highly dynamic setting, where a formation shape undergoes constant modification as a result of accommodating human presence.

Summary

In this short chapter we have introduced the three case studies of this thesis aiming at validation of our institutional framework and showcasing its main advantages. The primary objective of the proposed selection of case studies is to encompass a number of social contexts the robots moving in groups might be situated in, namely navigating among uncooperative humans, mixed groups where the humans serve either as leaders or followers. Consequently, we introduce institutions tailored for such contexts.

We will frequently revisit this chapter, and in particular Table 14.1, to always withhold the global objective of each case study when describing the implementation details.

CASE STUDY	$\mathcal{C}_I$	$\mathcal{C}_{II}$	$\mathcal{C}_{III}$
INSTITUTIONS, NORMS, DOMAINS			
Institution Name	NAVIGATING-AMONG-HUMANS	STEERING&TUTORING	ROBOT-GUIDANCE
Focus	Social Navigation	Robust & Social Navigation	Robust Navigation
Norms Institutions Domains Behaviors	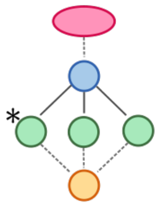	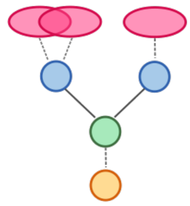	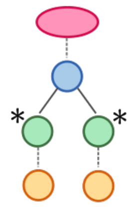
EXPERIMENTAL SETTINGS			
Mixed Societies	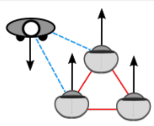	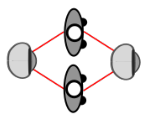	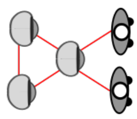
Formation | Flocking	● ○	● ○	● ●
Norm Adaptation	○	○	●
Participative Study	●	●	○
Tracking	●	○	○
Facility	Jordils	Örebro	Jordils

Table 14.1 – Overview of the case studies. INSTITUTIONS, NORMS, DOMAINS. The graphs illustrate the high-level overview of the norms (shown in magenta), institutions (blue), domains (green), and behaviors (yellow), operating in each case study. In $\mathcal{C}_I$, the NAVIGATING-AMONG-HUMANS institution is evaluated in three different domains. In $\mathcal{C}_{II}$, we deploy two different institutions, STEERING and TUTORING. The norms of the STEERING institution operate over two possible cases, therefore we can view them as two sets of norms for one institution. In $\mathcal{C}_{III}$ the ROBOT-GUIDANCE institution is evaluated with two different behaviors, and so, also two different domains. A star * indicates identical domain as in $\mathcal{C}_I$. In terms of experimental settings, $\mathcal{C}_I$ focuses on navigation among humans, while the studies $\mathcal{C}_{II}$ and $\mathcal{C}_{III}$ are designed to operate over mixed groups of humans and robots. In $\mathcal{C}_I$ and $\mathcal{C}_{II}$ the robots perform formation control, in $\mathcal{C}_{III}$ we compare both formations and flocks. In $\mathcal{C}_I$ and $\mathcal{C}_{II}$ we perform participative studies, while in $\mathcal{C}_{III}$ we demonstrate norm adaptation. Finally, in $\mathcal{C}_I$ we evaluated the multi-robot cooperative localization algorithm from Chapter 8.

15 Case Study: Social Navigation Among Humans

SOCIAL navigation among humans requires that the robots, on top of respecting navigational constraints, satisfy relevant social norms. While basic navigation methods are easily generalizable to a large variety of environments, introducing social aspects to robot behaviors requires a good understanding of the social context and a mapping between the context and the manner in which the norms are applied. In the institutional formalism, such mapping is accomplished through grounding and norm realization.

In the following case study, we propose an institution called NAVIGATING-AMONG-HUMANS for governing the social aspects of navigation among humans, designed for use when humans and robots share the same environment, but perform independent missions. We employ this institution in three different social contexts and show that the same institutional abstraction can yield very different results for different domains, without the need of redesigning it. With this plug-and-play property of the institutional formalism, we perform the norm-to-context mapping without the need to resort to heuristics or full re-parametrization for each new domain.

We carry out extensive evaluation in controlled settings with up to two humans. Subsequently, we perform participative studies to obtain subjective assessment of robots' sociability. Finally, we leverage the high complexity of the experimental dynamics to challenge the cooperative localization method presented in Chapter 8. With this last set of experiments we consolidate all elements of this thesis – adaptive formations, cooperative tracking, and sociality of robot behaviors.

HIGHLIGHTS

- **Reusability.** The NAVIGATING-AMONG-HUMANS institution is applied to three different domains.
- **Modularity.** The domains represent diverse social contexts, each stipulating disparate balance between team's efficiency in completing the mission and team's sociability. The

norms however remain the same and are applied to one behavior (see Table 14.1), which does not need to be readjusted for each domain.

- **Versatility.** In Chapter 17 we use one of the three domains presented in this chapter, but different institutions (and norms), demonstrating the plug-and-play principle of the formalism.
- **Customization.** In each domain, we distinguish individual human roles with distinct personal preferences to which robot behaviors are tailored to.
- **Generability.** The selection of the employed norms targets a large variety of social aspects, ranging from human comfort achieved through robots' navigational compliance, to understandability of robots' intentions reflected through gestures, expressions, and sounds.
- **Participative study.** We analyze subjective performance measures collected in the form of questionnaires with the aim of understanding individual human preferences with regard to robot sociality and their perception of the norm-following robot.
- **Consolidation.** We perform experiments that consolidate all elements of this thesis presented in Part II and Part III, namely adaptive formations with dynamically changing topology, cooperative tracking, and sociality of robot behaviors achieved through the institutional framework.

15.1 Representation of Social Context

The types of social contexts are diverse in their nature, and so, the levels of accommodation of robot behaviors required to satisfy social norms in these contexts vary. To represent such levels of conformation we draw a line between two extremities: on the one side of the spectrum, we consider a *social* context, where the priority of the robots is to accommodate human needs, on the other side, we consider an *efficient* context, where the priority of the robots is to accomplish a given task by minimizing disruption engendered by the human activities. Finally, we look into a the *socially-efficient* context between the two extreme cases.

We propose three domains representing the above social contexts with their basic characteristics in Table 15.1.

SOCIAL CONTEXT	CHARACTERISTICS	DOMAIN	HUMAN NAME
social	accommodating, human-oriented, safe, calm	SCHOOL	teacher, child
efficient	task-oriented, cost-effective, fast, productive	WAREHOUSE	worker, visitor
social-efficient	a combination of the two above	HOSPITAL	staff, patient

Table 15.1 – Social contexts explored in this case study and the domains that represent them. In each domain, we distinguish two human functions, for which robots personalize their behaviors with regard to human's preferences through norm realization.

So defined three domains, SCHOOL, WAREHOUSE and HOSPITAL, will be governed using the same institution – NAVIGATING-AMONG-HUMANS, which, however, given different groundings and norm realization, will yield different results, adjusted to the domain.

In each domain, we introduce two human agents, each playing a different function and with different social preferences regarding the robots (e.g., how far away a robot should stay, how fast it should move). Each of the humans will be treated accordingly, on the basis of the distinct role the institution assigns to him or her. Note that although names such as teacher or child, are general, they are part of the domain, and they refer to a specific person. Without loss of generality, we could have used labels such as teacher-Steven or child-Ariana instead, but we shorten them to simplify the notation.

15.2 The NAVIGATING-AMONG-HUMANS Institution

The institutional formalism allows for a large flexibility in terms of the granularity of institutions, the choice of which is typically dictated by the application requirements. In this case study, the institutional elements of NAVIGATING-AMONG-HUMANS must be able to encompass a large variety of domains that the institution applies to and so, this institution is more general than the ones presented in the case studies $\mathcal{C}_{II}$ and $\mathcal{C}_{III}$ (see Figure 11.3). As an alternative to the solution we propose, one could have defined three different institutions, one per each class of the domains: one institution designed to govern all different instances of SCHOOL, another all instances of WAREHOUSE, and a third one governing those of HOSPITAL; such institutions could *inherit* part of their definition from a related, generic institution.

In this section, we characterize the NAVIGATING-AMONG-HUMANS institution and describe how to apply the institutional abstraction to the distinct domains, while ensuring adequate customization for each domain and personalization tailored to the specific human role mentioned above.

15.2.1 Actions and Roles

At the institutional level, we define two generic human roles: 1) Acquainted-Human – a person that is assumed to be more familiar and experienced with the robots (such as teacher, worker, or staff) and Unacquainted-Human – a person with less exposure to the robots and thus little understanding of their behaviors (such as child, visitor, or patient). Human actions are not relevant in this case study, and they will not be defined formally in the institution.

Robot roles are dictated by the function a robot plays in the group movement, in this case a navigation in formation. We define a Leader role for a robot that guides the formation through an environment with a Guide action, and a Follower role, for a robot that follows with a Follow action.

15.2.2 Norms

Social navigation literature [31] offers a large variety of norms applicable to robots navigating in human-populated environments. In the following case study, we choose a selection of norms along two axes: a *spatio-temporal* axis, constituting the operational space of the robot, and *human-oriented* axis. The norms are listed in Table 15.2.

SPATIO-TEMPORAL NORMS		
Spatial (static)	n_1	*"Robots should not enter activity-critical areas"*
Spatial (dynamic)	n_2	*"Robots must respect human personal spaces"*
Temporal	n_3	*"Robots should not disturb during quiet activity times"*
HUMAN-ORIENTED NORMS		
Naturalness	n_4	*"When navigating among humans, robots should assume appropriate speed"*
Friendliness, appropriateness	n_5	*"Robots should indicate their intentions through facial expressions, gestures and sounds"*
Clarity, transparency	n_6	*"Robots should offer informative announcements"*
Cultural convention	n_7	*"When interacting with humans, robots should direct their gaze towards the interaction partner"*

Table 15.2 – Norms of the NAVIGATING-AMONG-HUMANS institution.

While the aforementioned norm selection is not exhaustive, we believe that it represents well the spectrum of social norms applicable to robots. Spatial norms n_1 and n_2 represent spatial constraints that should be respected by a robot. While n_2 is inherently dynamic, i.e. the constraint is fixed with respect to a dynamic point in the environment (in this case, a human), n_1 can be also used for temporarily constant environment constraints, such as permissions to go through some areas that can be granted or elevated during run time. Temporal norms similar to n_3 do not allow for plan composition (see [9] for examples of temporal norms that can be used for constraint satisfaction planning), but they provide means to introduce elements that require an interpretation of time. For example, norm n_3 can be used for scheduling robot behaviors when a general plan is known a priori. Note that each robot interprets the social norms individually, and given the role it is assigned to, determines whether it is obliged to comply to the given norm.

The human-oriented categories, including naturalness, appropriateness and clarity of robot behavior as well as cultural conventions are intended to improve acceptance of the robots. Norms n_5 and n_6 do so by enhancing human-robot interactions with basic interactive features, including gestures, facial expressions, sounds that convey information about robot's state and spoken dialog. Norms n_4 and n_7 introduce simple cues designed to give impression of the robot having a social understanding from the point of view of the observing human.

As explained in Chapter 11, interpretation of the above norms is based on the institutional components, namely *Roles*, *Conditions*, *Knowledge* and *Actions* according to the Definition 11.2:

$$\mathcal{N}: \textit{Conditions} \rightarrow \text{deontic}\,(\,\textit{Roles} \times \textit{Actions} \times \textit{Knowledge}\,).$$

Norms $n_1 - n_7$ rewritten according to this syntax are:

n_1 : ∅ → should not(*robots*, *enter*, *activity-critical areas*)

n_2 : (when close to human) → must(*robots*, *respect*, *personal spaces*)

n_3 : (during quiet activity times) → should not(*robots*, *disturb*, *quiet activity times*)

n_4 : when navigating among humans → should(*robots*, *assume*, *appropriate speed*)

n_5 : (when appropriate) → should(*robots*, *indicate their intentions*, *through expressions, gestures, sounds*)

n_6 : (when appropriate) → should(*robots*, *offer*, *informative announcements*)

n_7 : when interacting with humans → should(*robots*, *direct gaze towards the interaction partner*)

With the above, we complete Stage I of norm realization at the institutional level (see Definition 13.1).

15.2.3 Conditions and Knowledge

Institutional *Conditions*, including *Activation Conditions* for norm activation and *Outcome Conditions* for norm verification are presented in Table 15.3:

	Activation Conditions	*Outcome Conditions*
n_1	-	IN_ACTIVITY_CRITICAL_AREA
n_2	CLOSE_TO_HUMAN	RESPECTED_PERSONAL_SPACE
n_3	QUIET_ACTIVITY_TIME	RESPECTED_QUIET_ACTIVITY_TIME
n_4	IN_HUMAN_POPULATED_ENVIRONMENT	BELOW_SPEED_FACTOR
n_5	INTERACTING_WITH_HUMAN REACHED_GOAL CLOSE_TO_HUMAN QUIET_ACTIVITY_TIME	-
n_6	FIRST_TIME_PERSON_ENCOUNTERED INVADED_INTIMATE_SPACE	-
n_7	INTERACTING_WITH_HUMAN	-

Table 15.3 – *Conditions* of the institution NAVIGATING-AMONG-HUMANS.

Conditions strongly depend on institutional *Knowledge*, e.g. in order to evaluate a condition INVADED_INTIMATE_SPACE a robot must understand the concept of intimate space. Furthermore, when grounding the conditions to the state variables R of the domain, the *actual value* of the intimate space for the particular person must be known. Providing such information is the main purpose of the institutional *Knowledge*.

Knowledge	***Grounded Knowledge K***	**Symbol**
ActivityCriticalArea	area_x	Φ_{ACS}
SocialSpace	social_distance_x	$\Delta_{S,h}$
PersonalSpace	$\text{personal\_distance}_x$	$\Delta_{P,h}$
IntimateSpace	$\text{intimate\_distance}_x$	$\Delta_{I,h}$
QuietActivityTime	time_interval_x	T_{QAT}
AppropriateSpeed	speed_factor_x	s_A
InteractionMode	$\text{interaction\_mode}_x$	IM
SpeechMode	speech_mode_x	SM
FormationShape	formation_shape_x	FS
FormationBias	formation_bias_x	b

Table 15.4 – *Knowledge* and knowledge grounding K of the institution NAVIGATING-AMONG-HUMANS. We use the subscript $(\cdot)_x$ to abbreviate variables with multiple instances x.

Knowledge of the institution NAVIGATING-AMONG-HUMANS includes the elements listed in Table 15.4 (first column). SocialSpace, PersonalSpace and IntimateSpace represent the proxemics [63] – preferences for the distance relation between two agents during interactions we described in Chapter 5. ActivityCriticalArea is an affordance space where humans habitually engage in activities that should not be disturbed by robots, for instance, the area in front of a blackboard in schools. Similarly, QuietActivityTime represents time interval when, habitually, humans should not be disturbed, for instance, classroom time in school. InteractionMode and SpeechMode allow a robot to choose appropriate expressions, sounds and gestures given the circumstances. The circumstances and the mapping between the mode and the circumstance are determined during the grounding procedure. Finally, AppropriateSpeed, FormationShape and FormationBias allow the robots to choose the appropriate formation parameters.

15.2.4 Formal Definition

The NAVIGATING-AMONG-HUMANS institution is formally represented as:

$$\begin{aligned}
&\textbf{NAVIGATING-AMONG-HUMANS} = \Big\langle \\
&\quad \textit{Norms} = \{n_1,\ n_2,\ n_3,\ n_4,\ n_5,\ n_6,\ n_7\} \\
&\quad \textit{Roles} = \{\text{Leader, Follower, Acquainted-Human, Unacquainted-Human}\} \\
&\quad \textit{Actions} = \{\text{Guide, Follow}\} \\
&\quad \textit{Conditions} = \{\text{CLOSE_TO_HUMAN, QUIET_ACTIVITY_TIME}, \ldots\} \\
&\quad \textit{Knowledge} = \{\text{ActivityCriticalArea, SocialSpace, PersonalSpace}, \ldots\}\Big\rangle
\end{aligned}$$

With this general definition of institution, we proceed to describe how this single institution is applied to three different domains.

15.3 Domain and Grounding

Each domain presented in Section 15.1 includes the following elements:

$$A = \{\mathsf{mbot}_1,\ \mathsf{mbot}_2,\ \mathsf{mbot}_3,\ \mathsf{human}_\mathsf{x}\}$$
$$B = \{\mathsf{MoveOnTrajectory},\ \mathsf{MoveInFormation},\ \mathsf{HumanBehavior}_\mathsf{x}\}$$
$$R = \{\mathsf{time},\ \mathsf{pose}_\mathsf{x},\ \mathsf{speed}_\mathsf{x},\ \mathsf{distance}_\mathsf{x},\ \mathsf{location}_\mathsf{x},\ \mathsf{sound}_\mathsf{x},\ \mathsf{expression}_\mathsf{x},\ldots\}$$
$$K = \{\mathsf{social\_distance}_\mathsf{x},\ \mathsf{speed\_factor}_\mathsf{x},\ \mathsf{interaction\_mode}_\mathsf{x},\ \mathsf{formation\_shape}_\mathsf{x},\ldots\}$$
$$C = \{C_1,\ C_2,\ C_3,\ C_4,\ C_5,\ C_6,\ C_7\}$$

where we use the subscript $(\cdot)_x$ to abbreviate variables with multiple instances x. Elements in green are specified for each domain separately. According to the initial description in Table 15.1 the element $\mathsf{human}_\mathsf{x}$ is $\{\mathsf{teacher},\ \mathsf{child}\}$ for the SCHOOL domain, $\{\mathsf{worker},\ \mathsf{visitor}\}$ for the WAREHOUSE domain, and $\{\mathsf{staff},\ \mathsf{patient}\}$ for the HOSPITAL domain.

Since the human actions are undefined at the institutional level, we will leave the element $\mathsf{HumanBehavior}_\mathsf{x}$ unspecified. It is, however, entirely possible that what humans do is institutionally relevant, for example, in the case study in Chapter 17, where humans are the leaders of the formation. If so, $\mathsf{HumanBehavior}_\mathsf{x}$ would become playing, working, visiting, etc. Furthermore, depending on the design choices, some conditions, such as QUIET_ACTIVITY_TIME could be directly linked to the current human behavior, for example human activity is not to be disturbed when a teacher is teaching.

Grounded conditions C_1 - C_7 are the institutional conditions presented in Table 15.3 and evaluated over the state variables of the domain. For example, the activation condition CLOSE_TO_HUMAN, assessed by a robot mbot_i, is true if a distance between mbot_i and one of the humans $\mathsf{human}_\mathsf{j}$ falls below a threshold $\mathsf{personal\_distance}_\mathsf{j}$ dictated by the grounded knowledge (see Table 15.4). While the meaning of most of the conditions listed in Table 15.3 is unambiguous, we would like to point out that the condition INTERACTING_WITH_HUMAN is true when a human is facing the robot and the human and the robot are within interaction distance from each other (interaction distance is proportional to the personal_distance).

The full list of grounded knowledge K is presented in Table 15.4. QuietActivityTime is grounded to a time interval chosen a priori. Grounded InteractionMode and SpeechMode are lookup tables mapping a triggering condition (such as that the robot has REACHED_GOAL or the robot is INTERACTING_WITH_HUMAN) to the corresponding desirable facial expression, gesture and sound. The mapping is performed in the process of norm realization, and will be discussed in the next section.

In summary, in our case study we distinguish three domains, which leads to having three institution-to-domain groundings: $\mathcal{G}^S$ for the SCHOOL domain, $\mathcal{G}^W$ for the WAREHOUSE domain,

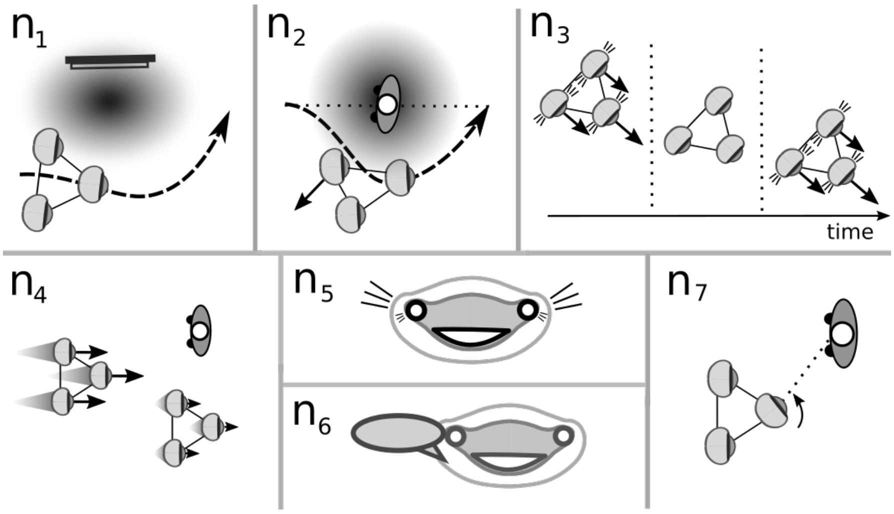

Figure 15.1 – Overview of the norms and their realization in the concrete system.

and $\mathcal{G}^H$ for the HOSPITAL domain. Customized results achieved with a single institution in such diversity of domains are accomplished by a combination of a) the role-to-agent grounding $\mathcal{G}_A$, b) the particular values that the grounded knowledge K assumes and c) the norm realization. While details on the latter will be given in the next section, the general description of the grounding is as follows.

$$
\begin{aligned}
\mathcal{G}_A &= \{(\mathsf{Leader,\ mbot_1}),(\mathsf{Follower,\ mbot_2,\ mbot_3})\}\\
\mathcal{G}_A^S &= \mathcal{G}_A \cup \{(\mathsf{Acquainted\text{-}Human},\ \mathbf{teacher}),\ (\mathsf{Unacquainted\text{-}Human},\ \mathbf{child})\}\\
\mathcal{G}_A^W &= \mathcal{G}_A \cup \{(\mathsf{Acquainted\text{-}Human},\ \mathbf{worker}),\ (\mathsf{Unacquainted\text{-}Human},\ \mathbf{visitor})\}\\
\mathcal{G}_A^H &= \mathcal{G}_A \cup \{(\mathsf{Acquainted\text{-}Human},\ \mathbf{staff}),\ (\mathsf{Unacquainted\text{-}Human},\ \mathbf{patient})\}\\
\mathcal{G}_B &= \{(\mathsf{Guide,\ MoveOnTrajectory}),\ (\mathsf{Follow,\ MoveInFormation})\}\\
\mathcal{G}_C &= \{(\text{CLOSE\_TO\_HUMAN},\ c_{11}^a),\ (\text{BELOW\_SPEED\_FACTOR},\ c_{41}^o),\dots\}\\
\mathcal{G}_K &= \{(\mathsf{PersonalSpace,\ personal\_distance_x}),\dots\}
\end{aligned}
$$

with $c_{ab}^x \in C_a$, where a denotes the index of a norm the condition pertains to, b is the index of the grounded condition within set C_a, and $x \in \{a, o\}$ allows to distinguish between the activation and the outcome conditions, as shown in Table 15.3.

After all elements of the domain are defined and matched with the institutional components through grounding, the domain-level stage of norm realization can take place.

15.4 Norm Realization

Steps of norm realization at the domain level presented in Definition 13.2 include A) norm activation, B) selection of behavior parameters to be modified, C) assignation of the respective parameter values, D) application of the parameters with the chosen values to the behavior, and E) evaluation of norm satisfaction. In this section, we provide a high-level description necessary for understanding the norm realization and we list the key elements that allow for adaptation of the institution to the social context at the grounding level[1]. The norms are illustrated in Figure 15.1.

Norm n_1: Activity-critical areas

In our implementation, norm n_1 is always active, as the affordance places are known a priori. One can however imagine that affordances can be bound to human activities, for example a robot should not move in front of a TV when someone watches it, but it can do so otherwise. In that case, the activation condition for that particular affordance would be TV_IN_USE.

In order to avoid an affordance area, the position and shape of which is given by Φ_{ACS} in K, the formation leader modifies the *TrajectoryShape* parameter of the MoveOnTrajectory behavior by adding social cost to the speed map used by the FMM path planning method (as explained in Section 5.5.1).

Norm n_2: Personal spaces

Norm n_2 is active when CLOSE_TO_HUMAN, i.e. when robot is within a distance from human smaller than threshold proportional to the personal_distance $\Delta_{S,h} \in K$ of person h. The threshold must be larger than personal_distance so that the robot can act *before* entering the space and to prevent it.

The Leader robot modifies the *TrajectoryShape* parameter by shifting its path away from the human, proportionally to the distance between them, until the condition is false again (see Section 5.5.1). A Follower, by utilizing the *RepulsionWeights* $\mathcal{W}_{\mathcal{R}}$, generates a repulsive force, driving it away from the human. Further details are provided in Section 5.2.2.

Norm n_3: Quiet activity times

Condition QUIET_ACTIVITY_TIME is triggered at known-in-advance time intervals provided to the robots by the time_interval element $T_{QAT} \in K$. For this norm the Leader robot stops the formation by setting the *TrajectorySpeed* to zero throughout the duration of the activity time.

Norm n_4: Appropriate speed

The activating condition IN_HUMAN_POPULATED_ENVIRONMENT is always true if a human takes part in an experiment. The overall formation speed is adjusted adequately to the social

1 We provide an example of algorithmic implementation of norm realization in the case study $\mathcal{C}_{II}$ in Chapter 16 and a detailed description of the norm realization for the case study $\mathcal{C}_{III}$ in Appendix E.

context (i.e. to the domain) by moderating the *TrajectorySpeed* parameter of the Leader, while the Follower speeds are mitigated by adjusting the *ControlGain* parameter K_u (see Section 5.2) proportionally to the speed_factor $s_A \in K$. Note that we truncate the speeds for safety purposes, as explained in Section 12.2.1, but decided to not limit them any further, as the robots need enough momentum to return to the formation after large disturbances (such as the one created as a result of the human presence) and for securing adequate response to the repulsion forces.

Norm n_5: Expressions, gestures, sounds

We distinguish four conditions that trigger relevant facial expressions, gestures and sounds, otherwise a robot assumes a default interaction mode. The mapping between condition and interaction_mode IM$\in K$ is shown in Table 15.5. Note that the same conditions that trigger norms n_2 and n_3 activate an interaction mode.

DEFAULT	INTERACTING_ WITH_HUMAN	REACHED_ GOAL	CLOSE_ TO_HUMAN	QUIET_ ACTIVITY_ TIME
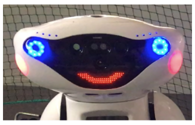				
no sound	recognition sound	success sound	recognition sound	no sound
no gesture	no gesture	arms up and down repeated	arms up and down	no gesture

Table 15.5 – Mapping between grounded conditions (top row) and robot interaction mode IM $\in$ K, which includes facial expressions, sound and gestures.

Norm n_6: Informative announcements

We define two conditions activating robot speech. First, a robot utters "Hello" upon triggering of the FIRST_TIME_PERSON_ENCOUNTERED activation condition. Second, it says "Sorry" when INVADED_INTIMATE_SPACE, i.e. when it intrudes the intimate space of a human. In some social contexts, robot speech might be undesirable - as an example of that we disable it in the WAREHOUSE domain with a suppressing condition NOT_IN_WAREHOUSE. The mapping between the condition and the utterance, similarly as in the case of n_5, is given by a lookup table SM$\in K$ embedded in robot's code.

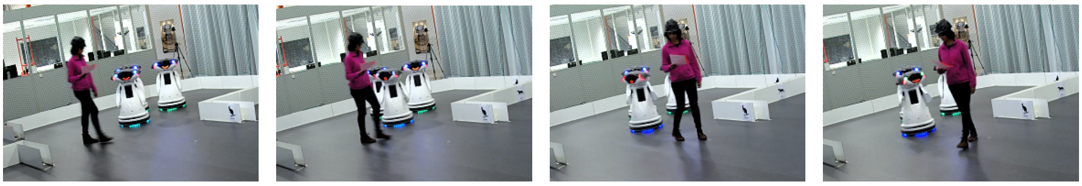

Figure 15.2 – Example of the gaze control the Leader robot engages in.

Norm n_7: Gaze control

To convey the impression of robots awareness of human presence, the Leader robot turns its head towards the human when INTERACTING_WITH_HUMAN, i.e. when passing close to the human and the human gaze is directed at the robot. We choose to direct the gaze of only one robot, as having all the robots in the formation directing their gaze at a human could be intimidating. Example of gaze control of the Leader robot is shown in Figure 15.2.

Summary

In the above examples the activation conditions are composed using an OR operator, i.e. only one condition in C_i is necessary to trigger norm activation of n_i. In our implementation, grounded knowledge forms a lookup table, the particular values of which will be provided in the experiments section, in Table 15.6. Ideally, knowledge would be encoded in open source databases, and be filled in by researchers with their experimental findings and shared publicly. In our work, knowledge is loosely based on the literature and empirical tests that yield the most promising results.

15.5 Experimental Campaign Overview

Experiments are performed in the Jordils arena, described in Chapter 2. A team of three robots moves in a formation on a path between three waypoints. Experiments involve up to two humans, each acting independently of the robots and of the other people. A participant plays a game involving a number of waypoints delineated on a map and to be visited in a particular order (randomized among the runs). Such design of the game allows us to maximize the number of interactions between the robots and the humans, while retaining comparative value across runs.

We perform three sets of experiments:

I. **One Institution for Three Domains**. In the first set of experiments, we analyze the results objectively with numerical performance metrics. Our hypothesis is that the SOCIAL context, represented by the SCHOOL domain, will yield the most human-friendly results, the EFFICIENT context, represented by the WAREHOUSE domain will yield an efficient robot behavior, while still taking human comfort into account, and the *social-*

efficient context will reside in-between. Additionally, we perform NON-NORMATIVE experiments where robots are not conforming to norms and do not reside in institutional environment, and are expected to be less human-friendly, as indicated by the metrics.

II. **Participative Study**. In the second set of experiments we evaluate subjective assessment of human participants. We ask a number of volunteers to assess the degree of robots' sociality through a set of questionnaires in the SCHOOL, WAREHOUSE and NON-NORMATIVE scenarios, to determine whether introducing normative aspects influences human perception of the robots.

III. **With Cooperative Localization**. The objective of the final set of experiments, is to assess the performance of the FI-GM-PHD filter proposed in Chapter 8 in scenarios where the dynamics of the formation are affected by the interactions with humans. In this chapter we provide a brief summary of the results, while detailed analysis is provided in Appendix C.

15.6 Experiment Set I

We distinguish four scenarios: NON-NORMATIVE, SCHOOL, WAREHOUSE and HOSPITAL, each performed with A) one human and B) two humans. Each person is represented by the domain function (e.g., teacher, visitor, etc.) and assigned an institutional role (i.e. teacher $\in \mathcal{G}^{\mathbf{S}}_{A}$(Acquainted-Human)), directly linking a person with its speed_factor and social, personal and intimate distances – elements of grounded knowledge K that allow for customization of the robot response to a particular human. For each scenario we perform 12 consecutive runs. In experiments with one human, one of the two human functions is selected for half of the runs. In experiments with two humans, both human functions are present simultaneously. An overview of the key elements of the grounded knowledge is provided in Table 15.6, while the remaining implementation details are provided in Appendix A.

15.6.1 Results

We analyze the experimental results with regard to three objective measures: human-robot distances, formation error, and formation speed. A summary of these results is presented in Table 15.6.

Human-Robot Distances

The largest impact on the motion of the formation is exerted by the norm n_2 for avoiding human personal space. Upon encountering a human, robots prioritize the avoidance component achieved through the trajectory change in the case of the Leader, or repulsion weights in the case of the Followers, with the intensity of the reaction being a function of the human-robot distance and the personal_distance$_h$ of the encountered human H_h, known to the robots as

SOCIAL CONTEXT		*Social*		*Efficient*		*Social-Efficient*	
DOMAIN	NON-NORMATIVE	SCHOOL		WAREHOUSE		HOSPITAL	
ROLE	*default*	child	teacher	worker	visitor	staff	patient
EXPERIMENT PARAMETERS							
PersonalSpace personal_distance	1.2 m	1.5 m	1.2 m	0.8 m	1.2 m	1.2 m	1.2 m
AppropriateSpeed trajectory s_A	1.0	0.75 (slower)		1.25 (faster)		1.0	
near-agent s_A	-	0.3	0.7	0.9	0.7	0.7	0.7
FormationBias formation_bias 	1.0	1.1		0.8		1.0	
EXPERIMENT RESULTS							
Human-Robot Distance Single-human experiments; Experiments with two humans 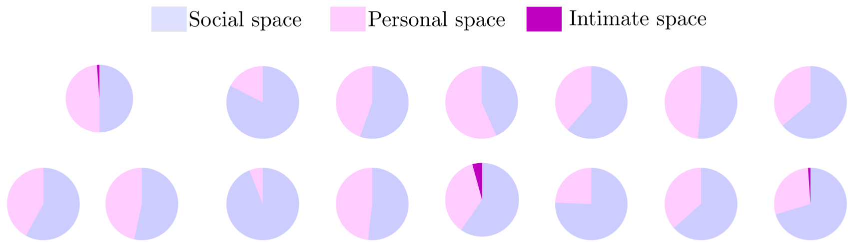							
Formation Speed Av. run duration	≃ 1000 s	5% longer		5% shorter		≃ 1000 s	
Formation Error Average e_F	0.33	0.49		0.53		0.50	
Std e_F	0.1	0.29		0.31		0.31	
QuietActivityTime Av. robot speed	-	-		-		94% slower	
ActivityCriticalArea Time in Φ_{ACS}	-	-		-		87% less	
Asocial Behaviors Offensive %; Unacceptable %; Acceptable % 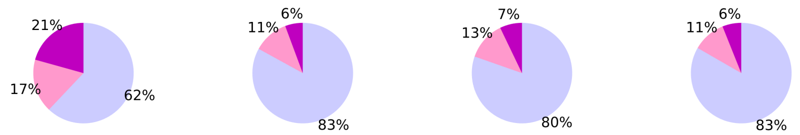							

Table 15.6 – Summary of the parameters included in robot's grounded knowledge K tailored to each domain and to the preferences of the humans within (top rows) and the summary of the experimental results highlighting the outcomes of the aforementioned customization (bottom rows). The pie charts of the human-robot distances show the distances falling into the intimate space, personal space and social space categories, for the human roles that label the respective columns.

part of grounded knowledge K. The element personal_distance, shown in Table 15.6, differs among human agents, with the largest space assigned to the child of 1.5 m, and smallest to the worker of 0.8 m.

After classification of the resulting human-robot distances during encounter according to three categories - in intimate space, in personal space, and in social space, the results confirm that, on average, robots keep larger distance to humans with larger personal_distance. As shown through the pie charts in Table 15.6, robots rarely intrude personal space of the child agent, in both single-human and multi-human experiments. As expected, the opposite holds for the worker, for whom the proportion of intrusion is the largest in both single-human and multi-human experiments.

With regard to the static affordance space, the robots visit Φ_{ACS} 87% less frequently than when the norm n_1 is not applied (although more strict constraints could potentially achieve 100% successful restriction).

An analysis of agent trajectories and recorded videos, shown in Table 15.7, provides further insights into understanding of the difference between the non-normative domain and the domains controlled by the institutions. In the non-normative case, upon encountering a human, robots continue their movement in the formation without any change[2]. In the WAREHOUSE domain, for a worker with the personal_distance of 0.8 m, robots minimally modify the formation to let the person pass. On the contrary, in the SCHOOL domain, the reaction to a human child with the personal_distance of 1.2 m occurs much earlier and the robots leave a large space for the human to pass, modifying the formation considerably.

While distance control has a substantial impact on the sociability of interactions, one should note that in our experiments, where robots act as a result of a balance between the attraction and the repulsion forces, short-term dynamics might outweigh any other factor that influences human's perception of robot sociability. Large repulsion forces and, consequently, considerable robot speeds occurring at a small distance from a human can appear as more threatening than similar forces taking place at larger distance, especially when large robots are involved. For this reason it is important to introduce subjective performance measures, as we will do later in this chapter.

Formation Error

As it can be expected, the degree of formation modification in presence of humans has a significant impact on the formation error. Summarized in Table 15.6, the formation error is significantly lower, in the case of non-normative experiments than in all the other scenarios. As shown in Table 15.7, in the NON-NORMATIVE scenarios robots remain largely undisturbed by the human presence. One should not, however, see the formation error as inability of the robots to maintain a desirable formation, but as a degree of compliance of the robots to the

2 For safety purposes, reactive obstacle avoidance is always enabled to prevent immediate collisions

NON-NORMATIVE

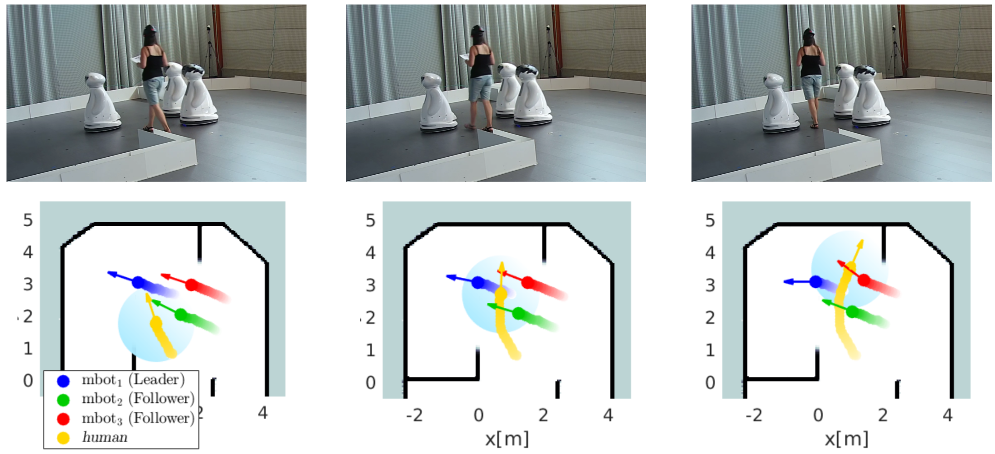

Robots continue their movement with little to no disturbance upon encountering a human.

SCHOOL

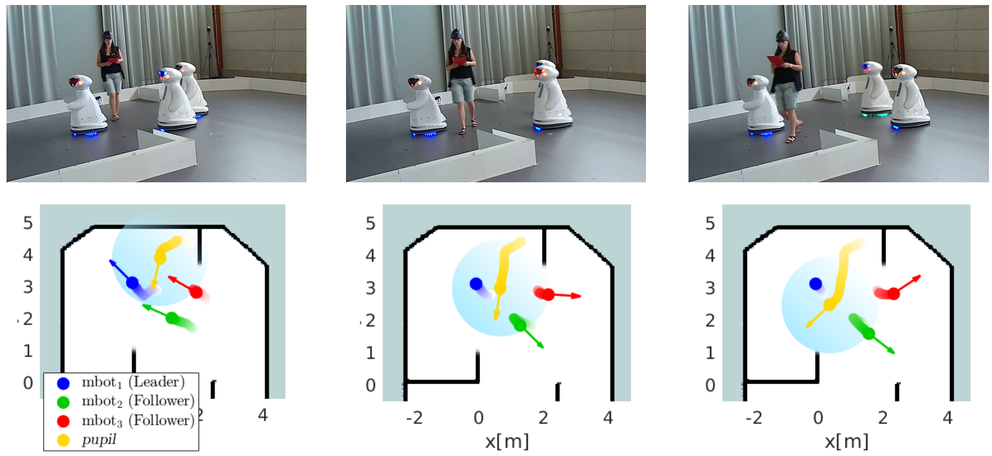

When a human designated as child with personal_distance of 1.5 m passes between the robots, followers temporarily break the formation to make a space for it.

WAREHOUSE

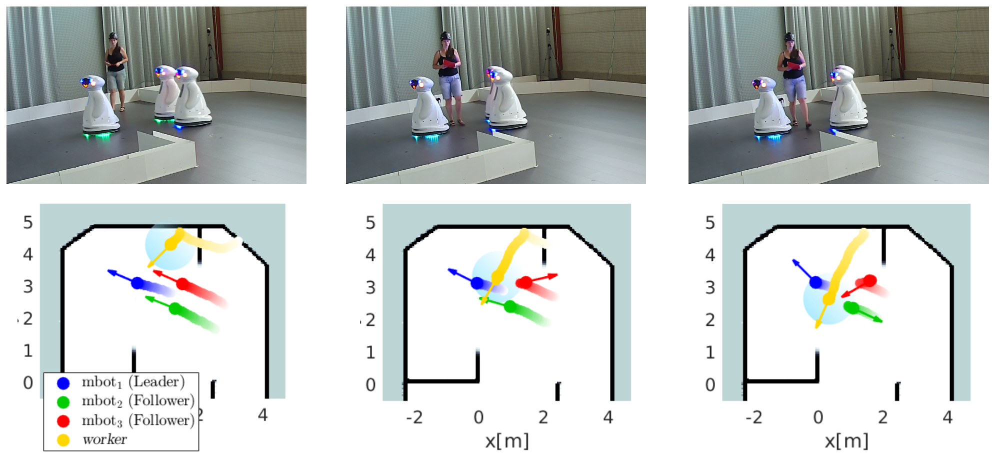

Similarly as in the SCHOOL domain, robots give way to the human, but the degree of compliance is smaller due to the smaller personal_distance of 0.8 m of the human worker.

Table 15.7 – Pictures taken from the recorded videos of the experiments and the respective trajectories of humans and robots. Light blue sphere around human position indicates the extent of the custom personal space.

presence of humans.

Formation Speed

Lastly, the control of the formation speed through speed_factor leads to changes in run duration, but does not result in differences of robot speeds around humans. On average, in the SCHOOL domain with the smallest trajectory speed_factor of 0.75, the formation reaches the goal in 5% longer time than in the default non-normative case with speed_factor of 1. Analogously, in the WAREHOUSE domain with the largest trajectory speed_factor of 1.25 formation completes the run in 5% less time.

Additionally, in a set of experiments in the HOSPITAL domain, where the QUIET_ ACTIVITY_TIME condition is true, during the time_interval $T_{QAT} = [80, 100]$ s, robots adapt 94% lower speed than outside of that time interval.

Video Annotations

All experiments were recorded and annotated in post-processing. Only experiments with two humans involve participants with no knowledge of the differences between the domains, and only these experiments are annotated. We classify human-robot encounters in the following three categories: *acceptable*, where human motion is not disturbed by the robots, *unacceptable*, where one or more robots interrupts human motion, and *offensive*, where a participant is visibly distressed as a result of robot's behavior. Note that although care was taken to replicate consistency among the scenarios, annotating videos is a process prone to subjectivity of the observer.

Altogether, we annotated 29 interactions in the NON-NORMATIVE domain, 53 in the SCHOOL domain, 47 in the WAREHOUSE domain and 66 in the HOSPITAL domain, the summary of which is shown in Table 15.6. The institutional scenarios generally perform much better than the NON-NORMATIVE domain, with 6-7% of the encounters being offensive and further 11-13% being unacceptable, versus 21% and 17% for offensive and unacceptable interactions respectively of the NON-NORMATIVE domain. Note that *acceptable* interactions in the NON-NORMATIVE domain include the encounters during which the robots ignore human presence but the dynamics of the situation do not result in the disturbance of human motion, hence the number of the offensive encounters is bounded.

15.6.2 Discussion

The experimental scenarios proposed in this chapter are characterized by high stochasticity and brisk dynamics, where sudden behavior changes on the human side result in large forces being exerted on the robots, which consequently lead to rapid changes of the formation. It is a challenge to devise relevant performance metrics for the evaluation of normative robot behaviors in such uncontrollable social environments. Although norm activation is deterministic (as dictated by the institutional conditions), the dynamic situations in which the robots operate are affected by a multitude of factors that in turn influence their behaviors. Moreover, the action-to-reaction time does not correspond well to the norm activation interval. For example, one may expect that the lower near-agent speed_factors would predominantly yield lower speeds next to the corresponding humans. However, the interplay between the repulsion forces, the formation forces that bring the robots back together after large disturbance, and the dynamics of the particular encounter might result in the opposite. For this reason, our analysis is focused on the power of the formalism to yield perceivable differences in formation behavior tailored to the domains and to the participating humans. While a subjective assessment is undertaken in the next section, the results of the objective evaluation presented in this first set of experiments provide an insight into the advantages of using the institutional formalism.

One of the most perceivable differences lies in the distance control – the element with the highest impact on the human comfort, closely followed by the speed the robots keep near humans. Our analysis of the human-robot distance with regard to the proxemics model clearly indicates the contrast in the robot's interpretation of the norm n_2 when encountering humans with different functions: for instance, a the robot keeps the largest distance to a child and smallest distance to the worker. The formation speed is tailored to the domain rather than to a particular human, and is the lowest in the domain where the knowledge dictates that is should be so, i.e. in SCHOOL. Such customization is achieved through the grounded knowledge, but no differences are made at the institution level, and neither in the behaviors themselves – it is only the parameters that change the behavior modality automatically, through norm realization.

As a final remark of this section, we would like to emphasize that there are many improvements that can be done to the above implementation of the norms. Specifically, the behaviors would yield better results if the parametrization embedded in institutional knowledge was based on empirical findings, tailored to specific humans. Furthermore, elements of deliberative planning and human prediction such as in [68] [88] are critical for achieving optimal results. The latter elements increase the complexity of the system manifold and could overshadow the main focus of experiments; therefore, we intentionally leave them out of the scope of this thesis. Although our methodology might not result in better physical performance of the robots when compared to the state-of-the-art approaches to norm-following social robots, it also not our objective.

15.7 Experiment Set II

The results presented in the previous section show how the differences across domains are manifested through objective performance metrics and how adding norms on top of standard, non-normative multi-robot behaviors shapes these behaviors to meet a desired objective.

To complete the evaluation, in this section we report the outcome of a participative study that was designed to analyze the impact of introducing social norms on human perception. From among the social domains introduced in this chapter, we select the SOCIAL domain – SCHOOL and the EFFICIENT domain – WAREHOUSE, as the two most marginal cases of the normative approaches. Furthermore, we compare the social methods with the NON-NORMATIVE case to achieve an unbiased estimate.

We endeavor to test the following hypotheses:

- $\mathbf{H}_0$ - (Null hypothesis) There is no relation between the three robot behaviors and their respective perception by the participants.
- $\mathbf{H}_1$ - The participants will perceive a difference between the three robot behaviors and will rate the NON-NORMATIVE behavior as the least acceptable, social, and legible.
- $\mathbf{H}_2$ - The participants will perceive a difference between the three robot behaviors and will rate the SOCIAL behavior as the most acceptable, social, and legible.

15.7.1 Experimental Protocol

Participants were recruited among EPFL staff, EPFL students, and the surrounding community through a combination of online announcements and emails, fliers, and word of mouth. In return for participation they received pizza coupons; no monetary incentive was provided. A requirement for participation was to be at least 18 years of age.

A within-subjects design was adopted – each participant experienced all three robot behaviors, NON-NORMATIVE, SOCIAL, and EFFICIENT. To minimize the bias that could have been introduced by the order of trials, we randomized their sequence.

Two participants were simultaneously engaged in each trial. In a case we did not find a pair of participants, a person already familiar with the experiments was asked for an assistance, but was not taking part in the evaluation.

Initially, each pair of participants was shown the robots and the arena, and was given a description of their task. Participants were informed that the robots are engaged in their own duty, but they will be sharing the same environment. Each participant was to take part in a simple game, the rules of which have been explained before the experiments. Each person was given a map of the arena populated with waypoints that were labeled with animal shapes and indicated as crosses on the floor. The map was accompanied with a list of waypoints to visit (see Appendix A for additional details on the experimental setup). Additionally, each participant was assigned a position to start and finish the game and given a cap with a set of MCS markers for pose tracking and a unique identification number. The designation of a waypoint list and a specific cap to a participant was directly linked to the assignment of an institutional role – of a child or an adult in the SOCIAL (SCHOOL) domain, or a worker or a visitor in the EFFICIENT (WAREHOUSE) domain. To simplify the notation, we will denote the

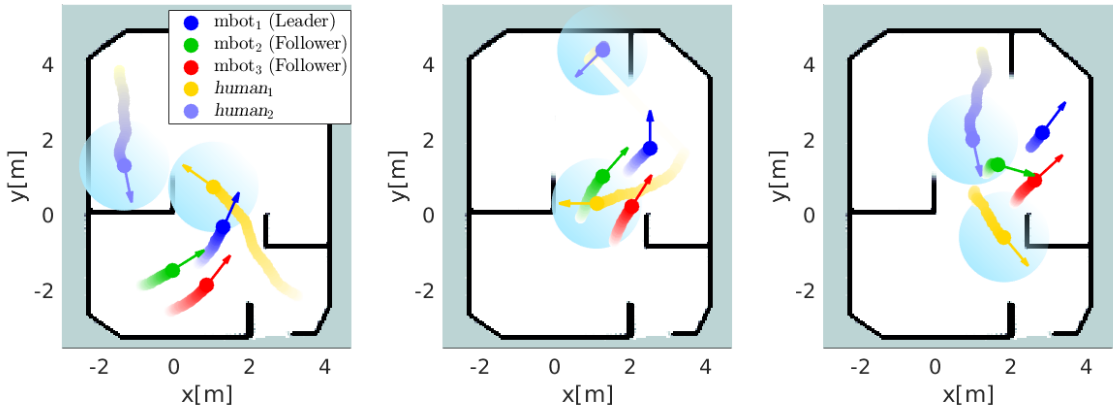

Figure 15.3 – Example of trajectories of the participants and the robots in the NON-NORMATIVE scenario. All three snapshots have been taken in the same run, at times $t = 21$ s (left), $t = 32$ s (middle) and $t = 37$ s (right).

participants with these roles as H_A (child or visitor) and H_B (teacher or worker). Neither the differences between the roles nor the differences in robot behaviors were explained to the participants. Participants were instructed to move naturally, in the same manner for each experiment, following the same path. After each trial we distributed a survey form and then asked to return to the arena for the next test. Finally, after completion of the three trials, the participants were asked to answer a set of comparison questions. Additionally, each part of the survey was accompanied with a free-response question to solicit comments on each robot behavior.

15.7.2 Results

A total of 38 people participated in our study (18 female and 20 male) in a set of 23 experiments. Participants were ranging in age from 21 to 61 years (mean = 31, SD = 10.31), and represented a variety of education levels and occupations, including administration, art, engineering, as well as a wide range of self-reported prior experience with robots (mean 1.18, SD 1.16, on a scale of 1 to 4). Videos of the experiments are available at the link provided in the footnote[3].

Figure 15.3 and Figure 15.4 illustrate typical encounters of the participants and the robots in the NON-NORMATIVE and the SOCIAL trials, respectively. In the NON-NORMATIVE case the robots change their default formation behavior only marginally, when the DWA obstacle avoidance is activated to avoid collisions. Otherwise, human presence is ignored and the robots do not make space for the passing humans. In contrast, in the SOCIAL trial, the formation undergoes significant modifications to accommodate the human presence. This can be clearly seen in Figure 15.4 (left), where the Leader adopts a path circumventing the participant with the role of a child, and in Figure 15.4 (middle, right), where the Followers break the formation to move away from the participants. The trials of the EFFICIENT domain are similar to the SOCIAL ones,

3 http://disalw3.epfl.ch/research/alicja/Chapter_15-7.mp4

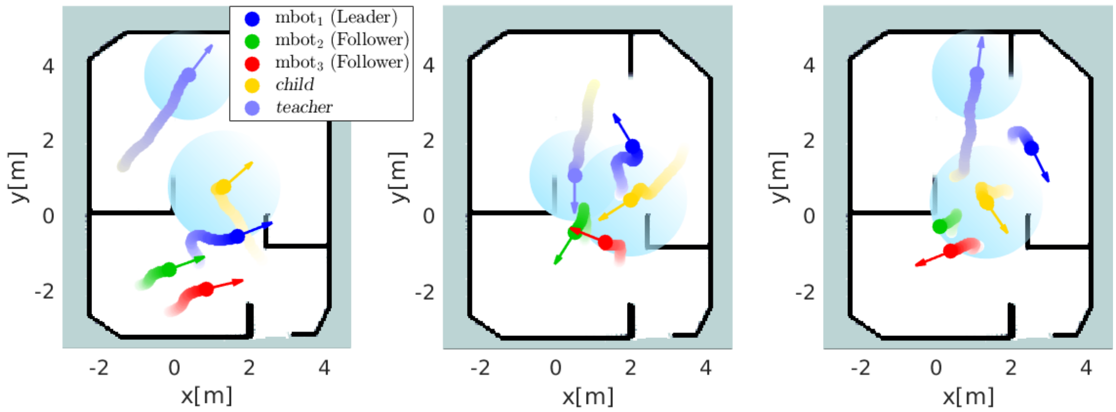

Figure 15.4 – Example of trajectories of the participants and the robots in the SOCIAL scenario. All three snapshots have been taken in the same run, at times $t = 21$ s (left), $t = 32$ s (middle) and $t = 37$ s (right).

but robots are less reactive to the human presence and leave less space for them.

Although the aforementioned characteristics are typical for the chosen domains, the experimental runs were highly stochastic; each participant assumed different speeds, some remained at the waypoints for longer times, some stopped to look around in search for the next step. Reaction of the participants to the encounter with the robot team also varied a great deal – some participants chose to stop and wait to see how to pass next to the robots, some chose to force their way, others adopted a mixture of compliance and confrontation. Consequently, we observed a great deal of variation between the trials. In some experimental runs, the robots were able to successfully adhere to social norms and respect the personal spaces of the participants. In most of the trials, however, we noted infringements of such spaces resulting from a situation the robots simply could not have dealt with given the current reactive nature of the behaviors – for example when two participants moved on either side of a robot, the resulting forces might have made it move towards one of them. Perception uncertainty combining noise in both the self-localization position and the human position obtained from the MCS also had a strong effect on the system – during some encounters humans were passing very close to the robots, so even errors that are usually negligible could have resulted in a robot believing that the human is on the opposite side of its body than where he or she was in reality.

Thereupon our analysis encompasses an evaluation of the survey results juxtaposed with the performance metrics and the behavior annotations.

Metrics: Human-Robot Distances

As in Section 15.6.1, we have classified human-robot distances into three zones *in-intimate-space* (< 0.45 m), *in-personal-space* (< 1.2 m) and *in-social-space* (< 1.68 m) and determined the amount of time the robots fall into these zones. Again, we do not distinguish between the human roles and the prescribed to them values of personal_space embedded through the

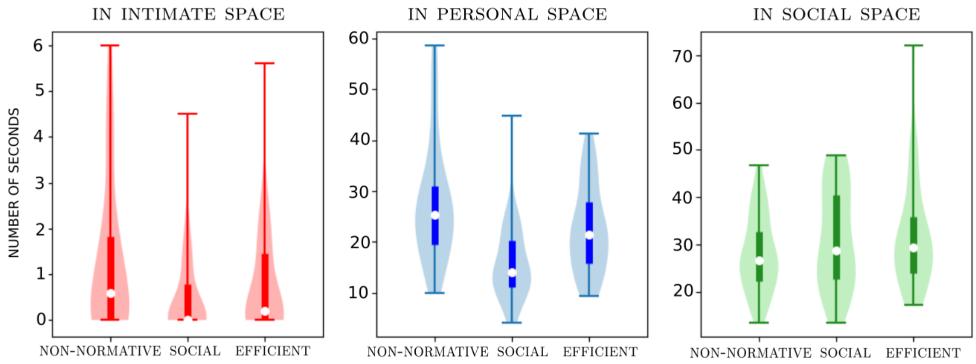

Figure 15.5 – Distribution of the time robots spent in the three human comfort zones. White marker indicates the median, extent of the bars is the 25 and the 75 percentiles.

institutional knowledge, as in this type of analysis we are interested in the final differences between the trials visible through the unprocessed data and not how well the robots adhere to the specifications.

Figure 15.5 outlines the distribution of the time robots spent in the three human comfort zones, with the y-axis showing the number of seconds with a precision of 0.1 s, averaged across all participants and trials and normalized per participant. In the NON-NORMATIVE scenario, the intimate space has been invaded frequently, with a median of 0.55 s per participant, and the 75 percentile reaching up to 1.8 s. The result is slightly better in the EFFICIENT case, where the median of 0.18 s is much lower. The least disturbance of the intimate space has been recorded in the SOCIAL case, with the median of zero and the 75 percentile being half that of the NON-NORMATIVE case. A similar trend is visible when considering the personal space, although we can also note one interesting phenomenon – in the EFFICIENT case a significant number of participants experienced a high number of seconds a robot remained in their personal space. We can explain it by the fact that by design, one of the human roles – that of the worker, has been prescribed personal_space of 0.8 m. Therefore robots had no incentive to move out of that zone. The same is true for the SOCIAL case, where robots largely remained in the social zone for a significant amount of time.

The Distance Cost

For further analysis of the impact the distances might have exerted on the participants' answers in the questionnaire, we developed a score capturing their performance with respect to that aspect. For a participant H_h and a scenario $s \in \{\text{NON-NORMATIVE}, \text{SOCIAL}, \text{EFFICIENT}\}$ the *distance cost* $\mathcal{M}^D$ is computed as follows:

$$\mathcal{M}^D_{s,h} = 2\, m^D_I + m^D_P - 0.5\, m^D_S \tag{15.1}$$

ACCEPTABLE

UNACCEPTABLE

OFFENSIVE

Table 15.8 – Snapshots from the videos recorded during the experiments, showing examples of *acceptable*, *unacceptable* and *offensive* robot behaviors before, during, and after an interaction.

where m^D is the portion of time a robot remained in a particular zone, i.e. a number of seconds in that zone normalized over the total interaction time. To this end, metric m_I^D stands for portion of time in intimate space, m_P^D for a personal space and m_S^D for social space. The metric penalizes close encounters and promotes interaction that occur at a social (norm-satisfying) distance. The resulting average costs are 6.06×10^{-2} for the NON-NORMATIVE case, 0.51×10^{-2} for the SOCIAL case, and 3.08×10^{-2} for the EFFICIENT case. Note that the lower the distance cost, the better the performance.

Behavior Annotations

We classified all encounters of humans and robots into the three categories: *acceptable*, *unacceptable*, and *offensive*, following the same procedure as in Section 15.6.1. Categorization has been done with no knowledge of the scenario label to avoid potential subconscious bias. Altogether, we annotated 152 interactions in the NON-NORMATIVE trials, 164 in the SOCIAL trials, 158 in the EFFICIENT trials. Examples of such interactions are shown in Table 15.8, where we show a situation before, during and after the encounter.

Figure 15.6 provides a detailed overview of the annotation results. Offensive type of encounters

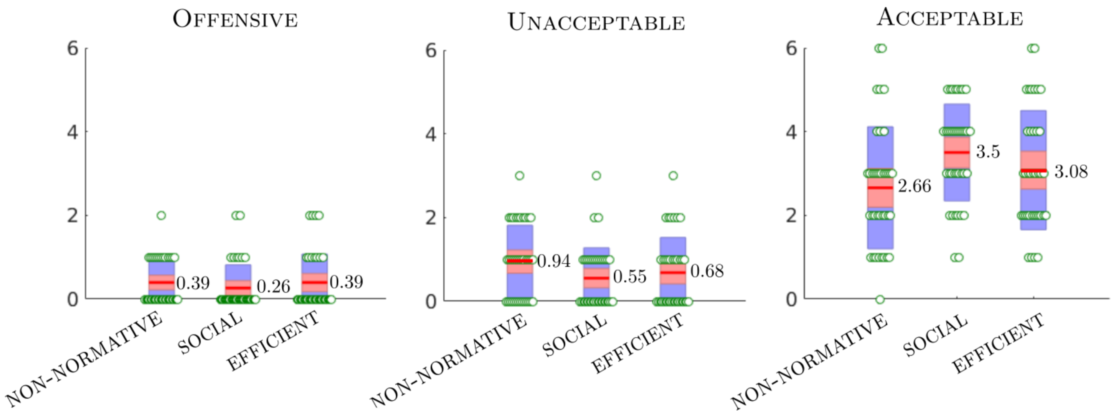

Figure 15.6 – Result of video annotations. Red line indicates the mean, light red area delineates 95% confidence interval and blue area the standard deviation. Raw data, for clarity jittered along the x-axis, is shown in green.

occurred in all three scenarios, with the lowest average in the SOCIAL case, and a comparable result between the NON-NORMATIVE and the EFFICIENT scenarios. A number of participants, in particular in the EFFICIENT case was subject to as many as two offensive encounters. With regard to unacceptable behaviors the tendency is much clearer; the least number of such encounters occurred in the SOCIAL case, closely followed by the EFFICIENT case. On average, there was less than one unacceptable encounter per participant per trial in all the scenarios. Acceptable interactions occurred much more frequently than negative encounters. On average each person experienced at least three acceptable interactions, with the lowest number in the NON-NORMATIVE case, where the distribution is skewed towards a smaller number (< 4) of acceptable encounters, and the best result in the SOCIAL case, with the average number reaching 3.5 positive encounters per participant, and a high number of participants experiencing 4 or 5 acceptable interactions.

The Annotations Score

Similarly as negative interactions might negatively impact the perception of robots' sociality, positive robot encounters have a potential to improve their acceptability. The effects of the types of interactions unique to each participant will be accounted for when analyzing the survey results. Similarly as for the distance metric, we propose a score capturing the impact that the quality of interactions might have exerted on the participants' answers in the questionnaire. For a participant H_h and a scenario $s \in \{$NON-NORMATIVE, SOCIAL, EFFICIENT$\}$ the *annotations score* $\mathcal{M}^A$ is computed as follows:

$$\mathcal{M}^A_{s,h} = m^A_A - m^A_U - 2\, m^A_O \tag{15.2}$$

where m^A_A, m^A_U and m^A_O is the number of interactions of the offensive, unacceptable and acceptable type, respectively. The metric rewards positive encounters, and penalizes negative ones, with the weight of the offensive being twice that of the unacceptable ones. The resulting average scores are 0.92 for the NON-NORMATIVE case, 2.42 for the SOCIAL case, and 1.61 for the EFFICIENT case. Note that higher scores indicate better performance of the robot team.

QUESTIONNAIRE, PART I-III. RATING OF THE INDIVIDUAL TRIALS

GENERAL ACCEPTANCE			CRONBACH'S ALPHA 0.85
Q_1	*"What is your overall feeling towards the robots?"*	Highly reject	Highly accept
Q_2	*"Would you like to see these robots helping people at their work places?"*	Not at all	Absolutely
Q_3	*"Were you comfortable around the robots?"*	Not at all	Absolutely
Q_4	*"How would you rate the robots in terms of safety?"*	Very dangerous	Very safe
Q_5	*"How would you rate the robots in terms of naturalness?"*	Artificial	Human-like
Q_6	*"How would you rate the robots in terms of confusing?"*	Not at all	Absolutely
SOCIABILITY			CRONBACH'S ALPHA 0.75
Q_7	*"Was it enjoyable to interact with the robots?"*	Not at all	Fully
Q_8	*"How would you rate the robots in terms of friendliness?"*	Very aggressive	Very friendly
Q_9	*"How would you rate the robots in terms of likability?"*	Very unpleasant	Very pleasant
LEGIBILITY			CRONBACH'S ALPHA 0.74
Q_{10}	*"Were you able to understand the behavior of the robots?"*	Not at all	Absolutely
Q_{11}	*"Were you able to anticipate the actions of the robots?"*	Not at all	Absolutely
Q_{12}	*"How would you rate the robots in terms of intelligence?"*	Very stupid	Very intelligent

Table 15.9 – Survey questions of Parts I-III, for analysis grouped according to three categories (the actual order of the questions was different in the survey provided to the participants, for explanation see Section 15.7.2). The third column lists the possible response on the Likert scale of 1, and the fourth column on the Likert scale of 5.

Survey Analysis

Participants were asked to complete a four-part survey, Parts I-III evaluating the individual three trials (NON-NORMATIVE, SOCIAL and EFFICIENT) and Part IV for their comparison. Parts I-III were identical, and administered after each trial. Trial labels were not provided so as to avoid any bias they could introduce. The questions, shown in Table 15.9 were answered on a Likert scale from 1 to 5, with the exact labels of the marginal lowest and highest scores provided in the table. Part IV included a set of questions regarding the differences perceived among the trials and a set questions for choosing the trial during which the robots' behavior was maximizing or minimizing one of the characteristics (e.g., friendliness, safety, intelligence). At the end of Part IV, we asked the participants to assign each trial a label from a set of available labels, namely *non-normative*, *social* and *efficient*. Questions in Part IV are listed in Table 15.10.

At the end of each part we invited the participant to provide their comments ("Any other

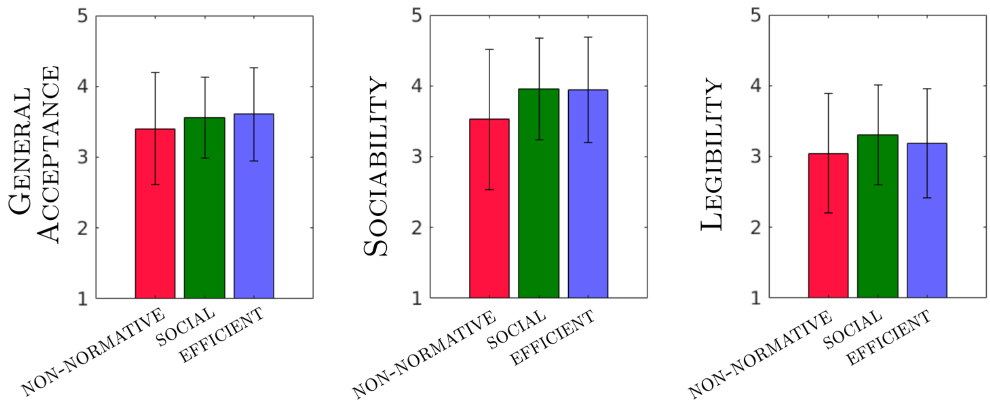

Figure 15.7 – Average score the participants have given to the robots after a particular trial, grouped according to three categories of questions, *general acceptance*, *sociability* and *legibility*.

comments or suggestions?"). The comments, although not codified, provide further insight into the perception of the robot behaviors by participants. A detailed analysis of the comments is provided in Appendix B.

Each of the Parts I-III consisted of nine questions, shown in Table 15.9. For analysis we grouped these questions into three categories[4], *general acceptance*, evaluating the overall perception of the robot team, *likeability*, measuring how enjoyable it was to interact with the robots, and finally *legibility*, for determining whether it was easy for the participants to understand robots' intentions. The adequacy of the proposed categories was confirmed by performing a Cronbach's alpha test, which surpassed the commonly-used 0.7 level of reliability[5]. For each participant, we computed the score in a given category by averaging the response to the questions comprising that category.

Statistical Analysis of Parts I-III

In this section, we aim to understand the impact that the robots' behavioral differences across the three trials (NON-NORMATIVE, SOCIAL and EFFICIENT) exerted on the participants' perception of the robots. To this end we conduct a one-way ANalysis Of VAriance (ANOVA)[6]. The results of ANOVA presented in this section will allow us to test the null hypothesis given in Section 15.7.

With ANOVA we test the effects a robot behavior in a particular trial had on the response in the survey[7]. Our independent variables are the specifications that lead to the differences of

4 Similar methodology has been adopted in [171] and [172].

5 Cronbach's alpha is a test used to determine an internal consistency of a composite score, providing an estimate of the reliability of a psychometric test. Values higher than 0.7 indicate that the questions are closely related and evaluate the same aspect, and therefore their results can be combined for analysis.

6 ANOVA is a statistical test for determining whether there are any statistically significant differences among group means in a sample.

7 ANOVA is considered to be robust against the normality assumption. It can be used with Likert data, even with

robots behaviors in the three trials. The dependent variables are the participants' acceptance of the robots, their perceptions of the robots' sociability and legibility, and other social and intellectual characteristics.

First, we perform a one-way ANOVA over the three categories of survey questions. In the *general acceptance* category no effects were significant (F = 0.92, p = 0.4). The average rating on a scale of 1 to 5, shown in Figure 15.7 was 3.4 for the NON-NORMATIVE behavior, 3.56 for SOCIAL behavior and 3.6 for the EFFICIENT behavior. Similarly, no statistically significant difference were found in the *legibility* rating (F = 1.1, p = 0.336), where the SOCIAL behavior achieved the highest score of 3.3, the EFFICIENT behavior scored 3.18, and the NON-NORMATIVE 3.04. In contrast, an analysis of the *sociability* category indicated a statistically significant effect of the behavior differences (F = 3.3, p = 0.04). Consequently, the null hypothesis $\mathbf{H}_0$ stating that all behavior score means are equal is rejected. The ANOVA test was carried out with a significance level of 0.05, hence the results can be stated with a confidence of 95%[8]. Furthermore, we performed a multiple comparison test to determine which trials are statistically significantly different than the others. Multiple comparisons with the Fisher's least significant difference procedure revealed that the score for the NON-NORMATIVE behavior is statistically significantly lower than SOCIAL ($p < 0.05$) and EFFICIENT ($p < 0.05$). Consequently, we partially confirmed hypothesis $\mathbf{H}_1$ – we verified that the participants rate the NON-NORMATIVE behavior as the least *social*, but we cannot do so with the aspects of *acceptability* and *legibility*. Lastly, no statistically significant difference was found between the SOCIAL and the EFFICIENT behaviors (p = 0.926). The average rating for the SOCIAL and the EFFICIENT trials was similar (3.96 and 3.94 respectively), while the NON-NORMATIVE trials received the lowest score of 3.53.

To understand which particular questions had the most influence on the response in the *sociability* category we further analyzed the three questions that it comprises. A further one-way ANOVA procedure carried out on individual questions indicated that there is a statistically significant difference in two of the questions, Q_8 *"How would you rate the robots in terms of friendliness?"* (F = 3.22, p = 0.044) and Q_9 *"How would you rate the robots in terms of likability?"* (F = 3.71, p = 0.028). In question Q_8, the NON-NORMATIVE behavior with the mean of 3.61 was rated significantly lower ($p < 0.05$) than the SOCIAL and EFFICIENT behaviors with means of 4.16 and 4.11 respectively on a scale 1 to 5. A similar pattern was visible in question Q_9, where the NON-NORMATIVE behavior with the mean of 3.5 was rated significantly lower ($p < 0.05$) than the SOCIAL and EFFICIENT behaviors with means of 3.97 and 4.02 respectively. The question Q_7 *"Was it enjoyable to interact with the robots?"* showed no significant effects, with an average response of 3.63 across the behaviors.

small sample sizes, with unequal variances, and with non-normal distributions [173].

8 The p-value represents the lowest significance level at which a null hypothesis can be rejected. The lower the p-value, the stronger is the evidence in favor of an alternative hypothesis.

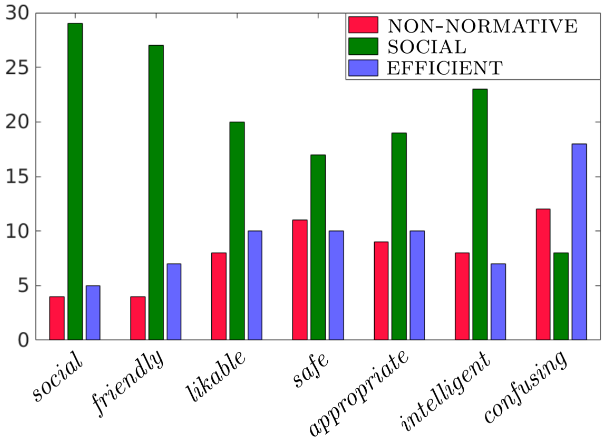

Figure 15.8 – Average response to question Q_{C3} of Part IV, composed of seven sub-questions determining the participants' preferences between the three trials NON-NORMATIVE, SOCIAL and EFFICIENT.

Comparison of the Trials in Part IV

Part IV of the survey was handed out to the participants once all trails have been completed and evaluated in Parts I-III. As shown in Table 15.10, Part IV consists of four questions, aiming at verifying whether the participants in retrospect perceived differences in robot behavior (Q_{C1}), what aspects these differences concerned (Q_{C2}), and how at the end of all trials they evaluate their preferences (Q_{C3}).

Participants' responses to questions Q_{C1} and Q_{C2} indicate that there are perceivable differences between the trials. In fact, all participant responded affirmatively to question Q_{C1} *"Did you see any difference in robot behavior?"*. Responses to question Q_{C2} determining in which aspect of the behavior the variability resided indicates that the vast majority of the participants understood that the robots reacted differently to human presence – 37 out of 38 participants positively answered question Q_{C2a} (*"how the robots reacted to your presence"*). With this result we once again confirm that the null hypothesis $\mathbf{H}_0$ can be rejected. The response was less decisive in questions Q_{C2b} (*"how the robots interacted with each other"*) and Q_{C2c} (*"how the robots interacted with the environment"*). Only 37% of the responders affirmed Q_{C2b}, while only two participants saw a difference in how robots interacted with the environment (Q_{C2c}).

Average response to question Q_{C3} composed of seven sub-questions determining the participants' preferences is shown in Figure 15.8. Retrospectively, the majority of participants judged the SOCIAL behavior most positively. As high as 76% of the responders indicated that the SOCIAL trials were the most social; the same behavior type received very high scores in terms of friendliness (71%), likability (52%), appropriateness (50%) and intelligence (61%). It was also judged safer (45%) and less confusing (only 21% indicated the SOCIAL behavior as the most confusing) than the other behaviors. The NON-NORMATIVE and the EFFICIENT behaviors received comparable scores throughout questions $Q_{C3a} - Q_{C3f}$, but the EFFICIENT behavior was slightly more confusing than the NON-NORMATIVE one (47% versus 32%). However, as the individual encounters with the robots as well as other environmental factors could potentially have a strong influence on the above response, we will analyze such effects later on in the next section.

Similarly as for the distance cost in Section 15.7.2 and for the annotations score in Section 15.7.2 we compute a score capturing the participant's perception of the given behavior. For a participant H_h and a scenario $s \in \{$NON-NORMATIVE, SOCIAL, EFFICIENT$\}$ the *comparison score* $\mathcal{M}^C$ is computed as follows:

$$\mathcal{M}^C_{s,h} = \sum_{q \in Q_{C3a:f}} m^C(q) - m^C(Q_{C3g}) \tag{15.3}$$

where m^C is the binary value indicating whether the option was selected (1) or not (0). Note that *confusing* was assigned a negative score. The metric rewards positive characteristics, and penalizes negative ones. The resulting average scores are 0.84 for the NON-NORMATIVE case, 3.34 for the SOCIAL case, and 0.82 for the EFFICIENT case.

The final question in Part IV asked for assigning a label to the trial. The correctness of the labeling assigned to the experiments by the participants was given a score from 0 to 2, where a score of 2 means correct answer, score of 1 means partially correct answer, when either the NON-NORMATIVE and EFFICIENT labels were swapped, or when the EFFICIENT scenario was labeled as SOCIAL. Zero score was awarded for labeling the NON-NORMATIVE experiment as SOCIAL or vice versa. Half of the participants (19 out of 38) labeled the trials correctly. Further 45% (17 out of 38) achieved partially correct answer. Only one person achieved zero score in the labeling. One answer was invalid.

The comparison score and the labeling score will form dependent variables of the statistical analysis in the next section.

Statistical Analysis of Part IV

In our analysis, we separated the effects of the experimental setup from the effects caused by the actual robot performances. First, we run a set of statistical tests on the effects of the sequence of trials and the role assigned to the participant. Second, we analyzed the effects of the distance score DS and the annotations score AS.

The comparison scores of the scenarios and the labeling score achieved by the participant formed our dependent variables. The independent variables we considered included age; gender; prior experience with robots, on a scale of 1 to 4 (none, low, medium, high), that was grouped into two categories, from-none-to-low and from-medium-to-high; role assigned to the participant (either H_A or H_B, as described in Section 15.7.1); sequence of the trials, classified into three groups: first, when the NON-NORMATIVE trial was carried out at the beginning, second, when the SOCIAL trial was first, and third, when the EFFICIENT trial was first. We also analyzed the impact of the distance score DS described in Section 15.7.2, averaged across the trials and encoded as a continuous predictor; and the average annotations score AS described in Section 15.7.2, classified into two groups, first for a satisfying score (≥ 1) and second for otherwise.

QUESTIONNAIRE, PART IV. COMPARISON OF THE TRIALS		
PERCEIVED DIFFERENCES		
Q_{C1}	*"Did you see any difference in robot behavior?"*	
Q_{C2}	*"If the answer to the above question was yes, the difference was in:"*	
	Q_{C2a}	*– "how the robots reacted to your presence"*
	Q_{C2b}	*– "how the robots interacted with each other"*
	Q_{C2c}	*– "how the robots interacted with the environment"*
COMPARISON OF BEHAVIOR CHARACTERISTICS		
Q_{C3}	*"Indicate during which experiment the robots were:"*	
	Q_{C3a}	*– "the most social"*
	Q_{C3b}	*– "the most friendly"*
	Q_{C3c}	*– "the most likable"*
	Q_{C3d}	*– "the most safe"*
	Q_{C3e}	*– "the most appropriate"*
	Q_{C3f}	*– "the most intelligent"*
	Q_{C3g}	*– "the most confusing"*
LABELING		
Q_{C4}	*"Each experiment was assigned a label:* ***non-normative***, ***social***, *and* ***efficient***. *Try to match the label to the experiment you participated in."*	

Table 15.10 – Survey questions of Part IV. The order of the questions was the same in the survey provided to the participants. In question Q_{C3} one had to indicate the number of the trial that is best described with one of the provided adjectives.

The first set of tests evaluating the effects of the sequence of trials and the human roles on the comparison score, in the case of the EFFICIENT scenario resulted in one statistically significant factor, namely that of a combination of both factors (F = 4.52, p = 0.019). Multiple comparison indicated that the response of the participants who first witnessed the EFFICIENT trial and had a role of H_A is significantly different from the ones with the same role but different sequences, namely mean of -1 for groups who first witnessed the EFFICIENT trial versus 1.38 for those who first experienced the NON-NORMATIVE behavior (p = 0.0348), and 1.6 for for those who first experienced the SOCIAL behavior (p = 0.0407). There was no significant difference between those who first experienced the NON-NORMATIVE behavior and those who first experienced the SOCIAL one. There were also no significant effects of the experimental setup on the labeling score.

In the second test, we evaluated the effects of the distance cost $\mathcal{M}^D$ and the annotations score $\mathcal{M}^A$. One significant effect was found in the SOCIAL scenario. Since the $\mathcal{M}^D$ metric was measured on a continuous scale, its effects were tested using ANalysis Of COVariance (ANCOVA) method. We discovered with that there is a correlation between the average distance score $\bar{\mathcal{M}}^D$ and the comparison score $\mathcal{M}^C_{\text{SOCIAL}}$, as shown in Figure 15.9 (middle). One can

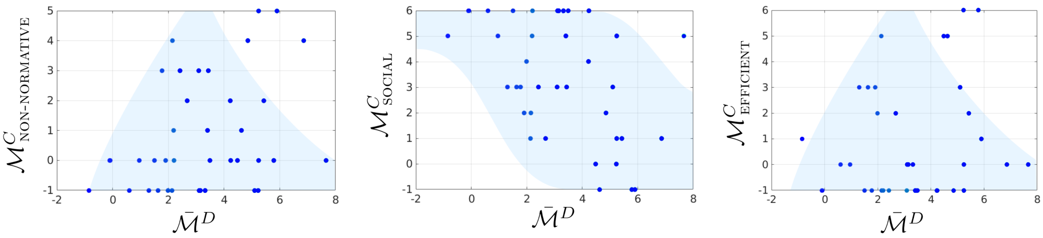

Figure 15.9 – Correlation between the average distance cost $\bar{\mathcal{M}}^D$ and the comparison score $\mathcal{M}^C$. Blue shaded area is shown for visualization of the general data trend.

observe that for lower $\bar{\mathcal{M}}^D$ (better robot performance in terms of distance control), no low $\mathcal{M}^C_{\text{SOCIAL}}$ scores have been assigned to the robots (bottom-left corner of the plot is empty). Similarly, for higher $\bar{\mathcal{M}}^D$ (worse performance), the tendency was as for the lower scores $\mathcal{M}^C_{\text{SOCIAL}}$. For comparison, we additionally visualized data for the NON-NORMATIVE case (left) and for the EFFICIENT case (right). In contrast to the SOCIAL scenario, satisfactory distance score $\bar{\mathcal{M}}^D$ have not been awarded with high values of $\mathcal{M}^C$ (bottom-left corner is populated with data points while top-left corner is empty).

Finally, we analyzed the comparison score for the effects of age, gender and robot experience using a three-way ANCOVA with the age being encoded as a continuous predictor. For none of the scenarios significant effects were found (all $p > 0.05$). Similarly, the labeling score was not affected by the personal factors.

Comments of the Participants

The participant comments provided at the end of each part of the survey are analyzed in details in Appendix B, while in this section we provide a brief overview of the main conclusions.

The majority of comments regarding the NON-NORMATIVE trial was rather negative, the robots were seen as aggressive and it was understood that they ignored human presence. Some participants used decidedly negative wording, such as “scary” or “creepy”. Interestingly, a small number of participants indicated their preference for the NON-NORMATIVE behavior, despite of having an understanding that the robots were not social. The SOCIAL trial received the highest number of general positive comments. The robots were described as “cute”, “polite” and “pleasant”. However, some of the participants disliked the particular implementation of the norms, such as the color of the eyes, or the volume of the sounds. Finally, the motion of the robot during the EFFICIENT trials was judged as asocial and illegible.

Finally we would like to bring attention to the fact that participants had a tendency to ascribe personality traits to the robot behaviors, a phenomenon that has been noticed in similar studies [171]. During the NON-NORMATIVE trials robots were judged as more or less “aggressive”, the SOCIAL behavior was “polite” and the robots had “kind of feelings”, while during the EFFICIENT trials they were “selfish”. We believe that this further demonstrates the ability of the

the institutional framework to produce different robot behaviors and impressions through simple modifications of the behavioral parameters in the process of norm realization.

15.7.3 Discussion

The results of the participative study we conducted allowed us to better understand the impact of introducing social norms to multi-robot behaviors.

The first and foremost finding is the confirmation that every human being has his or her personal preferences and expectations regarding robot's behaviors or other characteristics. Our approach in the presented experiments was to show the diversity of social behaviors we can achieve with our framework, but we did not encode the actual individual preferences of the participants. We believe this would be the natural next step of the validation of the institutional formalism – a human-in-the-loop design process, where initial input of the participants would iteratively affect the process of norm realization. Alternatively, one could try to learn or predict participant's preferences based on age, gender or other standard statistics, however our study did not indicate any of the a-priori human characteristics having statistically significant impact on their preferences.

The analysis of the participants' perception of the individual scenarios in Section 15.7.2 through the rejection of hypothesis $\mathbf{H}_0$ allowed us to confirm that although individual preferences vary, the greater majority of people is able to perceive differences between non-normative and social behaviors, as well as differences in the degree of sociality. Furthermore, we confirmed hypothesis $\mathbf{H}_1$ in the *sociability* category, where the score of the NON-NORMATIVE behavior was significantly lower than the two normative behaviors. With regard to the final hypothesis $\mathbf{H}_2$, stating that the SOCIAL behavior is the most acceptable, social and legible, results of Part IV of the survey shown in Figure 15.8 indicate strong preferences of the vast majority of the participants for the SOCIAL behavior. As high as 76% of the participants judged it as the most *social*, and 71% as the most *friendly*. Therefore we believe that in human-populated spaces, where collection of data with regard of individual preferences is impossible because of a short interaction time, such as at the airports or in museums, highly social behaviors can yield the best results, as people are generally willing to forgive robots' mistakes, as long as they are perceived as friendly. This conclusion can be deducted from the comments of the participants listed in Appendix B, in particular the comments regarding the SOCIAL behavior.

15.8 Experiment Set III

In this section we validate the FI-GM-PHD filter in highly dynamic environments where formation is frequently interrupted by the presence of a human affecting the dynamics of the norm-following formation. We simulate temporary communications failure and analyze its impact on the performance of the robot team. We distinguish two cases: A) where robots rely on communications only (labeled as NOTRACKING), and B) where each robot runs the FI-GM-

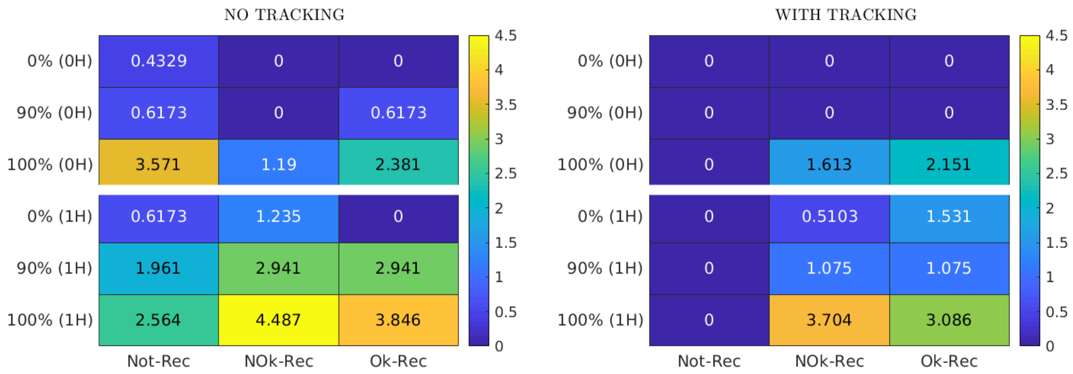

Figure 15.10 – Formation failure rates per second (multiplied by $\times 10^{-3}$), classified into three categories, NOT-REC, NOK-REC and OK-REC, calculated based on the video annotations. The higher the value, the worse the formation performance in the respective category. The categories are: NOT-REC, where robots are not able to recover the formation, NOK-REC, where the formation breaks substantially, but is able to recover, and OK-REC, where the formation separation is visible, but inconsequential. The row labels provide the percentage of dropped messages, 0H stands for experiments with no human, 1H for experiments with a single person.

PHD filter described in Chapter 8 (labeled as WITHTRACKING). Details of the experimental setup and a thorough analysis of the results are provided in Appendix C, while in this section we summarize the main conclusions.

The tracking performance in the WITHTRACKING case is shown to deteriorate gracefully with the degradation of the communication quality – a drop of 90% of messages for 20% of the experiment duration increases the error on average by 14%, while 100% message drop results in 56% increase. Human presence has a clear impact on the tracking performance, with OSPA being on average 5.2% higher when a human is present.

As an indicator of the formation performance we use failure rates, calculated based on the video annotations. We distinguish three categories of formation failures: NOT-REC, where robots are not able to recover the formation, NOK-REC, where the formation breaks substantially, but is able to recover, and OK-REC, where the formation separation is visible, but insignificant. Failure rates shown in Figure 15.10 confirm that the use of FI-GM-PHD filter in the WITHTRACKING scenarios yields much better performance than when tracking is not used (NOTRACKING). The formation failure occurs far less frequently in the case when the tracking is employed, and with lesser criticality – we recorded no cases when the formation breaks with no ability to recover.

With the above experiments we achieved two objectives. First, we have shown that the FI-GM-PHD filter performs well also in highly dynamic situations, where the formation is subject to frequent and significant changes resulting from a social response to the presence of humans. Second, we consolidated all the elements of this thesis in a single experimental study, showing complementarity of the algorithmic solutions to practical challenges of multi-robot deployment in human-populated spaces proposed in Part II, and the institutional framework providing a method for achieving robot sociality proposed in Part III.

Summary

The objective of the case study presented in this chapter was to highlight the advantages of our framework when designing complex normative behaviors. We demonstrated agility of the methodology to encompass a large variety of norms, ranging from navigation to expressions and HRI features. By grounding the institution to three different domains we confirmed the generalizability of the abstract norm formulation, while retaining customization. The latter was achieved through the grounded knowledge, but neither at institutional level nor at behavioral level modifications are inserted: – it is only the parameters that change the behavior modality automatically, through norm realization. This stands in a high contrast to what is proposed in the literature on normative navigation as discussed in Chapter 10, where the norms are embedded directly in the behavior design. As we emphasize in the same chapter and repeat here, our methodology does not aim directly at improving the physical performance of the robots but rather at facilitating the design of socially aware complex robot behaviors that can be reusable and understandable for the humans sharing the same environment.

We completed the analysis of the social aspects of the behaviors through a participative experimental campaign. Our study targeted a comprehensive evaluation of the subjective effects that have an impact on human perception of robots' sociality. Our findings confirm the power of the institutional framework to yield significantly different robot behaviors through the process of norm realization. We were able to show that the non-normative robot behaviors are generally seen as significantly less friendly than normative. Furthermore, we confirmed the need of personalization of robot responses to fit the individual robot preferences. At the same time we identified the fact that, in general, highly social behaviors are preferable by the majority of people, so if personalization of interactions is impossible in a given application, such behaviors have the potential to already yield highly desirable results.

The final set of experiments consolidated all the major technical elements of this thesis, namely formation control with adaptive topology controlled by norms, cooperative localization based on the FI-GM-PHD tracking filter, and sociality of robots behaviors created through our institutional formalism. The results show that complementing faulty communication with the FI-GM-PHD filter tracking significantly reduces the chance of the formation breaking irreversibly, confirming that the proposed tracking method performs well not only in controlled environments such as the ones tested in Chapter 8, but also in highly dynamic missions where human presence induces rapid changes of the formation shape.

16 Case Study: Mixed Formations For Human Guidance

We have seen in the previous case study how an abstract institution can be modeled, how it can be grounded to a concrete domain, and how a generic norm can be designed. In particular, we can use norms to *prescribe* the behavior of the controllable agents, the robots, so that they adhere to the institutional norms. In other words, institutions can be used for both control and for deliberation in multi-robot systems, as we will demonstrate in this chapter, where a different set of norms is automatically activated depending on the social context. More precisely, in this case study a first set is active in situations, when humans are following the robots, and a second set in case humans disobey. Deliberation between the cases is transparent and embedded through the rules of norm realization.

The social context we adopt in this case study is similar to a museum, where visitors do not know the surroundings and the objective of the robots is to guide them through the environment while simultaneously providing information. In contrast to the case study of Chapter 15, humans and robots *cooperate*, i.e. to achieve the objective they explicitly interact and coordinate their behaviors, while social norms make allowances for establishing mixed human-robot formations.

We illustrate our approach in a case study, where we translate abstract norms into concrete constraints on cooperative behaviors of humans and robots. We investigate the feasibility of our approach and quantitatively evaluate the performance of our framework in 30 real experiments with user-based evaluation with 40 participants.

Highlights

- **Composition.** Two institutions are deployed sequentially, with the second being activated as a result of the termination of the first. We provide a preliminary discussion on how institutions can be activated and terminated based on the social context.
- **Flexibility.** The institutional framework is independent of the tools, methods or systems governed by it. In this case study the institutional evolution coexists with a centralized

planner, but the relation is transparent at the institutional level and occurs fully at the domain level.

- **Norm deliberation.** One of the two dominant sets of norms is activated depending on the social situation – the first set in case humans are following the robots, the second set in case humans disobey. Deliberation between the cases is transparent and embedded through the rules of norm realization.
- **Feasibility relations.** The MBot and the Pepper robots coexist in the same institutional environment, but because of different capabilities, they cannot assume the same roles and perform the same actions. We show how to resolve the issue of linking agents to the roles they can perform and the respective actions through feasibility relations.
- **Computational Protocol.** Norm realization at the domain level is represented in a form readily implementable in robot's programming language.
- **Mixed Formations.** To the best of our knowledge, ours is the first study where multiple humans and multiple robots engage in coordinated navigational behaviors in real settings. The results are analyzed with the aim of understanding the social preferences in the context of human guidance in mixed formations.
- **Participative Study.** An analysis of the subjective assessment of participants taking part in the study indicates that there is a need for empirical evaluation of norms before they can be operated in real world applications and that perception of the same normative behavior of a robot varies from person to person, as a consequence of different expectations, background, or experience with robots.

16.1 The STEERING and TUTORING Institutions

Our case study consists of two phases. In the first phase, a group of humans is guided by a number of robots to a destination. The formation includes one leader robot, human visitors, and the follower robots[1]. The second phase takes place once the group reaches the destination, in an exhibition room, where two cooperative robots take the visitors on a tour. One of the robots, referred to as a tutor, describes the objects, while the other robot – an assistant, moves around the room while showing the objects to the visiting humans. The social context is similar to a museum, where the visitors do not know the surroundings and the objective is to guide them through the environment while simultaneously providing information.

The case study is governed by two institutions. The STEERING institution in phase one, and the TUTORING institution in phase two. The choice of institutions, the granularity of the activities and their roles are one of the many possible design choices that could be made. With reference to our discussion on the granularity of institutions in Section 11.3, STEERING and TUTORING fall under the category of institutions for human guidance in Figure 11.3.

1 In our case study we adopted a single robot as follower but the definition of the institutions and related norms support multiple robots in such role.

The STEERING institution is defined as follows:

$$
\begin{aligned}
\textbf{STEERING} = \Big\langle \\
& \textit{Norms}_{(S)} = \{n_{S1}, n_{S2}, n_{S3}, n_{S4}\} \\
& \textit{Roles}_{(S)} = \{\mathsf{Leader}, \mathsf{Follower}, \mathsf{Visitor}\} \\
& \textit{Actions}_{(S)} = \{\mathsf{Guide}, \mathsf{AssistGuidance}\} \\
& \textit{Conditions}_{(S)} = \{\text{VISITORS\_NOT\_FOLLOWING}, \text{LEADER\_STOPPED}, \ldots\} \\
& \textit{Knowledge}_{(S)} = \{\mathsf{PersonalSpace}, \mathsf{SocialForce}, \mathsf{FormationShape}, \ldots\}\Big\rangle
\end{aligned}
$$

The Tutoring institution is defined as follows:

$$
\begin{aligned}
\textbf{TUTORING} = \Big\langle \\
& \textit{Norms}_{(T)} = \{n_{T1}, n_{T2}\} \\
& \textit{Roles}_{(T)} = \{\mathsf{Tutor}, \mathsf{Assistant}, \mathsf{Visitor}\} \\
& \textit{Actions}_{(T)} = \{\mathsf{Teach}, \mathsf{AssistTutor}\} \\
& \textit{Conditions}_{(T)} = \{\text{IN\_EXHIBITION\_ROOM}, \text{NEXT\_TO\_OBJECT}\} \\
& \textit{Knowledge}_{(T)} = \{\mathsf{ObjectLocations}, \mathsf{ObjectDescriptions}\}\Big\rangle
\end{aligned}
$$

16.1.1 Norms

The norms of the STEERING and the TUTORING institutions are listed in Table 16.1.

STEERING		
Prevention	n_{S1}	*"When guiding humans in a formation, the followers should adopt a configuration designed to retain the humans behind the leader."*
Comfort	n_{S2}	*"The followers should maintain a comfortable distance from the humans."*
Waiting	n_{S3}	*"When humans fail to follow behind the leader, the leader should wait for them."*
Re-engaging	n_{S4}	*"When humans are not following, the followers should encourage them to return."*
TUTORING		
Accompanying	n_{T1}	*"The tutor robot should accompany the visitors to indicate the locations of the demonstrated objects."*
Instructing	n_{T2}	*"The assistant robot should provide informative descriptions of the demonstrated objects."*

Table 16.1 – Norms of the STEERING and the TUTORING institutions.

Human Visitors are uncontrollable agents that in general follow the instruction of following behind the Leader. When not, we assume that human behavior can be influenced by the robots. The norms n_{S1} and n_{S2} determine the social behavior of the robots in terms of the formation geometry in the ideal case humans agree to follow, chosen so as to take into account human comfort zone. The norms n_{S3} and n_{S4} operate in the case participants stop following; the leader robot simply waits for the humans to re-engage, and the followers move closer to the humans in an attempt to influence their motion based on the assumption that humans try to maintain a comfort space between themselves and the others. The motivation behind the norms n_{S3} and n_{S4} has been discussed in Section 12.2.2.

For completeness, we rewrite the norms in terms of the institutional components according to the syntax of the institutional norm given in Definition 11.2:

$$\mathcal{N} : \textit{Conditions} \rightarrow \text{deontic}\left(\textit{Roles} \times \textit{Actions} \times \textit{Knowledge}\right).$$

$$\begin{aligned}
n_{S1} &: \text{(when guiding humans)} \rightarrow \text{should}(\textit{followers}, \textit{adopt}, \textit{configuration to retain humans})\\
n_{S2} &: \emptyset \rightarrow \text{should}(\textit{followers}, \textit{maintain}, \textit{comfortable distance from humans})\\
n_{S3} &: \text{when humans fail to follow} \rightarrow \text{should}(\textit{leader}, \textit{wait}, \emptyset)\\
n_{S4} &: \text{when humans not following} \rightarrow \text{should}(\textit{followers}, \textit{encourage}, \textit{to return})\\
n_{T1} &: \emptyset \rightarrow \text{should}(\textit{tutor}, \textit{accompany the visitors}, \textit{indicate locations of objects})\\
n_{T2} &: \emptyset \rightarrow \text{should}(\textit{assistant}, \textit{provide}, \textit{descriptions of objects})
\end{aligned}$$

16.1.2 Conditions and Knowledge

The elements of institutional *Conditions* and *Knowledge* constituting the above decomposition of norms are listed in Table 16.2 and Table 16.3 respectively.

	Activation Conditions	***Outcome Conditions***	***C***
n_{S1}	VISITORS_FOLLOW	ROBOTS_IN_FORMATION	C_{S1}
n_{S2}	VISITORS_FOLLOW	RESPECTED_PERSONAL_SPACE	C_{S2}
n_{S3}	VISITORS_NOT_FOLLOWING	LEADER_STOPPED	C_{S3}
n_{S4}	VISITORS_NOT_FOLLOWING	ROBOTS_IN_FORMATION	C_{S4}
n_{T1}	IN_EXHIBITION_ROOM	-	C_{T1}
n_{T2}	NEXT_TO_OBJECT	-	C_{T2}

Table 16.2 – *Conditions* of the two institutions STEERING and TUTORING.

The conditions for norms n_{S3} and n_{S4} are the negative of the conditions for n_{S1} and n_{S2}, which results in the deliberation between two social situations – first, for the case when humans are following the rules of the mixed formation; second for when humans stop following and are to be encouraged to return. Further details are provided for grounded conditions.

	Knowledge	*Grounded Knowledge K*	Symbol
STEERING			
	PersonalSpace	personal_distance$_x$	$\Delta_{P,h}$
	IntimateSpace	intimate_distance$_x$	$\Delta_{I,h}$
	SocialForce	force$_x$	K_o, Δ_a, Δ_c
	FormationShape	formation_shape$_x$	FS
	FormationBias	formation_bias$_x$	b
	Destination	position$_x$	p_x
TUTORING			
	ExhibitionArea	area$_x$	EA
	ObjectLocations	position$_x$	p_x
	ObjectDescriptions	speech_mode$_x$	SM

Table 16.3 – *Knowledge* and its grounding *K*. We use the subscript $(\cdot)_x$ to abbreviate variables with multiple instances x.

The institutional knowledge is based on two models of human behavior described in Chapter 3, i.e. PM and SFM. The distances of PM are used for the case humans move in the formation (norms n_{S1} and n_{S2}) to determine the optimal parametrization of the formation geometry, given the positions of humans and their personal spaces. While the PM is used for parametrization, it is the FormationShape element that defines the high-level structure of the formation, with the locations of the followers placed so as to retain humans in the mixed formation. The SFM is used for forecasting the amount of virtual repulsive force a robot is to exert to affect human motion. As we explain in Section 5.6.2, we use the SFM to map the distance the robot keeps from the human to the resulting motion of the human aimed at increasing the uncomfortable distance, which results in the human moving away in the desired direction. As we point out in Chapter 10, a similar strategy is popularly used for human guidance in robotics.

16.1.3 Linking Institutions

In Section 11.2, we mentioned that an agent joins an institution upon recognition of a social context (the immediate physical and social setting of the environment) and leaves an institution when the context is no longer perceived. Correspondingly, we can specify the social context of the two institutions introduced in this case study. The STEERING institution is to be activated when humans are to be guided to a known destination by the robots. When humans are to be introduced to a selection of objects in the environment it is the TUTORING institution that is activated. Such abstract descriptions can be further concretized in a manner similar to that of a norm realization, for example through sets of conditions that must be satisfied so that an institution is operational. In our implementation, the STEERING institution is activated once the presence of the humans is detected and terminated once the Leader has reached the Destination location defined in institutional knowledge. The latter is the triggering condition for activation of the TUTORING institution, which terminates once all the objects listed in knowledge of TUTORING have been presented to the Visitors.

16.2 Domain and Grounding

The description so far provides specification for the abstract institutions. We now ground our institutions in the domain of the case study involving real robots and humans. The domain includes the following components:

$$\begin{aligned}
A &= \{\mathsf{mbot}_1,\ \mathsf{mbot}_2,\ \mathsf{pepper},\ \mathsf{participant}_\mathsf{x}\} \\
B &= \{\mathsf{MoveOnTrajectory},\ \mathsf{MoveInFormation},\ \mathsf{Speak}\} \\
R &= \{\mathsf{time},\ \mathsf{position}_\mathsf{x},\ \mathsf{distance}_\mathsf{x},\ \mathsf{location}_\mathsf{x},\ \mathsf{object_state}_\mathsf{x}, \ldots\} \\
K &= \{\mathsf{personal_distance}_\mathsf{x},\ \mathsf{force}_\mathsf{x},\ \mathsf{speech_mode}_\mathsf{x}, \ldots\} \\
C &= \{C_{S1},\ C_{S2},\ C_{S3},\ C_{S4},\ C_{T1},\ C_{T2}\}
\end{aligned}$$

with the subscript $(\cdot)_x$ abbreviating variables with multiple instances x. There are three heterogeneous robots, mbot_1, mbot_2 and pepper, and a number of human participants. We discern the diverse capabilities of the robots through the feasibility relations described in the next section.

The STEERING institution is grounded with $\mathcal{G}^{(S)}$:

$$\begin{aligned}
\mathcal{G}^{(S)}_A &= \{(\mathsf{Leader},\ \mathsf{mbot}_1),\ (\mathsf{Follower},\ \mathsf{mbot}_2),\ (\mathsf{Visitor},\ \mathsf{participant}_1,\ \mathsf{participant}_2)\} \\
\mathcal{G}^{(S)}_B &= \{(\mathsf{Guide},\ \mathsf{MoveOnTrajectory}),\ (\mathsf{AssistGuidance},\ \mathsf{MoveInFormation})\} \\
\mathcal{G}^{(S)}_C &= \{(\textsc{visitors_not_following},\ c^a_{S3}),\ (\textsc{leader_stopped},\ c^o_{S3}), \ldots\} \\
\mathcal{G}^{(S)}_K &= \{(\mathsf{PersonalSpace},\ \mathsf{personal_distance}_\mathsf{x}), \ldots\}
\end{aligned}$$

while the TUTORING institution is grounded with $\mathcal{G}^{(T)}$:

$$\begin{aligned}
\mathcal{G}^{(T)}_A &= \{(\mathsf{Tutor},\ \mathsf{pepper}),\ (\mathsf{Assistant},\ \mathsf{mbot}_1),\ (\mathsf{Visitor},\ \mathsf{participant}_1,\ \mathsf{participant}_2)\} \\
\mathcal{G}^{(T)}_B &= \{(\mathsf{Teach},\ \mathsf{Speak}),\ (\mathsf{AssistTutor},\ \mathsf{MoveOnTrajectory})\} \\
\mathcal{G}^{(T)}_C &= \{(\textsc{in_exhibition_room},\ c^a_{T1}),\ (\textsc{next_to_object},\ c^a_{T2})\} \\
\mathcal{G}^{(T)}_K &= \{(\mathsf{ObjectLocations},\ \mathsf{position}_\mathsf{x}),\ (\mathsf{ObjectDescriptions},\ \mathsf{speech_mode}_\mathsf{x})\}
\end{aligned}$$

Since the robots are heterogeneous, it is important to discern which robots can take which roles and their respective actions. First, we specify the feasibility relations between the institutional role and the action that is assigned to that role according to Definition 11.3. The relations include:

$$\mathcal{F}^{(S)} = \{(\mathsf{Leader},\ \mathsf{Guide})(\mathsf{Follower},\ \mathsf{AssistGuidance})\}$$
$$\mathcal{F}^{(T)} = \{(\mathsf{Tutor},\ \mathsf{Teach})(\mathsf{Assistant},\ \mathsf{AssistTutor})\}$$

Second, we assure that an agent can take a role only if it can carry out the linked action and the behavior the action is grounded to, i.e. an agent must have capabilities (in terms of sensors, actuators etc.) to perform the behavior. An agent is linked to a role through grounding $\mathcal{G}_A$.

When synthesized, the two above steps result in the relation between the agent, through the role it takes and the action assigned to that role, to the behavior the action is grounded to:

AGENT	$\xleftarrow{\mathcal{G}_A}$	ROLE	$\xrightarrow{\mathcal{F}}$	ACTION	$\xrightarrow{\mathcal{G}_B}$	BEHAVIOR
mbot_1	$\xleftarrow{\mathcal{G}_A^{(S)}}$	Leader	$\xrightarrow{\mathcal{F}^{(S)}}$	Guide	$\xrightarrow{\mathcal{G}_B^{(S)}}$	MoveOnTrajectory
mbot_2	$\xleftarrow{\mathcal{G}_A^{(S)}}$	Follower	$\xrightarrow{\mathcal{F}^{(S)}}$	AssistGuidance	$\xrightarrow{\mathcal{G}_B^{(S)}}$	MoveInFormation
pepper	$\xleftarrow{\mathcal{G}_A^{(T)}}$	Tutor	$\xrightarrow{\mathcal{F}^{(T)}}$	Teach	$\xrightarrow{\mathcal{G}_B^{(T)}}$	Speak
mbot_1	$\xleftarrow{\mathcal{G}_A^{(T)}}$	Assistant	$\xrightarrow{\mathcal{F}^{(T)}}$	AssistTutor	$\xrightarrow{\mathcal{G}_B^{(T)}}$	MoveOnTrajectory

According to the above relation, in the TUTORING institution, mbot_1 assigned the role of an Assistant performs a MoveOnTrajectory behavior, while a pepper robot assigned the role of a Tutor performs the Speak behavior. The robot capabilities are taken into account through the process of agent grounding $\mathcal{G}_A$. The mbot_1 robot is capable of robust navigation, thus it shows objects to the participants, while the humanoid pepper robot with enhanced interaction potentials but poor maneuvering capabilities explains these objects from far away.

The behaviors MoveOnTrajectory and MoveInFormation of the first phase of $\mathcal{C}_{II}$ are described in details in Chapter 5. The Tutor and the Assistant in phase two synchronize execution of their behaviors. The Tutor describes an object o_i with the Speak(o_i) behavior, while the Assistant navigates to that object with the behavior MoveOnTrajectory(o_i). The Speak behavior involves performing a speech act, with a pre-defined string, such as "The object is a yellow sphere". The MoveOnTrajectory behavior involves navigation using the FMM method to the pre-selected waypoints, such as nextToYellowSphere. Object descriptions and their locations are included the institutional *Knowledge* of ObjectDescriptions and ObjectLocations, respectively (see Table 16.3). Synchronization is assured by a communication protocol between the robots. Human behaviors are not represented formally.

The temporal course of the experiments is established by an offline plan devised by a temporal planner described in Chapter 2. The planner determines relations between time intervals

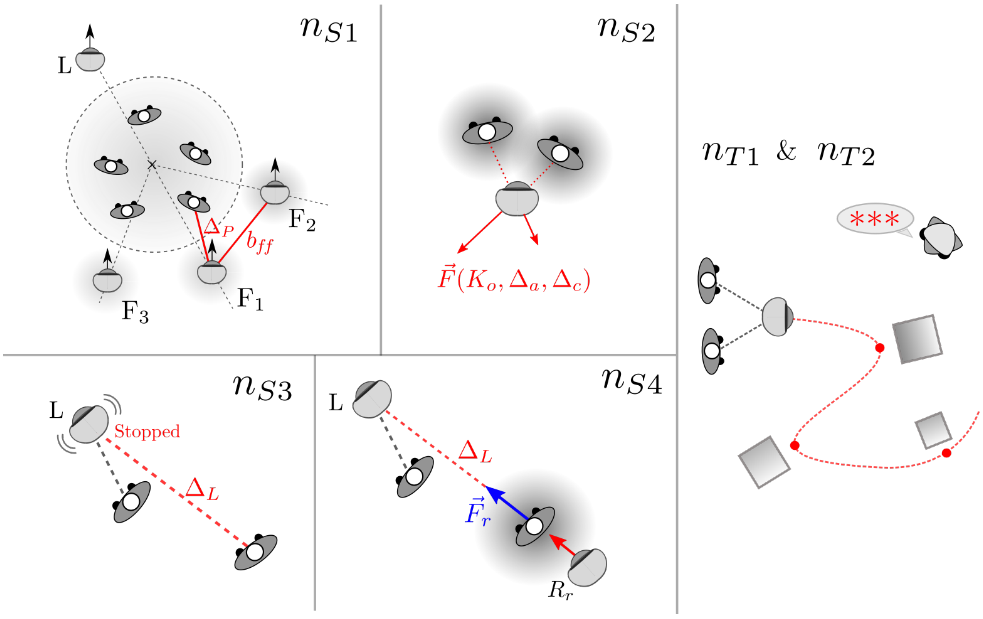

Figure 16.1 – Overview of the norms and their realization in the concrete system.

when the behaviors execute, based on institution specification, domain, and grounding, which are used to automatically generate a planning scheme [13]. The planner enforces the *temporal norms*, while the norm realization enforces the *spatial norms*.

16.3 Norm Realization

In the case study of Chapter 15, we have only provided high-level description of how norms are realized at the domain level. In this case study, we propose a concrete computational protocol of the steps A) - E) of norm realization presented in Definition 13.2. An illustration of the necessary computational steps with a simple example is presented in Appendix D, while here we provide a brief account of the adopted protocol.

16.3.1 Computational Protocol

The steps of the norm realization are captured in Algorithm 4. The same protocol is used in all the case studies in this thesis $\mathcal{C}_I$ - $\mathcal{C}_{III}$.

In the protocol, we can observe two dominant patterns. First, the flow of control resembles a traditional sensing-deliberation-actuation loop of robotic systems, namely we test the conditions defined over world state variables, then plan the next actuation command (behavioral modality), and finally, execute it. Second, the richness of the resulting normative behaviors is achieved through the depth provided by the r^P - r^V - r^B rules.

Our implementation rules are realized as input-output functions. Rules r^N and r^O (lines 5 and

Algorithm 4: NORM REALIZATION

1: Norms $\mathcal{N}$ of the institution are realized at time t by an agent performing behavior B_b given the grounded knowledge K and the values of the state variables R at time t as follows:

2: $\mathcal{N}_t^A \leftarrow \emptyset,\ \mathcal{N}_{t+1}^T \leftarrow \emptyset$ (reset the set of active and satisfied norms)

3: .... NORM ACTIVATION

4: FOR $n_k \in \mathcal{N}$:

5: IF $r_k^N(\Phi_k(C^a)) \equiv \text{True}$: (if the formula over the activation conditions is satisfied)

6: $\mathcal{N}_t^A \leftarrow n_k$ (norm is added to the set of active norms)

7: .... SPECIFICATION OF BEHAVIORAL MODALITY

8: FOR $n_k \in \mathcal{N}_t^A$: (for each active norm)

9: $P_k \leftarrow r_k^P(B_b)$ (find the set of parameters to be modified)

10: FOR $p \in P_k$:

11: $v_p \leftarrow r_k^V(K, R_t)$ (determine the parameter value)

12: .... APPLICATION OF BEHAVIORAL MODALITY

13: $\lambda_b \leftarrow r_k^B(p, v_p)$ (apply the resulting behavioral modality)

14: .... NORM VERIFICATION AT $t+1$

15: FOR $n_k \in \mathcal{N}_t^A$:

16: IF $r_k^O(\Phi_k(C^o)) \equiv \text{True}$: (if the formula over the outcome conditions is satisfied)

17: $\mathcal{N}_{t+1}^T \leftarrow n_k$ (norm has been satisfied at t)

16 in Algorithm 4) evaluate prepositional formulas over boolean conditions. Rule r^P (line 9), given the type of norm and the type of behavior, provides a list of parameters to be modified. Rule r^V (line 11) determines the specific value of the parameter given K and R; it is typically a complex function implemented as a part of the behavior semantics. Finally, rules r^B (line 13) are applied directly as a part of the behavioral procedure.

16.3.2 Realization of Norms

An overview of the norms of the STEERING and the TUTORING institutions is shown in Figure 16.1.

Norm n_{S1} specifies that the configuration of the followers should be chosen so as to retain the humans in the formation. The norm is active when the condition VISITORS_FOLLOW given in Table 16.2 is satisfied. The norm is complied with when ROBOTS_IN_FORMATION, i.e. when robots are close to the desired places in the formation and the formation error meets $e_F \leq e_F^{MAX}$. The rules r^P and r^V determine the exact value for the bias matrix b conceptualized in Figure 16.1. With such configuration the Visitors are placed between the Leader and the Followers and can be closely monitored for deviation from the desired path. This bias is calculated dynamically. It depends on the human positions and is calculated so that the desired place in the formation of the follower closest to a human is at the limit of the personal_distance $\Delta_{P,h} \in K$ of the human, as shown in Figure 16.1. The remaining followers choose the desired positions accordingly, taking into account the extent of the human group (simplified to a circular region). For guidance of a single human, the bias is chosen so as the desired robot-human distance is twice the $\Delta_{P,h}$.

Realization of a norm similar to n_{S2} has been already provided in Chapter 13, while in Appendix D we detail the steps necessary to apply this norm with our computational protocol. In summary, to maintain a comfortable distance from a human according to norm n_{S2}, the robot dynamically creates a repulsive field around the human to prevent interference with his comfort zone. The parameters K_o, Δ_a, Δ_c for a specific human H_h are extracted from the grounded knowledge of the particular domain. The resulting value of $\mathcal{W}_{\mathcal{R}}$ is applied by the MoveInFormation behavior according to Equation (5.8). The norm is active when a condition VISITORS_FOLLOW is true, and satisfied when RESPECTED_PERSONAL_SPACE, i.e. when the resulting human-robot distance d_{ih} at time $t+1$ is higher than the value of the personal distance $\Delta_{P,h}$, for all human Visitors H_h.

Norms n_{S3} and n_{S4} are active under the condition VISITORS_NOT_FOLLOWING, specified for the distance between the Leader and the human furthest from the Leader, $d(L, H_h)$ exceeding a threshold Δ_L. Realization of n_{S3} sets the desired *TrajectorySpeed* parameter of the Leader's MoveOnTrajectory behavior to zero for the duration when the norm is active. Norm n_{S3} is satisfied when LEADER_STOPPED, i.e. $d(p_L(t), p_L(t+1)) \to 0$. Realization of n_{S4} changes the bias matrix in the case when a human is not following. The new bias is chosen so that the desired place in the formation of the follower is at the limit of the intimate_distance $\Delta_{I,h} \in K$ of the human, as shown in Figure 16.1. Norms similar to n_{S3} and n_{S4} appear in the literature. In [93] the robot re-engages a strayed human by exerting a repulsive social force that urges the person towards a goal. In [170] a learning-based approach adapts robot trajectory to slow down or stop when the human is not following.

16.4 Experimental Campaign

We distinguish two scenarios with the aim of demonstrating the effect of introducing normative behaviors in real environments with human participants. Both scenarios consist of a steering phase with the STEERING institution and a tutoring phase with the TUTORING institution.

In Scenario I two human participants are asked to follow behind the Leader robot, commencing the steering phase. A single Follower robot accompanies the participants, while attempting to act in accordance with the norms n_{S1} and n_{S2}. Once the mixed formation enters the exhibition room, the second phase begins (as described in Section 16.1.3), the Follower robot mbot_2 moves aside and the Leader robot mbot_1 assumes the role of an Assistant for a third stationary agent, a pepper robot, that serves as a Tutor. The Tutor robot describes objects present in the scene, while the Assistant robot leads the participants to these objects.

Scenario II is essentially similar to Scenario I but involves only one human Visitor, who is asked to stop following the Leader for around 5 s during the steering phase of the experiment, triggering activation of norms n_{S3} and n_{S4}.

Figure 16.2 – Snapshots of the experiments during the steering phase (a-b) and the tutoring phase (c-d), for Scenario I (a,c,d) and Scenario II (b). Picture b) is taken at the time the participant stops following the Leader.

The experiments are carried out at the Örebro facility (for details see Chapter 2.6.3), formed out of a corridor that leads to a robotics laboratory. Further implementation details are provided in Appendix A.

16.4.1 Scenarios

For Scenario I, we performed 20 experiments with altogether 40 participants (2 participants per run), aged between 19 and 46 years, recruited among researchers and students with various backgrounds. The participants are given the instruction of following the Leader robot (for distinguishing between mbot_1 and mbot_2 we dress the Leader with a scarf). For half of the control group (20 participants) no further information is given. The remaining half of the group is informed about the course of the events (the participants are guided through a corridor and then inside the laboratory), including an explanation about robot roles and their behaviors (one robot is leading, another following, a third robot provides explanations). The participants are asked to carry a box with the UWB tags, the purpose of which is explained to them. After the experiments we hand out a survey with questions listed in Figure 16.7, evaluated on five-level Likert scale.

For Scenario II we perform 10 experiments with 10 participants (1 participant per run). In this scenario we inform the participants about the robot behavior and ask explicitly to stop following, so that the norms n_{S3} and n_{S4} are activated. The evaluation of Scenario II is based on performance metric, no subjective assessment is carried out.

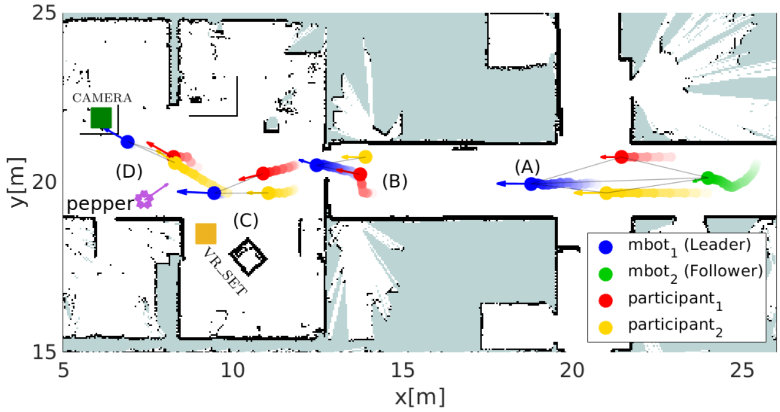

Figure 16.3 – Trajectories of the robots and the participants in Scenario I, shown for the times $t = 11\ s$ (A), $t = 35\ s$ (B), $t = 59\ s$ (C) and $t = 72\ s$ (D).

16.4.2 Results

The objective performance measures and the subjective assessment obtained through the participative study are analyzed with the aim of understanding the characteristics of, and the human preferences for the type of social interactions in the mixed human-robot formations. Representative pictures of the experiments are shown in Figure 16.2. Videos of the experiments are available at the link provided in the footnote[2].

Scenario I

The execution of a representative run of Scenario I is illustrated in Figure 16.3. The scenario starts once the human participants assume their initial positions. This triggers the activation of the STEERING institution with the mbot$_1$ robot assuming the Leader role and the mbot$_2$ robot the role of the Follower, guiding two adult Visitors, participant$_1$ and participant$_2$ through a corridor (A), to the Destination of the STEERING institution at the entrance of the room (B), when the institution becomes dormant.

Figure 16.4 – Average distances for Scenario I. The area between the two human (adult) comfort zones and the Leader or Follower robot – the personal space $\Delta_{P,p_x} = 1.6$ m (shaded in yellow) and the intimate space $\Delta_{I,p_x} = 1.0$ m.

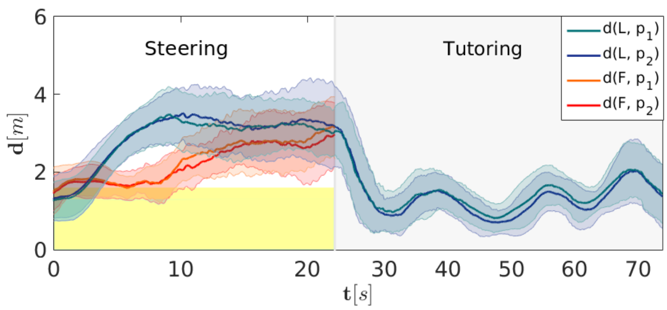

The arrival at the Destination of the STEERING institution coincides with entering the ExhibitionArea of the TUTORING institution. The TUTORING institution becomes active, with mbot$_1$ adopting

2 https://www.epfl.ch/labs/disal/research/InstitutionalRoboticsFormations

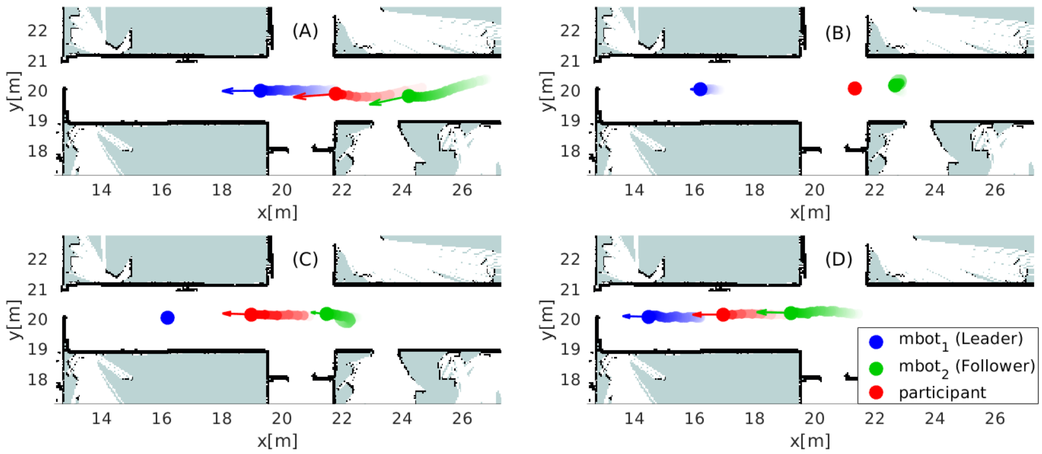

Figure 16.5 – Trajectories of the robots and the participants in Scenario II, shown for the times $t = 16\ s$ (A), $t = 26\ s$ (B), $t = 34\ s$ (C) and $t = 39\ s$ (D).

the role of the Assistant, and the pepper robot taking on the Tutor role. In each experimental run, two objects are chosen from the set of given in Section A.2, and presented by the robots to the Visitors (labels C and D in Figure 16.3). The Tutor and the Assistant synchronize their behaviors – the mbot$_1$ robot guides the Visitors to an object in the scene located at ObjectLocations while the pepper robot utters the description provided through the institutional knowledge ObjectDescriptions.

The norms n_{S1} and n_{S2} pertaining to the MoveInFormation behavior of the Follower robot mbot$_2$ are active throughout the experiments because the condition VISITORS_FOLLOW is always true (and VISITORS_NOT_FOLLOWING is false). The outcome condition of n_{S1} evaluated over the formation error, is not always satisfied. On average, after initialization ($t < 8$ s), e_F stabilizes at around 1.5 m. Log data suggests that the increase of e_F is caused by a saturation of the speed of mbot$_2$ to $\dot{p}_{max}$, illustrating the tradeoff between safety and performance. The evaluation of n_{S2} is based on the distances $d(F, p_x)$ between the adult participants and the Follower, where $p_x = \{p_1, p_2\}$ denotes the participants. The results shown in Figure 16.4 indicate that for the majority of the trials the Follower robot respects the personal_distance $\Delta_{P,p_x} = 1.6$ m, and comply with the norm n_{S2}.

Scenario II

Robot trajectories of a representative experimental run are shown in Figure 16.5. After the team attains an initial steady state at $t = 16$ s (A), the participant stops following (B at $t = 26$ s), as instructed. Once the distance between the Leader and the participant exceeds the threshold $\Delta_L = 3.2$ m, the condition VISITORS_FOLLOW becomes false (and VISITORS_NOT_FOLLOWING becomes true), which causes an activation of norms n_{S3} and n_{S4}. According to the norm n_{S3}, the Leader robot sets its speed to zero and due to the norm n_{S4} the Follower robot reduces the distance to the human to exert a force that is assumed to encourage the human to start following again. After the participant re-joins the Leader (at C, $t = 34$ s), the conditions reverse back, the norms n_{S3} and n_{S4} are no longer active and the team attains a steady state defined by n_{S1} and n_{S2} (at D, $t = 39$ s).

How the value of $d(L, p)$ affects the robot behavior can be further analyzed in Figure 16.6. The plot shows the correlation between $d(L, p)$ and $d(F, p)$. According to n_{S1}, the Follower bias regulates the robot-human distance to reach the set point $2\Delta_{P,p}$ for a single human guidance, so for $d(L, p) \leq \Delta_L$, $d(F, p)$ tends to the value 3.2 m (blue region in Figure 16.6). For $d(L, p) \geq \Delta_L$, $d(F, p)$ (red region) tends to the value $\Delta_{I,p} = 1.0$ m given by intimate_distance $\in K$, and enforced by n_{S4}.

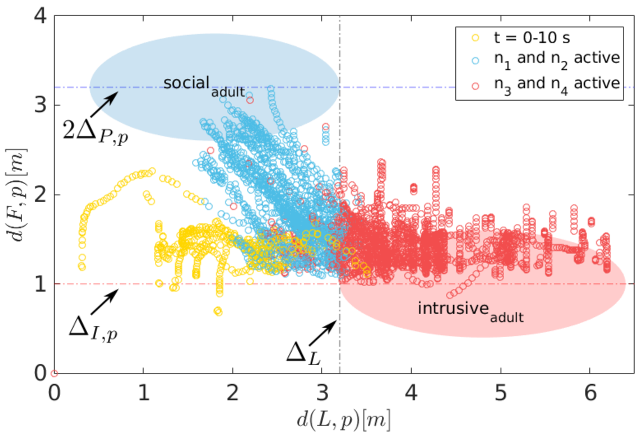

Figure 16.6 – Scenario II. Human-robot distances illustrating how norms are shaping robot behaviors. Blue region indicates activation of n_{S1} and n_{S2}. The value of the bias of the Follower (F) is chosen so as the desired robot-human distance is twice the $\Delta_{P,p}$ (3.2 m). After the participant (p) stops following, the red region is when the Leader (L) stops and waits according to the norm n_{S3}, and the Follower tends to a value of the bias at the limit of the intimate_distance $\Delta_{I,p}$ (1.0 m) given by n_{S4}.

Results of the Participative Study

Results of the survey, presented in Figure 16.7 show that in general, the Follower is perceived as the least understandable and natural among the robots. Moreover, the participants do not observe whether the Follower robot adheres to human social norms or not. In general, the Tutor robot is given the highest scores. Whether or not the participants were given a priori information about the course of the experiment does not have a significant effect on the perception of the robots. Among the most notable differences between the groups we note that the informed group judges the Leader robot as slightly less understandable, less natural and less comfortable to be around. On the other hand, the same group understands the role of the Follower better. Results suggest that robot acceptance does not depend on whether the behavior of the robot is explained. It is possible that the robot's actions are perceived on a subconscious level, and providing a priori explanation has little effect on the perception of robot's friendliness.

Since during the trials the participants always follow the Leader, we do not discern the correlation between the different modalities of the norms the Follower adheres to and the subjective assessment of the participants (we always only observe the norms n_{S1} and n_{S2} being active).

The analysis of the survey results leads us to two main conclusions. First, norms of human societies are not always suitable to be directly applied to robots, or can be applied in some

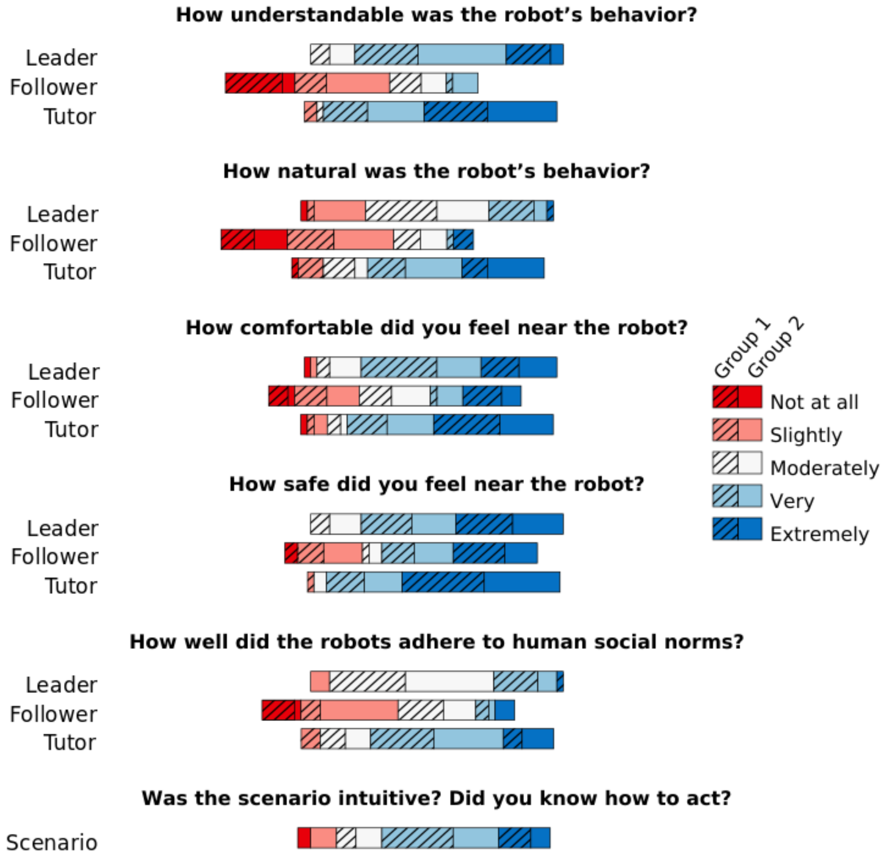

Figure 16.7 – Results of the user-based evaluation on five-level Likert scale. Group 1 involves uninformed participants and Group 2 informed users.

situations but not in others. In our example, the Follower robot, being the only agent that explicitly acted according to social rules, is not perceived as more social than the other robots. Some results reported in the literature are in accordance with our findings [67], [154] [174], while some convey the opposite [175]. Results of our experiments evince that norms should always be verified against user expectations for the particular application. Second, our user studies indicate that the chosen norms and their realizations are not well accepted by the users during the particular case of human guidance in the formation. However, the fact that the norms are encapsulated and abstracted from a system allows us to focus of our future attention on design and re-design of norms, whereas thanks to norm realization, norms can be easily modified or exchanged in a plug-and-play manner.

Summary

With the second case study we continued our course of validation of the institutional framework. In particular, we provided further analysis on three aspects of the formalism. First, we provided a preliminary solution for activation of institutions and building institutional relations through the context recognition in a concrete domain. Second,

we showed how within one institution norms can be conflicting, but the resolution of such conflicts boils down to the selection of appropriate formulas over institutional conditions. This stands in contrast to what is currently proposed in the literature, where conflicts are either never addressed or resolved heuristically (see our discussion in Chapter 10). Third, we demonstrated the flexibility of the framework in terms of embracing heterogeneous robotic platform with contrasting capabilities the institutions reside over, and tools and systems their execution coexists with. Simultaneously, we presented a computational protocol for norm realization, readily implementable in robot's programming language.

The experimental validation carried out in real settings with human participants provided us with insights into acceptance mechanisms for robots that follow human-defined norms and human preferences when in the mixed human-robot groups. The analysis of the user-based evaluation lead us to the conclusion that depending on the socio-cultural context, the same robot behavior can be perceived very differently by the onlooker. Only a systematic evaluation and engagement of the end users in the design process can lead to universally acceptable robot behaviors, but such behaviors must be tailored to the specific person with his level of experience with robotic systems.

17 Case Study: Mixed Teams For Robot Guidance

WITH the final case study we show that institutions can dynamically accommodate the evolution of a given domain by the means of adaptation rules for norm realization. The adaptation rules capture the fact that that robots gain experience, humans gain confidence in robots' actions and the environment undergoes continuous changes. We show that with adaptive norm realization it is possible for the robots to track and adjust to these changes through a self-regulated process, by which the robots can modify their institutional interpretation based on the experience gathered in a given domain.

We demonstrate adaptive norm realization in a new case study, where two humans guide a group of robots through a structured indoor environment. First, we show how conventional non-normative methods can be used for such scenarios and discuss the advantages and disadvantages of introducing the institutional approach with its additional deliberate planning layer. Second, we show how by using institutional norms we can reduce uncertainty of interactions due to norm-induced assumptions regarding the conduct of the other agents. The latter has far-reaching consequences in mixed human-robot teams, as in our example, where robots assume that the human leaders will guide them through safe (obstacle-free) areas of the environment and, based on this assumption, the robots simplify their decision-making routines for selecting the appropriate behavior modality. Third, we compare the performance of the robotic system with and without the capability to adapt and analyze how the introduction of adaptation rules affects the behavior of the robots.

HIGHLIGHTS

- **Adaptation.** With norm adaptation, a cognitive planning layer encompassing the retained experience of a robot presides over the purely reactive behaviors. Institutional evolution through self-regulated adaptation is shown to enhance the performance of the individuals.
- **Generality.** Through the power of institutional grounding and norm realization, the

same norms are leveraged onto two algorithmically different multi-robot behaviors – flocking and navigation in formation.

- **Reliance.** By applying the norms, robots take advantage of the fact that they can make assumptions regarding the behavior of the other agents.
- **Plug-and-Play Institution.** The ROBOT-GUIDANCE institution is applied to one of the domains of $\mathcal{C}_I$, demonstrating that institutions can be freely exchanged within one concrete system as long as the grounding is admissible.
- **Plug-and-Play Domain.** The same institution with the same norm realization is applied to two diverse environments. This is complementary to the institution plug-and-play principle – here one institution is grounded into multiple domains.
- **Norm Reuse.** A number of the norms appears across the institutions in $\mathcal{C}_I$, $\mathcal{C}_{II}$ and $\mathcal{C}_{III}$ and is accommodated by the means of norm realization, which not only concretizes them for a given domain, but also associates the norms to other institutional components making it possible to reuse them for different institutions.
- **Simulation to Reality.** The two aforementioned experimental environments are one in simulation, and one in physical world. The transition from one to another is smooth and does not require parameter tuning.

17.1 The ROBOT-GUIDANCE Institution

In contrast to the previous case studies $\mathcal{C}_I$ and $\mathcal{C}_{II}$, where the norms emphasize human comfort and human view of robots' sociability, the case study $\mathcal{C}_{III}$ described in this chapter is focused on robust multi-robot navigation that can be facilitated through human-robot cooperation mediated by an institution. In particular, we propose that human Leaders guide a group of robot Followers through an environment. Humans as agents with surpassing perception capabilities provide navigation through areas that are safe to pass, while the robots keep a formation. We achieve such cooperation through a ROBOT-GUIDANCE institution. It can be used in applications where the capabilities of the robots are insufficient for them to pass safely to a destination, for instance, when operating in unknown environments or environments with unmapped obstacles.

Human Leaders, $L_H = \{L_1, L_2, ...\}$ with $|L_H| = N_L$, as agents with superior perception capabilities, guide a group of robots through the environment. The space between the Leaders is assumed to be obstacle-free at all times, so the outermost paths of the Leaders delineate areas that can be safely passed by the robot team, i.e. *safe areas*. Robot Followers, $F = \{F_1, F_2, ...\}$ with $|F| = N_F$ follow the humans as a team using one of the two multi-robot behaviors, B_1 = MoveByFlocking or B_2 = MoveInFormation.

The ROBOT-GUIDANCE institution takes the form:

$$
\begin{aligned}
&\textbf{ROBOT-GUIDANCE} = \Big\langle \\
&\qquad \mathit{Norms} = \{n_1, n_2, n_3, n_4, n_5\} \\
&\qquad \mathit{Roles} = \{\mathsf{Leader},\ \mathsf{Follower}\} \\
&\qquad \mathit{Actions} = \{\mathsf{Guide},\ \mathsf{Follow}\} \\
&\qquad \mathit{Conditions} = \{\text{LEADER\_POSITIONS\_KNOWN},\ \text{FOLLOWER\_LOST}, \ldots\} \\
&\qquad \mathit{Knowledge} = \{\mathsf{PersonalSpace},\ \mathsf{SafeAreas},\ \mathsf{VirtualLeader}\} \Big\rangle
\end{aligned}
$$

The norms of the ROBOT-GUIDANCE institution are listed in Table 17.1.

ROBOT-GUIDANCE		
Safe Areas	n_1	"*Followers must follow leaders within the safe areas*"
Following	n_2	"*In complex spaces, followers must ensure to follow paths of the leaders*"
Help Request	n_3	"*When lost, the follower must notify the guide, and the guide must wait for that follower*"
Narrow Spaces	n_4	"*The followers should negotiate passing though narrow spaces*"
Social Spaces	n_5	"*Personal spaces must be respected*"

Table 17.1 – Norms of the ROBOT-GUIDANCE institution.

The color code highlights the components of the norm as specified by the syntax proposed in Definition 11.2, namely:

$$\mathcal{N} : \mathit{Conditions} \rightarrow \mathsf{deontic}(\mathit{Roles} \times \mathit{Actions} \times \mathit{Knowledge}).$$

With the above, the components of norms n_1 to n_5 can be summarized as follows:

$$
\begin{aligned}
n_1 &: \emptyset \rightarrow \mathsf{must}(\mathit{followers},\ \mathit{follow},\ \mathit{safe\ areas}) \\
n_2 &: \mathit{in\ complex\ spaces} \rightarrow \mathsf{must}(\mathit{followers},\ \mathit{follow},\ \mathit{leaders\ paths}) \\
n_3 &: \mathit{when\ lost} \rightarrow \mathsf{must}((\mathit{follower},\ \mathit{notify}), (\mathit{guide},\ \mathit{wait\ for\ follower})) \\
n_4 &: \mathit{narrow\ spaces} \rightarrow \mathsf{should}(\mathit{followers},\ \mathit{negotiate\ passing},\ \mathit{safe\ areas}) \\
n_5 &: \emptyset \rightarrow \mathsf{must}(\mathit{be\ respected},\ \mathit{personal\ spaces})
\end{aligned}
$$

Note that proposed set of norms is not exhausting all the possibilities. Moreover, the norms listed above pertain to the Follower agents, which, as it will be shown later, will be grounded to the robots. It would be natural to define a norm stating that n_6 =“*Leaders should always guide only through safe areas*”, but since the human agents, to whom the Leaders will be grounded to do not use norm realization to interpret the norms, we will only focus on the norms that concern the robots. Norm n_6, however, has important implications for the norm n_1, which allows the Follower robots to act based on a belief that the human Leaders are obliged to always guide through safe areas. Because of introducing such mutual norms, agent actions no longer rely on assumptions regarding the behavior of other agents, but rather on understanding and trust in their conduct.

The full list of institutional *Conditions* and *Knowledge* is given in Table 17.2 and Table 17.3, respectively. We will revise their elements when explaining how they are grounded in the domain.

17.2 Domain and Grounding

The domain includes the following components:

$$A = \{\mathsf{human}_1,\ \mathsf{human}_2,\ \mathsf{mbot}_1,\ \mathsf{mbot}_2,\ \mathsf{mbot}_3\}$$
$$B = \{\mathsf{MoveOnPath},\ \mathsf{MoveInFormation},\ \mathsf{MoveByFlocking}\}$$
$$R = \{\mathsf{time},\ \mathsf{pose}_\mathsf{x},\ \mathsf{path}_\mathsf{x},\ \mathsf{virtual\_point},\ \mathsf{environment}, \ldots\}$$
$$K = \{\mathsf{personal\_distance}_\mathsf{x},\ \mathsf{safe\_area},\ \mathsf{virtual\_point}\}$$
$$C = \{C_1,\ C_2,\ C_3,\ C_4,\ C_5\}$$
$$L = \{\text{IN\_SAFE\_AREA},\ \text{IS\_FOLLOWING},\ \text{IN\_SOCIAL\_SPACE}\}$$

where we use the subscript $(\cdot)_x$ to abbreviate variables with multiple instances x.

The robots engage in either of the two behaviors: MoveInFormation, or MoveByFlocking. We assume that all robots in one instance of the domain adopt the same behavior, therefore we distinguish two corresponding groundings: first $\mathcal{G}^{FL}$, where the robots perform the MoveByFlocking behavior, and second $\mathcal{G}^{FR}$, where the robots perform the MoveInFormation behavior:

$$\begin{aligned}
\mathcal{G}_A &= \big\{(\mathsf{Leader},\ \{\mathsf{human}_1,\ \mathsf{human}_2\}),\ (\mathsf{Follower},\ \{\mathsf{mbot}_1,\ \mathsf{mbot}_2,\ \mathsf{mbot}_3\})\big\}\\
\mathcal{G}_B^{FL} &= \big\{(\mathsf{Guide},\ \mathsf{MoveOnPath}),\ (\mathsf{Follow},\ \mathsf{MoveByFlocking})\big\}\\
\mathcal{G}_B^{FR} &= \big\{(\mathsf{Guide},\ \mathsf{MoveOnPath}),\ (\mathsf{Follow},\ \mathsf{MoveInFormation})\big\}\\
\mathcal{G}_C &= \big\{(\text{LEADER_POSITIONS_KNOWN},\ c^a_{11}),\ (\text{FOLLOWER_LOST},\ c^a_{31}), \ldots\big\}\\
\mathcal{G}_K &= \big\{(\mathsf{PersonalSpace}, \mathsf{personal_distance}_\mathsf{x}), \ldots\big\}
\end{aligned}$$

We will use the notation $\mathcal{G}_x(\cdot)$ with $x \in A, B, C, K$ to list the grounded counterpart of $(\cdot)$, e.g. $\mathcal{G}_A(\mathsf{Leader}) = \{\mathsf{human}_1, \mathsf{human}_2\}$.

	Activation Conditions		Outcome Conditions	
n_1	LEADER_POSITIONS_KNOWN	c^a_{11}	FOLLOWER_IN_SAFE_AREA	c^o_{11}
n_2	LEADER_POSITIONS_ALWAYS_KNOWN IN_COMPLEX_ENVIRONMENTS	c^a_{21} c^a_{22}	FOLLOWER_FOLLOWS_LEADERS_PATHS	c^o_{21}
n_3	FOLLOWER_LOST	c^a_{31}	LEADERS_WAIT	c^o_{31}
n_4	IN_NARROW_PASSAGE	c^a_{41}	FOLLOWER_FOLLOWS_LEADERS_PATHS	c^o_{41}
n_5	HUMAN_POSITIONS_KNOWN	c^a_{51}	RESPECTED_PERSONAL_SPACES	c^o_{51}

Table 17.2 – *Conditions* of the institution ROBOT-GUIDANCE.

Each element in *Conditions* is grounded to $c^x_{ab} \in C_a$, where a denotes the index of a norm the condition pertains to, b is the index of the grounded condition within set C_a, and $x \in \{a, o\}$ allows to distinguish between the activation and the outcome conditions. The full grounding of conditions can be found in Table 17.2.

Institutional *Knowledge* includes the following elements:

- PersonalSpace for determining a social_distance $\Delta_{P,h}$ to a human H_h,
- SafeAreas, the extent of which is provided by the Leaders at a run time,
- VirtualLeader and ObstacleCharacteristics, including real O^R, virtual O^V and human-induced O^H obstacles, necessary for understanding the shared navigation concepts of the MoveByFlocking behavior (see Section 5.3.2 for details of the algorithm),
- HumanSocialForce denoted HF and TeamSocialForce denoted TF for allowing each agent to participate in the collective motion of the team of humans and robots, and
- FollowerLostDistance and NarrowPassageMark for agreeing on common thresholds used by the norm conditions.

The full grounding of institutional knowledge is provided in Table 17.3. Note that in this particular institution we consider obstacles as mechanisms of social order – artifacts upon which robots negotiate their safe passage. Knowledge about the obstacles is shared so that all the robots know when deliberation is necessary for passing next to them.

We assume that basic capabilities of the robots, such as ability to perform a speech act or stop in place, can be categorized as behaviors. Such behaviors do not involve interactions and thus are not grounded to institutional *Actions*. They are referred to as *individual behaviors*.

The experience measures L used for norm adaptation are further explained in Section 17.3.2.

Knowledge	Grounded Knowledge K	Symbol
PersonalSpace	personal_distance$_x$	$\Delta_{P,h}$
SafeAreas	safe_area	Ω
VirtualLeader	virtual_point	VL
	former_virtual_point	FVL
HumanSocialForce	social_force$_x$	$\mathrm{HF}_h = \{K_o, \Delta_a, \Delta_c\}$
TeamSocialForce	team_force	$\mathrm{TF} = \{w_{ff}, w_{fl}\}$
FollowerLostDistance	lost_threshold	T_L
NarrowPassageMark	narrow_threshold	T_C
ObstacleCharacteristics	obstacle$_x$	$o_x \in \{O^R, O^V, O^H\}$

Table 17.3 – *Knowledge* and knowledge grounding K of the institution ROBOT-GUIDANCE.

17.3 Norm Realization

In this chapter, we perform the additional *adaptation* step of norm realization. In our implementation, the adaptation is applied continuously throughout the execution of robot behaviors. Whether to apply norm adaptation or not is a design choice, and depends on whether the performance of the robot team calls for adjustments, or whether the environment or humans within are expected to undergo any changes (e.g., large, structural obstacles are moved around, humans adjust to robot presence, etc.).

17.3.1 Realization of Norms

The detailed description of the norm realization is provided in Appendix E, while below we lay out an intuitive explanation of how norms are implemented in our system. The realization of norms n_1-n_5 refers to the Follower robots – all the robots of this case study and is illustrated in Figure 17.2.

Norm n_1: Safe Areas

This norm is active when LEADER_POSITIONS_KNOWN i.e. when the follower has access to the positioning information of the leaders.

- For the MoveByFlocking behavior, we include *virtual obstacles* O^V at the limits of the safe area so that the robots are kept inside it (see Figure 17.2 for illustration). Norm adaptation further affects the weights of the repulsive forces generated by O^V. Details of the computational aspects of the flocking algorithm are provided in Chapter 5.
- For the MoveInFormation behavior, the width d_Ω of the safe area drives computation

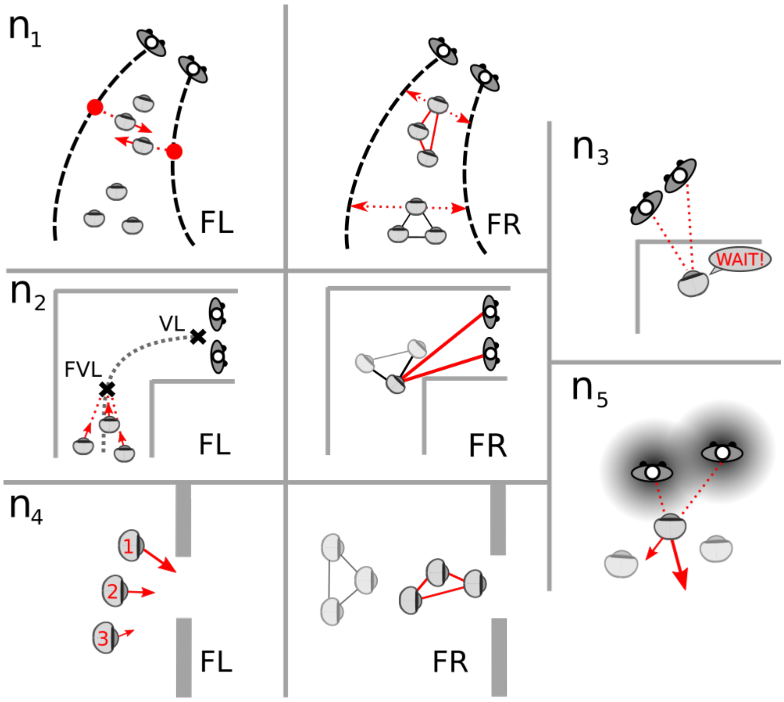

Figure 17.1 – Illustration of norm realization of the ROBOT-GUIDANCE institution. Label FL stands for the MoveByFlocking behavior, while FR for the MoveInFormation behavior.

of the formation bias using the LFT method presented in 7, scaling the original bias of F_i towards the column-like shape. Norm adaptation accelerates the rate at which the modification is done.

The norm is complied with if FOLLOWER_IN_SAFE_AREA.

Norm n_2: Reinforcement of Following

To facilitate navigation in *complex spaces* (i.e. confined areas with narrow spaces often requiring sharp turns) norm n_2 provides the Followers with the means to reinforce the Leader-Follower edges. It is active when LEADER_POSITIONS_ALWAYS_KNOWN (because it is important that the robot has access to the history of positions of the Leaders), and the robot is IN_COMPLEX_ENVIRONMENTS.

- For the MoveByFlocking behavior, instead of following the immediate virtual_point VL, the follower robot migrates towards a *former* FVL, the access to which is not obstructed by structured obstacles that can create local minima (see Figure 17.2, where F_i moves to the closer FVL(t) point instead of moving to VL(t)). By being more immediate than VL, FVL allows F_i to follow the path of the Leaders even when it has lost their line of sight. Norm adaptation amplifies the attraction weights of the flocking migration component.
- For the MoveInFormation behavior weights of the Laplacian Leader-Follower edges w_{lf}

are reinforced with respect to the Follower-Follower edges with weights w_{ff}. Norm adaptation further exacerbates the w_{lf} to w_{ff} ratio.

The norm is complied with if FOLLOWER_FOLLOWS_LEADERS_PATHS, i.e. the Follower is able to follow the general route provided by the Leaders, even if at a large distance.

Norm n_3: Help Request

This norm is active when FOLLOWER_LOST, i.e. when a Follower falls behind the Leaders further than a threshold T_L. The norm parametrizes an individual robot behavior – a speech act, not represented in the institutional formalism. The Follower asks the Leaders to stop and wait until it can catch up. The norm is satisfied if LEADERS_WAIT. We assume that the human Leaders stop each time a Follower requests it.

Norm n_4: Narrow Spaces

Norm n_4 facilitates passing of the Followers through narrow spaces by assigning who is to pass next. It is active when IN_NARROW_PASSAGE, i.e. when the passage that a Follower traverses is narrower than a threshold T_C.

- For the MoveByFlocking behavior, the control gain K_c of each individual F_i is tuned according to its ID i, so that some robots speed up when passing, while others slow down, letting the faster robots to proceed. Norm adaptation amplifies the gain further if the previous value did not yield the desired performance.
- For MoveInFormation, the coordination when passing a narrow opening is achieved by a length-wise change of the formation geometry with the LFT method, which spreads the bias along the direction of motion. Norm adaptation accelerates the rate of formation shape modification.

The norm is complied with when FOLLOWER_FOLLOWS_LEADERS_PATHS.

Norm n_5: Social Spaces

The last norm is similar to n_2 of $\mathcal{C}_I$ and n_2 of $\mathcal{C}_{II}$. For both behaviors n_5 introduces repulsive forces that drive the robots away from human personal spaces. It is active when HUMAN_POSITIONS_KNOWN.

- For the MoveByFlocking behavior humans are modeled as the β-agents of the flocking algorithm described in Chapter 5, i.e. a virtual obstacle $o_h \in O^H$ is placed at the position of the human H_h, while the strength of the repulsive force is proportional to the value $\Delta_{P,h}$ of the PERSONAL_DISTANCE $\in K$. Norm adaptation increases the relative strength of the force and increases the range it operates within.
- For the MoveInFormation behavior we change the parameter *RepulsionWeights* $\mathcal{W}_{\mathcal{R}}$ as

described in detail in Section 13.1.5. Norm adaptation further reinforces the gain of the repulsive weight.

The norm is complied with when the condition SOCIAL_SPACES_RESPECTED is satisfied.

17.3.2 Norm Adaptation

The experience measures L are collected throughout the course of the experiments and are used for online adaptation of the norm values. At the domain level, we define the following experience measures L to evaluate outcomes of the norms:

IN_SAFE_AREA	F_i was located within the safe area Ω
IS_FOLLOWING	F_i was able to follow within the distance T_L
IN_SOCIAL_SPACE	F_i respected human personal spaces Δ_P

In our implementation, we normalize the elements of L so that $l \in [0,1]\ \forall\ l \in L$. As explained in Chapter 13, the adaptation rules r^L take the current parameter value v_p, determined through the choice and the value rules of norm realization, and modifies it in a way that is thought to emphasize the enforcement of the norm.

As an example, consider the case of norm n_5. When the repulsive forces that are designed to keep the robot out of human comfort spaces do not yield the desired results (i.e. the experience measure IN_SOCIAL_SPACE indicates recurrent violation of the norm n_5), the adaptation rules amplify the strength of the forces as a function of the degree of noncompliance. In our implementation the strength of the forces decreases linearly with the experience measure IN_SOCIAL_SPACE, which evaluates whether the Follower is able to stay away from human personal space. In a marginal case, when the measure always evaluates to true, no norm adaptation is necessary and the original force strength is applied. On the contrary, if the robot experiences frequent, severe violations, the linear function of the measure IN_SOCIAL_SPACE reinforces the original strength of the force. Comprehensive details of the adaptation rules are provided in Appendix E.

17.4 Experimental Campaign

Our case study focuses on cooperative navigation of humans and robots in structured, indoor environments. In this case study humans and robots roles are reversed in respect with $\mathcal{C}_{II}$, where robots guide a group of humans.

First, we introduce the notation. A *path* $\mathcal{P}$ of an agent i is the collection of all the points the agent visited up to time t, i.e. $\mathcal{P}_i(t) = \{p_i(t_0), p_i(t_0+1), ..., p_i(t)\}$, where $p_i(t)$ is position of agent i at time t. A *safe area* Ω is an obstacle-free zone between the outermost paths of the Leaders. The distance $d_\Omega(p)$ is the width of the safe area at a position p (see Fig. 17.2

for illustration). In the case of the MoveByFlocking behavior, the Follower robots follow the Leaders by the means of a *virtual point* - a point calculated based on the Leaders' positions. In our implementation, we set the virtual point to be the center of positions of human Leaders L_H, i.e. $\mathrm{VL}(t) = |L_H|^{-1}\sum_{h\in L_H} p_h(t)$. The *virtual path* is $\mathcal{P}_{\mathrm{VL}}(t) = \{\mathrm{VL}(t_0), \mathrm{MP}(t_0+1), \ldots, \mathrm{VL}_i(t)\}$. The concepts of path and virtual point are shown in Figure 17.2.

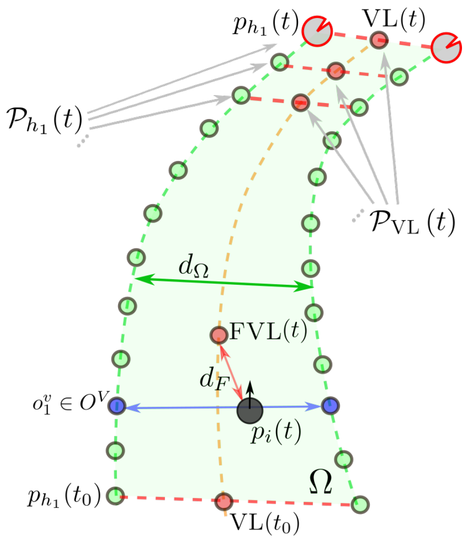

Figure 17.2 – Illustration of the definitions used in this chapter. A path $\mathcal{P}$ of an agent i is the collection of all the points the agent visited up to time t. A safe area Ω is an obstacle-free zone between the outermost paths of the Leaders. The width of the safe area at a position p is denoted as $d_\Omega(p)$. A virtual point VL is located at the center of the Leaders positions.

We perform two sets of experiments. The first series is performed in simulation, the second in real settings in the Jordils facility. Experiments are performed with two human Leaders and three Follower robots, unless stated otherwise. For safety reasons, we do not disable robots' self-localization functionalities. The robots communicate their positioning information to all other team members and have access to the human state, but the methods only rely on relative positioning information, therefore they are readily implementable in environments unknown to the robots, under the condition that robots can measure the relative localization of the other team members. The implementation details of this case study are provided in Appendix A, while the thorough account of norm adaptation is given in Appendix E.

17.5 Experiments Set I

All experiments are performed in a simulated indoor environment with narrow passages and sharp turns (see Figure 17.3).

17.5.1 Scenarios

We distinguish the following scenarios.

- Scenario I (NonNormative), where the robots perform a non-normative behavior without the institutional layer;

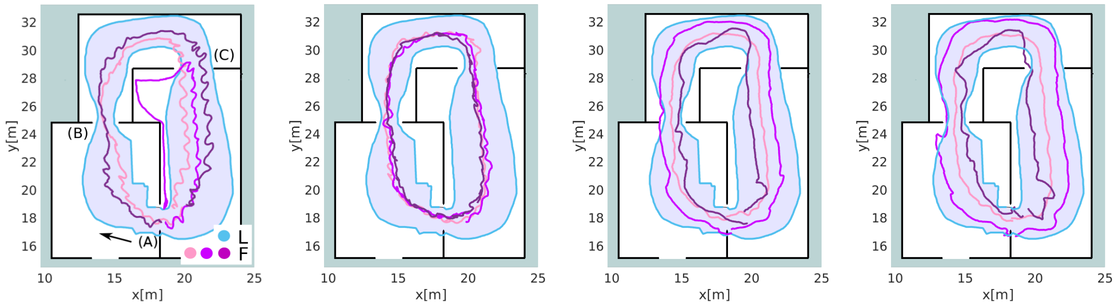

(a) (b) (c) (d)

Figure 17.3 – Trajectories of the Follower robots and the Leader humans for a) NonNormative, MoveByFlocking, b) Institutional, MoveByFlocking, c) NonNormative, MoveInFormation, and d) Institutional, MoveInFormation. Shaded blue area is the safe area Ω. Label (A) indicates the beginning of the run, the same for each scenario; labels (B) and (C) indicate the first and the second narrow passages; arrow indicates the direction of motion.

- Scenario II (Institutional), where the behaviors are governed by the institutional norms;
- Scenario III (Adaptive), where the measures $L(t)$ are used for value adaptation in the next time step $t+1$.

The goal of Scenario II is to understand the effect of introducing the institutional norms on top of a non-normative baseline system. Scenario III tests whether norm adaptation can improve the performance of the institutional approach.

Each scenario is completed with the Follower robots performing one of the two behaviors, MoveByFlocking, or MoveInFormation. For each scenario, we perform 10 sequential runs, each lasting approximately 100 s. Paths of the humans are obtained from an operator input, recorded and replayed during each run to achieve consistency.

17.5.2 Results

Figure 17.3 shows typical trajectories of the agents, robot Followers and human Leaders, navigating through the environment, with the starting location of the run at (A), and with the clockwise direction of motion. Subfigures a) and b) exemplify the MoveByFlocking behavior, for the NonNormative and the Institutional cases, respectively; subfigures c) and d) represent the NonNormative and the Institutional cases of the MoveInFormation behavior, respectively.

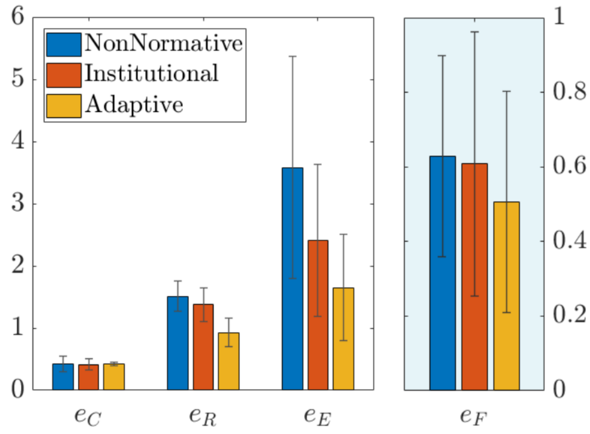

Figure 17.4 – Performance of the robots in Scenarios I to III. Connectivity e_C, cohesion e_R and deviation e_E are the metrics for the flocking behavior, while formation error e_F evaluates the formation behavior. Error bars show the standard deviation over the runs.

The MoveByFlocking Behavior

Adding the augmented navigational competences on top of the NONNORMATIVE baseline in the case of the MoveByFlocking behavior brings advantages visible at the first glance in Figure 17.3 a and b. A Follower mbot_3, the third robot to pass though narrow door at (A), disconnects from the rest of the team, as it is not able to pass the narrowing before the Leaders progress substantially on their paths. After the Leaders leave the LOS of mbot_3, the VL (centre of Leader positions) shifts further upwards, making the lost robot assume that direction of motion, but remain behind the wall (note that according to the Γ_α component of flocking, the Follower is not only attracted to the VL, but also accelerates in the direction of VL's velocity). It only recovers when the rest of the team enters the same room at (C) and the overall flocking force directs it through the obstacle-free area. Similar phenomenon is observed in all of the 10 runs of NONNORMATIVE MoveByFlocking, as the original flocking algorithm has no means to escape local minima.

By adding the normative layer onto the baseline flocking, we deal with the above problem by encouraging the robot through the realization of norm n_2 to use FVL, a former virtual point closer to the robot and close enough to be within LOS, which allows the robot to follow the general route of the Leaders, instead of rigidly following the most recent (and not always feasible) VL. Note that it would be entirely possible to augment the navigational behaviors with the additional features such as FVL without resulting to the use of norms, as we have emphasized many times over the course of this work. However, the primary purpose of employing institutional, norm-based approach is generalizability, encapsulation and reusability of heuristic, case-dependent methods.

The average performance for Scenarios I-III is shown in Figure 17.4. The INSTITUTIONAL method yields lower cohesion e_R and deviation e_E than the NONNORMATIVE method for two reasons. First, in the INSTITUTIONAL case no robot is trapped in local minima. Second, introducing virtual obstacles at the boundaries of the safe area in the realization of norm n_1 makes the robots attain a column-like shape in contrast to the well-separated lattice-shaped

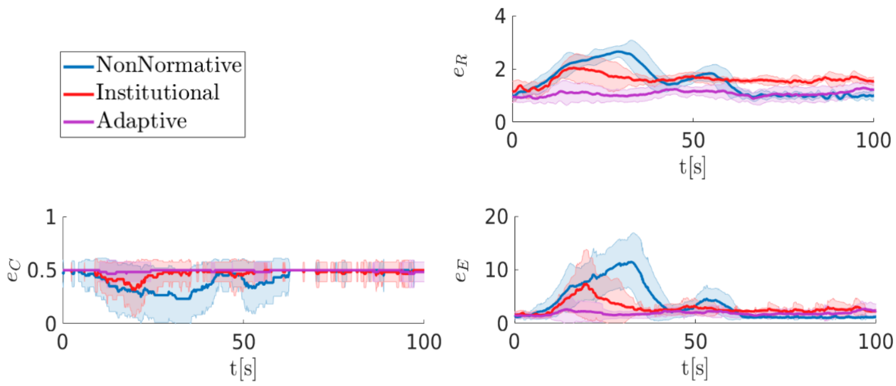

Figure 17.5 – Average flocking performance, i.e. connectivity e_c, cohesion e_R and deviation e_E, throughout the runs. Shaded areas show standard deviation.

pattern achieved in the NonNormative case.

The Adaptive method always leads to performance improvement over the Institutional case as indicated by all the metrics. Time-wise plots of the performance metrics shown in Figure 17.5 indicate that the norm-induced navigational capabilities smoothen the spikes of the performance errors, while adaptation levels them out completely, making the errors consistent irrespective of whether the robots pass through complex narrow passages with sharp turns or thorough straight open spaces.

The MoveInFormation Behavior

Compared to flocking, formation control leads to less oscillatory movement of the robots, with well-defined separation between all the agents, as it can be seen in Figure 17.3 c and d. All three versions of the MoveInFormation behavior perform well, with all the Followers being able to keep their desired places in the formation. In the case of the Institutional approach the Followers visibly modify the formation shape in a reaction to the environment governed by the concurrent operation of norms n_1 and n_4, increasing the formation size when navigating through open space and reducing it when encountering a narrow passage. In our understanding such reaction is social because it involves norm-based regulation of how the robots mediate their actions. Modification of the bias involves agreement on the mutuality of the formation change – robot F_i cannot change its bias b_{ij} to F_j without F_j doing the same. Additionally, modification of the formation shape in this case study is based on the input from the Leaders [1], which delineate obstacle-free safe areas, so that the robots act relying on trust in the Leaders complying to their obligations. Agent actions no longer rely on assumptions regarding the conduct of the other agents, but rather on understanding and trust introduced by such mutual norms.

The formation error e_F shown in Figure 17.4 is similar for both the NonNormative and the

1 In the LFT method presented in Chapter 7 the Followers alter the bias based on local sensing instead.

	MoveByFlocking			MoveInFormation		
	NonNormative	Institutional	Adaptive	NonNormative	Institutional	Adaptive
FOLLOWER_LOST	0.21	0.05	0.0	0.0	0.0	0.01
RESPECTED_PERSONAL_SPACES	0.87	0.97	0.97	0.99	1.0	1.0

Table 17.4 – Average fraction of time a condition is satisfied in Scenarios I-III.

INSTITUTIONAL methods. Norm adaptation results in a decrease of e_F of 23% compared to the INSTITUTIONAL case. One should note that during the transient time, an alteration of the formation shape unavoidably increases the formation error, as the dynamic system needs some time to adjust to the change of configuration.

Norm satisfaction

In Table 17.4 we present two critical results with regard to the rate of norm satisfaction, one for whether the Followers are able to keep up with the Leaders, one for sociability of the human-robot interactions. The results confirm that the INSTITUTIONAL and the ADAPTIVE methods result in the Followers complying with the norms. Norm adaptation results in an improvement of norm satisfaction, albeit a slight one, as the INSTITUTIONAL method already yields close-to-ideal results.

A baseline NONNORMATIVE flocking algorithm leads to breaking of the formation, as indicated by the FOLLOWER_LOST condition. Similarly, the same algorithm results in personal spaces of human Leaders being intruded 13% of the time. The latter is a result of the humans being treated simply as β-agents, with the strength of the repulsive forces ample for avoiding collisions, but insufficient for respective human comfort zones. The aforementioned conditions are far less frequently breached by the baseline NONNORMATIVE formation control method, due to the rigidity of the agent-agent edges. No formation failure is noted and the personal spaces of human Leaders are rarely intruded.

17.6 Experiments Set II

The goal of the following set of experiments is twofold. First, we validate the performance of the methods in the real world, but with a *different* grounding than in the simulation. The difference lies mostly in the shape and the size of the environment, which for the real experiments is almost three times smaller. Second, we use the *same* grounding as the one of a WAREHOUSE in Case Study $\mathcal{C}_I$ presented in Chapter 15. While in Chapter 15 we have shown that it is possible to apply the same institution to different domains, in this chapter we

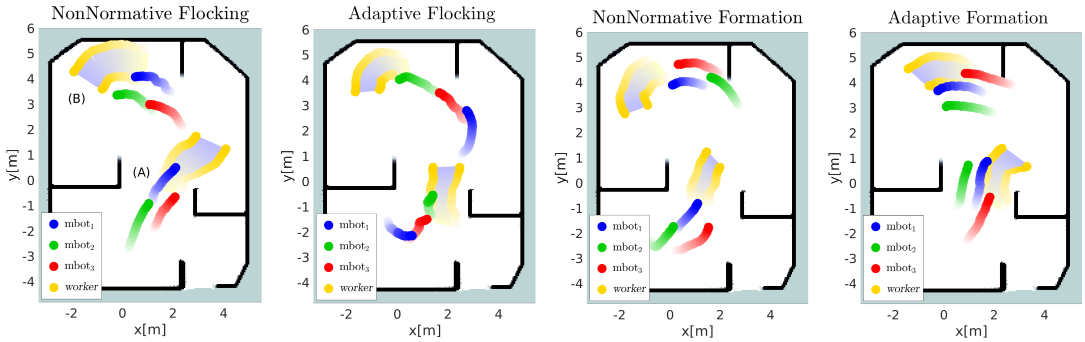

Figure 17.6 – Trajectories of the Follower robots and the human Leaders. The shaded blue area is the safe area Ω.

demonstrate that the opposite also holds: it is possible to ground different institutions to the same domain. And since the domain in the real experiments is different than the domain in the simulation, we once again demonstrate the institution plug-and-play property.

17.6.1 Scenarios

Experiments are carried out for three scenarios: NonNormative, Institutional and Adaptive. For each scenario, we perform 10 consecutive runs, each lasting approximately 50 s. The paths of the Leaders are provided by one human worker who carries two hand-held reference points, tracked by the MCS. By controlling the relative spacing between the reference points, the worker delineates the safe_area.

The agents of the domain are now $A = \{\mathsf{human}_1, \mathsf{mbot}_1, \mathsf{mbot}_2, \mathsf{mbot}_3\}$ where human_1 has the function of a worker and the role-to-agent grounding now includes the reference points the human provides $(\mathsf{Leader}, \{\mathsf{reference\_point}_\mathsf{A}, \mathsf{reference\_point}_\mathsf{B}\}) \in \mathcal{G}_A$. The size of the worker's personal_distance is the same as in the Case Study $\mathcal{C}_{III}$ and amounts to 0.8 m (see Table 15.6). Real experiments are carried out in an area of approximately a third of the size of the simulation setup and characterized by a relatively open space (see Figure 17.6), with wall constraints that force the Leaders to take sharp turns, or even to turn back on the spot, as compared to the simulated area, dominated by narrowings and ninety-degree turns (shown in Figure 17.3).

17.6.2 Results

Figure 17.6 shows representative snapshots of the NonNormative and Adaptive scenarios at the time of around 15 s (A) and 30 s (B). In the case of the MoveByFlocking behavior, in the NonNormative baseline, the flock is compact and keeps a lattice-like shape. Introducing norms in the Institutional and the Adaptive methods results in a stretch of the flock, which attempts to stay within the safe_area delineated by the reference points provided by the worker. This phenomenon is reflected in the metrics shown in Figure 17.8, where the cohesion e_R and

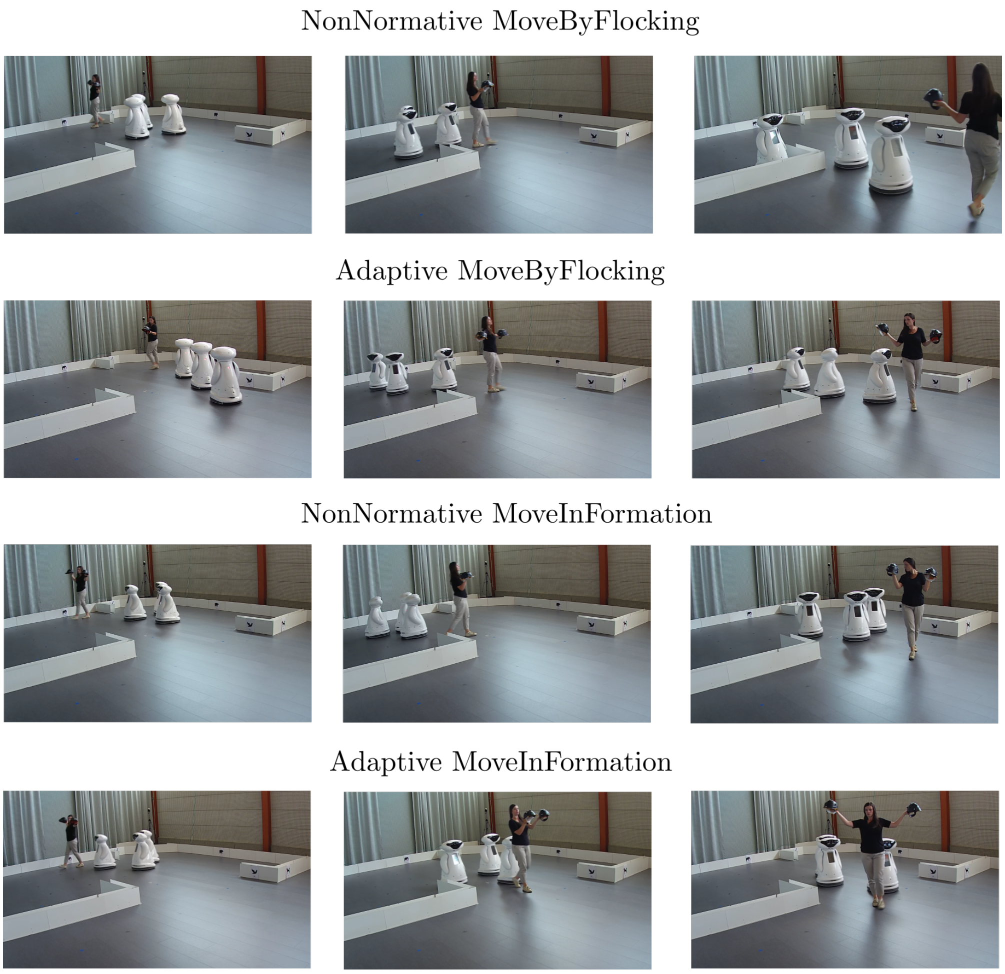

Figure 17.7 – Snapshots taken during the experiments. The third picture of the NonNormative MoveByFlocking behavior is an example of the team failure.

deviation e_E metrics show the smallest values for the NonNormative case.

In the case of the MoveInFormation behavior, in the NonNormative baseline the Followers lag behind the Leader, but keep the length of the Follower-Follower edges close to the desired value. Introducing norms that reinforce the Leader-Follower edges leads to an increase of formation compactness and a reduction of the formation error e_F (shown in Figure 17.8) at the expense of intruding human personal space, which in the case of a worker with personal_distance of 0.8 m is breached over 30% of the experimental time in the Institutional case, and over 40% of the time in the Adaptive case, but only around 5% of the time in the NonNormative case[2].

Shown in Table 17.5, the failure rate is defined as the portion of runs when at least one of the robots stops following and does not recover before the end of the run. In the case of

2 Here we refer to the role of the Leader grounded to a reference_point as a human, as for the robots the implementation details are transparent. In the next section we provide an additional explanation of how this grounding is interpreted by the robots.

	MoveByFlocking			MoveInFormation		
	NonNormative	Institutional	Adaptive	NonNormative	Institutional	Adaptive
SUCCESS RATE	50%	100%	100%	80%	90%	90%

Table 17.5 – The success rates in the real experiments. Failure is defined as the portion of runs when at least one of the robots stops following and does not recover before the end of the run.

the NONNORMATIVE MoveByFlocking experiment the success rate is only 50%, while for the NONNORMATIVE MoveInFormation method it amounts to 80%. No failures are recorded during the INSTITUTIONAL and ADAPTIVE scenarios for MoveByFlocking behavior, while for the MoveInFormation, we note one failure (90% success rate) during each of INSTITUTIONAL and ADAPTIVE scenarios.

Finally, as noticeable in Figure 17.8, adaptation helps improving e_R and e_E metrics in comparison to the non-adaptive institutional counterpart. For the MoveByFlocking behavior, e_R is 11% smaller and e_E is almost 40% lower for the ADAPTIVE method, indicating that the flock is more compact, while retaining the excellent success rate (100%) of the INSTITUTIONAL method. For the MoveInFormation behavior, the formation error e_F is 7% smaller in the ADAPTIVE case.

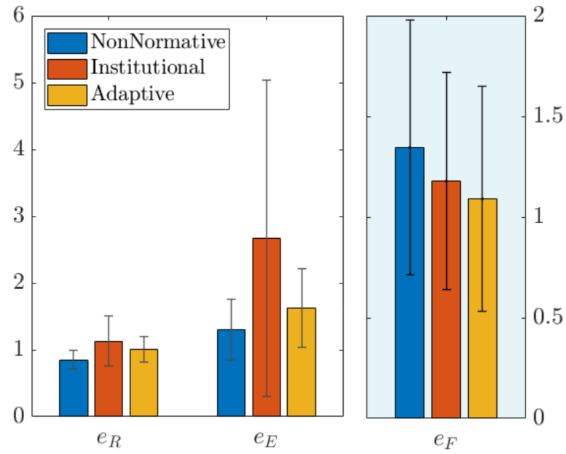

Figure 17.8 – Performance of the robots in scenarios I-III.

17.7 Discussion

The ROBOT-GUIDANCE institution presented in this chapter has been grounded to two different domains – one in the simulated experiments, one in the physical experiments. In the domain representation, the most tangible difference lies in the agents that are assigned the institutional roles of Leaders. In simulations, those agents are human_1 and human_2. In the physical experiments only one human_1 agent is present, but by providing two reference points, it grounds the role of two Leaders. The experimental area that can be formally represented by a state variable $\mathsf{environment} \in R$, on the other hand, is very different for the two sets of experiments.

From an institutional point of view, the only change that has to be applied when using the two domains happens in their respective groundings. In simulation, each of the two humans is assigned a role of a Leader and $(\mathsf{Leader}, \{\mathsf{human}_1, \mathsf{human}_2\}) \in \mathcal{G}_A$, while in the physical experiments the element A of the domain is $A = \{\mathsf{human}_1, \mathsf{mbot}_1, \mathsf{mbot}_2, \mathsf{mbot}_3\}$, where human_1 has the function of a worker, and the role-to-agent grounding includes the element $(\mathsf{Leader}, \{\mathsf{reference\_point}_\mathsf{A}, \mathsf{reference\_point}_\mathsf{B}\}) \in \mathcal{G}_A$, where each of the $\mathsf{reference\_point}_\mathsf{x}$ is provided by human_1. There is absolutely no difference in the definition of the institution (and in this case also very little difference in grounding).

The domain we used in the physical experiments is that of the WAREHOUSE presented in the Case Study of Chapter 15. All basic domain elements, including agents, behaviors and state variables (such as the environment with the obstacles) are the same. In order to introduce the ROBOT-GUIDANCE institution (instead of the SOCIALLY-AMONG-HUMANS institution from Chapter 15), only the grounding is adjusted: the complementary state variables and elements in C, K and L are determined based on the basic state variables; for instance, $\mathsf{path} \in R$ is based on the pose of an agent, $\mathsf{safe\_area} \in K$ based on path of the Leader agents, which is further used to determine $\text{IN\_SAFE\_AREA} \in L$. Note that the values of the above domain elements are not directly sensed or extracted from the domain representation, but calculated at run time. Therefore, effectively no changes to grounding are done except that in the case of SOCIALLY-AMONG-HUMANS the Leader role is grounded directly to human_1, in the case of ROBOT-GUIDANCE, Leader is grounded to $\mathsf{reference\_point}_\mathsf{A}$ and $\mathsf{reference\_point}_\mathsf{B}$. Both the interchangeability of the grounding mentioned above as well the easiness in substituting an institution with another outline once more the plug and play feature of our institutional formalism.

Summary

In the final case study we have shown how institutions can dynamically accommodate to the peculiarities of a domain by regulating behavioral adaptation. Experimental results confirm that the robotic system with the capability to adapt consistently outperforms the non-adaptive institutional approach, although the rules of adaptation are very simple. This promising outcome suggests that the incorporation of more sophisticated machine learning techniques in the institutional framework might lead to outstanding performance in systems characterized by high behavioral complexity, deployed in dynamic, stochastic environments, where interactions of multiple individuals, humans and robots, are influenced by numerous uncontrollable factors.

We have shown how by using institutional norms we can reduce uncertainty of interactions due to norm-induced assumptions regarding the conduct of the other agents. This phenomenon has far-reaching consequences in mixed human-robot teams. For robots, the reliance on humans can simplify the decision-making routines, while the consent to ask for assistance relaxes the actuation and algorithmic requirements. Inversely, when humans know the rules of robot conduct, they are more likely to accept the robots in their environment.

With the third case study we further demonstrated of the institutional modularity and reusability by exploiting the institutional plug-and-play property, where diverse institutions can apply to the same concrete system, one institutional abstraction can be reused across multiple domains, and one norm can be repurposed by several institutions. Furthermore, by applying the norms of the ROBOT-GUIDANCE institution to the flocking and formation behaviors we have demonstrated that the same norms can be leveraged onto two algorithmically different multi-robot routines. With the institutions, norms, domains and behaviors freely exchangeable, the modularity of the institutional framework is similar to that of the building blocks, where elements can be shuffled, matched and linked, as long as the connection is compatible.

This chapter concludes the presentation of our institutional formalism. Nonetheless, throughout our studies we have identified a number of developments towards the completeness of the framework, which we will discuss in the conclusive part of the thesis.

Part IV
Conclusions and Future Work

"Were I to attempt to be good to everyone, to the entire world and to all the creatures living in it,
it would be a drop of fresh water in the salt sea.
In other words, a wasted effort.
Thus, I decided to do specific good; good which would not go to waste.
I'm good to myself and my immediate circle."

Andrzej Sapkowski, Baptism of Fire

18 Conclusions and Outlook

DURING our endeavor to design normative multi-robot behaviors we have identified a key milestone necessary for the integration of robotic systems in human society, namely that of interpretation of generic norms formulated in human language to robot-understandable terminology. With an aspiration of developing a holistic theory for the translation of abstract norms to robotic systems, this thesis explores the notion of institution – a paradigm for reducing uncertainty, simplifying decision-making and promoting cooperation. Inspired by well-versed findings of research ranging from social studies to the early attempts in institutional robotics, we have formulated a framework that assures conceptualization, organization and reusability of social norms, so that they can be applied in different social contexts without resorting to the use of heuristics. Our institutional framework is validated through three diverse case studies, each highlighting one or more desirable properties of the formalism.

To address the practical aspects of the deployment of cooperative multi-robot systems in physical, human-populated environments, this thesis meets the three main challenges we identified during our preliminary experiments with a robot formation deployed in a real hospital environment: a) the complexity of collective navigation in structured indoor spaces, b) the robustness to communication failures and c) the sociability of robot behaviors. To this effect, our research effort is resting on three main pillars: an approach to agile group movement control in complex indoor environments, a multi-robot localization method robust to communication failures, and enhanced navigation methods with elements of HRI. All these methods in diverse combinations are leveraged and validated in the three case studies reported in the thesis.

We hope that the discussion about the need for a holistic theory of robot sociability we engage in this thesis will inaugurate the interest of the robotics community in the generic approaches for unification of different methods and solutions. To this end, we trust that the vision presented in this thesis, the insights gained through our exposure to the research on natural and artificial societies, and the lessons learned during experimental validation will constitute a valuable source of information for future ventures within this field.

18.1 Summary

Our work presents a number of contributions along two lines of research, the first one addressing the challenge of multi-robot deployment in social environments, the second concerning the institutional framework.

With regard to the multi-robot deployment in social environments, our contributions detailed in Part II can be summarized as follows:

- **Adaptive Multi-Robot Navigation.** We developed a Local Formation Transformation (LFT) approach for realizing adaptive multi-robot formations in structured environments that yields local and gradual change of formation shape with the level of alteration proportional to the density of obstacles ahead. The results confirmed that the LFT enables the formation to navigate as a unit through demanding environments with complex building features and uncertainties arising from the presence of static and dynamic obstacles as well as sensor and actuator errors. Motivated by deployments in human-populated environments, the LFT algorithm achieves desirable properties of smoothness of motion and aesthetic negotiation of obstacles.
- **Cooperative Localization.** We presented a strategy for providing reliable robot state estimates to be used for formation control when communication is sporadic or suffers from short-term outage. Our method called Formation Information GM-PHD (FI-GM-PHD) filter combines absolute positions exchanged by the robots, information about the formation geometry and sensory detections in an extension of the GM-PHD filter. The experimental results demonstrated that our approach is capable of maintaining the state estimates even when long-duration occlusions occur, and it allows for sustaining formations in cluttered environments with high measurement uncertainty and low quality of communication. The proposed method not only outperforms canonical multi-target algorithms stressed under the same communication conditions, but also delivers performances competitive with methods relying on perfect communication. Finally, in experiments with elevated dynamics caused by a norm-driven reaction to the presence of humans, introducing our method significantly reduces the chance of formation breaking and in turn the probability of mission failure.
- **Social Awareness.** Our contribution to the field of human-aware navigation is three-fold. First, we adopted well-established single-robot methods in the context of multi-robot systems and through the experiments presented in Part III we analyzed the impact of the presence of multiple robots on their acceptance by humans. Second, in contrast to the state-of-the-art research on multi-robot teams deployed in human-populated environments that is limited to individual, uncooperative robots or cooperative solutions that fail to consider realistic situations, we deploy cooperative robot teams that are socially

aware in scenarios with real human participants in authentic indoor spaces. Finally, by virtue of representing social norms in abstract and general form, easily understandable for humans, we have taken the first steps of a long-term process aimed at unifying the existing approaches to norm-following robot behaviors.

- We accomplished a unique set of multi-robot experiments in a real hospital setting. The analysis of the results provided valuable insight into the benefits and disadvantages of multi-robot deployments in highly sensitive environments. The lessons learned at the hospital served as the main motivation for this thesis, and we strongly believe that they can convince the robotics community about the need for integration of social norms into robot behaviors.

With regard to the institutional framework our contributions detailed in Part III can be summarized as follows:

- **Institutional Formalism.** We developed a model-based approach for abstraction, encapsulation, and formalization of generic social norms into reusable structures, called institutions. Our institutional framework brought together insights gained from the research in economics, multi-agent systems, the original principles of institutional robotics, and the approaches derived thereof. Through our case studies we have shown that our abstract representation of institutions allows for the governance over miscellaneous robot behaviors and integration of social constraints of diverse nature. The properties of institutions allow us to seamlessly reuse the same institution across the domains with a variety of agents capable of performing different behaviors.
- **Norm Realization.** By bridging the research on robot planning, where the focus lies on clear semantics and abstraction, and the normative navigation, which addresses the question of "*how?*", we have identified the key elements necessary for interpretation of social norms and putting them into practice. Our main contribution, norm realization, is a mechanism for translation of generic, language-defined norms in terms of robot-understandable language, making such norms readily implementable onto concrete restrictions of robot behaviors and executable in real physical systems. To this effect, we have introduced low-level sophistication to embrace the complexity of continuous multi-robot behaviors, at the same time retaining the desirable high-level properties, including abstraction, encapsulation, and modularity. The power of norm realization has been showcased through a number of norms related to a large variety of social aspects, ranging from human comfort achieved through robot navigational compliance, to understandability of robot intentions reflected through gestures, expressions, and sounds.
- **Norm Abstraction.** The proposed institutional formalism provides a simple tool for encoding behavior specifications in a plug-and-play way instead of programming hard-wired social compliance in ad-hoc behaviors, as it is done currently. To this end, we

pioneer a number of worthwhile properties lacking in the state-of-the-art normative robotics approaches, namely high reusability of institutions and norms, which we freely allot over different domains and behaviors without the need of redesigning, and modularity and scalability, where the encapsulation of norms allows to decouple their operation from the behavior design and where introducing a new norm does not require heuristics on how to merge it with the current solution. Instead of the design being driven by norms, norms are imposed as constraints operating over the parametrization of already existing behaviors.

18.2 Discussion and Outlook

Our work constitutes a first step towards integration of social norms in multi-robot systems. Although we have addressed some aspects towards this end, several remain to be taken into consideration.

To begin with, we have identified a number of practical challenges of cooperative navigation in indoor, human-populated spaces, but the selection is by no means comprehensive and there are ample opportunities to improve our solutions, which we will discuss next. Second, in spite of our best effort towards achieving completeness of the institutional formalism, we have identified a number of key research thrusts that can be undertaken for leveraging the current approach.

The lessons learned during our multi-robot experiments in Part II and in the case studies in Part III lead us to the following conclusions:

- **Suitability of Norms.** Bearing in mind that although the development of adequate social norms and methods is imperative for achieving robot acceptance in human societies, it was not our primary objective. Instead, we attempted to reuse existing single-robot approaches and extend them to the multi-robot context, with varied levels of success. This is partially a result of the combination of highly dynamic scenarios with multiple heterogeneous agents – humans and robots, and so, environments characterized by a low level of controllability. From our perspective, however, even if a norm is designed with the best intentions and has proven to be adequate in one social context, it may not necessarily be easily ported to another application. Norm validation necessitated extensive participative studies tailored for the particular context and performed at various stages of the effort. Furthermore, as we have emphasized throughout this manuscript, there is a clear need of an open source database where results of such studies can be collected and shared among the researchers.
- **Challenges of Real Environments.** Soaring dynamics, high stochasticity, noisiness and uncontrollability of the real social environments make deployments of multi-robot systems profoundly challenging. Although we have addressed a number of such challenges, there are still several improvements that can be made. Firstly, through our evaluation

in Chapter 15 we recognized the need for augmentation of the purely reactive formation control methods with deliberative components. In all our case studies, we have considered only established formations and flocks. However, it should be still investigated how group initiation, leaving and joining a team, and other behaviors necessary for dealing with re-grouping and emergency handling would fit into our framework. Although the methods we described performed well in most of the social interactions, the case studies involved also situations where it was simply impossible for the robots to meet the social requirements (e.g. when two humans approached a robot from two sides). Incorporating a prediction for the human motion into the graph and propagating such information in order to accommodate formation structure correspondingly could represent a promising next step. Secondly, throughout our experiments we relied on external systems (e.g., a MCS) to obtain estimates of human positions. However, in real environments, the deployment of similar systems can prove difficult or impossible, and the positioning information would have to be retrieved from robots' onboard sensing. To this end, effort towards dealing with robot perception, its limitations and inherent uncertainty must be made.

The key research thrusts that can be undertaken for leveraging the current institutional formalism and broaden its perspective can be summarized as follows:

- **Completeness of the Framework.** Although our institutional framework accommodates all the necessary steps of grounding and norm realization, we left part of the procedures in the hands of the designer. First, the identification of the elements of the sentence that correspond to the institutional components is trivial, so we foresee that it would be possible to perform this step with a language-based compiler. More prominently though, the step of formulating the rules based on the above opens the door for many possibilities, and allows for ad hoc methods to be developed. As correctly diagnosed in [147], this could possibly hamper the progress of the institutional approach by not providing the means for combination of distinct works. Unfortunately, the formulation of the rules is nontrivial. To the best of our judgment, the development of a fully automatic fit-for-all tool would be close to impossible, as the rules rely on the behaviors they operate upon, and not on the norms that dictate them. Nonetheless, this thesis offers comprehensive explanations and plenty of examples that would support the next step, namely that of open sourcing a compiler, which based on provided behavior specifications would guide the process of building the rules. Because of modularity of our methods, norm realization of typical robot behaviors could be open sourced in a form of building blocks, reusable across applications.
- **Outlook on the Institutional Formalism.** There remains a number of open problems in the institutional formalization that are yet to be addressed. First, we only provided a preliminary solution to the question of social context recognition and activation of the institution operating upon it. Second, a rigorous method for defining relations

within a network of institutions operating concurrently, their hierarchy and resolution of conflicts is yet to be developed. Third, to relax our assumption of admissible grounding an automatic process akin to that of institution recognition is yet to be established. This would provide declarations of grounding relations that in our examples we have assumed to be given.

- **Outlook on Norm Realization.** There are two further aspects related to social norms and norm realization yet to be tackled. The first refers to detection of norms existing in human societies, their automatic application to the robot behaviors and norm cognition, i.e. how to generate new norms based on robot's experience. From our perspective, however, these problems lay at an intersection of sociology, neurology and machine learning techniques, as proposing a solution requires a good understanding of how individuals interpret norms, how a society agrees on common norms, and how these solutions can be adopted by an artificial system. For this reason, it is out of our expertise to suggest how they should be approached. The second refers to introducing priorities between norms, necessary for conflict resolution as well as addressing ethical concerns. For example, in Case Study I, one can imagine that a robot should prioritize respecting human personal space (norm n_2) before complying to the norm preventing to enter spaces where humans might perform an activity (norm n_1). Therefore introducing priorities is a topic for further study.

- **Limitations.** Throughout our work on the institutions we always assumed full robot immersion in an institutional environment, i.e. that once a robot belongs to an institution, it has access to all the information and services it provides. We did not address the issue of incomplete institutions, i.e. institutions with limited knowledge whose components only partially comprise the elements of a social context. Although we proposed adaptation of norms at the domain level, in our approach institutions are static structures that do not change over time. Institutional adaptation is still to be addressed. Our formalism has been designed with the objective of scaling well with the complexity of the norms and behaviors they apply to. However, scalability with respect to the number of robots has not been validated as of now. To our best judgment, the institutional framework is not purely restricted to only small teams, or only heterogeneous teams of robots that apply the same norms, which we have proven on our case studies. However, the design of institutions and norms for large swarms or for multiple diverse teams operating in the same environment must be adapted to achieve the desired end results, similarly as we have demonstrated in the Case Study II where the institution embraces the complementary capabilities of the two robots, MBot and Pepper.

As a final conclusion of this thesis we would like to bring attention to the fact that an interdisciplinary discussion among researchers representing the fields of social sciences, economics, multi-agent systems and robotics is imperative for drawing full benefits from frameworks such as the one presented in this manuscript. For once, we have shown that drawing inspiration from other fields of study allows us, roboticists, to build our work upon well-established solutions. On the other hand, formal methods developed for concrete, physical and distributed robotic systems can provide invaluable insights on representation, modeling and control of socio-economic phenomena.

A Case Studies: Implementation Details

This chapter provides additional implementation details for the Case Studies $\mathcal{C}_I$-$\mathcal{C}_{III}$, presented in Chapter 15, Chapter 16, and Chapter 17.

A.1 Case Study I

In all the experiments of $\mathcal{C}_I$, the human participants played a game involving a number of waypoints to visit, delineated on a map. The waypoints were labeled with animal shapes and indicated as crosses on the floor. The basic experimental setup is illustrated in Figure A.1. Such design of the game allows to maximize the number of interactions between the robots and the humans, while retaining systematicity between the runs.

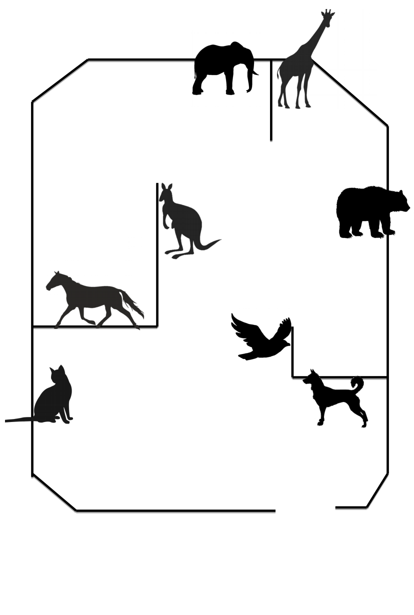

Figure A.1 – Arena set-up of the experimental campaign for Case Study I. Each person was asked to follow a path driven by animal-labeled waypoints. For instance, for human H_A the path was 1) Dog, 2) Kangaroo, 3) Bear, 4) Kangaroo, 5) Bird, 6) Dog, 7) Cat and 8) Bear, and for human H_B the path was 1) Horse, 2) Elephant, 3) Kangaroo, 4) Elephant, 5) Bear, 6) Kangaroo, 7) Giraffe, 8) Horse and 9) Bird. The roles and the paths were retained for all the runs. Each person was asked to move with their natural pace and keep a similar speed during each trial.

The experiments have been parametrized as follows. The virtual size of $\mathsf{personal\_distance}$ $\Delta_{P,h}$ and relative $\mathsf{speed\_factor}_\mathsf{h}$ for each human H_h are presented in Table 15.6. And so, the robots are implored to keep a larger distance to a child than to any other human (as driven by norm n_2), while also lowering the speed the most (as driven by norm n_3). The *TrajectoryShape* parameter of the Leader in n_2 is modified by adding a vector perpendicular to original path, directed away from the human, with the size of the vector being a sigmoid function of the Leader-to-human distance and the size of personal_distance $\Delta_{P,h}$. Also in n_2, the parameter Δ_a of *RepulsionWeights* of the Follower robot is proportional to personal_distance $\Delta_{P,h}$, with additional margin to prevent the robot entering the personal space. The speed_factor$_x$ coefficient plays two roles. First, for a given domain (where x stand for SCHOOL, WAREHOUSE or HOSPITAL), a *trajectory* speed_factor allows the Leader to adjust the *TrajectorySpeed* relatively to the default speed. Second, for the Followers, a *near-agent* speed_factor adjusts the *ControlGain* as in $K_C \leftarrow K_C * s_{A,h}$, where H_h is the human during the avoidance of whom the speed is being lowered. This leads to proportional reduction of the actuation control input, and consequently, lowers the robot speed. All discussed values are listed in Table 15.6.

The *FormationShape* is a triangle, while the *FormationBias* is scaled according to the domain – it is smallest for the WAREHOUSE domain, where efficiency is important, and largest for SCHOOL, where it might be desirable that the inter-robot spaces are large enough to allow a running child to pass through. Values of *FormationBias* are given in Table 15.6. Lastly, in a subset of experiments carried out in the HOSPITAL domain, we simulate QuietActivityTime to take place during time_interval $T_{QAT} = [80,\ 100]s$. We also tested the robot response to a static affordance area Φ_{ACS} at a location [0,0] and with radius of 1.3 m. These experiments are tested separately so as to evaluate norms n_1 and n_3 in an isolated context.

A.2 Case Study II

The institutional knowledge of the STEERING and the TUTORING institution is grounded as follows. The element PersonalSpace takes the values personal_distance$_{\text{child}}$=2.0 m and personal_distance$_{\text{adult}}$=1.6 m, and the element IntimateSpace the values intimate_distance$_{\text{child}}$=1.2 m and intimate_distance$_{\text{adult}}$=1.0 m. Values for K_o, Δ_a and Δ_c were obtained empirically in simulations in [9], where the human reaction was modeled with the SFM (for more details refer to Section 5.6.2). Robot behaviors are distributed, while their activation is regulated by a central planner.

The Destination of the STEERING institution is grounded to the entrance of the laboratory, which in turn is labeled as the ExhibitionArea of the TUTORING institution. The objects are chosen from the set {VR_SET, ROBOT_STATION, CAMERA}, their positions are provided to the robots in the grounding of ObjectLocations, while a short, one sentence description is supplied a priori via ObjectDescriptions.

We chose $e_F^{MAX} = 0.65$ m, equal to the diameter the MBot robot. The threshold determining whether VISITORS_FOLLOW is $\Delta_L = 3.2$ m. For safety reasons, we limit the maximum velocity of the Followers to $\dot{p}_{max} = 1$ m/s.

A.3 Case Study III

The experiments of the case study $\mathcal{C}_{III}$ are parametrized as follows. The values of the thresholds used in norm realization are: $T_L = 5.0$ m for determining if a follower is lost, and $T_C = 2.5$ m for MoveByFlocking and $T_C = 2.25$ m for MoveInFormation for triggering the condition IN_NARROW_PASSAGE. The value of the personal_distance is $\Delta_P = 1.2$ m for all humans.

The flocking coefficients are $K_{\alpha 1} = 2.0$, $K_{\alpha 2} = 1.0$, $K_{\beta 1} = 150$ and $K_{\beta 2} = 1.5$ for obstacles and $K_{\beta 1} = 200$ and $K_{\beta 2} = 2.0$ for β-agents that represent humans, $K_{\gamma 1} = 5.0$ and $K_{\gamma 2} = 2.0$. The interaction ranges for agents and for obstacles are $\Delta_A = 1.8$ m and $\Delta_O = 0.36$ m respectively. The desired agent-agent or agent-obstacle distance of flocking are $b_A = 1.2$ m and $b_O = 0.24$ m. For the β-agents that represent humans the values are $\Delta_{O^H} = 0.72$ m and $b_{O^H} = 0.48$ m. The flocking gain is $K_c = 0.15$.

Formation control is parametrized with the gain $K_u = 1.0$, bias that forms an “X”-shaped mesh of edge size 1.5 m with two human Leaders in front of the three followers behind. The values for calculating the *RepulsionWeights* $\mathcal{W}_{\mathcal{R}}$ are $K_o = 1.0$, $\Delta_a = 1.0$ m and $\Delta_c = 0.55$ m. The default Laplacian weights are $w_{ff} = w_{fl} = 1.0$.

B Case Study I: Comments of the Participants

This appendix provides a detailed summary of the additional comments provided by the participants in Case Study I, Experiment set II of Chapter 15. The comments, although not codified, provide further insight into the perception of robot behaviors by the participants.

B.1 The NON-NORMATIVE Trials

After the NON-NORMATIVE trial, some participants commented on the fact that their presence was ignored, for example: "We didn't interact much in this experiment so mostly neutral feeling"; "I feel like robots are ignoring me and they were too close during navigation"; "In this experiment, I've felt they were not eager to move out of my way". In two cases robot aggression was compared to the previous trial (in both cases it was preceded by the EFFICIENT trial): "In this experiment robots looked a bit less aggressive than in experiment 1 [EFFICIENT]"; "The robots seem a bit more aggressive [than in the EFFICIENT trial] in a sense that I don't really understand if they really care about human as for the 1st experiment". One person remarked on robots friendliness: "They are less friendly [than in experiment 1: SOCIAL trial]". Finally, two participants used decidedly negative wording to express their feelings: "A little bit scary"; "Creepier experiment". In contrast, a number of participants indicated their preference for the NON-NORMATIVE behavior, despite it being the non-social one, e.g. "They 3rd experiment [NON-NORMATIVE], while being the least social, was the most natural one for robots (...), therefore I could predict their behavior easier"; "It is difficult to say what was the difference between the 3 experiments, I just liked more number 2 [NON-NORMATIVE]".

B.2 The SOCIAL Trials

The SOCIAL trial received a number of general positive comments, such as: "The robots are cute [drawing of a smile]. I liked when the robots had their *hands* and *arms* moving" [emphasis in original]; "Robots during experiment 3 [SOCIAL] have according to me the best behavior & interaction with the human. In terms of comfortable feeling with descending order: 3>1>2

(SOCIAL > EFFICIENT > NON-NORMATIVE)"; "Polite"; "I felt like they are too slow to work in a real environment with humans, but it is interesting to work with them". Their movement was preferred against the preceding trials, i.e. "Far more receptive to my movement [than in the other trials]"; "It is better when they avoid us [compared to the previous NON-NORMATIVE trial]". Other positive comments regarded the spoken apologies and facial expressions, for example: "The robots are very pleasant now, they have kind of *feelings*, excuse themselves if they disturb us" [emphasis in original]; "The fact that they say 'excuse me' makes them a bit more friendly and pleasant"; "It is pleasant to see a little bit more a human face on them (mouth changing, language)"; "Nice how they excused themselves and changed facial expression and color on the bottom". Interestingly, the same characteristics were perceived negatively by a number of participants: "Robot apologizing all the time may become annoying in the long term"; "The volume of the robots' voice is a bit too loud. (...) The red eyes look like an *evil robot*" [emphasis in original]; "Sounds are a bit annoying"; "I didn't really understand the sounds they are making, also they are quite dumsy". Finally, some participants pointed out some of the asocial aspects of the robot behavior, e.g. "A robot apologized, after which it collided with my foot"; "It seems they try *too hard* to avoid each other and the humans, feels a bit too chaotic. (...) I liked a lot the smiley [drawing of a smile] face!"

B.3 The EFFICIENT Trials

The EFFICIENT trials were generally provided negative comments, such as: "They are not likely to give priority to people to walk. They give me selfish feeling"; "They were slightly more aggressive [than in the social trial]"; "They didn't care about people around. (...) They also seemed to disturb [me] intentionally"; "Difficult to understand their roles". Furthermore, the comments suggested that their motion was judged as asocial and illegible, for example: "During a close-by meeting, the robots spread apart and we did not collide. However, I had to stop. Then they gave me way"; "The robots collided into each other when trying to avoid the human"; "It was hard to anticipate their movement. They were paying less attention. I had to bypass one because it was in my way"; "He walked on my foot! They are a bit too *sticky*" [emphasis in original].

B.4 The Overall View

Overall, a number of participants mentioned that they felt like they were biased among the runs, as they got used to the robot presence, for example: "I was more comfortable with the robots in the second experiment because I got used to walk with robots"; "My feeling towards the first experiment might be biased because I discovered the robots, whereas I was able to be critical on the next runs". Some participants focused on the robot characteristics rather than on their behaviors. The comments referred to the robot appearance, e.g. "I felt that the robots were too big for the arena", "I felt like I was being watched", or to their motion characteristics, e.g., "A bit too vibrating to be considered *natural*" [emphasis in original].

C Evaluation of the FI-GM-PHD Filter in Human-Populated Environments

This chapter provides a detailed account of the third set of experiments referred to as *With Cooperative Localization* introduced in Chapter 15.

C.1 Scenarios

In the first two sets of experiments formation control is carried out with sufficient quality of communications. In the experiments of this appendix, we simulate temporary communications failure and analyze its impact on the performance of the robot team. We distinguish two cases: A) where robots rely on communications only, and B) where each robot runs the FI-GM-PHD filter described in Chapter 8 to compensate faulty communication data with sensory detections and information about the desired formation geometry. We label these two cases as NOTRACKING and WITHTRACKING, respectively.

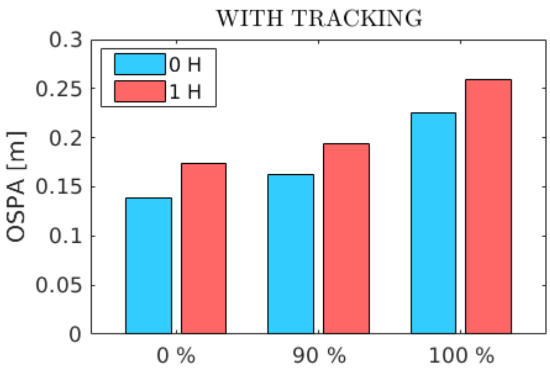

Figure C.1 – Summary of the OSPA errors for the scenarios with 0% message drop (i.e. ideal baseline communication), 90% drop and 100% drop, each evaluated in experiment with one human (1H) and no humans (0H).

Communication failure is simulated at 10 s intervals occurring every 40 s, adding up to 20% of the run time. Two types of communication losses are tested, 90% loss, where robots drop 9 out of 10 messages at random, and 100% loss, where there is no communication. For a baseline, we perform experiments with no losses (0% loss). Basic communication rate is 10 Hz. Communication failures are simulated for all robots simultaneously. All robots run the tracking system independently and onboard. All experiments are run with the SCHOOL domain

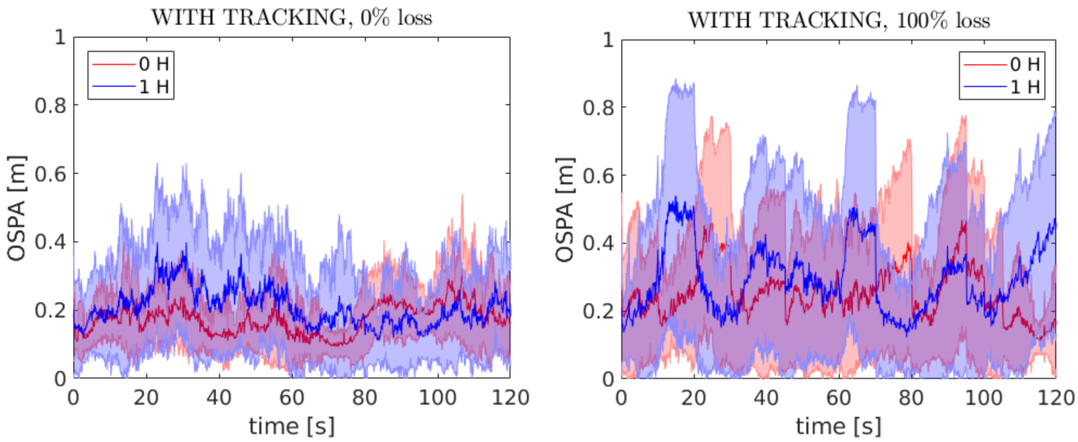

Figure C.2 – Time-wise plots of the OSPA metric of the the WITHTRACKING experiments and 0% and 100% message drop, averaged over the runs.

		0H			1H		
		0%	90%	100%	0%	90%	100%
NOTRACKING	*mean* e_F	0.22	0.25	0.23	0.38	0.34	0.3
	std e_F	0.07	0.12	0.12	0.26	0.24	0.22
WITHTRACKING	*mean* e_F	0.18	0.19	0.3	0.33	0.33	0.48
	std e_F	0.05	0.09	0.18	0.23	0.22	0.4

Table C.1 – Average formation error e_F.

and one human (labeled 1H) child, as that choice yields maximal social disturbance to the robot team. These are compared to experiments with no humans (labelled 0H). For each scenario, we perform 12 consecutive runs.

C.2 Results

The performance of the FI-GM-PHD tracking system in the WITHTRACKING experiments has been analyzed using the OSPA metric (described in Section 8.5.2), calculated between the ground truth poses obtained from a motion capture system and the estimates of the robots. The mean OSPA errors shown in Figure C.1 demonstrate that the tracking performance gracefully deteriorates with the degradation of the communication quality; a drop of 90% of messages increases the error on average by 14%, while a 100% message drop results in 56% increase. As shown in Figure C.2, the overall error mostly remains within around 0.4 m for 0% loss (left) and within around 0.6 m for 100% loss (right). Human presence has a clear impact on the tracking performance, with OSPA being on average 5.2% higher when a human is present and with around 40% larger standard deviation.

There is no significant difference in the formation error e_F between the NOTRACKING and

WITHTRACKING scenarios. Results summarized in Table C.1 show that tracking slightly improves e_F in the 0% and 90% cases for both 0H and 1H, bot not in the 100% case. Our intuition is that in the 100% case, where in the NOTRACKING scenario robots simply stop until communication resumes, that strategy leads to lower formation error than when robots compensate with other information sources that integrate also the sensing errors. On average, human presence leads to around 20% rise of e_F for both NOTRACKING and WITHTRACKING. The average error is bounded for all scenarios, with the highest being 0.48 m for 100% and 1H, which given the robot diameter of approximately 0.65 m can be considered acceptable. However, one should note that since it is desirable and driven by design that the human presence distorts the formation shape, e_F should not be considered an indicator of the formation performance, but a sign of the degree of compliance that the formation exhibits when navigating among humans.

As an indicator of the formation performance we use failure rates, calculated based on the video annotations. We distinguish three categories of formation failures: NOT-REC, where robots are not able to recover the formation, NOK-REC, where the formation breaks substantially, but is able to recover, and OK-REC, where the formation separation is visible, but insignificant. We consider situations where a failure is a result of insufficient data or a combination of missing data and other factors, such as the complexity of the environment or the presence of a human. Examples of the three categories are shown in Table C.2, where the snapshots from the videos recorded during the experiments show the formation before it breaks (left), at the time when formation of the most distorted (middle) and after recovery in the NOK-REC and OK-REC cases or in the case of no possibility to recover in the NOT-REC category (right).

Figure C.3 shows failure rates per second ($\times 10^{-3}$), where the higher the value, the worse the formation performance in the respective category. Data confirms that the use of FI-GM-PHD filter in the WITHTRACKING scenarios yields much better performance than when tracking is not used in the NOTRACKING case. As indicated by NOT-REC, no critical failures occur in the WITHTRACKING experiments and the rates of failures with recovery (NOK-REC and OK-REC) are on average much smaller than in the NOTRACKING case. Irrespective of whether the tracing system is used, presence of human increases the likelihood that the formation breaks in all scenarios.

C.3 Discussion

In Chapter 8, where we introduce the FI-GM-PHD filter, the evaluation is carried out in static environments with no humans. Although we challenge the tracking system in scenarios cluttered with obstacles in the arena and expose the formation to long-term communications outage, there are no external factors that affect the performance of the formation and of the filter.

With the presented experiments we validate the FI-GM-PHD filter in highly dynamic environ-

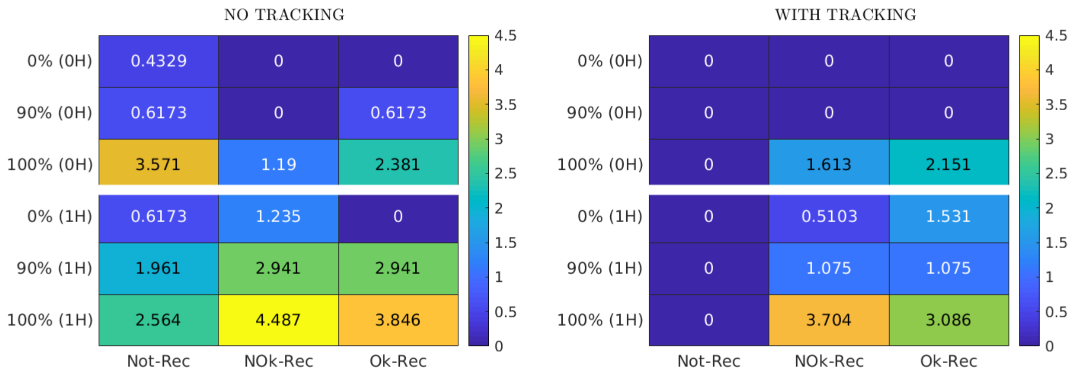

Figure C.3 – Formation failure rates per second (multiplied by $\times 10^{-3}$), classified into three categories, NOT-REC, NOK-REC and OK-REC, calculated based on the video annotations. The higher the value, the worse the formation performance in the respective category. The categories are: NOT-REC, where robots are not able to recover the formation, NOK-REC, where the formation breaks substantially, but is able to recover, and OK-REC, where the formation separation is visible, but inconsequential.

ments where the formation is frequently perturbed interrupted by the presence of human, which on top of causing sparse occlusions affects the dynamics of the norm-following formation. No changes in the filter implementation or models are made, neither did we perform parameter tuning. This confirms the robustness of the method to environmental changes. The results indicate that the performance of tracking degrades gracefully with the drop of communications quality. Although the presence of a human significantly affects the ability of the robots to track well, the performance remains stable. Finally, the formation failure occurs far less frequently when the tracking is employed, and with lesser criticality – we recorded no cases during which the formation broke and did not recover.

	Before breaking	Most distorted	After recovery / no recovery
Ok-Rec		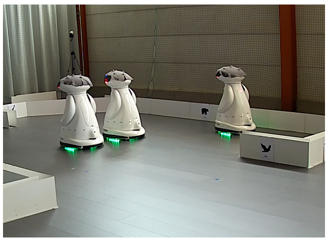	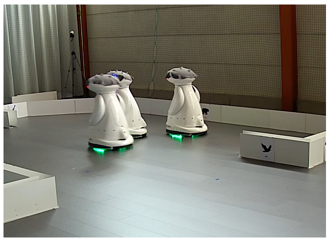
NOk-Rec			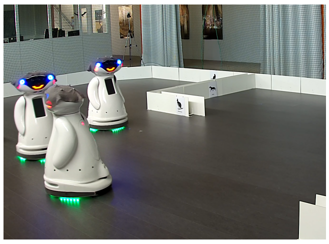
Not-Rec			

Table C.2 – Snapshots from the videos recorded during the experiments showing the examples of formation failures.

D Computational Procedure of Norm Realization: An Example

In this chapter, we illustrate the steps of the computational procedure described Chapter 16 with a simple example of the norm n_{S2} of Case Study $\mathcal{C}_{II}$. The realization of a similar norm has been already discussed in Chapter 13, but here we focus on the implementation details.

D.1 Overview of the Computational Procedure

The procedures for the specification and application of behavioral modality in our computational procedure presented in Chapter 16 take the following form:

B) A subset of relevant parameters $P_i = \{p_j\}$ of behavior B_k is selected: $P_i = r^P(\mathcal{N}_t^A, B_k) = \{p_j\}$.

C) Based on R and K, each parameter p_j takes value $v_{p_j} = r_j^V(p_j, R, K)$:

C') Value v_{p_j} is adapted based on the experience measure L:

$$v_{p_j} = r_j^L(r_j^V(p_i, R, K), L) = f|_{v=f_j(R,K)}(v, L)$$

D) Behavior B_k applies the value v_{p_j} of parameter p_j: $\lambda_k \leftarrow r^B(p_j, v_{p_j}) = B_k(p_j, v_{p_j})$

The labels B) - D) correspond to the steps of norm realization introduced in Chapter 13.

D.2 Computational Procedure: A Step-By-Step Implementation

Recall that completing the institutional level of norm realization resulted in the following partition of the n_{S2} sentence *"The followers should maintain a comfortable distance from the humans"*:

$$n_{S2}: \emptyset \rightarrow \mathsf{should}(\textit{followers, maintain, comfortable distance from humans})$$

Norm n_{S2} concerns robots with the roles of the Followers, which, through the relations described in the previous section, carry out the MoveInFormation behavior.

Although the sentence does not specify the conditions for activation of norm n_{S2}, from n_{S4} we deduce that a comfortable distance is only maintained when humans are following, otherwise the norm n_{S4} encourages the Follower robot to "push" the human to join back the formation. It is the role of the designer to interpret the missing elements of the sentence as we do in this case. Such condition denoted VISITORS_FOLLOW is the only activation condition of n_{S2} and grounded to $c^a_{S1} \in C_{S1}$ (see Table 16.2) and the outcome of its evaluation by the requirement rule r^N determines whether n_{S2} is active at time t:

$$\mathcal{N}^A_t \ = \ r^N(C^a) \ = \ \begin{cases} n_{S2} \in \mathcal{N}^A_t & \text{if } \ c^a_{S1} \equiv \mathit{True} \\ n_{S2} \notin \mathcal{N}^A_t & \text{otherwise} \end{cases}$$

The term *if* $c^a_{S1} \equiv \mathit{True}$ can be replaced by any propositional formula. We denote the generic equivalence of the formula for norm n_k with $\Upsilon(C^a_k)$. In our perspective every condition can be represented in terms of prepositional logic, even conditions defined on continuous variables. For instance, condition WHEN_WARM can be mapped to the state variable temperature being between two discrete values.

Once a norm is known to be applicable at the given time, rules r^P and r^V determine the socially appropriate behavior modality, and r^B applies them. In our example, to respect n_{S2}, a robot dynamically creates a repulsive field around the human to prevent interference with his comfort zone. As explained in detail in Section 13.1.5, according to the rule r^P one parameter – the *RepulsionWeights* $\mathcal{W}_\mathcal{R}$ of the MoveInFormation behavior is modified:

$$P_{S2} \ = \ r^P\big(\mathcal{N}^A_t, \mathsf{MoveInFormation}\big) \ = \ \{\mathcal{W}_\mathcal{R}\}$$

The value of $\mathcal{W}_\mathcal{R}$ is a function of three parameters, K_o, Δ_a and Δ_c, and the distance d_{ih} between the robot R_i and the human H_h. The parameters for a specific human H_h are extracted from the grounded knowledge of the particular domain. For example, in this case study we distinguish between the child Visitors and the adult Visitors as follows:

KNOWLEDGE	$\xrightarrow{\mathcal{G}_K}$	GROUNDED K	$\longrightarrow$	PARAMETER VALUE
SocialForce	$\xrightarrow{\mathcal{G}^{(S)}_K}$	$\mathsf{force}_{\mathsf{child}}$	$\longrightarrow$	$K_o = 1.2,\ \Delta_a = 2.2,\ \Delta_c = 0.65$
		$\mathsf{force}_{\mathsf{adult}}$	$\longrightarrow$	$K_o = 1.0,\ \Delta_a = 1.6,\ \Delta_c = 0.55$

Based on the above elements of knowledge, K_o, Δ_a, Δ_c, and the distance $d_{ih} \in R$, the value for the parameter $p_{S2} = \mathcal{W}_\mathcal{R}$ is chosen as follows:

$$v_{p_{S2}} \ = \ r^V\big(P_{S2},\ K,\ R\big) \ = \ f\big(K_o,\ \Delta_a,\ \Delta_c,\ d_{ih}\big)$$

with the function f given in Equation (5.9). The resulting value of $\mathcal{W}_{\mathcal{R}}$ is applied by the MoveInFormation behavior according to Equation (5.8), which when all other active norms have been applied, results in the behavior modality λ

$$\lambda \leftarrow r^B\left(P_{S2}, V_{P_{S2}}\right) = \mathsf{MoveInFormation}(p_{S2}, v_{p_{S2}})$$

The norm is satisfied when c^a_{S1} = RESPECTED_PERSONAL_SPACE is true, i.e. the resulting human-robot distance d_{ih} at time $t+1$ is higher than the value of the personal distance Δ_P, for all human Visitors:

$$\mathcal{N}^T_{t+1} = r^O\left(\mathcal{N}^A_t, C^o\right) = \begin{cases} n_{S2} \in \mathcal{N}^T_{t+1} & \text{if } c^a_{S1} \equiv \mathit{True} \\ n_{S2} \notin \mathcal{N}^T_{t+1} & \text{otherwise} \end{cases}$$

In our implementation Δ_a is equal to the value of the personal space Δ_P, the other values are determined empirically based on the model of human social forces discussed in Chapter 5.

E Norm Realization in Case Study III

This chapter provides details of norm realization of the Case Study $\mathcal{C}_{III}$ presented in Part III, Chapter 17.

E.1 Activation and Outcome Conditions

The requirement rules r^N, and so the activation and outcome conditions are independent from the behaviors agents engage in. The conditions, listed in Table 17.2 are determined by the Follower F_i as follows.

Activation Conditions

r_1^N c_{11}^a = LEADER_LOCATIONS_KNOWN
Satisfied when the state variables pose of the agents in $\mathcal{G}_A$(Leader), i.e. $\text{pose}_{\text{human}_1}$ and $\text{pose}_{\text{human}_2}$ are known. In this case study it is assumed that this is true at all times, i.e. $c_{11}^a \in \mathcal{N}_t^A \; \forall \; t$.

r_2^N c_{21}^a = LEADER_LOCATIONS_ALWAYS_KNOWN
Satisfied if $c_{11}^a \in \mathcal{N}_t^A, \forall \tau \in (t_0, t)$.
c_{22}^a = IN_COMPLEX_SPACES
We assume that $c_{22} \in \mathcal{N}_t^A \; \forall \; t$, because the environment consists of multiple narrow turns with few open, straight passages.

r_3^N c_{31}^a = FOLLOWER_LOST
The descriptor LOST is grounded to a distance threshold T_L between F_i and the Leaders, falling beyond which deems F_i to be lost.

r_4^N c_{41}^a = IN_NARROW_PASSAGE
Satisfied when the passage that F_i traverses is narrower than a threshold T_C.

r_5^N c_{51}^a = PERSON_LOCATIONS_KNOWN

Since the descriptor person is not a role, but rather an attribute, realization of c^a_{51} does not use the grounding $\mathcal{G}_A$ directly (as it is not expressed in terms on Leaders and Followers), but it addresses all the agents that are humans. In this case study person = {human$_1$, human$_2$}. Because of that, $c^a_{51} \equiv c^a_{11}$, and so $c^a_{51} \in \mathcal{N}^A_t\ \forall\ t$.

Conditions c^a_{11}, c^a_{21} and c^a_{51} ensure the information necessary for the robots to be able to follow the norms. These norms are always active when such information is available. In contrast, c^a_{22}, c^a_{31} and c^a_{41} are active only occasionally upon occurrence of an event. We can categorize the first type of conditions as the *capability conditions*, and the second type as the *triggering conditions*.

Outcome Conditions

The outcome conditions typically evaluate a number of relevant performance measures. Note that the state variable position of the Follower F_i is denoted position$_{\mathcal{G}(\text{Follower}_i)} = p_i$.

r^O_1 c^o_{11} = FOLLOWER_IN_SAFE_AREA
Evaluates whether $p_i \in \Omega$.

r^O_2, r^O_4 $c^o_{21} = c^o_{41}$ = FOLLOWER_FOLLOWS_LEADERS_PATHS
Satisfied by F_i when $\|p_i(t) - p_{\text{VL}}(t)\| < T_L$

r^O_3 c^o_{21} = LEADERS_WAIT
Satisfied when $\dot{p}_h(t) = 0,\ \ \forall h \in \mathcal{G}_A(\text{Leader})$.

r^O_5 c^o_{51} = SOCIAL_SPACES_RESPECTED
Gratified when $\|p_i - p_h\| > \Delta_{P,h},\ \ h \in \{\text{human}_1, \text{human}_2\}$.

E.2 Parameters and Values

The choice rules r^P determine which parameters of the given behavior are changed by the norm, while the value rules r^V set these parameters to concrete values. Since each behavior has different parameters, rules r^P, r^V and application rules r^B are designed specifically for that particular behavior. The choice and parameter rules for $\mathcal{N}$ and behaviors MoveByFlocking (FL) and MoveInFormation (FR) are outlined in Table E.1 and determined as follows. Note that n_1 and n_2 require the Followers to retain previous paths of the Leaders.

	MoveByFlocking		
	Parameter p	Value v_p	Adapted value v_p
n_1	$\mathcal{C}$	Add O^V, $o^V = f(\Omega)$	$c_k^\beta(o^V) \leftarrow g_1(l_1)c_k^\beta(o^V), k \in 1,2$
n_2	$\mathcal{C}$	Add FVL	$c_k^\gamma(\text{FVL}) \leftarrow g_1(l_{2,4})c_k^\gamma(\text{FVL}) \in 1,2$
n_4	K_C	$K_C = f(\text{ID})$	$K_C \leftarrow g_1(l_{2,4})K_C$
n_5	$\mathcal{C}$	Add O^H $(c_k^\beta, r_O, \delta_O) = f(r_s)$	$c_i^{\beta,h} \leftarrow g_1(l_5)c_k^{\beta,h}$, $k \in 1,2$ $r_{O,h} \leftarrow g_3(l_5)r_{O,h}$ $\delta_{O,h} \leftarrow g_3(l_5)\delta_{O,h}$
	MoveInFormation		
	Parameter p	Value v_p	Adapted value v_p
n_1	b^x	$b_i^x = f(d_\Omega)$	$d_\Omega = g_2(l_{2,4})d_\Omega$
n_2	$\mathcal{W}$	$w_{fl} > w_{ff}$, $w_{fl} = f(TF)$	$w_{fl} \leftarrow g_1(l_{2,4})w_{fl}$
n_4	b^y	$b_i^y = f(d_\Omega)$	$d_\Omega = g_2(l_{2,4})d_\Omega$
n_5	$\mathcal{W}_\mathcal{R}$	$w_x^\mathcal{R} = f(HF_x)$, $w_x^\mathcal{R} \neq 0$	$K_O \leftarrow g_1(l_5)K_O$

Table E.1 – Overview of the adaptation process of the ROBOT GUIDANCE institution.

$r_1^{P,V}$ **(Safe areas)** SafeAreas Ω are the obstacle-free areas between the outermost paths of the Leaders. The distance $d_\Omega(p)$ is the width of the safe area at a position p (see Fig. 17.2 for illustration).

FL: The connectivity parameter $\mathcal{C}_i$ of F_i includes *virtual obstacles* O^V at the points on the human paths that are closest to F_i at time t, i.e. $O^V = \{p_h(\tau) | \tau = \text{argmin}_\tau(\|p_i(t) - p_h(\tau)\|), \tau \in (t_0, t), h \in \text{Leaders}\}$. Note that the norm n_1 is designed to keep robots within Ω, but when F$_i$ falls outside of it, the repulsive weights might refrain the robot from entering it. The solution to this is provided by the adaptation rule of n_1.

FR: The safe area width $d_\Omega(p)$ drives computation of the formation bias using the LFT method presented in Chapter 7. The variable d_Ω drives the formation transformation, scaling the original bias of F_i to $b_i^x = f(d_\Omega)$ towards the column-like shape. The new bias is parametrized by s_x, which shrinks or expands the formation width. Note that s_y, which shortens or elongates the formation along the direction of motion is modified by norm n_4.

$r_2^{P,V}$ (**Reinforce following**) To facilitate navigation in *complex spaces* – confined spaces with narrow, sharp turns, norm n_2 provides the Followers with the means to reinforce the Leader-Follower interactions.

FL: The current virtual point VL of F_i is calculated based on the paths of the Leaders and not their immediate positions, and referred to as the *former* virtual point

FVL $\in \mathcal{P}_{\mathrm{VL}}$. FVL is computed as FVL $= \mathrm{argmin}_{\mathrm{VL}}(\|p_i(t) - \mathrm{VL}(p)\| - d_F)$, where indicator p allows to track and discard all the points in $\mathcal{P}_{\mathrm{VL}}$ already reached by F_i and d_F is the preferable distance of FVL being ahead of F_i (see Fig. 17.2). By being more immediate than VL, FVL allows F_i to follow the path of the Leaders even when it has lost their line of sight.

FR: To put larger emphasis on the following, rather than on keeping perfectly-shaped formation, the weights of the L-F edges, w_{lf} in the Laplacian are reinforced with respect to the F-F edges (with weights w_{ff}).

$r_3^{P,V}$ (**Help request**) Norm n_3 parametrizes an individual robot behavior - spoken help request. The value of spoken phrase is set to asking Leaders to stop and wait until F_i can catch up.

$r_4^{P,V}$ (**Narrow spaces**) Norm n_4 facilitates passing of the Followers through narrow spaces by assigning who is to pass next.

FL: The control gain K_c of each individual F_i is tuned according to its ID i, so that some robots speed up when passing, while the others slow down, letting the faster robots to proceed. This is achieved by multiplying K_c by a factor from a lookup table $[1.5, 0.5, 0.2]$, where the entry of the table corresponds to the robot ID, so the robot F_1 can render a maximum gain of $1.5K_c$, and the robot F_3 can render a maximum gain of $0.2K_c$, making F_1 arrive at the narrowing much earlier than F_3. Similar behavior can be observed when humans walk though door and let the ones with higher priority to pass first.

FR: Coordination is achieved in the length-wise change of the formation geometry b_i^y. Depending on the width of the available space d_Ω at the position of the Leaders, each F_i spreads its bias along the direction of motion (the y-axis of the formation). Note that norm n_1 assures that the Followers reduce the perpendicular bias components (x-axis). The result is a formation with a shape converging to a column for very narrow passages.

$r_5^{P,V}$ (**Social spaces**)

Norm n_5 for both behaviors introduces repulsive forces that drive the robots away from human social spaces.

FL: Humans are modeled as obstacles $O^H = \{o^h\}$, $h \in \{\mathsf{human}_1, \mathsf{human}_2\}$, with the parameter values $c_1^{\beta,h}, c_2^{\beta,h}, r_{O^H}$ and δ_{O^H} that depend on the $\mathsf{personal\_distance} \in K$.

FR: To respect social spaces of human Leaders, the follower F_i activates the repulsive forces described in Section 5.2.2, where the repulsion weight changes continuously as a function of the distance d_{hi}, gain K_o and ranges Δ_a, and $\Delta_c < \Delta_a$ that depend on $\mathsf{social\_force}_\mathsf{x} \in K$.

E.3 Value Adaptation

We define the following measures L to evaluate outcomes of the norms. For F_i:

IN_SAFE_AREA	$l_1 = 1$ if $\bar{x}_i \in \Omega$, otherwise $l_1 = 0$
IS_FOLLOWING	$l_{2,4} = 1 - \min(\frac{1}{T_L} \lVert \bar{x}_i - \mathrm{VL} \rVert, 1)$
IN_SOCIAL_SPACE	$l_5 = \min(\min_h(\frac{1}{2*\mathsf{social_h}} \lVert p_i - p_h \rVert), 1)$

Functions used for norm adaptation are: $g_1(x) = 2 - x$, $g_2(x) = x^2$, and $g_3(x) = \frac{3-x}{2}$. Note that $l \in [0,1]\ \forall\ l \in L$. Adaptation rules are outlined in Table E.1 and determined for F_i as follows.

r_1^L **(Safe areas)**

FL: Weights c_1^β and c_2^β of the virtual obstacles O^V are increased when IN_SAFE_AREA of F_i is false as $c_k^\beta \leftarrow g_1(l_1)c_k^\beta$, $k \in 1,2$.

FR: Bias b_i^x of F_i is a function of a safe area width scaled as according to $d_\Omega \leftarrow g_2(l_1)d_\Omega$. This narrows the formation when IN_SAFE_AREA of F_i is false.

r_2^L **(Reinforce following)**

FL: The attraction weights c_1^γ and c_2^γ towards FVLare increased if F_i stays behind according to $c_i^\gamma \leftarrow g_1(l_{2,4})c_k^\gamma$, $k \in 1,2$.

FR: If the Follower F_i stays behind, weights w_{fl} of the L-F edges are reinforced according to $w_{fl} \leftarrow g_1(l_{2,4})w_{fl}$.

r_3^L **(Help request)**

For norm n_3, we assume that the human Leaders stop each time F_i requests it.

r_4^L **(Narrow spaces)**

FL: The control gain K_C is increased if F_i falls behind while navigating narrow passages $K_C \leftarrow g_1(l_{2,4})K_C$.

FR: Bias b_i^y of F_i is a function of scaled $d_\Omega \leftarrow g_1(l_{2,4})d_\Omega$. It increases the bias lengthwise if F_i is failing to pass a narrowing.

r_5^L **(Social spaces)**

FL: Weights $c_1^{\beta,h}$ and $c_2^{\beta,h}$ of the repulsive forces are increased if the social space of a human has been invaded as $c_i^{\beta,h} \leftarrow g_1(l_5)c_k^{\beta,h}$, $k \in 1,2$. Moreover, the interaction range $r_{O,h}$ and the desired distance $\delta_{O,h}$ increase as $r_{O,h} \leftarrow g_3(l_5)r_{O,h}$ and $\delta_{O,h} \leftarrow g_3(l_5)\delta_{O,h}$.

FR: If social space is interrupted, repulsion weights w_{ih} are reinforced by increasing $K_O \leftarrow g_1(l_5)K_O$.

Glossary

AI	Artificial Intelligence	118
AMCL	Adaptive Monte Carlo Localization	20
ANCOVA	Analysis of covariance	181, 182
ANOVA	Analysis of variance	177, 178
DBSCAN	Density-Based Spatial Clustering of Applications with Noise	18
DWA	Dynamic Window Approach	14, 20, 60, 67, 68, 171
EI	Electronic Institutions	115
FI-GM-PHD	Formation Information GM-PHD	9, 12, 70, 77–79, 85–94, 96–100, 102–104, 164, 183–185, 226, 239–241
FISST	Finite Set Statistics	72
FMM	Fast Marching Method	13, 20, 49, 50, 99, 161, 193
FOV	Field of View	15, 60, 97
GM-PHD	Gaussian Mixture Probability Hypothesis Density	9, 35, 36, 69, 70, 72, 75, 79, 82, 85, 88–94, 97–100, 102, 103, 226
GNSS	Global Navigation Satellite System	30, 69, 70
HRI	Human-Robot Interaction	36, 136, 137
IAC	Institutional Agent Controller	119, 120, 129
IAD	Institutional Analysis and Development	112, 113, 117, 133, 134
ID	Identification number	34–36, 39, 69, 71, 85, 95, 210
IMU	Inertial Measurement Unit	15
IPOL	Instituto Português de Oncologia de Lisboa	5, 6, 12, 15, 17, 21–23, 33, 53, 55–58
IR	Institutional Robotics	10, 11, 118–120, 128
JPDA	Joint Probabilistic Data Association	72
LED	Light Emitting Diode	15, 16
LFT	Local Formation Transformation	9, 34, 59–68, 210, 226, 251
LIDAR	Light Detection and Ranging	35, 87
LOS	Line Of Sight	214

MAS	Multi-Agent Systems	4, 113–116, 118, 127
MCMCDA	Markov Chain Monte Carlo Data Association	72
MCS	Motion Capture System	14, 19, 21, 96, 170, 172, 217, 229
MHT	Multiple Hypothesis Tracking	72
MOnarCH	Multi-Robot Cognitive Systems Operating in Hospitals	12, 15, 57, 137
NLOS	Non Line Of Sight	19
NN	Nearest Neighbor	71, 72
OS	Operating System	15
OSPA	Optimal SubPattern Assignment	88, 91, 93, 94, 99–102, 184, 240
PHD	Probability Hypothesis Density	35, 69, 72–75, 79, 80, 95
PM	Proxemics Model	36, 51, 52, 135, 191
PMHT	Probabilistic Multiple Hypothesis Tracking	72
PN	Petri Net	119
RFS	Random Finite Set	72–75, 79, 87
RG	Relational Graph	129
ROS	Robot Operating System	13, 14, 16–20
SFM	Social Forces Model	51–53, 191, 234
SLAM	Simultaneous Localization and Mapping	20, 72
UWB	Ultra-Wide Band	14, 19, 22

Mathematical Symbols

Cooperative Multi-Robot Systems and Algorithms

Preliminaries	
R	Robot
L	Leader
F	Follower
H	Human
O	Obstacle
N	Number of robots
p_i	Position of robot R_i
$\dot{p}_i$	Velocity of robot R_i
u_i	Control input of robot R_i
$\mathcal{P}_i$	Path of robot R_i
$\mathcal{T}_i$	Trajectory of robot R_i
d_{ab}	Distance between p_a to p_b
α_i	Orientation of robot R_i
Δ_{ab}	Range from point p_a to point p_b
γ_{ab}	Bearing between p_a and p_b
$\mathbb{I}_{R_i}$	Coordinate frame of robot R_i
$\mathbb{I}_W$	Global coordinate frame
Γ	Motion vector
Evaluation	
e_F	Formation error
e_O	Orientation error
e_C	Connectivity metric
e_R	Cohesion metric
e_E	Deviation metric

Formation Control & Flocking Algorithms

Graph Definitions	
$\mathcal{G}$	Graph
$\mathcal{E}$	Set of edges
$\mathcal{V}$	Set of nodes
$\mathcal{C}$	Graph connectivity
Ξ	Set of neighbors
$\mathcal{H}$	Incidence matrix
$\mathcal{A}$	Adjacency matrix
$\mathcal{W}$	Weight matrix
$\mathcal{L}$	Laplacian matrix
Formation Control	
b	Bias matrix
α_D	Desired formation orientation
K_u	Gain of the formation distance control
K_ϕ	Gain of the orientation control
Flocking	
$\mathcal{V}_\alpha$	Set of α-agents (robots)
$\mathcal{V}_\beta$	Set of β-agents (obstacles)
$\mathcal{V}_\gamma$	Set of γ-agents (virtual leaders)
b_A	Desired agent-agent distance
b_O	Desired agent-obstacle distance
Δ_A	Agent-agent interaction range
Δ_O	Agent-obstacle interaction range
ϕ	Function for a smooth pairwise attractive or repulsive potential
$\mathbf{n}_{ab}$	Vector from p_a to p_b
VL	Virtual leader agent at position p_{VL}
K_C	Gain of flocking forces
K_x	Positive gain, $x \in \{\alpha, \beta, \gamma\}$
Adaptive Formations	
η	Local density of obstacles ahead estimated by robot R_i
S_{LFT}	Specification of the formation shape adaptation
ψ	Transformation function local to robot R_i
$\tilde{b}$	Transformed bias matrix
e_A	Measure of LFT reactivity

Cooperative Localization

Random Finite Sets Notation	
$\mathcal{X}$	State space
$\mathcal{Z}$	Observation space
$\mathcal{F}(\mathcal{X})$	Collection of all finite subsets of $\mathcal{X}$
$\mathcal{F}(\mathcal{Z})$	Collection of all finite subsets of $\mathcal{Z}$
X_k	Multi-target state (finite set) at time k
Z_k	Multi-target observation (finite set) at time k
$M(k)$	The number of the targets at time k
$N(k)$	The number of the measurements at time k
$\mathbf{x}$	Single-target state vector
$\mathbf{z}$	Single measurement
$\Theta_k(\mathbf{x})$	RFS generated by target with a state $\mathbf{x}$ at time k
$S_{k\|k-1}(\zeta)$	RFS of the target state ζ at time k
Γ	RFS of spontaneous birth
K	Clutter RFS
$f_{k\|k-1}(\cdot\|\cdot)$	Transition density
$g_k(\cdot\|\cdot)$	Likelihood function
p_S	Survival probability
p_D	Detection probability

Gaussian Mixture Probability Hypothesis Density Filter Notation	
v	Intensity
$v_{S,k\|k-1}(\mathbf{x})$	Survival intensity
$v_T(\mathbf{x})$	Missed-detection term of posterior intensity
$v_D(\mathbf{x};\mathbf{z})$	Detection term of posterior intensity
$\gamma(\mathbf{x})$	Birth intensity
$\kappa(\mathbf{z})$	Clutter intensity
$\mathcal{N}(\cdot;m,P)$	Gaussian density with mean m and covariance P
J	Number of Gaussian components
$w^{(i)}$	Weight of Gaussian component i
J_{max}	Max number of components after selection
T_S	Weight threshold used in selection step
U_S	Distance threshold used in selection step
T_{SE}	Weight threshold used for state extraction

Cooperative Localization

Problem Description	
F	State transition matrix
Q	Process noise covariance
H	Observation matrix
U	Observation noise covariance
σ_f^2	Standard deviation of the process noise
σ_ϵ^2	Standard deviation of the measurement noise
δ	Time step
r_s	Sensing range
$[\beta_1, \beta_2]$	Intersection of occlusion regions
$\bar{o}_p^{(c)}$	OSPA metric
e_L	Self-localization error
e_M	Measurement error
$p_{D,s}$	Sensor-dependent missed detection probability
p_{md}	Message drop probability

Formation Information GM-PHD Filter Notation	
v_ζ	Coalition intensity
$o^{(\cdot,\cdot)}$	Measure for sorting in the coalition step
$\Phi_{\zeta,0}$	Initial budget in the coalition step
h_i	Projected formation state with respect to R_i
$\hat{e}_S$	Estimated association error
A_Δ	Best role assignment

SOCIAL BEHAVIORS

METHODS FOR MIXED HUMAN-ROBOT FORMATIONS	
$\Delta_{S,h}$	Proxemics: range of social space of human H_h
$\Delta_{P,h}$	Proxemics: range of personal space of human H_h
$\Delta_{I,h}$	Proxemics: range of intimate space of human H_h
Γ_R	Force generating a repulsive field around humans
$\mathcal{W}_{\mathcal{R}}$	Collection of the human repulsion weights
K_o	Gain of the repulsive weight
Δ_a	Activation range of the human repulsive force
Δ_c	Imminent collision range
SOCIAL FORCES MODEL	
Γ_d	Attractive force driving human towards destination d
Γ_o	Force preventing human colliding with obstacles $o \in O$
$\Gamma_{h'}$	Force for repulsive effects from the other humans $\mathrm{H}_{h'} \in \mathrm{H}$
Γ_r	Force for repulsive effects from the robot $\mathrm{R}_r \in \mathrm{R}$
K_{fx}	Gain of the force, $x \in \{d, o, h, r\}$
γ_f	Area of influence of the force Γ_d
Δ_x	Areas of influence of the SFM forces $x \in \{o, h, r\}$
σ_f	Standard deviation for human motion fluctuation
$\dot{p}_d$	Desired human velocity
τ_f	Relaxation time in for calculating Γ_d

Institutional Formalism

General Definitions	
$\mathcal{I}$	Institution
$\mathcal{D}$	Domain
$\mathcal{G}$	Grounding
$\mathcal{N}$	Set of norms
$\mathcal{N}^A$	Active norms
$\mathcal{N}^T$	Satisfied norms
$n \in \mathcal{N}$	Norm
$\mathcal{F}$	Feasibility relation
$\mathcal{K}$	Norm realization
$\mathcal{C}$	Case Study
Domain Description	
A	Set of actions
B	Set of behaviors
R	Set of state variables
C	Grounded conditions
K	Grounded knowledge
L	Set of experience measures
Behavior Specification	
Λ_k	Set of modalities of behavior B_k
$\lambda \in \Lambda_k$	A behavior modality
P	Set of parameters
$p \in P$	Parameter
V_p	Set of values of parameter p
$v \in V_p$	Value of parameter p
Rules	
r^N	Requirement rules
r^P	Choice rules
r^V	Value rules
r^L	Adaptation rules
r^B	Application rules
r^O	Outcome rules

CASE STUDIES

CASE STUDY $\mathcal{C}_I$	
Φ_{ACS}	Activity Critical Area
T_{QAT}	Quiet Activity Time
s_A	Appropriate Speed
IM	Interaction Mode
SM	Speech Mode
FS	Formation Shape
CASE STUDY $\mathcal{C}_{III}$	
EA	Exhibition Area
Ω	Safe Areas
VL	Virtual Point
FVL	Former Virtual Point
TF	Team Social Force, TF = $\{w_{ff}, w_{fl}\}$
T_L	Follower Lost Distance
T_C	Narrow Passage Mark
o	Obstacle Characteristics, $o \in \{O^R, O^V, O^H\}$

Bibliography

[1] K. Dautenhahn, "Roles and functions of robots in human society: implications from research in autism therapy", *Robotica*, vol. 21, no. 4, pp. 443–452, 2003.

[2] A. Wasik, J. N. Pereira, R. Ventura, P. U. Lima, and A. Martinoli, "Graph-based distributed control for adaptive multi-robot patrolling using local formation transformation", in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 1721–1728.

[3] A. Wasik, P. U. Lima, and A. Martinoli, "A robust localization system for multi-robot formations based on an extension of a Gaussian mixture probability hypothesis density filter", *Autonomous Robots*, vol. 44, 395–414, 2019.

[4] A. Wasik, A. Martinoli, and P. U. Lima, "A robust relative positioning system for multi-robot formations leveraging an extended GM-PHD filter", *Proceedings of the First International Symposium on Multi-Robot and Multi-Agent Systems*, pp. 71–77, 2017.

[5] A. Wasik, R. Ventura, J. N. Pereira, P. U. Lima, and A. Martinoli, "Lidar-based relative position estimation and tracking for multi-robot systems", in *Robot 2015: Second Iberian Robotics Conference*, Springer International Publishing, 2016, pp. 3–16.

[6] P. Silva and P. U. Lima, "Institutional robotics", in *Advances in Artificial Life*. Springer Berlin Heidelberg, 2007, pp. 595–604.

[7] P. Silva, R. Ventura, and P. U. Lima, "Institutional environments", in *Proceedings of workshop AT2AI: From agent theory to agent implementation workshop, 7th International Conference on Autonomous Agents and Multiagent Systems*, 2008, pp. 157–164.

[8] A. Wasik, S. Tomic, A. Saffiotti, F. Pecora, A. Martinoli, and P. U. Lima, "Towards norm realization in institutions mediating human-robot societies", *2018 IEEE/RSJ International Conference On Intelligent Robots And Systems (IROS)*, IEEE International Conference on Intelligent Robots and Systems, pp. 297–304, 2018.

[9] S. Tomic, A. Wasik, P. U. Lima, A. Martinoli, F. Pecora, and A. Safiotti, "Towards institutions for mixed human-robot societies", in *Intternational Joint Conference on Autonomous Agents and Multiagent Systems*, 2018, pp. 2216–2217.

[10] A. Wasik, A. Martinoli, and P. U. Lima, "An institutional robotics approach to the design of socially aware multi-robot behaviors", *Proceedings of the RO-MAN 2017 Workshop on Towards Intelligent Social Robots: Social Cognitive Systems in Smart Environments*, pp. 2–7, 2017.

[11] J. Messias, R. Ventura, P. Lima, J. Sequeira, P. Alvito, C. Marques, and P. Carriço, "A robotic platform for edutainment activities in a pediatric hospital", in *IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, 2014, pp. 193–198.

[12] O. Michel, "Cyberbotics ltd. webots™: professional mobile robot simulation", *International Journal of Advanced Robotic Systems*, vol. 1, no. 1, p. 5, 2004.

[13] S. Tomic, F. Pecora, and A. Saffiotti, "Norms, institutions, and robots", *ArXiv preprint arXiv:1807.11456*, 2018.

[14] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with rao-blackwellized particle filters", *IEEE Transactions on Robotics*, vol. 23, no. 1, pp. 34–46, 2007.

[15] D. Fox, "Kld-sampling: adaptive particle filters", in *Advances in neural information processing systems*, 2002, pp. 713–720.

[16] J. A. Sethian, "Fast marching methods", *SIAM review*, vol. 41, no. 2, pp. 199–235, 1999.

[17] C. W. Reynolds, "Flocks, herds and schools: a distributed behavioral model", *SIGGRAPH Comput. Graph.*, vol. 21, no. 4, pp. 25–34, 1987.

[18] D. Helbing and P. Molnár, "Social force model for pedestrian dynamics", *Physical Review E*, vol. 51, no. 5, pp. 4282–4286, 1995.

[19] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: algorithms and theory", *IEEE Transactions on Automatic Control*, vol. 51, no. 3, pp. 401–420, 2006.

[20] N. Moshtagh and A. Jadbabaie, "Distributed geodesic control laws for flocking of nonholonomic agents", *IEEE Transactions on Automatic Control*, vol. 52, no. 4, pp. 681–686, 2007.

[21] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.

[22] J. M. Soares, A. P. Aguiar, A. M. Pascoal, and A. Martinoli, "A distributed formation-based odor source localization algorithm - design, implementation, and wind tunnel evaluation", in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 1830–1836.

[23] T. Balch and R. C. Arkin, "Behavior-based formation control for multirobot teams", *IEEE Transactions on Robotics and Automation*, vol. 14, no. 6, pp. 926–939, 1998.

[24] M. A. Lewis and K.-H. Tan, "High precision formation control of mobile robots using virtual structures", *Autonomous Robots*, vol. 4, no. 4, pp. 387–403, 1997.

[25] J. Fredslund and M. J. Matarić, “A general algorithm for robot formations using local sensing and minimal communication”, *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 837–846, 2002.

[26] S.-W. Hong, S.-W. Shin, and D.-S. Ahn, “Formation control based on artificial intelligence for multi-agent coordination”, in *IEEE International Symposium on Industrial Electronics Proceedings*, vol. 1, 2001, pp. 429–434.

[27] N. E. Leonard and E. Fiorelli, “Virtual leaders, artificial potentials and coordinated control of groups”, in *Proceedings of the 40th IEEE Conference on Decision and Control*, IEEE, vol. 3, 2001, pp. 2968–2973.

[28] R. Olfati-Saber, J. A. Fax, and R. M. Murray, “Consensus and cooperation in networked multi-agent systems”, *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[29] C. Ghedini, C. H. C. Ribeiro, and L. Sabattini, “Toward efficient adaptive ad-hoc multi-robot network topologies”, *Ad Hoc Networks*, 2018, published.

[30] J. N. Pereira, P. Silva, P. U. Lima, and A. Martinoli, “An experimental study in wireless connectivity maintenance using up to 40 robots coordinated by an institutional robotics approach”, in *IEEE International International on Intelligent Robots and Systems*, 2013, pp. 5073–5079.

[31] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, “Human-aware robot navigation: a survey”, *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726 –1743, 2013.

[32] C. Bartneck and J. Forlizzi, “A design-centred framework for social human-robot interaction”, in *RO-MAN 2004. 13th IEEE International Workshop on Robot and Human Interactive Communication*, IEEE, 2004, pp. 591–594.

[33] T. Fong, I. Nourbakhsh, and K. Dautenhahn, “A survey of socially interactive robots”, *Robotics and Autonomous Systems*, vol. 42, pp. 143–166, 2003.

[34] J. Krueger, “Extended cognition and the space of social interaction”, *Consciousness and Cognition*, vol. 20, no. 3, pp. 643 –657, 2011.

[35] S. Monteiro and E. Bicho, “Attractor dynamics approach to formation control: theory and application”, *Autonomous Robots*, vol. 29, no. 3, pp. 331–355, 2010.

[36] J. Vilca, L. Adouane, and Y. Mezouar, “Adaptive leader-follower formation in cluttered environment using dynamic target reconfiguration”, in *Distributed Autonomous Robotic Systems*, Springer Japan, 2016, pp. 237–254.

[37] W. L. Seng, J. C. Barca, and Y. A. Şekercioğlu, “Distributed formation control of networked mobile robots in environments with obstacles”, *Robotica*, vol. 34, no. 6, 1403–1415, 2016.

[38] R. Alur, A. Das, J. Esposito, R. Fierro, G. Grudic, Y. Hur, V. Kumar, I. Lee, J. P. Ostrowski, G. Pappas, B. Southall, J. Spletzer, and C. J. Taylor, “A framework and architecture for multirobot coordination”, in *Experimental Robotics VII*. Springer Berlin Heidelberg, 2001, pp. 303–312.

[39] C. R. Kuppan, M. Singaperumal, and T. Nagarajan, “Distributed planning and control of multirobot formations with navigation and obstacle avoidance”, in *2011 IEEE Recent Advances in Intelligent Computational Systems*, 2011, pp. 621–626.

[40] J. P. Desai, J. P. Ostrowski, and V. Kumar, “Modeling and control of formations of nonholonomic mobile robots”, *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 905–908, 2001.

[41] M. Deghat, I. Shames, and B. D. Anderson, “Safe formation control with obstacle avoidance”, *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 11252 –11 257, 2011.

[42] A. Santos Brandao, M. Sarcinelli-Filho, R. Carelli, and T. F. Bastos-Filho, “Decentralized control of leader-follower formations of mobile robots with obstacle avoidance”, in *2009 IEEE International Conference on Mechatronics*, 2009, pp. 1–6.

[43] F. Zhang and N. E. Leonard, “Cooperative filters and control for cooperative exploration”, *IEEE Transactions on Automatic Control*, vol. 55, no. 3, pp. 650–663, 2010.

[44] J. Henry, H. P. H. Shum, and T. Komura, “Interactive formation control in complex environments”, *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 2, pp. 211–222, 2014.

[45] R. Falconi, S. Gowal, and A. Martinoli, “Graph based distributed control of non-holonomic vehicles endowed with local positioning information engaged in escorting missions”, in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 3207–3214.

[46] M. Cao, C. Yu, and B. D. Anderson, “Formation control using range-only measurements”, *Automatica*, vol. 47, no. 4, pp. 776 –781, 2011.

[47] A. N. Bishop, “A very relaxed control law for bearing-only triangular formation control”, *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 5991 –5998, 2011.

[48] J. M. Soares, A. P. Aguiar, A. M. Pascoal, and A. Martinoli, “Joint asv/auv range-based formation control: theory and experimental results”, in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 5579–5585.

[49] D. Dias, P. U. Lima, and A. Martinoli, “Distributed formation control of quadrotors under limited sensor field of view”, in *Proceedings of the 15th International International on Autonomous Agents and Multiagent Systems*, ACM, 2016, pp. 1087–1095.

[50] M. Ye, B. D. O. Anderson, and C. Yu, “Bearing-only measurement self-localization, velocity consensus and formation control”, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 53, no. 2, pp. 575–586, 2017.

[51] M. M. Zavlanos and G. J. Pappas, “Dynamic assignment in distributed motion planning with local coordination”, *IEEE Trans. on Robotics*, vol. 24, no. 1, pp. 232–242, 2008.

[52] N. Michael, M. M. Zavlanos, V. Kumar, and G. J. Pappas, “Distributed multi-robot task assignment and formation control”, in *IEEE International International on Robotics and Automation*, 2008, pp. 128–133.

[53] Z. Sun and Y. Xia, "Consensus-based formation control with dynamic role assignment", in *The 26th Chinese Control and Decision International*, IEEE, 2014, pp. 3681–3686.

[54] Y. Zhang, G. Song, G. Qiao, Z. Li, Y. Li, and A. Song, "Strategy research of role assignment and formation control for multi-robot systems", in *2013 IEEE International International on Robotics and Biomimetics*, 2013, pp. 958–963.

[55] K. Panta, D. E. Clark, and B.-N. Vo, "Data association and track management for the gaussian mixture probability hypothesis density filter", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 3, pp. 1003–1016, 2009.

[56] P. Stegagno, M. Cognetti, L. Rosa, P. Peliti, and G. Oriolo, "Relative localization and identification in a heterogeneous multi-robot system", in *2013 IEEE International International on Robotics and Automation*, 2013, pp. 1857–1864.

[57] A. Franchi, G. Oriolo, and P. Stegagno, "Mutual localization in multi-robot systems using anonymous relative measurements", *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1302–1322, 2013.

[58] B. Jiang, M. Deghat, and B. D. O. Anderson, "Simultaneous velocity and position estimation via distance-only measurements with application to multi-agent system control", *IEEE Transactions on Automatic Control*, vol. 62, no. 2, pp. 869–875, 2017.

[59] Z. Han, K. Guo, L. Xie, and Z. Lin, "Integrated relative localization leader-follower formation control", *IEEE Transactions on Automatic Control*, pp. 1–1, 2018.

[60] B.-N. Vo and W.-K. Ma, "The gaussian mixture probability hypothesis density filter", *IEEE Transactions on signal processing*, vol. 54, no. 11, pp. 4091–4104, 2006.

[61] J. Burgess, "Interpersonal spacing behavior between surrounding nearest neighbors reflects both familiarity and environmental density", *Ethology and Sociobiology*, vol. 4, no. 1, pp. 11 –17, 1983.

[62] R. Gockley and M. J. Matarić, "Encouraging physical therapy compliance with a hands-off mobile robot", in *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-robot Interaction*, 2006, pp. 150–155.

[63] E. T. Hall, *The hidden dimension*. Doubleday & Co, 1966.

[64] J. Rios-Martinez, "Socially-aware robot navigation : combining risk assessment and social conventions", PhD thesis, Université de Grenoble, 2013.

[65] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: crowd-aware robot navigation with attention-based deep reinforcement learning", IEEE International Conference on Robotics and Automation ICRA, pp. 6015–6022, 2019.

[66] J. A. Kirkland and A. A. Maciejewski, "A simulation of attempts to influence crowd dynamics", in *IEEE International International on Systems, Man and Cybernetics*, vol. 5, 2003, pp. 4328–4333.

[67] G. Ferrer, A. G. Zulueta, F. H. Cotarelo, and A. Sanfeliu, "Robot social-aware navigation framework to accompany people walking side-by-side", *Autonomous Robots*, vol. 41, no. 4, pp. 775–793, 2017.

[68] A. Alahi, V. Ramanathan, K. Goel, A. Robicquet, A. Sadeghian, L. Fei-Fei, and S. Savarese, "Learning to predict human behavior in crowded scenes", *Group and Crowd Behavior for Computer Vision*, pp. 183–207, 2017.

[69] A. Vega, L. J. Manso, D. G. Macharet, P. Bustos, and P. Núñez, "Socially aware robot navigation system in human-populated and interactive environments based on an adaptive spatial density function and space affordances", *Pattern Recognition Letters*, vol. 118, pp. 72 –84, 2019.

[70] G. Diego and T. K. O. Arras, "Please do not disturb! minimum interference coverage for social robots", in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2011, pp. 1968–1973.

[71] K. L. Koay, E. A. Sisbot, D. S. Syrdal, M. L. Walters, K. Dautenhahn, and R. Alami, "Exploratory study of a robot approaching a person in the context of handing over an object.", in *AAAI Spring Symposium: Multidisciplinary Collaboration For Socially Assistive Robotics*, 2007, pp. 18–24.

[72] J. Kessler, C. Schroeter, and H.-M. Gross, "Approaching a person in a socially acceptable manner using a fast marching planner", in *Intelligent Robotics and Applications*, Springer Berlin Heidelberg, 2011, pp. 368–377.

[73] J. T. Butler and A. Agah, "Psychological effects of behavior patterns of a mobile personal robot", *Autonomous Robots*, vol. 10, no. 2, pp. 185–202, 2001.

[74] L. Takayama and C. Pantofaru, "Influences on proxemic behaviors in human-robot interaction", in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2009, pp. 5495–5502.

[75] A. K. Pandey and R. Alami, "A framework towards a socially aware mobile robot motion in human-centered dynamic environment", in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 5855–5860.

[76] E. A. Martinez-Garcia, Ohya Akihisa, and Shin'ichi Yuta, "Crowding and guiding groups of humans by teams of mobile robots", in *IEEE Workshop on Advanced Robotics and its Social Impacts*, 2005, pp. 91–96.

[77] J. Guzzi, A. Giusti, L. M. Gambardella, G. Theraulaz, and G. A. Di Caro, "Human-friendly robot navigation in dynamic environments", in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 423–430.

[78] A. Sgorbissa, I. Papadopoulos, B. Bruno, C. Koulouglioti, and C. Recchiuto, "Encoding guidelines for a culturally competent robot for elderly care", in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1988–1995.

[79] R. Kirby, R. Simmons, and J. Forlizzi, "Companion: a constraint-optimizing method for person-acceptable navigation", in *18th IEEE Int Symposium on Robot and Human Interactive Communication*, 2009, pp. 607–612.

[80] M. Shiomi, T. Kanda, D. F. Glas, S. Satake, H. Ishiguro, and N. Hagita, "Field trial of networked social robots in a shopping mall", in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 2846–2853.

[81] R. R. Murphy, "Human-robot interaction in rescue robotics", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 34, no. 2, pp. 138–153, 2004.

[82] R. Vaughan, N. Sumpter, J. Henderson, A. Frost, and S. Cameron, "Experiments in automatic flock control", *Robotics and Autonomous Systems*, vol. 31, no. 1, pp. 109 –117, 2000.

[83] M. A. Potter, L. Meeden, and A. C. Schultz, "Heterogeneity in the coevolved behaviors of mobile robots: the emergence of specialists", in *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2*, ser. IJCAI'01, Morgan Kaufmann Publishers Inc., 2001, pp. 1337–1343.

[84] J. Alonso-Mora, E. Montijano, T. Nägeli, O. Hilliges, M. Schwager, and D. Rus, "Distributed multi-robot formation control in dynamic environments", *Autonomous Robots*, vol. 43, no. 5, pp. 1079–1100, 2019.

[85] S. Zhang and Y. Guo, "Distributed multi-robot evacuation incorporating human behavior", in *2013 10th IEEE International Conference on Control and Automation (ICCA)*, 2013, pp. 864–869.

[86] B. Capelli, C. Secchi, and L. Sabattini, "Communication through motion: legibility of multi-robot systems", in *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, 2019, pp. 126–132.

[87] A. Garrell and A. Sanfeliu, "Cooperative social robots to accompany groups of people", *The International Journal of Robotics Research*, vol. 31, no. 13, pp. 1675–1701, 2012.

[88] Z. Talebpour and A. Martinoli, "Adaptive risk-based replanning for human-aware multi-robot task allocation with local perception", *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3790–3797, 2019.

[89] C. Godsil and G. Royle, *Algebraic Graph Theory*, ser. Graduate Texts in Mathematics. volume 207 of Graduate Texts in Mathematics. Springer, 2001, vol. 207.

[90] R. Falconi, "Coordinated control of robotic swarms in unknown environments", PhD thesis, Università di Bologna, 2009.

[91] Meng Ji, A. Muhammad, and M. Egerstedt, "Leader-based multi-agent coordination: controllability and optimal control", in *2006 American Control Conference*, 2006.

[92] S. A. Gowal, "A Framework for Graph-Based Distributed Rendezvous of Nonholonomic Multi-Robot Systems", PhD thesis, EPFL, Lausanne, 2013.

[93] A. K. Pandey and R. Alami, "A step towards a sociable robot guide which monitors and adapts to the person's activities", in *2009 International Conference on Advanced Robotics*, 2009, pp. 1–8.

[94] Z. Xie, Y. Long, and H. Cheng, “Distributed adaptive formation control of a team of aerial robots in cluttered environments”, in *International Conference on Intelligent Robotics and Applications*, Springer, 2019, pp. 544–558.

[95] K. Mease, S Bharadwaj, and S Iravanchy, “Timescale analysis for nonlinear dynamical systems”, *Journal of guidance, control, and dynamics*, vol. 26, no. 2, pp. 318–330, 2003.

[96] M. Vasic and A. Martinoli, “A collaborative sensor fusion algorithm for multi-object tracking using a gaussian mixture probability hypothesis density filter”, *Proceedings of the 2015 IEEE 18th International International on Intelligent Transportation Systems*, ITSC ’15, pp. 491–498, 2015.

[97] C. Qiu, Z. Zhang, H. Lu, and H. Luo, “A survey of motion-based multitarget tracking methods”, *Progress In Electromagnetics Research B*, vol. 62, pp. 195–223, 2015.

[98] Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association*, ser. Mathematics in Science and Engineering Series. Academic Press, 1988.

[99] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*, ser. Artech House radar library. Artech House, 1999.

[100] R. L. Streit and T. E. Luginbuhl, “Maximum likelihood method for probabilistic multi-hypothesis tracking”, in *SPIE’s International Symposium on Optical Engineering and Photonics in Aerospace Sensing*, 1994, pp. 394–405.

[101] Y. Bar-Shalom, “Multitarget-multisensor tracking: advanced applications”, *Norwood, MA, Artech House, 1990, 391 p.*, vol. 1, 1990.

[102] S. Oh, S. Russell, and S. Sastry, “Markov chain monte carlo data association for multi-target tracking”, *IEEE Transactions on Automatic Control*, vol. 54, no. 3, pp. 481–497, 2009.

[103] R. P. Mahler, “Multitarget bayes filtering via first-order multitarget moments”, *IEEE Transactions on Aerospace and Electronic systems*, vol. 39, no. 4, pp. 1152–1178, 2003.

[104] M. Adams, B. N. Vo, R. Mahler, and J. Mullane, “Slam gets a phd: new concepts in map estimation”, *IEEE Robotics Automation Magazine*, vol. 21, no. 2, pp. 26–37, 2014.

[105] D. Moratuwage, B. N. Vo, and D. Wang, “A hierarchical approach to the multi-vehicle slam problem”, in *2012 15th International Conference on Information Fusion*, 2012, pp. 1119–1125.

[106] P. Dames and V. Kumar, “Cooperative multi-target localization with noisy sensors”, in *International International on Robotics and Automation*, 2013, pp. 1877–1883.

[107] P. Dames, “Distributed multi-target search and tracking using the phd filter”, in *International Symposium on Multi-Robot and Multi-Agent Systems*, 2017, pp. 1–8.

[108] K. R. Sapkota, S. Roelofsen, A. Rozantsev, V. Lepetit, D. Gillet, P. Fua, and A. Martinoli, “Vision-based unmanned aerial vehicle detection and tracking for sense and avoid systems”, in *IEEE International Conference on Intelligent Robots and Systems*, 2016, pp. 1556–1561.

[109] J. Melo and A. Matos, “A phd filter for tracking multiple auvs”, in *2014 Oceans - St. John’s*, 2014, pp. 1–8.

[110] H. W. Kuhn, “The hungarian method for the assignment problem”, *Naval Research Logistics Quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[111] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, “A consistent metric for performance evaluation of multi-object filters”, *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3447–3457, 2008.

[112] R Ventura and A Ahmad, “Towards optimal robot navigation in urban homes”, in *RoboCup 2014: Robot World Cup XVIII*, ser. Lecture Notes in Computer Science ; 8992, Springer, 2015, pp. 318–331.

[113] W. Kasper and M. E. Streit, *Institutional Economics*, ser. Books 1582. Edward Elgar Publishing, 1999.

[114] G. Boella, L. Torre, and H. Verhagen, “Introduction to normative multiagent systems”, *Computational & Mathem. Organization Theory*, vol. 12, no. 2-3, pp. 71–79, 2006.

[115] G. M. Hodgson, “What are institutions?”, *Journal of Economic Issues*, vol. 40, no. 1, pp. 1–25, 2006.

[116] E. G. Furubotn and R. Richter, *Institutions and Economic Theory: The Contribution of the New Institutional Economics*. University of Michigan Press, 2005.

[117] M. A. Mahmoud, M. S. Ahmad, M. Z. Mohd Yusoff, and A. Mustapha, “A review of norms and normative multiagent systems”, *The Scientific World Journal*, vol. 2014, 2014.

[118] E. Ostrom, *Understanding Institutional Diversity*. Princeton University Press, 2005.

[119] V. Dignum and J. Padget, “Multiagent organizations”, *Multiagent systems*, vol. 2, pp. 51–98, 2013.

[120] F. Dechesne and V. Dignum, “No smoking here: compliance differences between legal and social norms”, in *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 3*, ser. AAMAS ’11, International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 1205–1206.

[121] M. Hechter and K.-D. Opp, *Social norms*. Russell Sage Foundation, 2001.

[122] D. C. North, “Institutions”, *Journal of Economic Perspectives*, vol. 5, no. 1, pp. 97–112, 1991.

[123] F. A. Hayek, “Notes on the evolution of systems of rules of conduct”, *Studies in philosophy, politics and economics*, pp. 66–81, 1967.

[124] M. Polanyi, *The tacit dimension*. University of Chicago press, 2009.

[125] P. Noriega, “Fencing the open fields: empirical concerns on electronic institutions (invited paper)”, in *Proceedings of the 2005 International Conference on Agents, Norms and Institutions for Regulated Multi-Agent Systems*, ser. AAMAS’05, Springer-Verlag, 2006, pp. 81–98.

[126] G. Boella and L. van der Torre, “An architecture of a normative system: counts-as conditionals, obligations and permissions”, in *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS ’06, ACM, 2006, pp. 229–231.

[127] J. Carmo and A. J. Jones, “Deontic logic and contrary-to-duties”, in *Handbook of philosophical logic*, Springer, 2002, pp. 265–343.

[128] C. D. Hollander and A. S. Wu, “The current state of normative agent-based systems”, *Journal of Artificial Societies and Social Simulation*, vol. 14, no. 2, p. 6, 2011.

[129] B. T. R. Savarimuthu and S. Cranefield, “Norm creation, spreading and emergence: a survey of simulation models of norms in multi-agent systems”, *Multiagent and Grid Systems*, vol. 7, no. 1, pp. 21–54, 2011.

[130] D. Grossi and F. Dignum, “From abstract to concrete norms in agent institutions”, in *Formal Approaches to Agent-Based Systems*, Springer Berlin Heidelberg, 2005, pp. 12–29.

[131] R. Malyankar, “A pattern template for intelligent agent systems”, in *Agents’ 99 Workshop on Agent-Based Decision Support for Managing the Internet-Enabled Supply Chain*, 1999.

[132] P. Noriega, “Agent mediated auctions: the fishmarket metaphor”, PhD thesis, Universitat Autonoma de Barcelona, 1997.

[133] M. Esteva, “Electronic institutions: from specification to development”, PhD thesis, Universitat Politecnica de Catalunya, 2003.

[134] A. García-Camino, J. A. Rodríguez-Aguilar, C. Sierra, and W. Vasconcelos, “Constraint rule-based programming of norms for electronic institutions”, *Autonomous agents and multi-agent systems*, vol. 18, no. 1, pp. 186–217, 2009.

[135] J. Vázquez-Salceda, “The role of norms and electronic institutions in multi-agent systems applied to complex domains. the harmonia framework”, *Ai Communications*, vol. 16, no. 3, pp. 209–212, 2003.

[136] M. Dignum, “A model for organizational interaction: based on agents, founded in logic”, PhD thesis, Utrecht University, 2004.

[137] J. Vázquez-Salceda, V. Dignum, and F. Dignum, “Organizing multiagent systems”, *Autonomous Agents and Multi-Agent Systems*, vol. 11, no. 3, pp. 307–360, 2005.

[138] M. Esteva, D. De la Cruz, and C. Sierra, “ISLANDER: an electronic institutions editor”, in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 2002, pp. 1045–1052.

[139] K. M. Carley and L. Gasser, “Multiagent systems”, in, MIT Press, 1999, ch. Computational Organization Theory, pp. 299–330.

[140] F. Zambonelli, N. R. Jennings, and M. Wooldridge, “Developing multiagent systems: the gaia methodology”, *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 12, no. 3, pp. 317–370, 2003.

[141] M. Hannoun, O. Boissier, J. S. Sichman, and C. Sayettat, "MOISE: an organizational model for multi-agent systems", in *Advances in Artificial Intelligence*, 2000, pp. 156–165.

[142] M. Amiguet, J.-P. Müller, J.-A. Baez-Barranco, and A. Nagy, "The moca platform", in *International Workshop on Multi-Agent Systems and Agent-Based Simulation*, Springer, 2002, pp. 70–88.

[143] J.-J. C. Meyer, "Deontic logic: a concise overview", in *Deontic Logic in Computer Science: Normative System Specification*, Wiley, 1993, pp. 3–16.

[144] H. Aldewereld, D. Grossi, J. Vázquez-Salceda, and F. Dignum, "Designing normative behaviour via landmarks", in *Coordination, Organizations, Institutions, and Norms in Multi-Agent Systems*, Springer Berlin Heidelberg, 2006, pp. 157–169.

[145] F. Dignum, "Abstract norms and electronic institutions", in *Proceedings of International Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA'02)*, 2002, pp. 93–104.

[146] P. Silva, J. N. Pereira, and P. U. Lima, "Institutional robotics", *International Journal of Social Robotics*, vol. 7, no. 5, pp. 825–840, 2015.

[147] J. N. Pereira, "Advancing social interactions among robots: an institutional economics-based approach to distributed robotic systems", PhD thesis, EPFL, 2014.

[148] J. N. Pereira, P. Silva, P. U. Lima, and A. Martinoli, "Formalization, implementation, and modeling of institutional controllers for distributed robotic systems", *Artificial Life*, vol. 20, no. 1, pp. 127–141, 2013.

[149] S. Tomic, F. Pecora, and A. Saffiotti, "Too cool for school - adding social constraints in human aware planning", in *Proceedings of the 9th International Workshop on Cognitive Robotics, CogRob (ECAI-2014)*, 2014.

[150] S. Alili, R. Alami, and V. Montreuil, "A task planner for an autonomous social robot", in *Distributed Autonomous Robotic Systems 8*. Springer Berlin Heidelberg, 2009, pp. 335–344.

[151] F. M. Carlucci, L. Nardi, L. Iocchi, and D. Nardi, "Explicit representation of social norms for social robots", in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015, pp. 4191–4196.

[152] E. Ferrera, J. Capitán, A. R. Castano, and P. J. Marrón, "Decentralized safe conflict resolution for multiple robots in dense scenarios", *Robotics and Autonomous Systems*, vol. 91, pp. 179 –193, 2017.

[153] B. Okal and K. O. Arras, "Formalizing normative robot behavior", in *8th International Conference on Social Robotics*, 2016, pp. 62–71.

[154] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon, "A human aware mobile robot motion planner", *IEEE Trans on Robotics*, vol. 23, no. 5, pp. 874–883, 2007.

[155] R. Triebel, K. Arras, R. Alami, L. Beyer, S. Breuers, R. Chatila, M. Chetouani, D. Cremers, V. Evers, M. Fiore, H. Hung, O. A. I. Ramírez, M. Joosse, H. Khambhaita, T. Kucner, B. Leibe, A. J. Lilienthal, T. Linder, M. Lohse, M. Magnusson, B. Okal, L. Palmieri, U. Rafi, M. van Rooij, and L. Zhang, "Spencer: a socially aware service robot for passenger guidance and help in busy airports", in *Field and Service Robotics: Results of the 10th International Conference*. Springer International Publishing, 2016, pp. 607–622.

[156] C. P. Lam, C. T. Chou, K. H. Chiang, and L. C. Fu, "Human-centered robot navigation - towards a harmoniously human-robot coexisting environment", *IEEE Trans on Robotics*, vol. 27, no. 1, pp. 99–112, 2011.

[157] S. G. J.V. Gomez Nikolaos Mavridis, "Social path planning: generic human-robot interaction framework for robotic navigation tasks", in *IROS workshop on Cognitive Robotics Systems: Replicating Human Actions and Activities*, 2013.

[158] É. Pacherie, "The phenomenology of joint action: self-agency vs. joint-agency", *Joint Attention: New Developments*, 2012.

[159] G. Knoblich, S. Butterfill, and N. Sebanz, "Chapter three - psychological research on joint action: theory and data", in *Advances in Research and Theory*, ser. Psychology of Learning and Motivation, vol. 54, Academic Press, 2011, pp. 59 –101.

[160] A. Reckwitz, "Toward a theory of social practices: a development in culturalist theorizing", *European Journal of Social Theory*, vol. 5, no. 2, pp. 243–263, 2002.

[161] A. Clodic, E. Pacherie, R. Alami, and R. Chatila, "Key elements for human-robot joint action", in *Sociality and Normativity for Robots: Philosophical Inquiries into Human-Robot Interactions*. Springer International Publishing, 2017, pp. 159–177.

[162] A. Clodic, J. Vázquez-Salceda, F. Dignum, S. Mascarenhas, V. Dignum, A. Augello, M. Gentile, and R. Alami, *On the Pertinence of Social Practices for Social Robotics*. IOS Press. Envisioning Robots in Society – Power, Politics, and Public Space, 2018, pp. 36–74.

[163] U. Köckemann, "Constraint-based methods for human-aware planning", PhD thesis, Örebro University, School of Science and Technology, 2016, p. 194.

[164] S. Lemaignan, M. Warnier, E. A. Sisbot, A. Clodic, and R. Alami, "Artificial cognition for social human–robot interaction: an implementation", *Artificial Intelligence*, vol. 247, pp. 45 –69, 2017, Special Issue on AI and Robotics.

[165] S. Devin, A. Clodic, and R. Alami, "About decisions during human-robot shared plan achievement: who should act and how?", in *Social Robotics*, Springer International Publishing, 2017, pp. 453–463.

[166] E. Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.

[167] M. Casper, "A definition of "social environment"", *American Journal of Public Health*, vol. 91, p. 465, 2001.

[168] C. I. Mavrogiannis, W. B. Thomason, and R. A. Knepper, "Social momentum: a framework for legible navigation in dynamic multi-agent environments", in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '18, 2018, 361–369.

[169] H. Khambhaita and R. Alami, "Viewing robot navigation in human environment as a cooperative activity", in *Robotics Research*, Springer International Publishing, 2020, pp. 285–300.

[170] D. Feil-Seifer and M. MatariĆ, "People-aware navigation for goal-oriented behavior involving a human partner", in *2011 IEEE International Conference on Development and Learning (ICDL)*, vol. 2, 2011, pp. 1–6.

[171] R. Kirby, "Social robot navigation", PhD thesis, Carnegie Mellon University, 2010.

[172] A. G. Zulueta, "Cooperative social robots: accompanying, guiding and interacting with people", PhD thesis, Universitat Politècnica de Catalunya, 2013.

[173] G. Norman, "Likert scales, levels of measurement and the "laws" of statistics", *Advances in health sciences education*, vol. 15, no. 5, pp. 625–632, 2010.

[174] J. E. Young, Y. Kamiyama, J. Reichenbach, T. Igarashi, and E. Sharlin, "How to walk a robot: a dog-leash human-robot interface", in *2011 RO-MAN*, IEEE, 2011, pp. 376–382.

[175] R. Gockley, J. Forlizzi, and R. Simmons, "Natural person-following behavior for social robots", in *2nd ACM/IEEE Intternational Conference on Human-Robot Interaction*, 2007, pp. 17–24.

[176] M. Teixidó, T. Pallejà, D. Font, M. Tresanchez, J. Javier Moreno, and J. Palacìn, "Two-dimensional radial laser scanning for circular marker detection and external mobile robot tracking", *Sensors (Basel, Switzerland)*, vol. 12, pp. 16 482–97, 2012.

[177] G. P. Huang, N. Trawny, A. I. Mourikis, and S. I. Roumeliotis, "Observability-based consistent ekf estimators for multi-robot cooperative localization", *Autonomous Robots*, vol. 30, no. 1, pp. 99–122, 2011.

[178] F. He, Z. Du, X. Liu, and Y. Ta, "Laser range finder based moving object tracking and avoidance in dynamic environment", in *The 2010 IEEE International Conference on Information and Automation*, 2010, pp. 2357–2362.

[179] J. Pugh, X. Raemy, C. Favre, R. Falconi, and A. Martinoli, "A fast onboard relative positioning module for multirobot systems", *IEEE/ASME Transactions on Mechatronics*, vol. 14, no. 2, pp. 151–162, 2009.

[180] L. L. Hokuyo Automatic CO. (2014). Scanning rangefinder urg-04lx-ug01 specifications, [Online]. Available: https://www.hokuyo-aut.jp/search/single.php?serial=166 (visited on 08/06/2019).

[181] Y. Okubo, C. Ye, and J. Borenstein, "Characterization of the hokuyo urg-04lx laser rangefinder for mobile robot obstacle negotiation", in *Unmanned Systems Technology XI*, ser. SPIE Proceedings, vol. 7332, 2009, p. 733 212.

[182] J. J. Moré, “The levenberg-marquardt algorithm: implementation and theory”, in *Numerical Analysis*, Springer Berlin Heidelberg, 1978, pp. 105–116.

[183] J. Messias, R. Ventura, P. Lima, and J. Sequeira, “Ros in the monarch project: a case study in networked robot systems”, in *Robot Operating System (ROS): The Complete Reference (Volume 1)*. Springer International Publishing, 2016, pp. 375–395.

[184] J. Alonso-Mora, R. Knepper, R. Siegwart, and D. Rus, “Local motion planning for collaborative multi-robot manipulation of deformable objects”, in *2015 IEEE International Conference on Robotics and Automation*, 2015, pp. 5495–5502.

[185] M. Saska, V. Krátký, V. Spurný, and T. Báča, “Documentation of dark areas of large historical buildings by a formation of unmanned aerial vehicles using model predictive control”, in *2017 22nd IEEE International Conference on Emerging Technologies and Factory Automation*, 2017, pp. 1–8.

[186] T. Nägeli, L. Meier, A. Domahidi, J. Alonso-Mora, and O. Hilliges, “Real-time planning for automated multi-view drone cinematography”, *ACM Trans. Graph.*, vol. 36, no. 4, 132:1–132:10, 2017.

[187] A. Ahmad, G. Lawless, and P. Lima, “An online scalable approach to unified multirobot cooperative localization and object tracking”, *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1184–1199, 2017.

[188] D. Saldaña, R. Assunção, M. A. Hsieh, M. F. M. Campos, and V. Kumar, “Cooperative prediction of time-varying boundaries with a team of robots”, in *2017 International Symposium on Multi-Robot and Multi-Agent Systems*, 2017, pp. 9–16.

[189] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, “A density-based algorithm for discovering clusters in large spatial databases with noise”, in *Proceedings of the Second International International on Knowledge Discovery and Data Mining*, AAAI Press, 1996, pp. 226–231.

[190] J. N. Pereira, P. Silva, P. U. Lima, and A. Martinoli, “Formalizing institutions as executable petri nets for distributed robotic systems”, in *European Conference on Artificial Life*, 2011.

[191] C. Hollander and A. Wu, “The current state of normative agent-based systems”, *Journal of Artificial Societies and Social Simulation*, vol. 14, no. 2, p. 6, 2011.

[192] G. M. Hodgson, “What is the essence of institutional economics?”, *Journal of Economic Issues*, vol. 34, no. 2, pp. 317–329, 2000.

[193] J. N. Pereira, A. L. Christensen, P. Silva, and P. U. Lima, “Coordination through institutional roles in robot collectives”, in *Proceedings of the 9th International International on Autonomous Agents and Multiagent Systems*, 2010, pp. 1507–1508.

[194] E. Ostrom, “An agenda for the study of institutions”, *Public Choice*, vol. 48, no. 1, pp. 3–25, 1986.

[195] R. David and H. Alla, *Discrete, continuous, and hybrid Petri nets*. Springer Science & Business Media, 2005.

[196] D. Gouaillier, V. Hugel, P. Blazevic, C. Kilner, J. Monceaux, P. Lafourcade, B. Marnier, J. Serre, and B. Maisonnier, “The NAO humanoid: a combination of performance and affordability”, *CoRR abs/0807.3223*, 2008.

[197] G. Therborn, “Back to norms! On the scope and dynamics of norms and normative action”, *Current Sociology*, vol. 50, no. 6, pp. 863–880, 2002.

[198] G. Weiss, *Multiagent Systems*. The MIT Press, 2013.

[199] N. Alechina, N. Bassiliades, M. Dastani, M. De Vos, B. Logan, S. Mera, A. Morris-Martin, and F. Schapachnik, “Computational models for normative multi-agent systems”, *Dagstuhl Follow-Ups*, vol. 4, pp. 71–92, 2013.

[200] A. Ricci, M. Viroli, and A. Omicini, “Give agents their artifacts: the A&A approach for engineering working environments in MAS”, in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 2007, p. 150.

[201] E. A. Martinez-Garcia, A. Ohya, and S. Yuta, “A multi-robot system architecture for trajectory control of groups of people”, in *Cutting Edge Robotics*, IntechOpen, 2005, ch. 28.

[202] A. Bogdanovych, S. Simoff, and M. Esteva, “Virtual institutions: normative environments facilitating imitation learning in virtual agents”, in *Intelligent Virtual Agents*, Springer Berlin Heidelberg, 2008, pp. 456–464.

[203] Y. Shoham and M. Tennenholtz, “On social laws for artificial agent societies: off-line design”, *Artificial Intelligence*, vol. 73, no. 1, pp. 231 –252, 1995, Computational Research on Interaction and Agency, Part 2.

[204] E. Martinson, “Acoustical awareness for intelligent robotic action”, PhD thesis, Georgia Institute of Technology, 2007.

[205] N. P. Cuntoor, R. Collins, and A. J. Hoogs, “Human-robot teamwork using activity recognition and human instruction”, in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 459–465.

[206] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, “Market-based multirobot coordination: a survey and analysis”, *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, 2006.

[207] C. R. Kube and H. Zhang, “Collective robotics: from social insects to robots”, *Adaptive behavior*, vol. 2, no. 2, pp. 189–218, 1993.

[208] V. Villani, L. Sabattini, C. Secchi, and C. Fantuzzi, “A framework for affect-based natural human-robot interaction”, in *27th IEEE International Symposium Robot and Human Interactive Communication (RO-MAN)*, 2018, pp. 1038–1044.

Curriculum Vitae

Alicja Roelofsen

Education

2014-2020	Ph.D. in IST-EPFL Joint Doctoral Initiative *Ph.D. program in Robotics, Control and Intelligent Systems,* École Polytechnique Fédérale de Lausanne (EPFL), Switzerland *Ph.D. program in Electrical and Computer Engineering,* Instituto Superior Técnico (IST), Portugal
2010-2014	M.Eng in Aerospace Engineering *University of Sheffield, United Kingdom*

Experience

2013	Internship at the Institute of Simulation and Software Technology *German Aerospace Center (DLR), Braunschweig, Germany*
2012	Internship within the SAMULET project with the BAE Systems *Advanced Manufacturing Research Centre (AMRC), Sheffield, UK*
2011	Engineering Internship in Quality Assistance *European Brakes and Chassis Components (EBCC), Wrocław, Poland*

Honors and Awards

2014	Royal Aeronautical Society University Prize *for best aeronautical engineering graduate of the academic year*
2012	Sir Basil Blackwell Prize *for academic performance*
2011	Hatfield Memorial Prize *for academic performance*
2011	Sir Harold West Award *for academic and personal promise*

Publications

Journal Articles

1. A. Wasik, P. U. Lima, and A. Martinoli, "A robust localization system for multi-robot formations based on an extension of a Gaussian mixture probability hypothesis density filter", *Autonomous Robots*, vol. 44, 395–414, 2019

Refereed Conference Proceedings

1. A. Wasik, S. Tomic, A. Saffiotti, F. Pecora, A. Martinoli, and P. U. Lima, "Towards norm realization in institutions mediating human-robot societies", *2018 IEEE/RSJ International Conference On Intelligent Robots And Systems (IROS)*, IEEE International Conference on Intelligent Robots and Systems, pp. 297–304, 2018
2. S. Tomic, A. Wasik, P. U. Lima, A. Martinoli, F. Pecora, and A. Safiotti, "Towards institutions for mixed human-robot societies", in *Intternational Joint Conference on Autonomous Agents and Multiagent Systems*, 2018, pp. 2216–2217
3. A. Wasik, A. Martinoli, and P. U. Lima, "A robust relative positioning system for multi-robot formations leveraging an extended GM-PHD filter", *Proceedings of the First International Symposium on Multi-Robot and Multi-Agent Systems*, pp. 71–77, 2017
4. A. Wasik, A. Martinoli, and P. U. Lima, "An institutional robotics approach to the design of socially aware multi-robot behaviors", *Proceedings of the RO-MAN 2017 Workshop on Towards Intelligent Social Robots: Social Cognitive Systems in Smart Environments*, pp. 2–7, 2017
5. A. Wasik, J. N. Pereira, R. Ventura, P. U. Lima, and A. Martinoli, "Graph-based distributed control for adaptive multi-robot patrolling using local formation transformation", in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 1721–1728
6. A. Wasik, R. Ventura, J. N. Pereira, P. U. Lima, and A. Martinoli, "Lidar-based relative position estimation and tracking for multi-robot systems", in *Robot 2015: Second Iberian Robotics Conference*, Springer International Publishing, 2016, pp. 3–16

Project Supervision

1. Michiaki HIRAYAMA, Internship (2019)
   Realization and Validation of the FI-GM-PHD Filter With Real Multi-Robot Formations
2. Lucas BURGET, Semester Project (2018)
   Normative Multi-Robot Navigation in Crowds
3. Emil BRYNGELSSON, Master Thesis (2015)
   Distributed Multi-Robot Coverage in Realistic Environments
4. Michael SPAHR, Semester Project (2015)
   Distributed Formations With Non-Cooperative Human Agents
5. Arnaud WALD, Semester Project (2015)
   Social Awareness in Multi-Robot Systems - An Institutional Approach to Robotic Formations

Languages

Polish	native
English	fluent
German	intermediate
French	basic

Personal Details

Date of birth:	25 July 1991
Place of birth:	Wrocław, Poland
Email:	alicja.roelofsen@gmail.com