



Combining proactive planning and situation analysis for human-aware robot navigation

Phani Teja Singamaneni

► To cite this version:

Phani Teja Singamaneni. Combining proactive planning and situation analysis for human-aware robot navigation. Robotics [cs.RO]. Université Paul Sabatier - Toulouse III, 2022. English. NNT : 2022TOU30226 . tel-04006782v2

HAL Id: tel-04006782

<https://theses.hal.science/tel-04006782v2>

Submitted on 27 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par : *l'Université Toulouse 3 Paul Sabatier (UT3 Paul Sabatier)*

Présentée et soutenue le 14/12/2022 par :

Phani Teja SINGAMANENI

**Combining Proactive Planning and Situation Analysis for
Human-Aware Robot Navigation**

JURY

MARIE BABEL

LUIS MERINO

RACHID ALAMI

SIMON LACROIX

ANNE SPALANZANI

AARON STEINFELD

THIERRY SIMÉON

Professeur

Professeur Associé

Directeur de Recherche

Directeur de Recherche

Professeur

Professeur

Directeur de Recherche

Rapporteur

Rapporteur

Directeur de Thèse

Président du Jury

Membre du Jury

Membre du Jury

Membre du Jury

École doctorale et spécialité :

MITT : Informatique et Télécommunications

Unité de Recherche :

LAAS-CNRS

Directeur de Thèse :

Rachid ALAMI

Rapporteurs :

Marie BABEL et Luis MERINO

Acknowledgments

I still remember the first time I landed in Toulouse, unprepared for the winter and arriving at LAAS the next day without knowing even a bit of French. From that day to now, it has been an incredible and challenging journey of over three years. It would not have been possible to be here without the support and help of my friends, family and colleagues.

I am very grateful to my family, especially my mother and sister, for believing and supporting me in all possible ways, even when times were hard. I would also like to thank my father and uncle, as I would not have been here without them. I know that you encouraged me to keep moving forward even when it was difficult for you to understand what I was doing, and you will continue to do so in any circumstances.

The most important part of preparing a PhD thesis is having an advisor who believes in you and gives you an opportunity to explore the field of your interest. I am very thankful to my thesis director Rachid Alami for all the ideas, discussions and guidance that led me towards a fruitful thesis. Thanks for answering all my last-minute emails, giving some very useful feedback on my work and providing me with more opportunities to excel further in the field.

Thanks to Anthony Favier for assisting me in the development of this thesis by providing tools and feedback to improve my framework. Thank you for all the help while writing the papers, taking up my last-minute calls and demands, and the constructive criticism of my results to make them better.

I like to thank Marie Babel and Luis Merino for accepting to be reporters for my thesis. Thank you for taking the time out for providing me with your valuable feedback. I thank Simon Lacroix, the president of the Jury, and the members of the Jury, Anne Spalanzani, Aaron Steinfeld and Thierry Siméon, for taking their time out to be a part of my PhD defence.

Although there were tough times, the journey has been a fun ride, thanks to all my teammates at LAAS. Thank you, Amandine, for helping me out in settling when I first moved here and bearing with all my queries over the course of time. Without Guilhem, I would have been lost in the beginning stages of my PhD. Thanks for helping me to set up things and for answering all the questions that followed. I am thankful to Guillaume and Yoan for all the ideas and discussions that helped me progress. I will remember all those after-lunch games and interesting coffee break discussions with my colleagues Yannick, Kathleen, Rafael, Amélie, Phillipe, Shashank, Antoine, Jérémy, Jérôme and Simon.

I would like to thank Aurélie and Matthieu for all their help with the robots and for setting up things to work properly on the project. Thanks to all the administration of LAAS that helped me with the documents and missions from time to time.

I am glad that I have met some interesting people at LAAS to discuss amazing things over coffee and beer. Thanks to Rohan, Melya, Sai, Abhishek, Aarushee, Vidushi, Suresh and Amit for all that. Thank you, Suresh and Kalyani, for sup-

porting me all through this journey and being there like a family when I needed something. I know that I could always crash at your place for some amazing parties and conversations over beer.

Not long after starting my PhD, Covid-19 hit the world, and we were all locked up in our homes for months. Staying alone for such long times in a small room is not an easy task, and thanks to all my friends who made it possible to get through the pandemic. Thanks to Yami for all those long conversations and supermarket visits which made the lockdown look very easy to handle. Online gaming was one more thing that helped me during the pandemic. Playing with friends is always fun, and thanks to Saif, Rachana, Samyutka, Ravi, Sandily, Puneeth and Poorna for all games of ‘Codename’ and ‘AmongUs’. Suresh Santhanam, Bharat Kota and Srikar Allu, thanks for all those ‘Catan’ games which were more of long conversations with little gaming.

Some friends are there for life, and I am very grateful to have such childhood friends. SaiKiran Samudrala, Chanakya and Varun, we have all grown together, supporting each other continuously and acknowledging each other’s flaws. All during the pandemic and after, you have been my support system with whom I can discuss anything without being judged. Thanks to all of you, I survived the lows of this journey without panicking.

Finally, I cannot forget the bizarre talks and some very helpful suggestions from Yashwanth. You had been patient with all my stupid talks and helped me make somewhat better decisions. Thanks for walking randomly into my hostel room and deciding to be a helpful friend. Thanks to Mahtab and Jerin for all that psychoanalysis, motivational talks about the gym and career advice. They definitely helped me understand things better and make the right decisions throughout this journey.

My old friends from RRC, Sri Harsha and Vignesh, are all on the same career paths and thank you for sharing your experiences and all the ideation sessions. I know this is just the beginning, and we will have more talks and collaborations in the future.

Abstract

The navigation of a mobile robot in human environments is referred to as Human-Aware Navigation (HAN) or sometimes as Social (Robot) Navigation. We present the evolution of HAN over the years and the challenges associated with it before we proceed with the explanation of our approach. With a growing interest in the field, new frameworks are required to address this problem better, and as a contribution, we propose a proactive multi-context HAN planning system.

The proposed proactive planning framework for HAN is based on the Human-Aware Timed Elastic Band (HATEB) approach. This framework adds a decision-making loop at the trajectory planning level and switches between different planning modes based on the assessed situation. Using this framework, we propose a ROS based HAN Stack called Cooperative Human-Aware Navigation (CoHAN) planner. CoHAN is designed to address multi-context robot navigation among static and dynamic humans with updated costmaps and new human-robot social constraints. These optimization constraints aim to make the robot’s motion legible and avoid surprise appearances from behind. CoHAN offers four different path prediction mechanisms for humans and three different planning modes, along with a new recovery mode. The qualitative analysis performed on simulated human-robot scenarios highlights the advantages of proactive and situation based planning. They also show how the newly added social constraints improve the robot’s legibility. Quantitative comparisons with other HAN planners show that CoHAN can solve more intricate scenarios in an acceptable manner.

In the next part, CoHAN is extended further to proactively address sudden human appearance from occluded regions on the map. We propose the concept of ‘invisible humans’ to estimate the locations of such human emergences. These estimations are added to CoHAN via another optimization constraint, and a new mode of planning to pass through doors and narrow passages cautiously is introduced. The experiments in simulation and on the real robot demonstrate the benefits of ‘invisible humans’ in HAN and show that the robot maintains a greater distance from occluded humans compared to other approaches.

The current metrics based on proxemic zone violations cannot do justice in evaluating intricate scenarios. Therefore, we propose a set of metrics taking into consideration the velocities and the visibility of the human, which may be pertinent to various human-robot navigation contexts. This part of the thesis presents the mathematical formulation of these metrics, followed by the evaluation of four different human-robot navigation contexts using them. These metrics, combined with the velocity profiles and paths, clearly distinguish a HAN planner from a simple navigation planner. We conclude this thesis with a discussion on the contribution followed by the challenges and the future perspectives of HAN.

Résumé

La navigation d'un robot mobile dans un environnement humain est appelée Human-Aware Navigation (HAN) ou parfois Social (Robot) Navigation. Nous présentons l'évolution de la HAN au fil des ans et les défis qui lui sont associés avant de procéder à l'explication de notre approche. Avec un intérêt croissant dans le domaine, de nouveaux schéma sont nécessaires pour mieux aborder ce problème, et comme contribution, nous proposons un système de planification proactive multi-contexte pour HAN.

Le schéma de planification proactive proposé pour le HAN est basé sur l'approche Human-Aware Timed Elastic Band (HATEB). Ce schéma ajoute une boucle de décision au niveau de la planification de la trajectoire et passe d'un mode de planification à l'autre en fonction de la situation évaluée. Sa cette base, nous proposons une pile HAN basée sur ROS appelée planificateur Co-operative Human-Aware Navigation (CoHAN). CoHAN est conçu pour aborder la navigation de robot multi-contexte parmi des humains statiques et en mouvement avec des cartes de coûts mises à jour et de nouvelles contraintes sociales humain-robot. Ces contraintes d'optimisation visent à rendre le mouvement du robot lisible et à éviter les apparitions surprises par derrière. CoHAN propose quatre mécanismes différents de prédiction de trajectoire pour les humains et trois modes de planification différents, ainsi qu'un nouveau mode de récupération. L'analyse qualitative effectuée sur des scénarios simulés humain-robot met en évidence les avantages de la planification proactive et basée sur la situation. Elle montre également comment les contraintes sociales nouvellement ajoutées améliorent la lisibilité du robot. Les comparaisons quantitatives avec d'autres planificateurs HAN montrent que CoHAN peut résoudre des scénarios plus complexes de manière acceptable.

Dans la partie suivante, CoHAN est étendu pour traiter de manière proactive l'apparition soudaine d'humains dans des régions masquées de la carte. Nous proposons le concept d'"humains invisibles" pour estimer l'emplacement de telles émergences humaines. Ces estimations sont ajoutées à CoHAN via une autre contrainte d'optimisation, et un nouveau mode de planification pour passer les portes et les passages étroits avec précaution est introduit. Les expériences en simulation et sur le robot réel démontrent les avantages des "humains invisibles" dans le HAN et montrent que le robot maintient une plus grande distance avec les humains occultés par rapport aux autres approches.

Les mesures basées sur les violations de zones proxémiques ne peuvent pas convenir à l'évaluation de scénarios complexes. Par conséquent, nous proposons un ensemble de métriques prenant en compte les vitesses et la visibilité de l'humain, qui peuvent être pertinentes dans différents contextes de navigation humain-robot. Cette partie de la thèse présente une formulation mathématique de ces métriques, suivie de l'évaluation de quatre différents contextes de navigation humain-robot en les utilisant. Ces métriques, combinées aux profils de vitesse et aux trajectoires,

distinguent clairement un planificateur HAN d'un simple planificateur de navigation. Nous concluons cette thèse par une discussion sur la contribution, suivie par les défis et les perspectives futures du HAN.

Contents

Introduction	1
1 Evolution of Human-Aware Navigation	7
1.1 Motion Planning and Navigation	8
1.1.1 The Motion Planning Problem	8
1.1.2 Mobile Robot Navigation	11
1.1.3 2D Navigation Stack in ROS	13
1.2 Human-Aware Navigation	14
1.2.1 State-of-Art over the years	14
1.2.2 Human-Aware Navigation at LAAS	22
1.2.3 Human-Aware Navigation as a part of HRI	24
1.2.4 General challenges in Human-Aware Navigation	25
1.2.5 Addressing the challenges	26
1.3 Background and Our Approach	28
1.3.1 Timed Elastic Band for Trajectory Planning	29
1.3.2 Human-Aware Timed Elastic Bands for Proactive Planning	31
1.3.3 Principles of Joint-Action in Human-Aware Navigation	33
1.3.4 Our Approach and Contribution	34
2 HATEB-2 for Legible Proactive Planning with Situation Analysis	37
2.1 Introduction	37
2.2 Related Work	38
2.3 Proactive Planning in HATEB	39
2.3.1 The Entanglement Problem	40
2.4 HATEB-2: Situation Analysis and Mode Shifting	41
2.4.1 Formulation and Implementation	41
2.4.2 Modes of Planning	41
2.4.3 Mode Shifting based on Situation Assessment	42
2.5 Improving the Legibility in HATEB-2	45
2.5.1 New Human-Aware Constraints	45
2.5.2 Modified Elastic Band	46
2.5.3 Better Human Predictions	47
2.6 Results in Simulation	48
2.6.1 Entanglement Resolution through Mode Shifting	48
2.6.2 Advantages of Proactive Planning	49
2.6.3 The Effect of New Human-Aware Constraints	51
2.6.4 Quantitative Comparison between HATEB and HATEB-2	51
2.7 Real-World Tests	53
2.8 Conclusion	55

3	A Tunable Human-Aware Navigation Planner for Multi-Context Navigation	57
3.1	Introduction	57
3.2	Related Work	59
3.3	Co-operative Human-Aware Navigation Planner	60
3.3.1	Architecture	60
3.3.2	Types of Visible Humans and Costmap Layers	62
3.3.3	Human Path Prediction Mechanisms	63
3.3.4	An updated HATEB local planner	64
3.4	Testing CoHAN under Different Conditions	69
3.4.1	Door Crossing with Static and Moving Humans	69
3.4.2	A Very Narrow Corridor Scene	71
3.4.3	Co-operative Navigation Scenarios	73
3.4.4	Robot in a Crowd	74
3.4.5	Comparison with Another Human-Aware Planner	75
3.5	Real-world Experiments with CoHAN	76
3.6	Multi-context HAN and CoHAN	78
3.6.1	Modality-based HAN Planning	78
3.6.2	Multi-Context HAN Planning using Mode Shifting	79
3.6.3	CoHAN in Multiple Human-Robot Navigation Contexts	79
3.7	Conclusion	80
4	Proactive Planning for Unseen Humans in the Environment	81
4.1	Introduction	81
4.2	Invisible Humans Detection	82
4.2.1	Corner Detection using Laser Contour	83
4.2.2	Estimation of Invisible Humans' Locations	84
4.3	Introducing Invisible Humans into CoHAN	87
4.3.1	The 'Invisible Humans Constraint'	87
4.3.2	Passage Detection and a New CoHAN Planning Mode	88
4.4	Results	89
4.4.1	The Effect of the Invisible Humans Constraint	89
4.4.2	Navigating in the Presence of Visible and Invisible Humans	93
4.4.3	Testing the Accuracy and Robustness	95
4.5	Real-World Tests	97
4.6	Discussion and Limitations	98
4.7	Conclusion	99
5	Evaluation of Human-Aware Navigation	101
5.1	Introduction	101
5.2	Existing Evaluation methods	102
5.2.1	Navigation Metrics	102
5.2.2	Naturalness Metrics	103
5.2.3	Discomfort Metrics	103

5.2.4	Surveys and Studies	104
5.3	Proposed Discomfort Metrics	105
5.3.1	Velocity based Metrics	105
5.3.2	Vision based Metrics	108
5.4	Analysis	111
5.4.1	Scenario 1: Cross	111
5.4.2	Scenario 2: Follow and Overtake	112
5.4.3	Scenario 3: Approach	113
5.4.4	Scenario 4: Appear	114
5.5	Discussion	116
6	Conclusions	117
A	Simulating Human Agents for Testing HAN	121
A.1	Planning and Control for Human Agents	121
A.1.1	Manual Control	121
A.1.2	Planning based Control	122
A.2	InHuS	123
A.2.1	Architecture	123
A.2.2	Supervisor and Geometric Planner	124
A.2.3	Human Behaviour Model	125
A.2.4	Logs and Metrics	125
A.2.5	Generating Different Behaviours	126
B	Effects of Social Constraints	129
B.1	TTCplus Constraint	129
B.1.1	Approach	129
B.1.2	Open Space Crossing	130
B.1.3	Corridor Crossing	131
B.2	Relative Velocity Constraint	131
B.2.1	Corridor Crossing	131
B.2.2	Open Space Crossing	132
B.3	Visibility Constraint	132
B.4	Updated Invisible Humans Constraint	133
B.4.1	Testing the Updated Constraint	134
	Bibliography	137

Acronyms

- CADRL** Collision Avoidance with Deep Reinforcement Learning. 13, 28
- CALU** Collision Avoidance with Localization Uncertainty. 13
- CoHAN** Cooperative Human-Aware Navigation. iii, 58, 60, 61, 62, 63, 64, 65, 67, 68, 71, 73, 74, 75, 76, 78, 79, 82, 87, 89, 91, 92, 93, 94, 95, 97, 110, 111, 112, 113, 114, 115, 116, 122, 124, 127
- CRL** Composite Reinforcement Learning. 13
- DWA** Dynamic Window Approach. 12, 15, 28
- EKF** Extended Kalman Filter. 11
- FRP** Freezing Robot Problem. 17, 18, 19, 28, 40, 42
- HAN** Human-Aware Navigation. iii, 1, 14, 38
- HATEB** Human-Aware Timed Elastic Band. iii, 24, 29, 31, 33, 37, 38, 39, 40, 41, 42, 45, 47, 48, 50, 51, 53, 55, 58, 60, 61, 62, 63, 64, 65, 67, 68, 69, 76, 79, 80
- HRI** Human-Robot Interaction. 1, 2, 7, 14, 18, 20, 22, 23, 24, 25, 28, 29, 33, 57, 79
- HRVO** Hybrid Reciprocal Velocity Obstacle. 18
- IRL** Inverse Reinforcement Learning. 17, 18, 20, 23
- LIDAR** Light Detection and Ranging. 12
- MPC** Model Predictive Control. 10, 13, 17, 19, 28
- ORCA** Optimal Reciprocal Collision Avoidance. 13, 27, 28, 121
- POMDP** Partially Observable Markov Decision Process. 17, 19, 28, 39, 60, 78
- PRM** Probabilistic Road Map. 9, 12
- RGL** Relational Graph Learning. 18
- RL** Reinforcement Learning. 18, 19, 20, 28
- ROS** Robot Operating System. iii, 8, 13, 29, 34, 60, 121, 122, 124, 127

RRT Rapidly exploring Random Tree. 9, 17, 18, 20, 23, 28

RVO Reciprocal Velocity Obstacle. 13, 28

SA-CADRL Socially Aware Collision Avoidance with Deep Reinforcement Learning. 13, 18, 28

SFM Social Force Model. 15, 18, 26, 27, 28, 38, 39, 121

SLAM Simultaneous Localization and Mapping. 11, 12, 16

TEB Timed Elastic Band. 13, 29, 30, 31, 33, 41

TTC Time-to-Collision. 33, 45, 46, 51, 53, 55, 68, 104, 107

Introduction

Several methodologies were proposed by roboticists for mobile robot navigation, and they are promising as long as humans are not considered. Treating humans as dynamic obstacles can solve the robot navigation problem but do not offer an acceptable and comfortable solution to the humans co-existing in the environment. As humans are social beings and have certain notions and expectations about the environment around them, they expect the robot to respect their personal space and preferences while navigating. Consequently, a new field of navigation called ‘*Human-Aware (Robot) Navigation*’ (or *Social Robot Navigation*) has emerged, combining studies on humans and robot navigation.

Human-aware navigation originated in the field of Human-Robot Interaction (HRI) rather than robot navigation, and hence it shares some properties with HRI. It was not an independent field of research until recent years and used to be a part of the overall task that the robot needs to do in order to assist or coordinate with a human or group of humans. As time progressed, robots have become cheaper and are now being used in many indoor [Guldenring 2020], as well as outdoor [Ferrer 2013b] settings, solely to move from one place to another to deliver things [Bogue 2016] or to accompany a person [Repiso 2017]. The advent of autonomous vehicles [Rasouli 2019] further soared the interest in this field. This thesis explores human-aware navigation and presents a cooperative framework for robot navigation that is built on the joint-action principles of HRI [Curioni 2019]. We also present some new ideas to improve robot navigation and propose some new metrics for evaluation.

Human-Aware Robot Navigation

Human-Aware Navigation (HAN) is a special case of robot navigation where path planning or trajectory planning or both integrate humans into planning to generate paths and/or trajectories for the robot such that it reduces the discomfort to the humans while navigating. This could mean that the navigation motion executed by the robot should be legible and acceptable to the humans sharing the environment with the robot. The term discomfort could be ambiguous and could refer to different things in different settings.

Situation assessment and proactive planning

When a robot navigates around humans, it is important for a HAN planning system to analyse the situation and take proper action that mitigates any deadlocks or the ‘freezing robot problems’ [Trautman 2010]. This requires the HAN system to have decision-making capabilities on top of legible motion generation. Once a situation is identified, it may be possible that the current version (or parameters) of the

navigation planning system cannot avoid the occurrence of a conflict. Therefore, it is required for a HAN to have different modalities and switch between them depending on the context.

Proactive planning could mean controlling and mitigating a situation at hand instead of responding to it when it happens. This allows the robot to respond quickly and minimizes the occurrence of the robot freezing problem. Moreover, this kind of planning complements situation analysis and lessens the burden on the decision-making system. Therefore, proactive planning offers a better framework to address HAN compared to reactive planning schemes like Social Force [Ferrer 2013b]. This thesis explores how proactive planning combined with situation analysis can benefit HAN planning and then develop a HAN system capable of handling multiple human-robot navigation contexts.

Thesis Contributions

This thesis has four main contributions, which are briefly explained below. The first three contributions can be seen as three different versions of a HAN planning system in chronological order, and the new version inherits almost all the properties of the previous version with some exceptions. All the versions of the proposed system treat HAN as a cooperative activity that obeys the following four principles of joint-action in HRI: **‘sharing a common perspective’**, **‘coordinating’**, **‘predicting others’ contributions’** and **‘communicating’**. The core idea of all these contributions is to offer a better solution to HAN planning, respecting the above principles and assuming humans as partners in navigation who can cooperate. The final contribution is a set of new metrics for HAN that may be more relevant to multiple human-robot navigation contexts than the ones based on proxemics.

Combining situation assessment with proactive planning in HAN

This is the first major contribution of this thesis that combines situation assessment with proactive planning. The idea of proactive planning in our HAN system is to actively plan for the robot and the other agents involved in the navigation, assuming a possible goal for each agent while controlling only the robot. The advantage of this kind of planning is that the planning system considers both the robot and the humans while planning for the robot, which makes the robot act proactively in many situations and avoid conflicts. In deadlock scenarios, this system elicits plans for all the agents, which, if followed, will resolve the deadlock. However, some of the deadlocks cannot be solved by this kind of proactive planning. So, we introduce a simple situation assessment to detect such deadlocks and switch the planning modality with a different set of parameters to resolve the deadlock.

A HAN system to address multi-context navigation

After introducing the situation assessment and modality shifting into HAN, the system has been extended to take care of the different types of visible humans in the environment. Numerous changes were made to make the system scalable and more pertinent to real-world applications. The most interesting thing about the proposed system is that the parameters of the system are highly tunable. Depending on the type of goal prediction (for humans) selected and the allowed thresholds of various human-aware constraints, the robot's behaviour could be adjusted to handle different human-robot navigation contexts.

Proactively addressing unseen humans in HAN

The final version of the proposed system introduces a concept called 'invisible humans' to improve the human awareness of the robot. The intention behind this work is to address the possible future appearances of humans from the occluded or hidden regions of the environment in the navigation scene. Firstly, an algorithm to detect the locations of 'invisible humans' in a 2D map is proposed. Then a new human-aware constraint and a new mode of planning for such unseen humans are proposed and added to our HAN system that makes the robot cautiously mitigate such locations and avoid collisions proactively.

New metrics for HAN

The last contribution of this thesis is a set of metrics that can be applied to several human-robot navigation contexts. Unlike the proxemics-based metrics, we believe that the proposed metrics are better suited for understanding and evaluating HAN systems in intricate scenarios. We present the mathematical formulation of these metrics and show how they can be used to evaluate and differentiate HAN planners from standard navigation planners.

Thesis Organisation

A major part of this thesis is based on the published work (core publications). The chapters based on the published work are elaborated compared to papers, including more details and discussions. The supportive publications include the work that is complementary during the development of this thesis, and these are detailed in the Appendix part.

This thesis has five chapters which can be grouped into three different parts. The first group consists of Chapter 1, which presents different aspects of robot navigation and the evolution of human-aware navigation. It also talks about the challenges in HAN and how they are addressed in the literature before presenting the mathematical background necessary to understand this thesis. The following chapters are based on the core publications and form the second group. These chapters are presented in chronological order of the development of the proposed

HAN system. Hence, Chapter 2 talks about combining situation assessment in HAN, followed by Chapter 3, in which a complete HAN system that can address multiple human-robot navigation contexts is presented. Next, Chapter 4 introduces the concept of ‘invisible humans’ to HAN and talks about proactive avoidance of potential future collisions. Throughout these three chapters, the advantages of proactive planning in combination with situation assessment are discussed in different settings. The last group has only one chapter, Chapter 5, that proposes and evaluates some new metrics for HAN.

The final remarks, lessons learnt, and future perspectives are discussed in the Conclusions chapter. The supportive work presented in Appendix A shows how an intelligent human agent is developed for the case of HAN. Further, different methodologies employed to simulate human agents for testing our HAN system are also presented. Throughout this thesis, whenever we refer to robot navigation, it is always a mobile robot with either differential or omnidirectional drive navigating on a 2D plane.

List of Publications

Published : Core Publications

- Singamaneni, Phani-Teja, and Alami Rachid. “**HATEB-2: Reactive Planning and Decision making in Human-Robot Co-navigation.**” 2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN). IEEE, 2020.*
- Singamaneni, Phani-Teja, Anthony Favier, and Rachid Alami. “**Human-Aware Navigation Planner for Diverse Human-Robot Interaction Contexts.**” 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2021.
- Singamaneni, Phani-Teja, Anthony Favier, and Rachid Alami. “**Invisible Humans in Human-aware Robot Navigation.**” Workshop on Social Robot Navigation: Advances and Evaluation in 2022 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2022.
- Singamaneni, Phani-Teja, Anthony Favier, and Rachid Alami. “**Watch out! There may be a Human Addressing Invisible Humans in Social Navigation.**” 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2022.

Published : Supportive Publications

- Favier, Anthony, Phani-Teja Singamaneni, and Rachid Alami. “**Simulating Intelligent Human Agents for Intricate Social Robot Navigation.**”

*Nominated for Best Paper award

2021 Workshop on Social Robot Navigation in Robotics: Science and Systems (RSS). 2021.

- Favier, Anthony, Phani-Teja Singamaneni, and Rachid Alami. “**An Intelligent Human Avatar to Debug and Challenge Human-aware Robot Navigation Systems.**” 2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, 2022.
- Hauterville, Olivier, Camino Fernández, Phani-Teja Singamaneni, Anthony Favier, Vicente Matellán, and Rachid Alami. “**IMHuS: Intelligent Multi-Human Simulator.**” 2022 Workshop on Artificial Intelligence for Social Robots Interacting with Humans in the Real World in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2022.
- Hauterville, Olivier, Camino Fernández, Phani Teja Singamaneni, Anthony Favier, Vicente Matellán, and Rachid Alami. “**Interactive Social Agents Simulation Tool for Designing Choreographies for Human-Robot-Interaction Research.**” 2022 Iberian Robotics conference, pp. 514-527. Springer, Cham, 2023.

Published : Other Publications

- Singamaneni, Phani-Teja, Amandine Mayima, Guillaume Sarthou, Yoan Sallami, Guilhem Buisan, Kathleen Belhassein, Jules Waldhart, and Aurélie Clodic. “**Guiding Task through Route Description in the MuMMER Project.**” Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, pp. 643-643. IEEE, 2020.
- Truc, Jérôme, Phani-Teja Singamaneni, Daniel Sidobre, Serena Ivaldi, and Rachid Alami. “**KHAOS: a Kinematic Human Aware Optimization-based System for Reactive Planning of Flying-Coworker.**” 2022 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2022.

Submitted

- Mayima, Amandine, Guillaume Sarthou, Guilhem Buisan, Phani-Teja Singamaneni, Yoan Sallami, Kathleen Belhassein, Jules Waldhart, Aurélie Clodic, and Rachid Alami “**Direction-giving considered as a Human-Robot Joint Action.**” Submitted to *User Modeling and User-Adapted Interaction (UMUAI) Journal*.

Evolution of Human-Aware Navigation

Contents

1.1 Motion Planning and Navigation	8
1.1.1 The Motion Planning Problem	8
1.1.2 Mobile Robot Navigation	11
1.1.3 2D Navigation Stack in ROS	13
1.2 Human-Aware Navigation	14
1.2.1 State-of-Art over the years	14
1.2.2 Human-Aware Navigation at LAAS	22
1.2.3 Human-Aware Navigation as a part of HRI	24
1.2.4 General challenges in Human-Aware Navigation	25
1.2.5 Addressing the challenges	26
1.3 Background and Our Approach	28
1.3.1 Timed Elastic Band for Trajectory Planning	29
1.3.2 Human-Aware Timed Elastic Bands for Proactive Planning	31
1.3.3 Principles of Joint-Action in Human-Aware Navigation	33
1.3.4 Our Approach and Contribution	34

Beginning with the topic of robot navigation, this chapter outlines the development of human-aware (robot) navigation (HAN). We briefly talk about navigation in the context of motion planning and explain the methodologies and tools for 2D navigation. Then we present the evolution of HAN over the years using the existing literature before presenting a small discussion on how HAN is closely associated with HRI. Further, we also present the chronological developments and contributions of our lab to this field. Next, we present the challenges faced by the researchers in this field and discuss how they were addressed in the literature. Finally, we present our approach to HAN along with the required mathematical background. In addition, we explain how the joint-action principles can be applied to the context of human-aware navigation.

1.1 Motion Planning and Navigation

We briefly introduce the motion planning problem and robot navigation in this section. We present a detailed explanation of navigation in 2D and the different components involved in it. Finally, we provide a brief overview of the ROS based 2D navigation stack for mobile robots.

1.1.1 The Motion Planning Problem

Motion planning has varied applications in several fields [Latombe 1999], but robotics remains the core field of development. The intention behind any motion planning problem in robotics (manipulation or navigation) is to plan a continuous motion sequence for a robot to reach the final goal configuration from its initial configuration in a given environment while avoiding collisions. Although motion planning can refer to either path planning or trajectory planning, in this thesis, we treat them as the two steps of the motion planning problem. The first step is path planning which generates a collision-free path from start to goal. In the second step, the generated path is used as a guide, and a feasible trajectory is planned to trace this path partly or entirely. Generally, a path is a sequence of configurations, whereas a trajectory is a sequence of configurations and time differences that satisfy all the kinodynamic constraints.

1.1.1.1 Path Planning

Path planning in robotics deals with the generation of a continuous path from an initial configuration of the robot to a final or goal configuration while avoiding obstacles in the given environment. Hence it requires the description of the environment (2D or 3D grids, maps, etc.) and a goal position as inputs for the path planning.

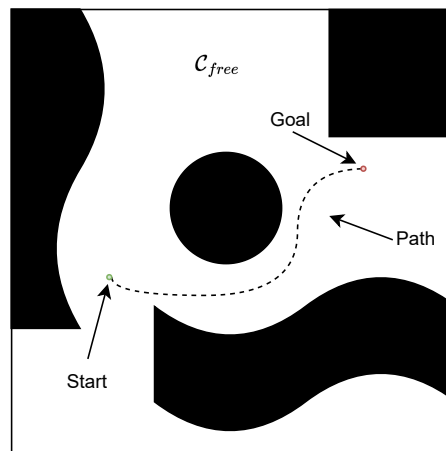


Figure 1.1: Motion Planning and Configuration Space

The configuration space, \mathcal{C} , for the path planning is the set of all possible configurations, and it depends on the type and degrees of freedom of the robot. For example, a mobile robot that can translate and rotate in 2D has a configuration space, $\mathcal{C} \in SE(2) = \mathbb{R}^2 \times SO(2)$, where $SE(2)$ is the special Euclidean group in 2D and $SO(2)$ is the special orthogonal group of 2D rotations. If the robot is a fixed base manipulator with n degrees of freedom (or joints), then \mathcal{C} is an n -dimensional vector. The subset of all configurations in \mathcal{C} that is collision-free comprise the free space, \mathcal{C}_{free} . Hence, any path from start to goal should be within the \mathcal{C}_{free} as shown in Fig. 1.1.

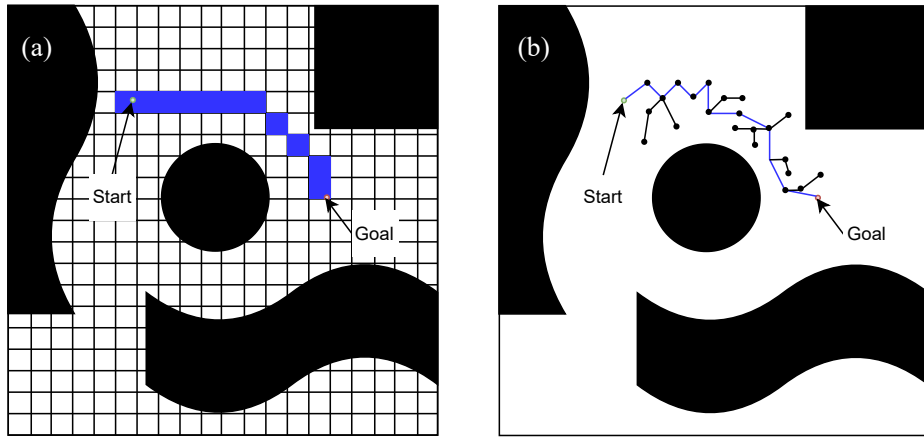


Figure 1.2: Path Planning. (a) Grid-based search (b) Sampling-based algorithm

Various methodologies can be employed to find a path, and the most famous ones are grid-based searches and sampling-based algorithms (Fig. 1.2). In grid-based methods, the entire configuration space is divided into small grids, and searching algorithms like A* [Hart 1968], D* [Stentz 1997] or Dijkstra are employed to search the connectivity from start to goal. A sample path generated using this method is shown in Fig. 1.2 (a) (blue region). However, these methods suffer from the curse of dimensionality as the grid size becomes smaller and the configuration space (or map) becomes larger. Sampling-based algorithms like RRT [LaValle 1998] or PRM [Kavraki 1996], on the other hand, can handle high dimensional spaces better as their run-time is not exponentially dependent on the size or dimension of the configuration space, \mathcal{C} . These methodologies grow a connectivity map or tree from the starting position based on the points sampled in small regions. The sampled points are connected and added to the tree only if the connecting line segment is completely in \mathcal{C}_{free} . A path is found when the start and goal are connected by this tree, as shown in Fig. 1.2 (b). The shortest connected path is shown in blue. These methods are probabilistically complete and will find a solution or a better solution as more time is spent searching but cannot verify if a solution exists or not. Other frequently used techniques involve artificial potential fields [Vadakkepat 2000], cell decomposition [Garrido 2006] (for e.g. Voronoi regions), interval-based search etc.

1.1.1.2 Trajectory Planning

Trajectory planning is used to generate the reference inputs for the robot's control system to execute the desired motion. Therefore, a trajectory is usually parameterised by time and can include the velocities and accelerations, apart from the configurations of the robot. The trajectory planning problem takes the path planned from the previous step as input, along with the kinodynamic constraints like the velocity, acceleration and kinematics of the robot to generate a continuous sequence of configurations and their associated time intervals.

The planned trajectory can either be in the operating space (world frame) or the joint space (local frame) of the robot. If the motion of the robot has to follow certain geometric characteristics defined in the operating space (e.g. navigation), the trajectory planning is done in the operating space. To generate the reference control commands from this trajectory, kinematic inversion has to be performed, mapping the trajectory in the world frame to the desired controller or joint frames. As the degrees of freedom of a robot increase, this kinematic inversion for a densely populated trajectory becomes computationally expensive. Hence, trajectory planning is largely done in the joint space. For this, a set of waypoints are sampled from the planned path and mapped to the joint space using kinematic inversion. Then, trajectories are planned, such that all these waypoints are connected while satisfying the kinodynamic constraints. This can be done by interpolating between the points using a polynomial function or spline satisfying the constraints. Planning in joint spaces generally avoids the problems with redundant joints and singularities. However, the issue with such trajectory planning is that the robot might have some undesired configurations between these waypoints, and depending on the granularity of sampling, the number of these configurations might increase or decrease.

No matter in which space the trajectories are planned, they should be continuous not only in terms of configurations but also in terms of velocity and acceleration to avoid jerks and vibrations of the robot while executing them. This means that the trajectory has to obey these additional constraints along with the kinodynamic constraints, and hence, it is often posed as an optimization problem. The most significant optimality criteria used for this are the minimum time, energy or jerk [Gasparetto 2015]. One can employ any optimization technique like minimax [Vathsal 1977], genetic algorithms [Tian 2004], multi-objective optimization [Oleiwi 2014], unconstrained optimization [Rösmann 2013], or model predictive control (MPC) [Ardakani 2015] to solve this problem, depending on the requirements. Re-planning times for trajectory are usually large as it involves solving a complex optimization problem. Therefore, trajectory planning can either be global, covering the entire path, or local, covering only a small part of the path. For dynamic environments, local planning is employed so that the trajectory can be dynamically re-planned in real-time. Global trajectory planning is employed when the robot's operating space is constant and does not change over time.

1.1.1.3 Navigation and Manipulation

Although both navigation and manipulation are motion planning problems and use similar algorithms, there are some basic differences. The navigation problem specifically deals with moving the entire robot from one place to another, whereas the manipulation problem deals with the motion of the end-effector and/or joints from the initial configuration to the goal configuration, while the base of the robot is assumed to be fixed. This thesis focuses only on the navigation problem, specifically mobile robot navigation on the ground using a 2D description of the world.

1.1.2 Mobile Robot Navigation

Mobile robot navigation is a well-studied field as it lays the foundation for autonomous vehicles and robots. Navigating in an unknown environment requires Simultaneous Localization and Mapping (SLAM) [Thrun 2007] along with some kind of global positioning system (GPS) to guide the robot or vehicle to the goal. As global path planning cannot be done in such environments, robot navigation is largely handled by local motion planning and collision avoidance techniques. However, when it comes to the known environments with prior maps, the navigation problem can be split into two discrete parts, namely, localization and motion planning. The localization part of the problem continuously locates and tracks the position of the robot in either a static or dynamic environment [Thrun 2006]. Numerous methods can be employed for this like EKF localization using landmarks, grid-based localization [Thrun 2006], Monte-Carlo or Adaptive Monte-Carlo localization [Fox 2003] using particles filters [Fox 2001], vision-based localization [Adorni 2001] etc. In this thesis, we do not address the entire navigation problem and only focus on the motion planning part, which is usually referred to as

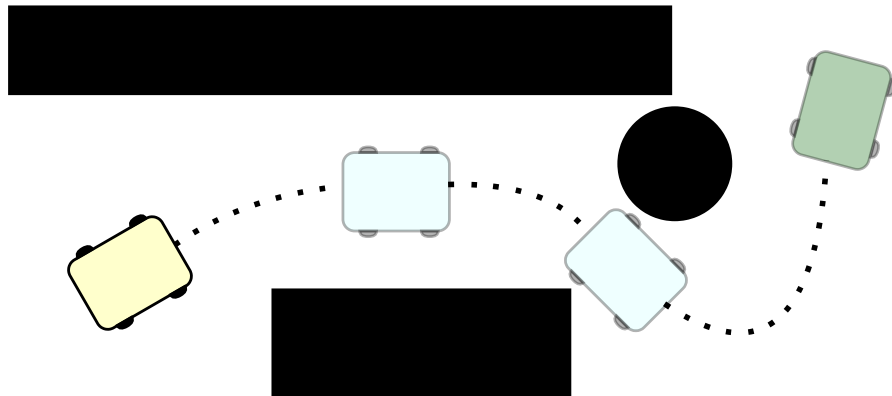


Figure 1.3: 2D mobile robot Navigation from start (yellow) pose to goal (green) pose. The path planning problem has to plan the path and the intermediate poses (blue) that the robot needs to follow to reach the goal without colliding with the obstacles. The obstacles are marked in black. The trajectory planning then has to generate the control commands for the robot to follow this path.

‘navigation planning’. As discussed in the previous section, given a goal configuration, in the case of 2D, a goal pose (position and orientation), the motion planning part (or navigation planning) has to solve the problem of the robot’s motion from the start to the goal. Fig. 1.3 shows an example of such navigation planning avoiding the static obstacles in the environment. The path planned from the start pose to the goal pose of the robot in the free space, \mathcal{C}_{free} is shown in this figure. The trajectory planner has to take this path as input and generate the control commands for the robot’s base controller.

The maps for mobile robot navigation are usually represented as 2D occupancy-grids [Thrun 2006] as they are easy to acquire. The fundamental concept behind occupancy grids is to visualise the map as an evenly spaced grid of binary random variables. The occupancy grid based maps are built by updating the occupancy value of each grid using different kinds of range sensors like Sonars or LIDARs. Cameras and other sensors are also used in fusion with the range sensors to improve the accuracy of this estimation. However, all these methodologies assume that the robot knows its position and path in the map, which is generally provided by SLAM techniques. Unless a blueprint of the environment is available, maps are built using SLAM and Occupancy Grid mapping algorithms [Thrun 2006] and are used for path estimation.

In such occupancy grids, search-based path planning algorithms like modified A* [Warren 1993] are much faster compared to sampling-based approaches like RRT-connect [Kuffner 2000] or PRM [Kavraki 1996]. Hence, in 2D navigation, path planning is done globally using grid-based search, which provides waypoints or a reference path for trajectory planning. The trajectory planning can be either local or global. Global trajectory planning will work in a setting where nothing moves or changes, and the environment remains in the same state as it was represented in a map. This is mostly true in the case of a manipulator, but for navigation, the environments usually change, and more obstacles like tables, chairs, beds etc., are added to the basic (static) environmental map. Hence, trajectory planning also requires obstacle information which is provided through the sensors. As the range of these sensors is limited, they can only provide information about the obstacles in a small region. Therefore, in a mobile robot navigation setting, the trajectory planning is mostly local, and it is interlinked with obstacle avoidance which modifies the trajectory of the robot to avoid collisions with obstacles found on its path. Some of the commonly used obstacle avoidance methods are Dynamic Window Approach (DWA) [Fox 1997], Elastic Bands [Quinlan 1993] and artificial potential fields. Consequently, trajectory generation algorithms combined obstacle avoidance methods with their optimization scheme to produce online trajectories in real-time. However, these obstacle avoidance methods did not consider the velocity of the obstacles [Cai 2020], and hence, in the case of an environment with a lot of dynamic obstacles, the trajectory optimization had to be run at a high frequency which increased the computational cost and made online re-planning infeasible.

Concurrently, collision avoidance strategies were developed for environments with a lot of dynamic obstacles or humans (treated as obstacles). These algorithms

generate instantaneous velocities or control commands for the robot's base that ensure collision-free motion with fewer computational resources. Therefore, many works in literature employ collision avoidance techniques with global path planning for robot navigation in highly dynamic environments and crowds, which are sometimes called 'reactive navigation' schemes. The trajectories were generated by accumulating a sequence of control commands provided by the collision avoidance strategies [Prassler 1999]. These methodologies include velocity-obstacle based methods like RVO [Van den Berg 2008], ORCA [Berg 2011] or CALU [Hennes 2012], social force and reinforcement learning methods like CADRL [Chen 2017b], SA-CADRL [Chen 2017a] or CRL [Ciou 2018]. This is still a very active research field as it helps to simulate and understand crowds. Although the collision avoidance methods can produce bounded velocities, the resultant trajectory is not guaranteed to be smooth in terms of acceleration and jerk. Consequently, researchers reverted to online trajectory planning, for example, MPC and Timed Elastic Band (TEB), as computational power has significantly improved over this time. The main contribution of this thesis is an online trajectory planner based on timed elastic bands [Rösmann 2013] while relying on grid-based path planning (Dijkstra or A*) to generate the path.

1.1.3 2D Navigation Stack in ROS

Robot Operating System (ROS) [Quigley 2009] provides an environment to develop different tools for robotic platforms. It has several built-in packages and software bundles that make this development easier. Being an open-source platform, there are several custom packages and large community support to discuss and resolve issues. ROS readily supports motion planning and has official packages for navigation as well as manipulation. As the navigation planning system presented in this thesis is developed using the ROS Navigation Stack¹, we briefly present its features. The main package of this stack is *move_base* which provides the plugins to include new planners.

Inputs: Apart from the 2D navigation goal (position and orientation), the navigation stack takes as input the odometry information, the sensory data like laser or point cloud to obtain the real-time obstacle information and the static map data (as occupancy grid) for planning the path. It also requires the necessary transformation to connect the 'map' and the 'base_link' of the robot, which usually is published by the localisation module. The navigation stack by default uses the *amcl* package in ROS for robot localization.

Architecture: The navigation stack has a 'global planner' that takes care of the path planning and a 'local planner' that deals with the trajectory planning. Both these planners are built as plugins and provide an easy way to use your own custom planners besides the existing ones. As it is built over ROS, the parameters and the planners can be updated easily using an 'yaml' configuration file.

Output: It outputs the command velocity for the robot's base controller.

¹<http://wiki.ros.org/navigation>

1.2 Human-Aware Navigation

Human-Aware Navigation (HAN) deals with the navigation planning of the robot in human workspaces. Therefore, it has to take into account several factors while planning. We believe that HAN needs to integrate human expectations, social norms and situation analysis into its planning while taking care of other static and dynamic obstacles. It should try to lessen the discomfort and make the robot's motion more legible to the humans in the environment.

In this section, we present various works related to HAN, distributed in four parts. We then proceed with a description of the research and evolution of HAN at our lab over the years and the proposed systems. Next, we briefly discuss how HAN is seen and used as a part of HRI before presenting the challenges faced by researchers in HAN. Finally, a short description of how these challenges are addressed is provided.

1.2.1 State-of-Art over the years

We briefly present the works that are related to HAN in this section. This section is divided into four parts to show for ease of explanation and also to show how the number of works increased over time. We start with some early works and move on to the most recent literature.

1.2.1.1 The initial works: Before 2000

Some of the early works in HAN were published before 2000. In the work by Tadokoro et al. [Tadokoro 1995], a trajectory planning algorithm was presented that takes human motion prediction into account and plans the robot's motion. As the interest in the sector of service robots grew, new design philosophies and architectures were proposed like the ones by Kawamura et al. [Kawamura 1996] and Wilkes et al. [Wilkes 1998]. During this period, one of the first robotic tour guides, RHINO [Burgard 1999b], was deployed in the “Deutsches Museum Bonn” for six days. According to the article by Burgard et al. [Burgard 1999a], it had a success rate of over 99%. Later, a second generation of the robotic tour guide, MINERVA [Thrun 1999], was deployed in the “Smithsonian's National Museum of American History” for two weeks and gave 620 hours of tours. Both these robots used similar probabilistic architectures as mentioned in [Burgard 1999a] and [Thrun 2000]. However, MINERVA had an additional interactive component and communicated its navigation intent through facial expressions and sounds [Schulte 1999]. The work by [Imai 1999] presents a migration system with a computer generated agent that shifts from PC to the Pioneer I robot, which guides the people in their labs.

Robotic wheelchairs also have to address HAN along with the comfort of the person in the wheelchair. Yanco [Yanco 1998a] presented a list of the then-existing projects for robotic wheelchairs and proposed a wheelchair for indoor navigation [Yanco 1998b]. In parallel, Prassler et al. [Prassler 1999] deployed their robotic wheelchair, MAid [Prassler 1998], in a railway station during rush hour and through

crowds to evaluate their navigation system. They used temporal maps with decaying observations over the grid for path planning and velocity-obstacle based evasive manoeuvres.

1.2.1.2 The early 2000's: 2000-2006

The number of works in the field started to rise over these years with an increased focus on robots for human-aid. Boy et al. [Boy 2002] proposed a collaborative wheelchair with guided paths for aiding the manoeuvres. This was later realised in work by Zeng et al. [Zeng 2006]. The wheelchair moves along the software-generated path between the start and the goal defined by the user. The user can modify this path online and also controls the speed. The path modification proposed by these works is similar to the elastic band approach. The works by Gonzalez et al. [Gonzalez 2006b] and Galindo et al. [Galindo 2006] also use the interactive path selection approach in their wheelchair, SENA [Gonzalez 2006a]. This system was tested mainly in indoor environments. The works presented in [Rao 2002] and [Parikh 2003] show the autonomous navigation of a wheelchair in the presence of obstacles, while [Morris 2003] uses a mobile robot with autonomous navigation capabilities to assist visually impaired elderly people. All these robotic systems have to navigate the human environment in addition to taking care of the comfort of the passenger. The work by Prassler et al. [Prassler 2002] presented a wheelchair accompanying a person side-by-side in a railway station. Here, the robot had to coordinate its motion with the accompanying person while avoiding collisions with the other humans in the environment.

Coming to the mobile robotic tour guides, researchers from *ETH*² designed and deployed a mobile robot, RoboX [Arras 2003] in a robotics expo and presented the complete architecture and their experiences [Siegwart 2003, Jensen 2005]. Their system used a modified local path planning combining the DWA and elastic band approaches [Philippsen 2003]. Another robotic tour guide is CiceRobot [Macaluso 2005] which used a cognitive architecture and was deployed in the “Archaeological Museum of Agrigento”. Garcia et al. [Martinez-Garcia 2005] studied the scenario in which several robots were guiding multiple people to steer the crowd without explicit communication using the Social Force Model (SFM) [Helbing 1995]. Most of the above works did not consider human motion prediction in the planning. Bennewitz [Bennewitz 2004] presented an approach for learning the motion patterns of humans [Bennewitz 2005] using Expectation-Maximisation methodology and integrated it into robot planning to have adaptive navigation strategies yielding a complete HAN system. Althaus et al. [Althaus 2004] integrated the navigation into the interaction loop and presented a methodology to approach a group of people and dynamically maintain its formation as the group changes. It used a potential field based approach with several constraints on the direction of the robot along with speed modulation around humans. A probabilistic extension of velocity-obstacle based collision avoidance was proposed

²<https://ethz.ch/en.html>

by [Kluge 2004], which tried to account for the uncertainty associated with human locomotion and the potential future behaviour of the agents. They modelled the obstacles (or humans) as intelligent decision-making agents with the same level of intelligence as the robot.

Some works studied the human-robot interaction in the context of navigation and published the observations. Butler et al. [Butler 2001] conducted several experiments with a mobile robot that is either approaching, avoiding or exploring the room while a person stands still in the same room. The mobile robot's speed, paths and bodies were different in all of these experiments. They concluded that slower approach velocities were preferred, and the larger distance was preferred for the robot with a bigger body while avoiding. The room exploration study did not yield any significant results. One of the early studies on communication through gaze was done by Kanda et al. [Kanda 2001], where the robot tried to communicate its intent and interest through gaze and body rotations. The conclusion was that the "gaze control (moving the camera in the direction of interest) made the robot impressions more enjoyable and active". The works by Pacchierotti et al. [Pacchierotti 2005, Pacchierotti 2006] studied a robot passing a human in a corridor. They concluded that people preferred higher speeds while crossing and early signalling of the direction the robot was going to take.

1.2.1.3 The late 2000's: 2007-2012

An increase in the number of works on person following and approach was observed during this period. In the context of the person following, Gockley et al. [Gockley 2007] reported that a robot following the direction of the person rather than the exact path was found to be more natural. The work by Hoeller et al. [Hoeller 2007] used a potential field based approach with virtual targets and motion prediction in the local planning that allowed the robot to evade the possible interferences by other humans while following the target person. In the works presented in [Zender 2007] and [Granata 2012], people tracking was combined with SLAM to determine where humans are in the environment and decide the action for the robot or the location to meet. Further, these works introduced situation assessment into the planning to shift between different robot behaviours like increased distance from doors, going to the person, following the person etc. However, these behaviours are on the top of the control layer and not inside it, as we present in this thesis. Seifer et al. [Feil-Seifer 2011] presented a goal-oriented robot behaviour when a human partner is involved. The modified trajectory planner presented in this work allowed the robot to slow down or stop and wait to let the human partner catch up with the robot.

Kessler et al. [Kessler 2011] and Satake et al. [Satake 2009] both presented the planning methodologies to approach a human for the mobile robot. While the work in [Kessler 2011] dealt with static humans using simple sampling-based planning, [Satake 2009] presented a complex framework to approach moving humans in the frontal direction and show the robot's intention to interact, based on the detected

state of the human. Similar work was presented by Hayashi et al. [Hayashi 2012] to naturally encounter people walking in malls using gaze and a predefined ‘friendly patrolling’ behaviour. In a contrasting setting, Hansen et al. [Hansen 2009] proposed an approach to predict whether a human intends to interact with a robot or not, based on motion pattern (trajectory) analysis of the human. Some studies regarding approach methodologies, like Shi et al. [Shi 2008] and Takayama et al. [Takayama 2009] reported some results relating the distance of approach to the speed and gaze, respectively. They concluded that slower speeds are preferred, and it is better to avert the gaze while approaching very close, to avoid the conveyance of threat. The works presented in [Lichtenthaler 2012b] and [Kruse 2012a] study the legibility of robot’s motion in the crossing scenario. They concluded that velocity modification is more legible than path modification in such scenarios.

As the interest grew in HAN, robotic wheelchairs were actively improving, and several works were published during this period. The works presented in [Demeester 2008, Taha 2008, Vanhooydonck 2010] studied the ideas of shared control and user intent recognition. Demeester et al. [Demeester 2008] presented a user-adapted plan recognition and shared control of wheelchair using the Bayesian approach and POMDPs by using explicit communication, whereas Vanhooydonck et al. [Vanhooydonck 2010] used artificial neural networks to learn the user intent implicitly. Taha et al. [Taha 2008] used POMDP based approach again, but instead of aiming at a semi-autonomous control, the system predicts the user’s intended goal in a given environment and takes the human there with minimal input from the joystick. Intention detection could be of use in autonomous vehicles as well. Fully autonomous HAN systems on wheelchairs were presented by Martinez et al. [Rios-Martinez 2012b] and Tomari et al. [Tomari 2012] in shared human spaces, considering the comfort of the other humans in the environment.

Guiding robots were still an active part of the HAN research, and some of them include TOOMAS [Gross 2009], a shopping guide robot in large stores, Urbano [Rodriguez-Losada 2008], a mobile robot guide deployed in several environments and a museum tour guide proposed by Han et al. [Han 2010]. These are just a few of them, and the list is not exhaustive. Huang et al. [Huang 2010] proposed a HAN system for the robots in the home based on proxemics. In a similar work by Lam et al. [Lam 2010], they not only consider humans but also other robots and propose a robot navigation system based on sensitive fields which are similar to proxemics. The work by Svenstrup et al. [Svenstrup 2010] uses an RRT based path planning combined with an MPC controller to navigate dynamic human environments. Henry et al. [Henry 2010] used Inverse Reinforcement Learning (IRL) to learn motion planners for the robot to navigate the crowds like humans. However, their system was tested only in the simulation, and no real-world results were reported. The work presented in [Muller 2008] proposes a strategy to navigate in populated environments by tracking and following people who move towards the robot’s goal.

One of the common problems that exist in crowd navigation is the Freezing Robot Problem (FRP) [Trautman 2010]. Trautman et al. [Trautman 2010] ex-

explicitly addressed the FRP problem by proposing an Interactive Gaussian Process (IGP) based motion model for the navigating agents. To improve the human model, Thompson et al. [Thompson 2009] used a mixture of annotated spaces and heuristically determined waypoints to infer the navigation intent of humans and predict their short-term navigation goals. Chung et al. [Chung 2010] proposed something similar by exploring the spatial effects of the environment on the behaviour of pedestrians. They combined their trajectory prediction model with the proposed spatial effects detection model and showed that a mobile robot navigates better by using them. Garrell et al. [Garrell 2010] proposed a new cost function for optimizing co-operative robot movements when a group of robots tries to guide a group of people. This makes the robots adapt their trajectories to avoid humans getting away from the group. It can be seen that more trajectory planning and complex optimization schemes were used during this period compared to simple collision avoidance methods. This can be attributed to increased computational power and improved studies in HRI and HAN. Kirby et al. [Kirby 2009] proposed an optimization scheme for person-acceptable navigation (COMPANION) where social norms or conventions were modelled as constraints on the robot’s navigation. The HAN system proposed in this thesis uses a similar approach.

1.2.1.4 In the past decade

Over the last decade, the research on HAN has gradually increased, and the field has expanded to include drones and autonomous vehicles. The literature is mainly spread among different categories like HAN in crowded or densely populated areas, indoor or structured environments, and service robots.

Crowd navigation is one of the challenging tasks for a HAN system, and some of the early works addressed this using SFM [Ferrer 2013a, Ferrer 2017, Patompak 2016] or learning approaches [Pérez-Higueras 2014, Vasquez 2014, Triebel 2016, Luber 2012]. The guiding robots, FROG [IDM 2011] presented in [Pérez-Higueras 2014] used IRL based controller while SPENCER, presented in [Triebel 2016] used an RRT* based planning over an IRL based costmap. Vasquez et al. [Vasquez 2014] have presented how features and selected IRL algorithms affected the learnt behaviour of the robot. The work by Ferrer et al. [Ferrer 2013a] is one of the first works to use SFM based mobile robot navigation in an outdoor setting. It was later extended to address the problem of accompanying a person in outdoor scenarios by Ferrer et al. [Ferrer 2017] and Repiso et al. [Repiso 2017]. Patompak et al. [Patompak 2016] proposed new strategies for path planning in populated environments based on an extended SFM and RRT. Truong et al. [Truong 2017a] combined the extended SFM with HRVO [Snape 2011] to generate a proactive social motion model that can navigate crowded environments taking care of human-human as well as human-object interactions. As deep Reinforcement Learning (RL) gained popularity in robotics and control, it was applied to learn human-aware crowd navigation strategies in works by Chen et al. [Chen 2017a, Chen 2019, Chen 2020] (SA-CADRL, SARL, RGL), Xie et al. [Xie 2021] and Liu et al. [Liu 2020]. These

works not only use deep RL but also use neural network based architectures to extract the features [Liu 2020] or learn salient relations [Chen 2020]. The work by Wang and Steinfeld [Wang 2020] uses 3D convolutional networks to predict the splitting and merging of people in groups. As mentioned previously, FRP is recurring in a crowd navigation setting, and some works address this issue precisely while proposing a robot navigation solution. Trautman et al. [Trautman 2015] extended their previous work on IGP and proposed multi-goal IGP (mgIGP) to develop a robot navigation algorithm that encourages cooperation between humans and the robot while avoiding the FRP. The studies conducted in real environments inferred that the proposed algorithm was comparable to the results of the teleoperation. Sathyamoorthy et al. [Sathyamoorthy 2020] have proposed a geometric approach to identify the freezing areas and combine it with a deep RL based collision avoidance [Long 2018] to move the robot among humans without getting frozen. Wang et al. [Wang 2022] use group-based motion predictions and a special form of MPC to mitigate the FRP. Nishimura et al. [Nishimura 2020] tried to learn the balance between safety and efficiency in crowd navigation. This approach is based on the assumption that the robot in the crowd can be seen as an agent in sequential social dilemmas (SSDs) [Leibo 2017] setting. The results in various simulated scenarios showed satisfactory navigation of the robot without taking long deviations and safely passing through the crowd. In a recent work by Dugas et al. [Dugas 2020], multi-behaviour navigation planning was proposed for robots in a crowd using some interactive actions. They have introduced three multi-modal behaviours called *Intend* (move to free space), *Say* (communicating the robot’s intention to pass through verbally) and *Nudge* (communicating verbally and gesturing using a hand to pass through). The *Nudge* action was specifically designed to unfreeze the robot in dense crowds.

As indoor navigation is one of the core areas of focus in HAN, there were numerous works in the previous decade which proposed several interesting methodologies to address the problem. One of the simplest approaches was proposed by Lu et al. [Lu 2013] in which they include a costmap layer around the detected human based on the proxemic zones. Then the path is planned in the free areas of the costmap and, hence, avoids humans. The works by Qian et al. [Qian 2013] and Cunningham et al. [Cunningham 2015] proposed POMDP based decision-making and modality shifting for robot navigation in constricted or crowded spaces. The modality shifting is not at the level of trajectory planning as presented in this thesis. One of the recent works by Buchegger et al. [Buchegger 2018] modified the robot’s trajectory at the control level based on human predictions. This work is similar to ours, but we include situation analysis and modality shifting as well at the control level. Some works like [Mavrogiannis 2018, Vasconcelos 2015] presented methods to produce legible robot navigation. The authors of [Mavrogiannis 2018] obtained legible motion for the robot by minimizing ‘social momentum’, a concept they introduce in their work. In [Vasconcelos 2015], the velocity of the robot is adjusted based on the distance from the human, similar to the proposal in [Kruse 2012a]. Truong et al. [Truong 2014] proposed the dynamic social zones based on proxemics and

published a series of works [Truong 2017c, Truong 2016, Truong 2017d] addressing various concerns in social navigation like approach, human-object interactions etc. Similarly, Vega et al. also published a series of works [Vega 2017, Vega-Magro 2018] that proposed HRI architectures for HAN with an adaptive spatial density function to address human needs based on the context. Very recent work from the same group [Manso 2020a] proposed the use of graph neural networks for learning a similar spatial density function based on context. The work by Kollmitz et al. [Kollmitz 2015] used time-dependent planning to address dynamic humans in home-like environments. Araujo et al. [Araujo 2015] proposed a potential field based approach for navigation in structured office-like environments. One of the recent works by Bruckschen et al. [Bruckschen 2020] proposed an idea of long-term goal prediction based on the possible goal locations for humans and designed a navigation strategy for the robot to arrive at the goal before or at the same time as the human for assisting. Human intention prediction was also used by some of the works [Ratsamee 2013, Peddi 2020, Park 2016] to predict the willingness of a human to interact with the robot or the navigation intent. Based on the detection, the robot takes appropriate actions to avoid smoothly or approach the human for interaction. Navigation in warehouses is a part of HAN, and the works presented in [Fernandez Carmona 2019, Dondrup 2016, Kenk 2019, Guldenring 2020] specifically address this. The works presented in [Fernandez Carmona 2019, Kenk 2019] used geometric approaches, whereas Guldenring et al. [Guldenring 2020] used deep RL to solve the same problem. As deep learning and RL have become significantly improved over the years, there are many works that use these to address HAN [Qiu 2022, Pérez-Higueras 2018b, Fahad 2020]. However, there are still some works that use classical approaches like RRT [Majd 2021] or optimization [Shin 2020, Banisetty 2018]. Depending on the requirements and the computational capacities of the system any of the above approaches can be chosen for implementing HAN. Note that none of the methods is perfect and has some limitations and drawbacks. The multi-context navigation, unlike the multi-behaviour navigation that is addressed in this thesis, was also proposed in one of the recent works based on multi-objective optimization and deep learning by Banisetty et al. [Banisetty 2020]. Shin et al. [Shin 2020] used optimization to follow a person, while Kollmitz et al. [Kollmitz 2020] used IRL to make the robot learn human preferences while navigating in an environment. Finally, one of the new developments in HAN is virtual reality (VR) based robotic control, and in this context, Becerra et al. [Becerra 2020] proposed a navigation system for a remote robot that takes into account the comfort of the human using the VR set.

The development of HAN in service robots continued, and Rioz et al. [Rios-Martinez 2012a] proposed intention-aware navigation for a robotic wheelchair using Risk-RRT [Rios-Martinez 2011] and face control [Escobedo 2012]. For the human motion prediction, they used the Growing Hidden Markov Model [Govea 2010] trained with collected trajectories from real-world experiments [Vasquez 2013]. Narayanan et al. [Narayanan 2016] proposed a semi-autonomous approach for wheelchair navigation with user intention prediction and

socially compliant planning. In the subsequent work, a transient goal driven approach [Narayanan 2018a] was proposed to make the robot socially compliant locally. Having such local goals improves the navigation behaviour of the robot among pedestrians when compared to global optimal planning. In the work presented in [Narayanan 2018b], this transient goal approach was extended such that these transient goals are communication aware. It is proposed for large environments to avoid the loss of wireless communication signals while navigating among humans. A series of works published by Pasteau et al. [Pasteau 2016, Pasteau 2013, Pasteau 2014] presented methods for corridor following and door passing using visual servoing for wheelchairs. Morales et al. [Morales 2017] proposed a motion planning framework for wheelchairs taking into consideration both passenger and pedestrian comforts. They performed a study to decide motion parameters for comfort and preferred location within a straight corridor (to one side) for passengers. A very recent work by Paez et al. [Paez-Granados 2022] demonstrated a control strategy that enables a mobile service robot to attain reactive control after a collision, allowing it to absorb some of the impacts and keep travelling by driving slowly around the pedestrian. They proposed it as an alternative solution to the general safe approach of freezing a robot upon contact.

There were numerous studies in HAN during these years, and we will talk about some of them here. The studies by Lichtenthaler et al. [Lichtenthaler 2012a, Lichtenthaler 2013b] on legibility concluded that a robot using human-aware motion planning systems was more legible and the robot changing speed while moving towards the goal was more legible than changing path. Communication of intent through signals was studied in three different works [May 2015, Hart 2020, Palinko 2020], and each of them reported some very interesting results. The works in [May 2015] and [Hart 2020] have a similar experimental setup and tested whether an expression-based turn indication was better than a traffic signal like indication. Both these studies conflict with each other as one suggested traffic indicators were better, and the other suggested that facial expressions were better. These differences could be due to differences in the timelines of these studies or the design of the signals. The take-home message from these studies would be some kind of indication is better than none. The study by Palinko et al. [Palinko 2020] also supported that blinking signals show the robot’s intention. This study showed that turning the whole robot in the direction of the turn was the best indicator, and combining it with other signals improves it further. They have also concluded that people need not follow the “right” or “left” lane rules while crossing the robot. Hetherington et al. [Hetherington 2021] studied different kinds of signals for a robot yielding to a human at a doorway. Almost all the signals were correctly interpreted, with retreat (stopping and moving back) scoring the highest in measures of trust, likeability, comprehension, comfort, and social compatibility. The work presented in [Senft 2020] studied the scenario of a robot giving way to humans in a narrow passage. The results showed that sliding to a side and rotating sideways was preferred by the people. Even though there was no extra space or clearance due to the robot’s actions, it was more acceptable than the one without such social behaviour,

indicating that HAN is not just about providing enough space for pedestrians. Mavrogiannis et al. [Mavrogiannis 2019] tested different navigation algorithms on the robot with real humans and reported that humans did not find any significant difference between the strategies. Sorrentino et al. [Sorrentino 2021] tried to model the personalities of the robot based on proxemics, and the study revealed that humans could notice the differences in the personalities and their perceptions were based on their prior experience with the robots. A recent work by Salvini et al. [Salvini 2022] presented the risks of deploying mobile robots in crowded human environments and suggested that psychological safety should also be considered by HAN than just physical safety.

With the growing interest in HAN, autonomous vehicles and drones also entered the field. Some of the works concerning HAN in drones and AVs are presented in [Evens 2022, Luo 2018, Garza Elizondo 2016, Garrell 2019]. There are many more works, and we will not go into details about these works as the main focus of this thesis is on mobile robot navigation. The popularity has also led to the creation of new datasets [Rudenko 2020, Karnan 2022, Manso 2020b, Othman 2020] and simulators [Mizuchi 2017, Tsoi 2020, Holtz 2021, Biswas 2022] for learning and studying HAN. The datasets are collected with robots navigating in real environments or under controlled laboratory settings. The presented simulators provide navigation behaviours based on realistic data or algorithmically generated behaviours. At the end of this thesis, we briefly present our contributions to simulating rational humans for HAN.

1.2.2 Human-Aware Navigation at LAAS

Human-aware navigation has been an active part of research at LAAS³ since the very early stages. One of the first works by Alami et al. [Alami 2000] presented an HRI architecture focussed on a human-friendly navigation task called Diligent. Instead of looking at HAN as a motion planning problem, a supervisor module supervises and controls the navigation task execution. It was based on the elastic band based path planning using nearness diagram [Minguez 2000] for collision avoidance. Later, in 2005, Sisbot et al. [Sisbot 2005] presented a human-aware path planning system that takes into account different factors apart from the proxemics like the visibility and hidden zones while planning the path. By introducing safety, visibility, and hidden zone costs into a grid-based planner based on the human states (position, orientation, standing or sitting), they solve for better and adaptive human-aware paths, mostly in the case of static humans. Some studies and trials were conducted [Dautenhahn 2006] that provided a basis for the grid modelling the preferred direction of approach while sitting. The proposed framework was tested in several scenarios [Sisbot 2006, Sisbot 2007a] to show its effectiveness. This framework was later extended to develop a path planner for complete motion planning tasks, including human-aware manipulation [Sisbot 2007b]. Similar to the case of navigation, human-aware manipulation takes safety, visibility and human

³<https://www.laas.fr/public/>

comfort into account while planning the path. A placement position is found based on the spatial location of the human using perspective placement [Sisbot 2007b], and then the HAN system plans a path to reach this position. Then the robot moves and places itself there before finally planning the manipulation to hand over the object. To determine the placement and the time to start the handover task, a user study was conducted [Koay 2007] before developing the framework. The study concluded that the preferred approach distance could change, but the manipulation should start as the robot arrives near the placement position, not too early or after reaching the position, according to the majority of humans. They also studied the approach direction, but it contradicted the results from [Dautenhahn 2006], which could be due to the different levels of familiarity of humans with the robot in the two studies. Later, this framework was combined with a trajectory planner [Sisbot 2010] and tested under several real-world scenarios, along with a user study [Dehais 2011], which inferred that the proposed planner was preferred to the standard motion planning frameworks.

In the subsequent years, attention shifted to HAN planning again based on the experiences of the museum guide robot, Rackham [Clodic 2006]. Pandey et al. [Pandey 2009a, Pandey 2009b] proposed a generic framework for incorporating various social norms at different stages of execution, like avoiding humans and groups or guiding a person to a goal. The initial path is generated using a set of milestones or waypoints that obey certain conventions, and these are modified whenever required, during the execution of motion, depending on the state [Pandey 2010]. Decision trees and splines were used to generate the path from these milestones. They also include intention-show into this framework by signalling early and selecting a path that conveys the robot's motion intention to the human. This is particularly useful in the case of a robot taking a human to a goal. The goal-oriented navigation [Pandey 2009b] also modifies its path to bring back the human if he/she stops moving. Following this, the HAN system proposed in [Sisbot 2005] was extended to the case of moving humans using velocity-based predictions and talked about the concept of cooperation in HAN [Kruse 2010b, Kruse 2010a]. In 2010, Mainprice et al. [Mainprice 2010] proposed a new RRT based planner for combined navigation and manipulation planning.

In recent years, the idea of humans as partners in HRI (and HAN) was studied at LAAS. This was studied in detail in the context of handover tasks. Mainprice et al. [Mainprice 2012a] included the mobility factor (whether the person can move more or less) of humans while planning a path for object handovers. This kind of combined planning, considering humans as partners, yielded very interesting results [Gharbi 2013]. However, this considered only the static humans and was implemented using a grid-based approach. Different kinds of approach mechanisms were also studied in HAN, and the work by Ramirez et al. [Ramírez 2016] used an IRL based methodology to learn the approach from the human demonstrations. Subsequently, a multi-agent setting was exploited, consisting of several humans and robots. Waldhart et al. [Waldhart 2015a] presented a planner that can find a sequence of handovers involving several agents to transfer an object from the initial

agent to the desired agent. It was built on a graph, with each node representing a possible object position and the agent holding the object, and the edges representing the various paths an object can travel between the nodes. These paths can be of either navigation or manipulation (handover). LAAS has also contributed to some of the studies in HAN. Specifically, the studies regarding the head motion [Gharbi 2015, Khambhaita 2016a] revealed that goal or path-oriented head motion is preferable in navigation, as well as object handover. Moreover, while navigating, the robot’s velocity modulation was found to be preferable, compared to path modulation as studied in [Kruse 2014a, Kruse 2012a]. Coming back to HAN planning alone, Khambhaita et al. [Khambhaita 2017c, Khambhaita 2017b] proposed the idea of proactive planning by considering the human as a cooperative agent and introduced the concept of Human-Aware Timed Elastic Band (HATEB) for trajectory planning. This idea is the basis for this thesis. Unlike the previously proposed frameworks, HATEB does not concentrate much on path planning and mainly deals with trajectory planning and control, combining all the ideas discussed above. Qualitative comparison of HATEB with other planners showed that it produced better trajectories [Khambhaita 2017a] than the other HAN planners. The proposed system also introduced the approach and head movements into HAN. Apart from works on mobile robots, a recent work by Truc et al. [Truc 2022] addressed the human-aware path planning problem in the case of a drone flying in human environments. This thesis is built over the previous work, HATEB, and proposes new ideas and architecture for HAN while improving on the previous ones.

1.2.3 Human-Aware Navigation as a part of HRI

HAN is not just a simple motion planning problem involving dynamic obstacles. It is navigation involving interaction with humans, and hence, it is essentially a special case of HRI. Interaction should be a part of HAN so that the robot can successfully communicate its intention or reach its goal with minimal discomfort to humans. In many HRI frameworks, HAN is treated as one of the many tasks the robot has to do to complete an interaction. For example, Kaiser et al. [Kaiser 1997] present an intelligent robot framework capable of acquiring knowledge and analysing the situation based on simple inputs from the users. However, some of the instructions implicitly involve navigation. Particularly, if the environment has humans in it, the navigation should be human-aware and intelligent. At LAAS, one of the first human-friendly navigation systems [Alami 2000] used HRI as the base and a modular system with multiple tasks. Later this idea was refined in different projects [Foster 2016], parallelly working on a better HAN system. Some recent works like [Vega 2018] and [Vega-Magro 2018] proposed cognitive architecture for social navigation involving decision-making and communication capabilities. Many works in HAN present how situation assessment and behaviour selection can generate acceptable and legible motions for the robot. Therefore, we can say that HAN is similar to HRI, involving decisions about the environment and the humans present in it. Hence, a robot navigating among humans must incorporate the principles of HRI into its design. In this

thesis, we see HAN as a cooperative activity involving humans and the robot, and hence, we apply the joint-action principles in HRI [Curioni 2019] while developing the framework.

1.2.4 General challenges in Human-Aware Navigation

Being a multifaceted problem, HAN planning faces several challenges, and a single framework may not be able to provide a general solution to the problem. Some of the many challenges faced by HAN are given below. This may not be the complete list of challenges, but it comprises the majorly discussed issues in HAN.

- **Human Modelling:** The most basic requirement for any HAN system is to have a model for a human navigating in the given environment. Treating humans as dynamic obstacles is not sufficient as a human have certain expectations and notions about other humans or agents in the environment. Therefore, a special model is required for the human that is independent of place, culture, gender etc., and built over some common ground.
- **Hard to Generalise:** The robot navigation must comply with the social norms of the environment, and one of the major issues is that they change rapidly with human density, geometric context (corridor, door, open area etc.), place (office, warehouse, street etc.) and many other factors. Furthermore, each social norm requires a different way of modelling, and sometimes they could conflict with each other. Hence, it is important to determine which social norms are relevant to the situation at hand. Lastly, people from different backgrounds react differently to the robot, and it adds more complexity to the planning.
- **Need for Decision-Making Capabilities:** Without proper situation assessment and handling, the robot could be contributing to the discomfort of humans rather than reducing it. Unless a HAN is designed to handle only a specific situation, it needs to evaluate the situation and take pertinent actions, which requires decision-making capabilities. The situation analysis also needs to anticipate possible actions and intentions of humans to avoid the occurrence of undesired situations and erratic behaviour of the robot. This requires good predictive models for human motion and intention detection.
- **Communication and Negotiation:** The robot motion planned by a HAN system should not only respect the social norms but also needs to be legible. It requires the robot to show its navigation intention to the humans with or without explicit communication through path or speed changes, gestures like waiting, turning to a side etc., signals and sometimes through voice or video. Explicit communication may also be needed to negotiate in intricate scenarios. The difficulty here is to know ‘when’, ‘what’ and ‘how’ to communicate.
- **No Standard Metrics:** As humans are social beings with different states of mind and backgrounds, it is difficult to come up with metrics that apply to every

situation. Therefore, different researchers in HAN use different sets of metrics for evaluation. Although there are some metrics that are widely used, they are not universally accepted and may be misleading in some contexts.

1.2.5 Addressing the challenges

As HAN has been a topic of research for quite some time, researchers have addressed the challenges in different ways. Human modelling was the first thing to be addressed. So, we start with it and then proceed to robot planning.

1.2.5.1 Human Modelling and Motion Prediction

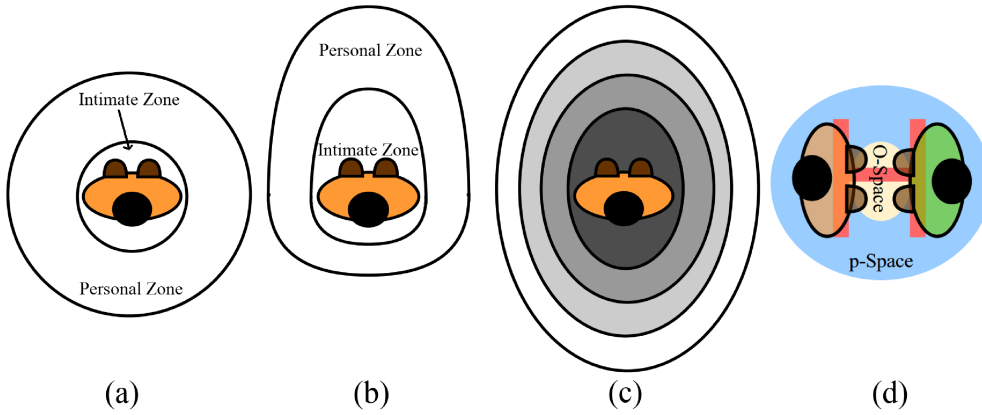


Figure 1.4: Zones and Spaces in Human and Group Modelling.

For modelling humans, proxemics theory [Hall 1966, Mishra 1983] provides some baselines that are common to most humans. Based on this, humans are designed as special obstacles with different zones around them. The very first modelling was done using concentric circles as shown in Fig. 1.4 (a) with four different zones [Rios-Martinez 2015], namely, 1) the public zone ($> 3.6m$), 2) the social zone ($> 1.2m$), 3) the personal zone ($> 0.45m$) and 4) the intimate zone ($\leq 0.45m$). These measures are not very strict and vary with age, culture, background etc. The robot can interact and enter the public and social zones while it is prohibited from entering the personal and intimate zones. Later different shapes are proposed for these zones [Rios-Martinez 2015] and the most commonly used one these days is the *Egg* shape shown in Fig. 1.4 (b) that makes the robot maintain a larger distance towards frontal side compared to the sides or back. This is specifically true in the case of moving humans, and some works [Kostavelis 2016] adapt the shape based on the estimated human velocity. As SFM gained popularity in robotics, human is assumed to have a repulsive potential field in the shape of an ellipse that is monotonically decreasing with a peak at the centre. Fig. 1.4 (c) shows this shape and the decaying potential field. The direction of this field is the direction of motion of the human. As research progressed, different concepts like Information Processing Space [Kitazawa 2010] and Object Affordance Spaces were intro-

duced into the human model. For groups, the concepts of *o-space*, *p-spaces*, and *f-formations* [Rios-Martinez 2015, Kendon 2010] were introduced, which allowed the robot to avoid, join or leave a group of humans appropriately. An example of *o-space* and *p-space* with *vis-a-vis* (or H) *f-formation* is shown in Fig. 1.4 (d).

It is necessary to predict the trajectory and motion of the human for the robot to act or react appropriately. As human detection and tracking is not the main focus of HAN, any off-the-shelf frameworks and systems can be used for this. After getting the necessary information about the humans' positions and velocities, a motion model is required to process these and predict the possible future trajectory. Thanks to the crowd navigation research, different methodologies like constant-velocity (or linear), ORCA, SFM, Social-LSTM [Alahi 2016] etc., were developed, and these are usually employed for the motion models. The trajectory prediction is performed by forward simulating all the agents (humans and the robot) that are currently involved in HAN planning using their motion models (can be the same or different). Generally, the motion model has built-in collision avoidance strategies, and hence, the predicted trajectories are collision-free. However, when simplistic models are used (for example, linear), the trajectory prediction has to take care of collision avoidance as well.

Human intention prediction is necessary to understand what humans might do in the future and take corrective actions. The meaning of intention may not be very clear as one can use it to refer to many things, like whether the human is going to cross the sidewalk or not [Köhler 2012], which direction the human will move or turn next [Peddi 2020], whether the human is willing to interact or not [Ratsamee 2013] and many others. Although intention prediction will greatly benefit HAN, it is not very easy to predict and generalise. The recent advancements in machine learning and computer vision provided tools that can be used to predict intentions, and some works [Peddi 2020, Ratsamee 2013] have already integrated it into HAN.

1.2.5.2 Robot Planning and Decision Making

In the early stages, HAN planning was done by adding the proxemic zones around the humans and then planning a path that avoids intrusions into the personal and intimate zones. This mostly involved grid search (like A*) or potential field based path planning with simple controllers that tracked the path. In most of these settings, the humans and the groups were static while the robot navigates around them. With progress in dynamic obstacle avoidance, different strategies evolved for HAN, like continuous re-planning of the path with updated proxemic zones [Truong 2014] and temporal planning [Kollmitz 2015]. However, scalability is an issue with such global planning strategies, which consider the entire 2D map and all possible states to re-plan or update the path. On the other hand, as discussed previously, trajectory planning was mostly local and used only a small portion of the map and states. This offered a better alternative for doing temporal or reactive planning in HAN, and consequently, a lot of work today focuses on complex trajectory planning using simple path planning. As mentioned previously, the trajectory

planning problem always involves some kind of optimization (minimization or maximization), and naturally, any kind of optimization technique could be employed to solve this problem. The ‘*human awareness*’ is included in the robot’s motion (or trajectory) through the constraints of optimization, and sometimes they are referred to as ‘social or human-aware constraints’. Some commonly used optimization techniques in HAN include MPC [Rösmann 2021], Pareto-optimality [Forer 2018], graph optimization [Rösmann 2013] etc. There are other ways as well in which this problem was addressed, and some popular approaches are SFM, velocity-obstacle based approaches (ORCA, RVO etc.), sampling-based approaches like DWA or RRT and RL based methodologies like CADRL, SA-CADRL etc.

Trajectory planning in HAN should do more than simple collision avoidance with humans and obstacles. Since it is highly dependent on the context, it should best serve the task in the context rather than just avoiding discomfort. Therefore, the system has to be tuned to make the robot’s motion legible and the interactions possible. One way to make robot’s motion legible is through intention show using velocity or path modulation [Kruse 2014b, Lichtenthäler 2013a], signals [May 2015] or visual projections [Shrestha 2018]. All these have to be done while considering and following the social norms in the context. Some norms like ‘passing on right (or left)’ could be employed using costmaps [Lu 2013], but some complex norms like waiting or advancing at a doorway to avoid blocking or giving way (moving to a side) for a human in a hurry require situation analysis and decision-making [Zender 2007, Granata 2012, Feil-Seifer 2011, Hayashi 2012] capabilities. Therefore, some of these works employ state machines that switch between behaviours, such as navigating in a crowd [Dugas 2020], approaching a human [Hayashi 2012], following a human [Zender 2007], standing in line etc., after analysing the situation. Some other recent works use POMDP [Qian 2013] or deep learning [Banisetty 2020] based approaches to include the decision-making capabilities and modality switching into HAN.

In HRI, if the robot needs to communicate something or interact with a human, it needs to estimate ‘where’ and ‘when’ it should approach or meet the human and ‘how’ it should plan its trajectory to be legible to all the humans in the environment (or context). Such things are usually addressed using Inverse Reinforcement Learning [Ramírez 2016] or potential fields [Hansen 2009]. Unlike the in-place communication that occurs in most HRI tasks, HAN planning has to communicate (through voice or gestures) during navigation to negotiate or prevent the occurrence of the FRP. This is usually done by maintaining a knowledge base and a model [Dugas 2020] that decides ‘what’ and ‘when’ to communicate. This again emphasises the need for decision-making in HAN.

1.3 Background and Our Approach

In this section, we provide details on the mathematical modelling and the background that is necessary to understand our approach to HAN planning. We present

the TEB formulation and its hypergraph modelling, followed by the formulation of HATEB and the modelling of social constraints. Following this, the application of the joint-action principles of HRI to HAN is presented. Finally, we discuss our ideology while developing a HAN system and briefly describe the approach.

1.3.1 Timed Elastic Band for Trajectory Planning

Timed Elastic Band (TEB) for trajectory planning was proposed by Rösmann et al. [Rösmann 2013] and is available as one of the local planners⁴ in ROS. This approach uses hypergraph representation to build the timed elastic band and then optimizes the graph to get the optimal command velocity. We briefly provide the mathematical details of this framework here.

1.3.1.1 TEB Formulation

The robot's trajectory in 2D navigation can be represented as a sequence of n poses, $p_i = (x_i, y_i, \theta_i) \in \mathbb{R}^2 \times S^1$ and with a set of $n - 1$ time intervals, $\delta t_i \in \mathbb{R}^+$ between two consecutive poses. Classically, in an elastic band [Quinlan 1993] approach, only the poses were considered along with the kinematic model of the robot, and therefore, the obtained trajectory may not satisfy the dynamic constraints of the robot. On the other hand, TEB considers both poses and time intervals and hence, the resultant trajectory is both kinematically and dynamically feasible. Let the sequence of poses, $\{p_i\}$, and the time intervals $\{\delta t_i\}$ be represented by:

$$\begin{aligned} P &= \{p_i\}_{i=0..n} \quad n \in \mathbb{N} \\ \tau &= \{\delta t_i\}_{i=0..n-1} \end{aligned}$$

Using these sequences, P and τ , TEB is defined as a tuple:

$$B := (P, \tau) \tag{1.1}$$

TEB is then optimized using a weighted multi-objective optimization to obtain the optimal poses and time intervals in real-time:

$$f(B) = \sum_k \alpha_k f_k(B) \tag{1.2}$$

$$B^* = \arg \min_B f(B) \tag{1.3}$$

where $f(B)$ is the objective function and B^* denotes the optimized TEB. The objective function $f(B)$ is a weighted (by α_k) sum of different components $f_k(B)$ that represent different kinodynamic constraints. The constraints are formulated as objectives in terms of a piecewise continuous, differential cost function that penalises the violation of a constraint [Rösmann 2013].

⁴http://wiki.ros.org/teb_local_planner

1.3.1.2 Hypergraphs and TEB Hypergraph

A hypergraph is a generalised graph with hyperedges [Bretto 2013] that can connect any number of vertices (or nodes), unlike a normal graph where an edge connects only two vertices. Mathematically, if H is a hypergraph and V is the set of vertices,

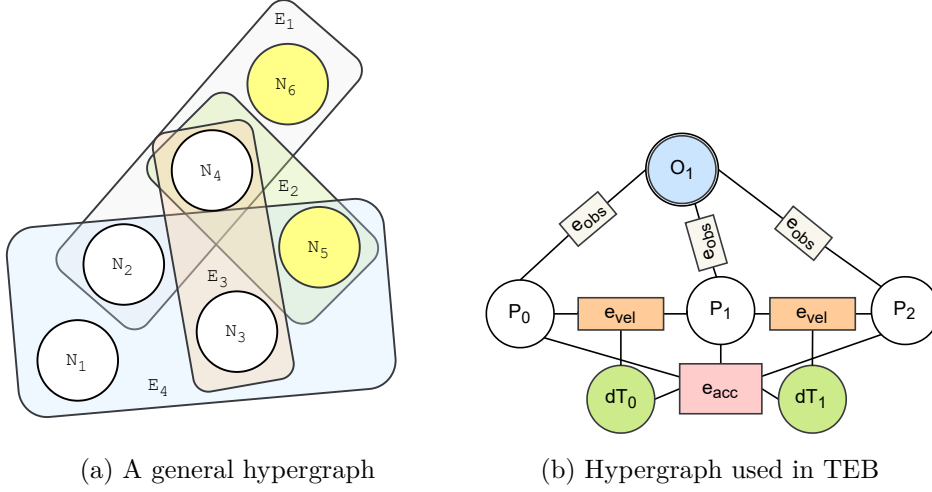


Figure 1.5: HyperGraph

and E is the set of hyperedges, then

$$H = (V, E) \implies E = \{e_i \mid e_i \subseteq V\} \forall i \in I \quad (1.4)$$

where I is a finite set of indices. An example of a hypergraph is presented in Fig. 1.5 (a) with six vertices and four edges. Hypergraphs can contain different kinds of vertices and edges. In Fig. 1.5 (a) each edge is of a different kind (different colours), and there are two different types of vertices (yellow and white).

Using hyperedges, any mathematical relationship between any of the vertices can be defined, and this plays an important role in the HAN system proposed in this thesis. Therefore hyperedges can naturally be employed to model the kinodynamic constraints of the robot, and so TEB is modelled using a hypergraph. The vertices of this hypergraph are the poses and the time differences, while the hyperedges represent different constraints or relations between these vertices. A small part of this hypergraph is presented in Fig. 1.5 (b), containing five vertices (three poses and two-time differences), and shows the velocity, acceleration and obstacle edges. The obstacle is shown in blue with concentric circles. From this figure, it can be seen that the constraints are mostly local, as they depend only on a few consecutive configurations. It is also the case for the majority of the components in the objective function, $f(B)$, of TEB, and hence, the TEB hypergraph has a sparse system matrix. Therefore, this hypergraph can be optimized online and in real time. TEB is optimized using an open-source graph optimization framework, g2o, which we present next.

1.3.1.3 g2o Optimization

g2o (or General Graph Optimization) [Grisetti 2011] is an optimization framework developed for solving graph-based nonlinear error functions. It offers solutions to the problems with the following structure:

$$\mathbf{F}(\mathbf{x}) = \sum_{k=\langle i,j \rangle} \mathbf{e}_k(\mathbf{x}_i, \mathbf{x}_j, \mathbf{z}_{ij})^T \boldsymbol{\Omega}_k \mathbf{e}_k(\mathbf{x}_i, \mathbf{x}_j, \mathbf{z}_{ij}) \quad (1.5)$$

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \mathbf{F}(\mathbf{x}) \quad (1.6)$$

where \mathbf{x} represents the parameters to be optimized, $\mathbf{x}_i, \mathbf{x}_j$ are the block parameters and \mathbf{z}_{ij} denotes the constraint between them. $\boldsymbol{\Omega}_k$ denotes the information matrix between the constraints and $\mathbf{e}_k(\mathbf{x}_i, \mathbf{x}_j, \mathbf{z}_{ij})$ is the error vector between the parameters and the constraint. The optimized set of parameters is represented by \mathbf{x}^* . Note that, the part inside the summation in Eq. (1.5) becomes $\Omega_k e_k^2$ if the error is scalar. In case of TEB [Rösmann 2013], the objective function, $f(B)$ takes the above form where B is the set of parameters to be optimized, $\mathbf{x}, \alpha_k = \Omega_k, e_k = \sqrt{f_k}$ and $\mathbf{x}_i = (p_i, \delta t_i)$. g2o uses the Levenberg-Marquardt method to optimize the objective function numerically and offers two Cholesky decomposition solvers, CHOLMOD and CSpase, to solve it efficiently.

1.3.2 Human-Aware Timed Elastic Bands for Proactive Planning

In this section, we present the idea of Human-Aware Timed Elastic Band (HATEB), proposed by Khambhaita and Alami [Khambhaita 2017c] that lays the basis for this thesis. We briefly discuss how the objective function and hypergraph are modified to include human-aware constraints in robot navigation.

1.3.2.1 HATEB Formulation and Hypergraph

The idea of HATEB is to add elastic bands to all the humans in the robot's vicinity in addition to the robot and jointly optimize the robot and the human trajectories under kinodynamic and human-aware (or social) constraints. This gives rise to a multivariate multi-objective optimization problem to solve, and the same weighted sum approach proposed in TEB is used for HATEB as well. Therefore, if the robot's band is represented by B_r and the human bands by $\{B_{h_k}\}_{k=0}^n$ when $n \in \mathbb{N}$ humans are present in the vicinity, the new objective function to solve is,

$$f(B_r, B_{h_0} \dots B_{h_n}) = \underbrace{\sum_i \alpha_i f_i(B_r)}_{F_R} + \sum_k \left(\underbrace{\sum_j \alpha_j f_j(B_{h_k})}_{F_H} + \underbrace{\sum_l \alpha_l f_l(B_r, B_{h_k})}_{F_S} \right) \quad (1.7)$$

$$B_r^*, B_{h_0}^* \dots B_{h_n}^* = \arg \min_{B_r, B_{h_0} \dots B_{h_n}} f(B_r, B_{h_0} \dots B_{h_n}) \quad (1.8)$$

where B_r^* , $\{B_{h_k}^*\}_{k=0}^n$ are optimized trajectories for the robot and the humans respectively. The initial paths of humans for this optimization is obtained from the human path prediction module of the system.

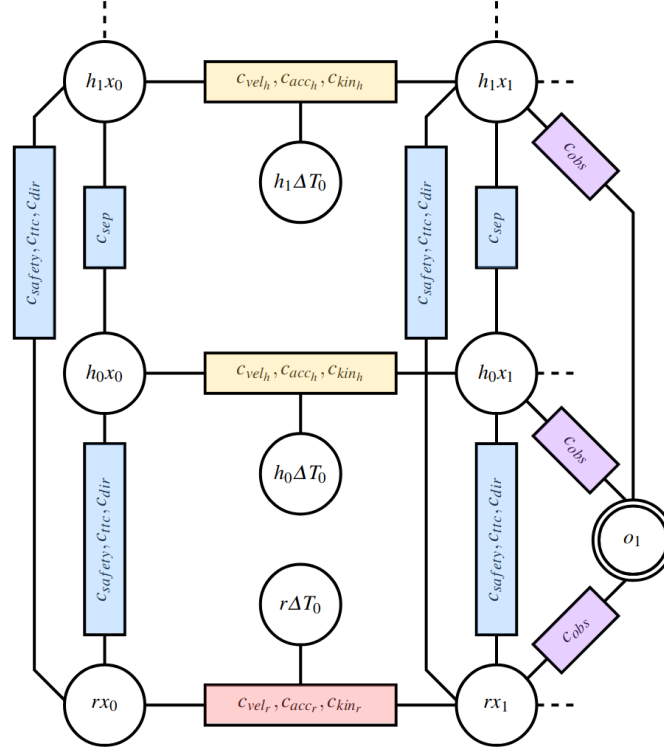


Figure 1.6: The updated hypergraph of HATEB [Khambhaita 2017c]. There are 3 bands, one for the robot and the other two for humans. The kinodynamic constraints are shown as horizontal edges (yellow, red), social constraints are shown as vertical edges (blue), and the obstacle avoidance constraints are shown as diagonal edges (purple).

In Eq. 1.7, F_R and F_H represent the objective functions for the robot and human trajectories, whereas F_S denotes the objective for human-robot social constraints. All these constraints are represented as hyperedges of an updated hypergraph, and a small part of this hypergraph is shown in Fig. 1.6. The social constraints between human-robot and human-human are highlighted in blue, while the kinodynamic constraints of the robot and humans are highlighted in red and yellow, respectively. Finally, the obstacle constraints are shown in purple. Even though the system matrix is still sparse, it is denser than before, and as we keep adding bands for humans, the density grows, and the trajectories cannot be obtained in real-time.

1.3.2.2 Proactive Planning and Social Constraints

The joint optimization produces the robot trajectory that obeys the social norms implemented in F_S . Further, the combined human-robot trajectory planning makes

the robot proactively estimate humans' plans and quickly adapt its trajectory before it is too late. We call this idea of adding double-sided (robot side, human side) bands as '*Dual Band*' throughout this thesis. This approach always elicits solutions (provided they exist) for all agents, which can solve the navigation scenario at hand if every agent follows its trajectory. Using the above architecture, three social constraints are defined, 1) *Human-Robot Safety*, which takes care of human safety and proxemics, 2) *Time-to-Collision (TTC)*, which adds early intention show to the robot's trajectory and 3) *Directional*, that tries to establish a trade-off between slowing down and changing the path. In HATEB, the objective function for safety clearance from obstacles is also updated to make the robot move towards obstacles rather than humans in confined scenarios. All these constraints are added to the objective function as piecewise continuous and differential error functions, like in TEB. We believe that this architecture is better suited to solve the HAN problem than the reactive approaches, especially in intricate indoor scenarios.

1.3.3 Principles of Joint-Action in Human-Aware Navigation

As mentioned previously, HAN can be seen as a part of HRI and therefore, we believe that it should possess some of the properties of HRI. Specifically, during the development of this thesis, the principles of joint-action [Curioni 2019] in HRI were applied to HAN. These principles include: sharing a common perspective, coordinating, predicting others' contributions and communicating. In terms of human-aware robot navigation, these can be seen as:

- *Share a common perspective*: The robot should know the social norms and expectations of a human in a given environment. For example, moving on the right, avoiding collisions with humans and other objects etc.
- *Coordination*: The robot and human should coordinate and cooperate with each other to reach their navigation goals. For example, the robot should give way for the human to pass in a door crossing scenario instead of blocking.
- *Predict others' contribution*: The robot should predict humans' motion (trajectory) and/or intentions to know how much the human is willing to contribute to the navigation task. For example, in a corridor, the robot facing a human can either continue with or change its path depending on whether the human is willing to change his/her path or not.
- *Communication*: The robot should be able to communicate to humans what action it is going to take or sometimes take permission or inform humans before taking an action. Communication in HAN can be split into two types, 1) Communicating navigation intention (for example, signals) and 2) Communication through speech or video to negotiate in a complex scenario.

These principles impose more restrictions on the robot's trajectory and require the inclusion of decision-making capabilities into robot navigation planning. This makes HAN a complicated problem to solve in the motion planning paradigm.

1.3.4 Our Approach and Contribution

One of the noted limitations with *Dual Band* approach is that it assumes that humans are always cooperative agents and will not cause any deadlock situations. So it always proposes solutions with the assumption that the human will be moving, which might not be the case. When the human stops moving or behave unconventionally, the robot gets stuck without moving or oscillates without making any progress towards the goal. Therefore, we introduce situation assessment over proactive planning as shown in Fig. 1.7. This situation analysis is local and happens at the trajectory planning level. After analysing a situation and determining if it's a

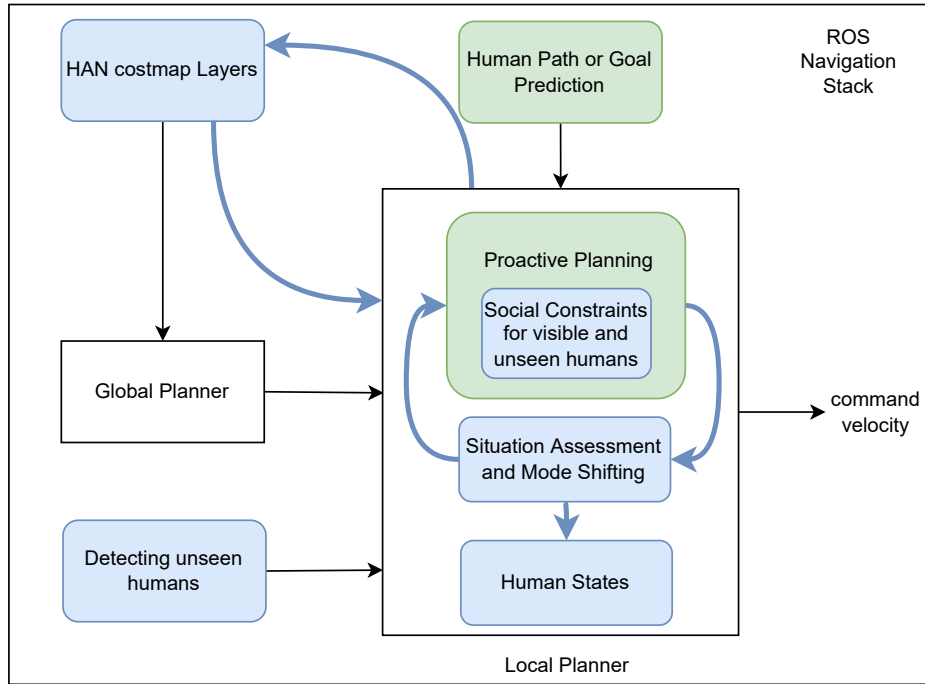


Figure 1.7: The situation assessment module is combined with proactive planning inside the local planner. The system can use any existing global planner.

deadlock or any other uncomfortable setting, we have to find a way to mitigate it. In this thesis, we do this by shifting between different navigation modalities. The blue arrows in Fig. 1.7 represent how proactive planning, situation assessment and mode shifting are interlinked with each other. The integration of situation analysis into trajectory planning and the process of mode shifting are presented in detail in the subsequent chapters.

The crux of this thesis is finding and mitigating the uncomfortable human-robot interactions in HAN. Note that this includes the humans in view and the humans the robot has not seen yet. The main contributions of this thesis in developing a human-aware ROS navigation stack are shown in coloured boxes in Fig. 1.7. The ones in blue colour are the new contributions, while the ones in green represent modifications over the previous work. The blue arrows stand for the new connections

we have made in our proposed HAN. Proactive planning has been significantly improved by proposing new human-robot social constraints for both visible and unseen humans. The updated path prediction module with new goal prediction methodologies further improves proactive planning, and hence, better navigation behaviours can be obtained. The new HAN costmap layers are proposed to specifically address the global planning around static humans. Yet, these layers update based on the states of humans (blue arrows) given by the local planner and play a role in trajectory planning as well. While analysing a situation, the state of the human (static, moving, stopped etc.) is necessary, and hence, situation assessment modules maintain a list of states for the humans it has encountered during its navigation. The choice of modality is based on the state of the human as well as the situation. The idea of mitigating collisions with unseen humans is quite new in the case of HAN and makes the robot ready to face any kind of situation when combined with the multi-context HAN system proposed in this thesis. We believe that the idea of mode shifting within local planning is significantly new and can be seen as a major contribution as well.

HATEB-2 for Legible Proactive Planning with Situation Analysis

Contents

2.1	Introduction	37
2.2	Related Work	38
2.3	Proactive Planning in HATEB	39
2.3.1	The Entanglement Problem	40
2.4	HATEB-2: Situation Analysis and Mode Shifting	41
2.4.1	Formulation and Implementation	41
2.4.2	Modes of Planning	41
2.4.3	Mode Shifting based on Situation Assessment	42
2.5	Improving the Legibility in HATEB-2	45
2.5.1	New Human-Aware Constraints	45
2.5.2	Modified Elastic Band	46
2.5.3	Better Human Predictions	47
2.6	Results in Simulation	48
2.6.1	Entanglement Resolution through Mode Shifting	48
2.6.2	Advantages of Proactive Planning	49
2.6.3	The Effect of New Human-Aware Constraints	51
2.6.4	Quantitative Comparison between HATEB and HATEB-2	51
2.7	Real-World Tests	53
2.8	Conclusion	55

2.1 Introduction

Human-Aware Timed Elastic Band (HATEB) [Khambhaita 2017c] includes human predictions by simultaneously planning for humans and the robot. This allows HATEB to handle intricate situations like narrow corridor crossing and door crossing, where human and robot cooperative motion is needed. With the increasing complexity of environments and the need to navigate robots in such environments,

decision-making has been introduced into planning [Qian 2013, Mehta 2016]. However, these frameworks might make the robot wait in confined spaces instead of proactively planning, hence, resulting in larger execution times. Therefore, in this chapter, we propose HATEB-2, a new modality based human-robot co-navigation framework that uses decision-making to solve semi-crowded as well as intricate scenarios. This is achieved by shifting between different modalities and simultaneous human-robot planning, similar to HATEB. This chapter discusses the following three things in detail: 1) HATEB-2, a new human-robot co-navigation planner comprising decision-making. 2) Improvements and modifications to HATEB. 3) Analysis of human-robot co-navigation in a variety of situations.

Note that this chapter discusses only the first version of HATEB-2 and its implementation. The subsequent chapters of this thesis discuss the evolution of this planner over time. Throughout this chapter, humans are represented as green cylinders with yellow arrows, where the direction of the arrows corresponds to the frontal direction of humans.

The chapter is organised as follows. Section 2.2 presents the related work corresponding to this chapter. Section 2.3 briefly presents the proactive planning in HATEB and the ‘*entanglement problem*’. Section 2.4 presents the architecture of HATEB-2, describes the different modes of planning and finally explains the modality-shifting process based on the situation assessment. Following this, section 2.5 presents the modifications to HATEB and the new human-aware constraints. The results and the analyses of the proposed architecture in various simulated experiments are presented in section 2.6, and the tests in the real world are presented in section 2.7. Section 2.8 finally concludes this chapter.

2.2 Related Work

Most state-of-art HAN planners use proxemics zones around humans in a grid-based map representation of the robot’s working environment [Kruse 2013]. Nonetheless, proxemics alone may not be sufficient to generate completely human-acceptable motion for the robot. Sometimes planning approaches like SFM [Helbing 1995, Ferrer 2013a] and velocity-obstacle models [Snape 2011, Berg 2011] are also used for HAN. As these approaches are reactive, they do not necessarily produce legible robot motion. In the human-aware navigation planner proposed by Sisbot et al. [Sisbot 2007c], other social criteria like visibility and hidden zones were considered along with the proxemics. Another framework proposed by Kruse et al. [Kruse 2012b] introduced the directional cost model, which attempts to solve the spatial conflict by adjusting velocity instead of the path whenever possible. The study conducted by Kruse et al. [Kruse 2014b] showed that humans prefer the robot to follow this strategy, especially in path-crossing scenarios. This model has also been shown to increase the legibility of the robot motions, and hence, in HATEB-2, we introduced some new constraints that restrict the path change and adjust the velocity based on the distance between the human and the robot.

Employing social constraints alone may not be sufficient to develop a socially acceptable navigation planner, and this raises the need for including human motion predictions into the framework [Kuderer 2012]. Many methods based on the SFM predict homotopically distinct trajectories for humans and design planners that learn the navigation policies for the robot based on human demonstrations [Kuderer 2012]. Although these methods involving independent human predictions work fine in large open spaces, they might require to re-learn the parameters to handle situations such as passing through a long corridor or a door, where coordination is needed between humans and the robot. Hence planning for humans along with the robot is required in such situations. The approach presented by Ferrer et al. [Ferrer 2015] uses the SFM, both for predicting human paths and controlling the robot’s motion. In this approach, human predictions based on the previously planned paths were used. Other approaches [Bordallo 2015, Nagariya 2015] try to predict the possible human goals based on some type of reasoning and generate locally optimal motion for the robot. One of the recent approaches [Fisac 2018] suggests the use of probabilistic human predictions to handle various uncertainties and plan the robot’s motion on top of these probabilistic predictions. This approach is particularly useful in systems with unreliable sensors. All these approaches are effective in densely crowded environments as a virtue of remaining purely reactive but can lead to needless detours in intricate situations. Our previous work [Khambhaita 2017c] is specifically developed to handle such intricate situations in semi-crowded environments. Such planning for humans along with the robot is usually required in robot-human handover scenarios to know where to perform a task and who performs a task [Mainprice 2012b, Waldhart 2015b].

The concept of modality shifting in human-aware navigation is discussed in works by Mehta et al. [Mehta 2016] and Qian et al. [Qian 2013], where POMDP is used for decision making. In both works, different modalities necessary for human-aware navigation are proposed, assuming that the robot takes all the load of the navigation process. Hence these methods may also suffer problems like purely reactive planners in complex situations leading to unnecessary detours or long halts. HATEB-2 includes HATEB as one of the modalities and hence can handle both intricate as well as semi-crowded scenarios by switching between different modalities when needed. In this chapter, however, we focus mainly on different intricate situations involving cooperative motion between the human and the robot.

2.3 Proactive Planning in HATEB

As presented in Chapter 1, the formulation of HATEB allows us to plan for the robot and the humans involved in the interaction simultaneously. As this approach plans for all the agents, the robot’s trajectory can be proactively updated before something happens. Hence, it can solve human-robot navigation scenarios in a better way when compared to reactive-only schemes. However, some assumptions made about humans during this formulation and the structure of its implementation

can lead to deadlock situations, as shown in Fig. 2.1. We call it the ‘*entanglement problem*’ throughout this thesis, and its details are presented below.

2.3.1 The Entanglement Problem

In HATEB, the robot’s trajectory is adapted based on the predicted humans’ (planned) trajectories. While planning, it considers a constant velocity for humans and assumes that they are always in motion. Hence, if the human stops moving, the system keeps anticipating the human movement in the immediate time step. One more issue is that the path prediction system for humans is called only once at the start of the planning, and it puts more burden on trajectory planning to handle the discrepancies. The assumptions about the agents and this planning strategy result in the ‘*entanglement*’ of trajectories when the human stops moving, and the robot is stuck, waiting for the human to move. These situations can occur commonly in corridors, and two such situations are shown in Fig. 2.1. The optimization scheme neglects the other possible solutions as these trajectories have the minimum cost. This leads to another case of the Freezing Robot Problem (FRP) even when there is enough space to move. We address this issue in HATEB-2 by detecting such situations and then taking mitigating actions.



Figure 2.1: Robot getting stuck due to *entanglement* of trajectories. The poses in blue correspond to the human’s trajectory and the ones in red to the robot’s trajectory. In the situations shown, there exists an alternate solution for the robot to solve the problem. However, the assumption about the human having non-zero velocity and the safety constraint makes the robot wait in the same entanglement, speculating the motion of the human. The picture on the left is of open space, whereas the one on the right is of a narrow corridor.

2.4 HATEB-2: Situation Analysis and Mode Shifting

The proposed framework, HATEB-2, combines decision-making and planning into a single framework and opens up new frontiers for HAN planning. This new framework can encompass a large variety of problems by allowing the transition between different modalities based on context. In this chapter, however, we study only the human-robot co-navigation problem. HATEB-2 introduces decision-making on top of planning, and this makes the planner adapt better to the situations at hand. These situations might be very different from each other and need to be handled differently. A single way of planning may not be sufficient in such situations, and hence, we introduce three different modalities into planning. All these modalities use Timed Elastic Band (TEB) [Rösmann 2013] as their base, and the transition between them is handled by the decision-making loop.

2.4.1 Formulation and Implementation

HATEB-2 uses the same hypergraph structure and formulation used by HATEB for proactive planning and introduces a situation assessment module over this. Depending on the planning modality chosen, the system uses HATEB with different settings and human predictions. In fact, ‘*Dual Band*’ is used as one of the modalities and encompasses a large part of human-robot co-navigation planning. The other modalities may possess the same human-robot social constraints as HATEB but differ significantly in their behaviour. We have also made a few modifications to HATEB to remove its drawbacks and improve legibility before using it in HATEB-2, which are presented in Section 2.5.

HATEB-2 is implemented in ROS and integrated as a local planner in the ‘*move_base*’ package¹ of Navigation Stack. HATEB-2 follows the same software architecture as HATEB and uses a global planner for human path prediction based on an assumed or predicted human goal. This version of our HAN planner is available on GitHub at https://github.com/sphanit/hateb_local_planner/tree/hateb_new. However, this only deals with robot navigation, and we still might have to depend on some other software package to move humans in simulation while testing the system. For this, a separate human navigation package is developed using ROS (see Appendix A) that allows a human agent (in simulation) either to directly execute the trajectory planned by HATEB-2 or follow the control command sent via a Joystick. We now proceed to the explanation of different modalities in HATEB-2 and the decision-making process involved.

2.4.2 Modes of Planning

HATEB-2 operates mainly in three modes of planning: 1) ‘**Single Band**’, 2) ‘**Dual Band**’ and 3) ‘**VelObs**’. However, an intermediate mode is present before the occurrence of **Dual Band** → **VelObs**² transition. The intermediate mode refers to

¹http://wiki.ros.org/move_base

²→ represents one-sided transition

the trajectory planning in the close vicinity of the human (for human to robot distances $\leq 2.5m$), where large velocity changes are restricted, and the elastic band is made tighter. More details about these changes are presented in the next subsection. Different modes of this framework below are explained below.

2.4.2.1 Single Band Mode

In the mode of planning, the elastic band is added only to the robot to avoid obstacles in the environment. This mode is computationally less expensive as it does not deal with human estimates and trajectory predictions (or proactive planning). This mode can be seen as purely reactive planning with social constraints on robot navigation. In this work, this mode is used only when there are no humans in the vicinity, or they are far from the robot. However, as robot planning still includes human-robot social constraints, this mode could be extended to semi-crowded situations with some modifications.

2.4.2.2 Dual Band Mode

This mode is the same as standard HATEB, where multiple elastic bands are added to humans and the robot, and the combined hypergraph is optimized to get the trajectories for all the agents. However, a few modifications are made before using it in HATEB-2, which are explained in the next subsection. This mode adapts trajectory planning according to the motion of the humans and the predicted goals. The main advantage of this mode is that it always proposes a possible solution from the current scenario besides being proactive. The main drawback of this mode of planning is the ‘*entanglement problem*’ discussed above. To overcome this drawback and continue the robot navigation, the ‘**VelObs**’ mode is defined, and a switching strategy is employed to allow this transition only under certain conditions.

2.4.2.3 VelObs Mode

This mode of planning also adds elastic bands to both humans and the robot, but the path prediction and the trajectory planning for humans are performed only when the humans are moving (have a non-zero velocity). This modality uses the linear (or constant velocity) prediction for humans under a predefined time window of prediction. Throughout this thesis, the duration of this prediction window is taken as 5s. Although this kind of human prediction makes the robot less proactive, it allows for an active re-planning when the human stops moving to mitigate the FRP in many situations. Hence, it can be used to resolve situations of *entanglement*, once they are identified.

2.4.3 Mode Shifting based on Situation Assessment

Now we move on to the explanation of the decision-making process involved in transitioning between the above modalities. Since HATEB-2 has three modes of

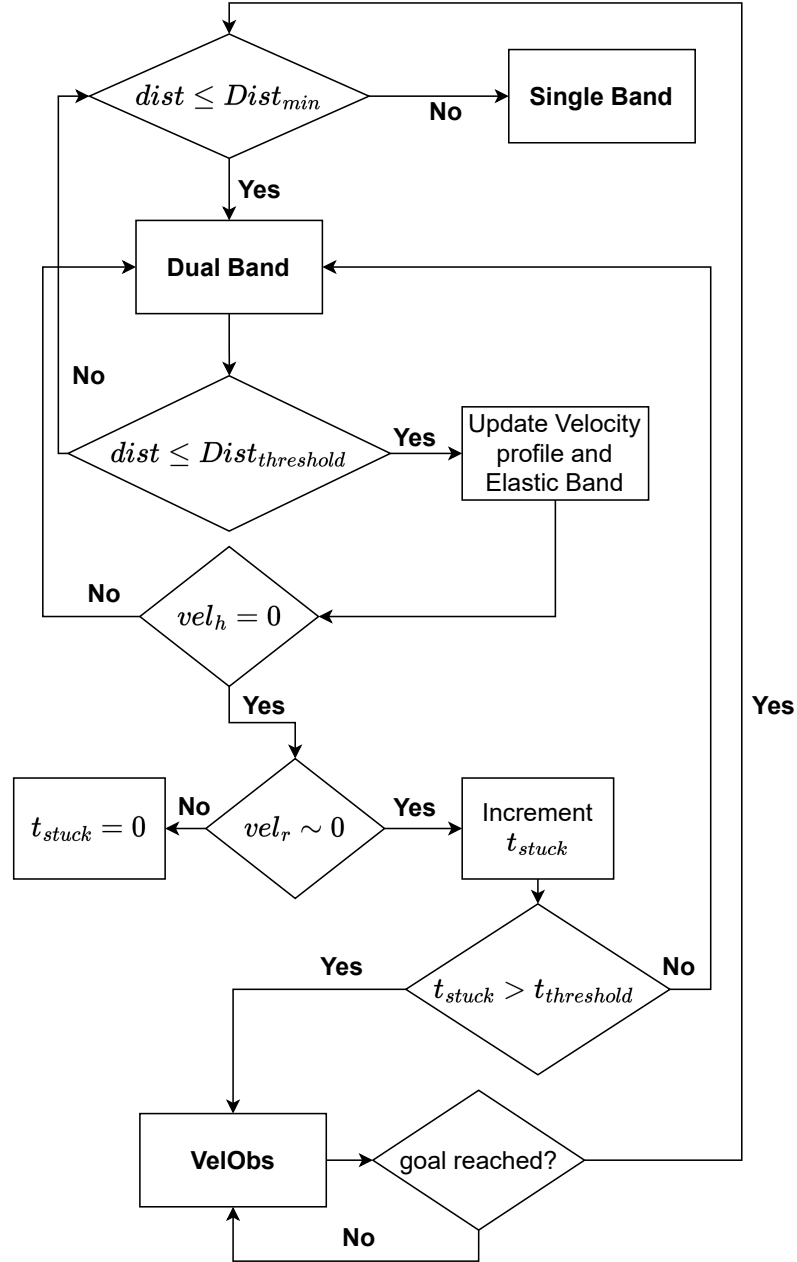


Figure 2.2: Mode transition procedure. $dist$ is the current distance between the closest human and the robot, $Dist_{min}$ is the minimum value of $dist$ to add a double band and $Dist_{threshold}$ is the minimum cutoff $dist$ to initiate transition between **Dual Band** and **VelObs**. vel_h , vel_r are the velocities of the human and the robot. t_{stuck} is the amount of time the robot's velocity is nearly zero, and $t_{threshold}$ is the time to wait before triggering the transition to **VelObs** mode. Note that, under $Dist_{threshold}$, the elastic band and velocity profile are changed irrespective of the mode of planning. The transition loop resets after reaching the navigation goal.

planning, it has the following two mode transitions:

1. **Single Band** \leftrightarrow **Dual Band**³
2. **Dual Band** \rightarrow **VelObs**.

The transition procedure and decision-making loop are presented in Fig. 2.2. The decision concerning the transition from **Single Band** to **Dual Band** is dependent on a cutoff distance, $Dist_{min}$. $Dist_{min}$ is the distance between human and robot, above which the influence of humans on the robot's trajectory is negligible. If the current distance of the robot from any human is less than $Dist_{min}$, the planning shifts from **Single Band** to **Dual Band** and vice versa. Therefore it is a double-sided transition, and it is indicated by ' \leftrightarrow '. $Dist_{min}$ can be chosen based on several factors and the present environment of the robot. During the development of this thesis, we have taken a distance of 10 m for $Dist_{min}$.

From Fig. 2.2, we can see that the second transition only occurs when humans and the robot are under a specified distance, called the $Dist_{threshold}$. Under this $Dist_{threshold}$, the robot's maximum velocity is reduced, and the homotopy class change is constrained to make the robot's motion more legible for humans. The weight of the proxemics constraint is also reduced under this distance to allow the planner to find a solution in near proximity to humans. $Dist_{threshold}$ is taken as 2.5m for the most part of this work. This mode of **Dual Band** with all these changes is the intermediate mode we mentioned previously. Under this intermediary mode, the transition from **Dual Band** to **VelObs** mode occurs when the following situation is detected:

"The human under co-navigation with the robot stops proceeding towards the predicted goal, and the robot either stops or oscillates near this human for more than a specified amount of time without any progress towards the goal".

The amount of time chosen is arbitrary and can be tuned depending on the context. As it can be seen from the transition loop, the above condition is tested using the human's velocity (vel_h), the robot's velocity (vel_r) and a time threshold ($t_{threshold}$). For all the experiments presented in this thesis, the waiting time to trigger this shift is taken as 2 s, i.e., $t_{threshold} = 2$ s. Therefore, the robot does not stay frozen for long and quickly resolves the '*entanglement problem*'. The transition from **Dual Band** to **VelObs** is one-sided, indicated by ' \rightarrow ' and does not happen the other way around. As this transition occurs mostly at the human-robot crossing, it is intuitive to assume that the human would be behind the robot and no longer interferes with the robot's trajectory after the transition. This assumption also reduces the cost of computation as we no longer plan for stationary humans. The mode transition loop resets once the robot reaches a navigation goal. If the goal change occurs in between, the robot continues to stay in the same mode as before the change, and the transitions may (**Single** or **Dual**) or may not (**VelObs**) occur.

An important point to emphasise here is that all the situation analysis and mode shifting happens within the local trajectory planning during the robot's navigation

³ \leftrightarrow represents two side transition

to a specified goal. Therefore, HATEB-2 inherits all the advantages of proactive trajectory planning while reducing its limitations. This method of planning is different from the ones proposed by [Qian 2013] or [Mehta 2016], where there is a higher level of decision-making that chooses between different possible navigation modes (or actions). Our system shifts between planning modes rather than navigation modes. HATEB-2 also incorporates some changes and modifications to the previous version of HATEB to improve legibility, and we present these in the next section.

2.5 Improving the Legibility in HATEB-2

In addition to the situation assessment at the planning level, HATEB-2 proposes some new human-robot social constraints for HAN and modifies some modules in HATEB. All these proposals and modifications aim to increase the legibility of the robot’s navigation and hence, can lead to an increase in acceptability. We present the new human-aware constraints (or human-robot social constraints) before moving on to the other modifications.

2.5.1 New Human-Aware Constraints

We propose two new human-aware constraints for HAN that increase the legibility of the robot’s navigation. The first constraint is inspired by the work presented in [Kruse 2014b]. This constraint restricts the rapid velocity changes in the human vicinity as concluded by [Kruse 2014b]. The second one is an improved *Time-to-Collision* (*TTC*) constraint that was introduced by HATEB.

2.5.1.1 Updated Robot Velocity Constraint

In the place of a constant velocity, we assigned a non-linear profile for the velocity of the robot, which slows down the robot up to 75% in the close vicinity of the human. This change reduces the rapid change in velocity around humans and leads to lesser confusion. The velocity function used in this constraint is given as follows:

$$v(d) = \min(1.0, \max(10^{d-2}, 0.25)) \quad (2.1)$$

where d is the distance between the human and the robot and v is the velocity of the robot. The above equation induces a rapid decrease in the maximum attainable velocity of the robot from 1 m/s at $d = 2$ m to 0.25 m/s at $d \leq 1.4$ m (personal zone in proxemics).

2.5.1.2 TTCPlus Constraint

Time-to-collision, as per its name, calculates the time the robot takes to collide with the human from the current position and velocity. The main advantage of including *TTC* constraint is better trajectory planning with early intention-show by making

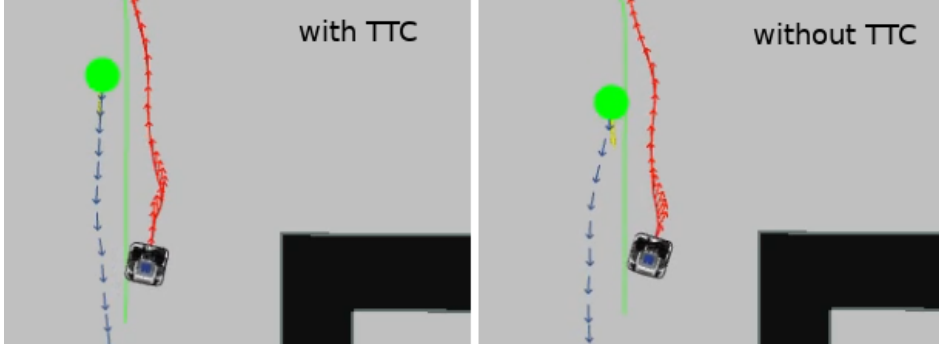


Figure 2.3: Intention show of the robot with and without *TTC*. Inclusion of *TTC* or *TTCplus* constraint results in an early intention show as can be seen from the picture on left. Even though the difference seems small from the pictures, this corresponds to more than half a meter in the real world. The robot’s trajectory is shown in red, while the human’s trajectory is shown in blue.

the robot move quickly towards the intended direction of its motion. This can be seen in Fig. 2.3. Although the difference seems small in the image, it is more than half a meter in the real world. The original implementation in [Khambhaita 2017c] computes the error at every time step and adds it to the optimization. This implementation results in many false negatives that affect the quality of the trajectory. To decrease the number of false alarms while maintaining the advantages of the constraint, we have to regulate the addition of the error to the optimization. The regulation is implemented in HATEB-2 as follows:

$$error = \begin{cases} tt_{error} & \text{if } t_a > t_d \text{ and } t_m < \kappa t_d \\ 0, & \text{otherwise} \end{cases} \quad (2.2)$$

where t_d is the threshold time, t_a is the cumulative time with positive tt_{error} and t_m is the cumulative time with zero tt_{error} . t_a is reset whenever $t_m \geq \kappa t_d$ and t_m is reset when a positive error is observed. κ is a constant determining how often the tt_{error} is added to the optimization, and in this work, we take $\kappa = 5$ based on empirical analysis. This new implementation shown in Eq. (2.2), called *TTCplus*, decreases the number of false negatives, avoiding unnecessary oscillations and improving the quality of trajectory. Hence, *TTCplus* eliminates the drawbacks of *TTC* while preserving its properties of early intention-show.

2.5.2 Modified Elastic Band

The default settings of the elastic band allow it to rapidly change the homotopy⁴ class. This change is good when the robot is at a large distance from humans. At closer distances, it may lead to the loss of legibility, as the study in [Kruse 2014b]

⁴In mathematics, two paths are homotopic if their endpoints are fixed, and one path can be continuously deformed into another within a specified region.

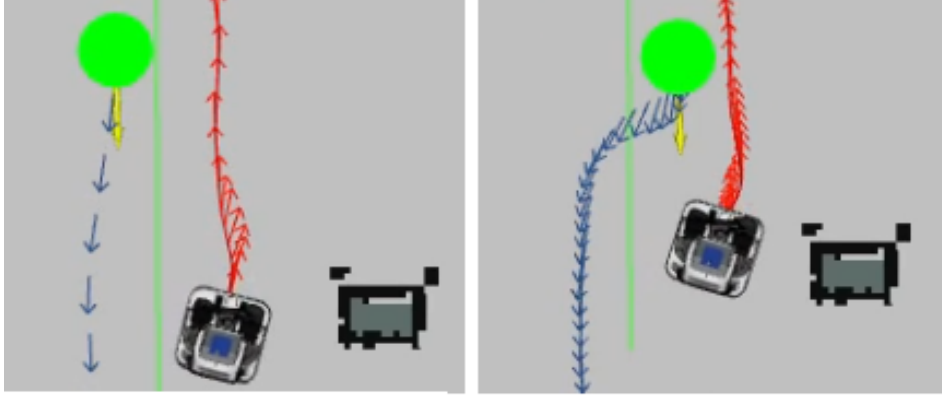


Figure 2.4: Modified Elastic band. The blue trajectory is of the human, and the red one is of the robot. Tightening of the elastic band results in very close pose predictions in the trajectories and slower velocities. The picture on the left shows the trajectories before the band tightening, and the one on the right shows the latter.

says. Hence, we address this issue by restricting the changing of the homotopy class under the threshold distance, *DistThreshold*. This restriction is implemented by decreasing the time interval between consecutive poses, which tightens the elastic band and decreases the possibility of homotopy class change, as shown in Fig. 2.4. This figure shows the snapshots of the trajectories before and after the online band tightness modification. The picture on the right side clearly shows a significant increase in the trajectory resolution, and hence, results in smoother path transitions.

2.5.3 Better Human Predictions

The better the estimates of humans, the better the proactive planning can plan the joint trajectories. Therefore, we have made the following improvements in HATEB-2 to have better estimates and path predictions for humans.

2.5.3.1 Human Velocity Estimate

The nominal velocity of a human was assumed to be constant in the previous work. However, this assumption leads to a trajectory plan that does not necessarily comply with the current human trajectory. Although HATEB-2 being a proactive planner, quickly re-plans and adapts, this wrong estimate can sometimes lead to unexpected behaviours of the robot. Therefore, a moving average filter-based estimation of velocity is added to the human prediction to avoid this and provide an adaptive velocity estimate for optimization.

2.5.3.2 Human Goal Prediction

In HATEB, human goals were assumed to be behind the robot for path prediction. However, this is not necessarily true in many cases, and the human destination

might even change during the navigation. Therefore, we try to include better human goal predictions into the system and address the changes in the human goal during the planning process. If we know a set of possible destinations for humans in an environment, we can predefine a set of goal positions for humans. Ferrer et al. [Ferrer 2014] proposed a methodology to estimate the possible goal among these predefined goals based on the history of motion of the human. HATEB-2 adopts this goal estimation framework to improve human estimations. It quickly adapts to the changes in the human goal and re-plans, as shown in Fig. 2.5, making it more adaptable to real-world scenarios.

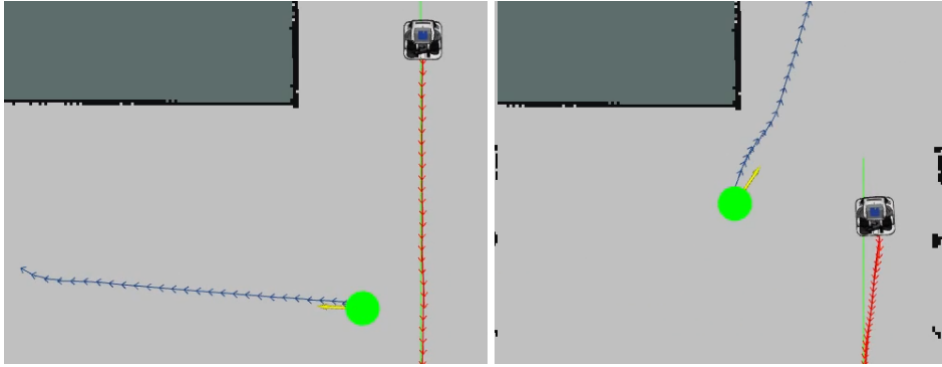


Figure 2.5: Human goal prediction in HATEB-2. On the left side, the initial predicted goal, and the calculated trajectory by HATEB-2 are shown. On the right, the human decides to move in a different direction, and HATEB-2 predicts a new possible goal and calculates the path. The robot’s trajectory is shown in red, while the human’s trajectory is shown in blue.

2.6 Results in Simulation

Various experiments are conducted using PR2⁵ robot and the simulated humans in MORSE [Echeverria 2011] to demonstrate the capabilities of HATEB-2. Two different environments are used to simulate different human-robot navigation scenarios. Both *Qualitative* and *Quantitative* analysis is performed, and the results are presented. The first three experiments presented in this section show the *Qualitative* analysis, highlighting the improvements in HATEB-2 and the roles of situation assessment and proactive planning in HAN. In all three experiments, the human agent is manually controlled using a Joystick.

2.6.1 Entanglement Resolution through Mode Shifting

One of the main drawbacks of HATEB is the *entanglement* issue presented previously. With the introduction of decision-making and mode transitioning in HATEB-2, this *entanglement problem* is resolved, and the robot finally reaches the goal with-

⁵<http://wiki.ros.org/Robots/PR2>

out getting stuck. The various stages of this *entanglement* resolution are illustrated in Fig. 2.6. When the human stops moving and block the robot, a new trajectory is planned, and the transition from **Dual Band** mode to **VelObs** mode occurs (Fig. 2.6 (a, b)). Since the human’s velocity is zero, the **VelObs** mode does not plan any trajectory for the human until he starts moving again. During this time, the robot escapes from the *entanglement* and starts following its trajectory to the goal, as seen in Fig. 2.6 (c) and (d).

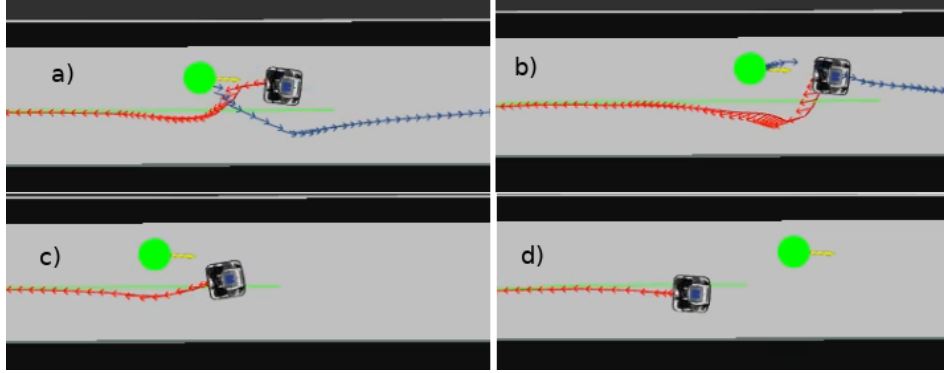


Figure 2.6: HATEB-2 solving the *entanglement*. The various stages of *entanglement* resolution are as shown: a) Detection of Entanglement: The *entanglement* is detected based on the human velocity and the current distance between human and robot. b) Re-planning: HATEB-2 tries to re-plan the trajectory before changing the mode c) Mode Transition: Mode transition occurs as the human is still and no longer moves. d) Execution of new plan: Finally, the robot executes the planned trajectory reaching the expected goal. The robot’s trajectory is shown in red, while the human’s trajectory is shown in blue.

2.6.2 Advantages of Proactive Planning

We present two experiments to show the advantages of proactive planning in HAN. The second experiment tests proactive planning and situation analysis in a complex setting. It shows how HATEB-2 solves the navigation problem in a more legible way.

2.6.2.1 Single Band vs Double Band

The single band refers to the addition of an elastic band only to the robot along with the human-aware constraints, whereas the double band (or dual band) includes the addition of an elastic band to the human as well. We test the following hypothesis to see if the double band has any advantage over the single band:

“The presence of an elastic band for human and co-planning allows the robot to predict the human motion better and adapt its trajectory accordingly.”

To test the hypothesis, we controlled the human manually, tried to block the robot’s trajectory, and observed the reactivity of the robot. We conducted two

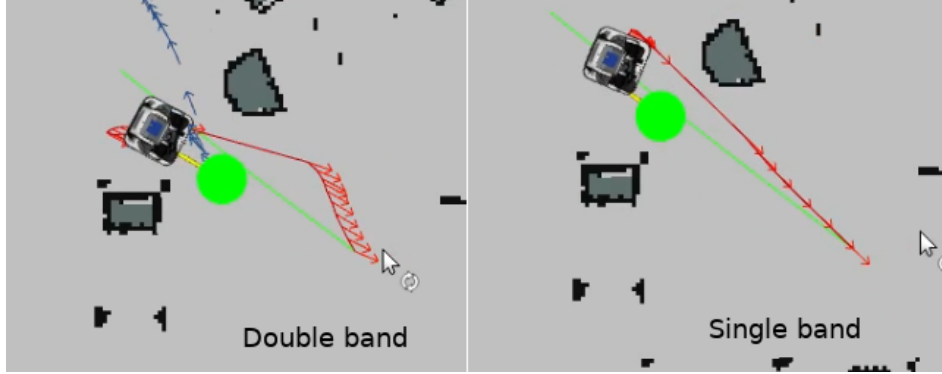


Figure 2.7: Double band and single band trajectories while passing through a narrow passage. As the red trajectory corresponds to the robot, it can be observed from the picture on the left that double band based planning makes the robot back off proactively and provide a way for the human. Whereas single band based planning shown on the right reacts slowly and moves sideways to clear the way for the human.

different experiments (in open space and narrow passage (Fig. 2.8)), and in both experiments, the robot reacted slowly while using a single band. However, while using the double band, the robot proactively backs off as the human moves towards it. Therefore, we can say that our hypothesis is correct, and the inclusion of an elastic band for humans is advantageous for HAN planning. These experiments can be seen clearly in the video⁶. Snapshots of this experiment in case of the narrow passage scenario are shown in Fig. 2.7

2.6.2.2 Pass through Narrow Opening

This experiment can be thought of as passing through a door where only a single person can fit. Suppose two persons arrive at the narrow opening at the same time, one of them has to back off and give way for the other to pass through. We tried to simulate this scenario⁶ in the human-robot co-navigation, and we want the robot to back off and give way to the human. To increase the complexity of this problem further, human crosses the opening and stops close to this opening. Enough space is present for the robot to pass through, but the trajectory might need re-planning. This scenario is shown in Fig. 2.8.

We have tested all three planners (HATEB, HATEB-2 and Single Band) in this scenario and snapshots of the trajectories at this crossing are shown in Fig. 2.7 for both double band (HATEB, HATEB-2) and single band planning. Both HATEB and HATEB-2 reacted, in the same way, to back off and provide a way for the human, which is shown in the left picture of Fig. 2.7. Although they reacted similarly at this instant, HATEB gets stuck in *entanglement* when the human stops moving after crossing the opening. HATEB-2 breaks this *entanglement* and re-plans to reach the given goal. Coming to the case of the single band, the human has to

⁶<https://youtu.be/xEG4e-Y9z8g>



Figure 2.8: Narrow passage passing scenario. In this scenario, the human and robot arrive at the common passage at the same time, and one should back off to provide a way for another. Otherwise, there exists no solution. This is one of the intricate situations addressed in this work, and the picture on the left shows the simulation of this scenario in MORSE. The right side picture shows the trajectories for the human (blue) and the robot (red), generated using the proposed framework, HATEB-2. It can be seen from the picture that the robot’s trajectory (red) is making the robot move backwards and hence providing a way for the human.

stop in front of the robot and wait for the robot to back off and re-plan. The single band also solves this case as it does not suffer from *entanglement* but with lesser reactive speeds. Also, note that the robot backs off in a double band scenario, whereas it tries to move aside in the single band case.

2.6.3 The Effect of New Human-Aware Constraints

TTC constraint in HATEB makes the robot very reactive and leads to unnecessary oscillations, as mentioned previously. Fig. 2.9 shows the trajectory of the robot in the same scenario using *TTC* and *TTCplus* constraints respectively. As we can see from the trajectory on the left, the plan using *TTC* constraint makes the robot back off further when it is already at a sufficient distance from the human. This exaggerated reaction results in long execution times apart from the oscillations. The trajectory planned using *TTCplus* in HATEB-2 is shown on the right, and it can be clearly seen from Fig. 2.9 that this trajectory results in faster execution as it removes the problem of oscillations. The change in the elastic band tightness can be seen in the right part of Fig. 2.9. This change, coupled with the reduced velocity in the close vicinity of the human, helps the human choose the path without confusion. Clear differences in the trajectory followed, and the effect of all these constraints can be seen in the video⁷.

2.6.4 Quantitative Comparison between HATEB and HATEB-2

To perform the *Quantitative* analysis, five different experiments were conducted, and each experiment was repeated 10 times using HATEB and HATEB-2. The

⁷<https://youtu.be/xEG4e-Y9z8g>

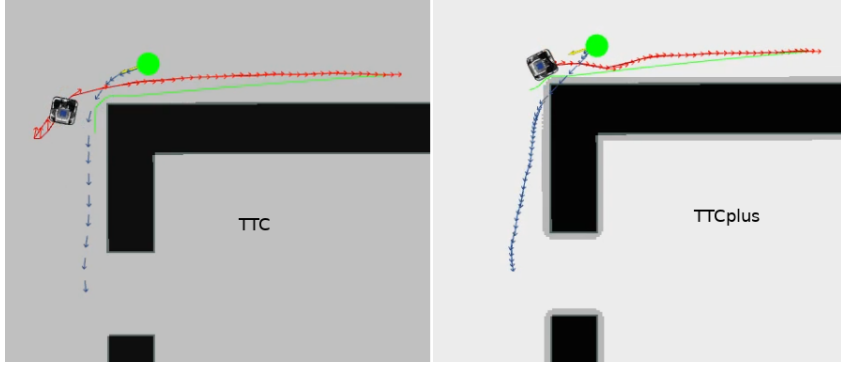


Figure 2.9: The original *TTC* constraint-based trajectory shown on the left results in unnecessary oscillations due to exaggerated constraint. However, *TTCplus* constraint regulates this exaggeration and results in a smooth trajectory as shown on the right. Note that the trajectory on the left is making the robot move backwards even when there is no necessity. The robot's trajectory is shown in red, while the human's trajectory is shown in blue.

list of the experiments performed is given in Table 2.1. The experiment ‘*Narrow opening 1*’ is the same scenario presented in Fig. 2.8, whereas ‘*Narrow opening 2*’ corresponds to a similar case with an opening that can be seen in Fig. 2.9. ‘*L-crossing*’ is the same experiment that is shown in Fig. 2.9 and finally ‘*Narrow corridor*’ and ‘*Wide space*’ represent the scenarios presented in Fig. 2.1.

In all these experiments, the goal of the human was assumed to be behind the robot, and the human executed the trajectory planned by the corresponding local planner. A set of five metrics are used to analyse the results: 1) Initial plan length, *ipl*, 2) Total time for completion, *ct*, 3) Traversed path length, *tpl* 4) Minimum distance from human, d_{min} and 5) Length deviation factor, α . The minimum distance from the human metric, d_{min} , refers to the closest distance between the human and the robot while executing a planned trajectory. The length deviation factor, α , is defined as follows:

$$\alpha = \frac{|tpl - ipl|}{ipl} \quad (2.3)$$

where $| \cdot |$ denotes the absolute value. After determining *ipl* and *tpl* from the experiments, α is calculated using Eq. (2.3).

All these metrics are calculated for each experiment, and the mean value over 10 experiments is presented in Table 2.1. The values highlighted in bold correspond to the best values in the given experiment. The evaluation of the best values for the metrics is done in the following manner. For *ipl*, the value closest to the *tpl* is taken as the best value, as it suggests that the initial plan is very close to the traversed path. In case of *tpl* and *ct*, smaller value represents the best value. The greater the distance of the robot from the human, the more the safety factor for the human, and hence, the larger distance is the best value. Finally, the lower value of α represents the lesser deviation from the initial plan and hence shows better performance of the planner. By observing the values of α from the table, it can be inferred that

Experiment	HATEB				
	$ipl(m)$	$tpl(m)$	$ct(s)$	$d_{min}(m)$	α
<i>Narrow opening 1</i>	5.26	5.94	13.63	0.72	0.13
<i>Narrow opening 2</i>	8.71	11.09	24.29	0.14	0.27
<i>L-crossing</i>	10.66	15.29	32.56	0.14	0.43
<i>Narrow corridor</i>	8.51	13.18	27.92	0.72	0.55
<i>Wide space</i>	8.51	9.54	19.81	0.72	0.12
	HATEB-2				
	$ipl(m)$	$tpl(m)$	$ct(s)$	$d_{min}(m)$	α
<i>Narrow opening 1</i>	5.36	6.34	17.63	0.54	0.19
<i>Narrow opening 2</i>	9.02	9.16	21.93	0.15	0.02
<i>L-crossing</i>	13.46	13.06	29.90	0.34	0.03
<i>Narrow corridor</i>	12.88	13.24	30.88	0.48	0.03
<i>Wide space</i>	9.86	9.68	22.25	0.83	0.02

Table 2.1: Mean values of the metrics over 10 repetitions. ipl : Initial path length, tpl : Traversed path length, ct : Completion Time of the experiment, d_{min} : Minimum distance between human and robot during the execution of the trajectory in the given experiment. α : Length deviation factor.

HATEB-2 performs better than HATEB in all the cases, except ‘*Narrow opening 1*’. In all these scenarios, it can also be seen that ipl differs less from tpl , and hence, we can say that HATEB-2 predicts a better plan than HATEB. The cause of this result can be directly associated with the improvements in human prediction and the new *TTCplus* constraint, thereby demonstrating the importance of human prediction in HAN. As the maximum allowed velocity decreases below $Dist_{threshold}$ in HATEB-2, an increase in ct is expected, and it is true in 3 out of the 5 cases. In ‘*Narrow opening 2*’ and ‘*L-crossing*’ cases, HATEB-2 has lesser ct than HATEB and this is because of the improved *TTC* constraint, *TTCplus*. Since HATEB uses the original *TTC*, the robot suffers from unnecessary oscillations and results in a longer path length, tpl as well as ct . It can be observed that HATEB-2 completely outperforms HATEB in these two scenarios. Although HATEB has better tpl values in the other three scenarios, the tpl values of HATEB-2 are very close to those of HATEB. Finally, it can be said that the overall performance of HATEB-2 is better than HATEB.

2.7 Real-World Tests

In this section, we present the experiments we have conducted using the proposed framework. We have ported the framework to Pepper⁸ robot and used it for this study. Since the main objective is to study the navigation framework, we used the OptiTrack⁹ motion-capturing system to track humans. The localization of the robot

⁸<https://www.ald.softbankrobotics.com/en/pepper>

⁹<http://www.optitrack.com/>

in the map is done using the localization technique based on ArUco¹⁰ markers.




Figure 2.10: Human follows his path without disturbing the robot. (a) Initial positions (b) Intermediate positions. (c) and (d) are the trajectories at (a) and (b) respectively. In (d), the band tightening can be seen as the robot is close to the human. Red: The robot's trajectory is shown in red, while the human's trajectory is shown in blue.

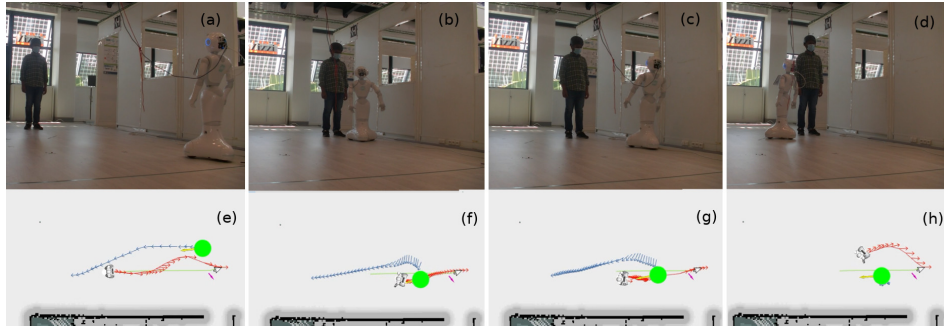


Figure 2.11: Entanglement resolution in a real-world scenario. Here the human goes out of his path and blocks the robot's planned path (f). While trying to block the path, he moves very close to the robot, as can be seen from (f). Therefore, the robot backs off (g) before resolving the entanglement and finding an alternative path (h). (a)-(d) represent different positions of the robot during the experiment, and (e)-(h) show the planned trajectories at these positions. The robot's trajectory is shown in red, while the human's trajectory is shown in blue.

We have conducted two experiments to check the capabilities and improvements

¹⁰<https://chev.me/arucogen/>

in HATEB-2. In the first experiment, shown in Fig. 2.10, the human moves along his path without blocking the way for the robot. It can be seen from Fig. 2.10 that the robot continues to navigate on the same side and follow a path similar to the initially planned path. We can also observe the band tightening in Fig. 2.10 (d) as the human is close to the robot. In the second experiment, the human goes out of his path and blocks the robot's planned path. Two kinds of scenarios are possible depending on how the human acts in this setup. If the human goes very close to the robot, the robot has to back up before resolving the *entanglement problem*. In the other case, where the human is at a nominal distance from the robot, only *entanglement* resolution happens. These two cases are presented in the video. However, we have presented only the first scenario in this section as it brings out more capabilities of the system. This is shown in Fig. 2.11. In Fig. 2.11 (f), the human goes very close to the robot and hence the robot backs off (Fig. 2.11 (g)) before resolving the *entanglement*. Finally, the *entanglement* is resolved (Fig. 2.11 (h)) and the robot proceeds to its goal.

2.8 Conclusion

In this chapter, we have proposed a new framework combining proactive planning and decision-making to handle human-robot co-navigation, called HATEB-2. This framework includes three different modes of planning, namely, **Single Band**, **Dual Band** and **VelObs**. Switching between these modes allows for solving many complex human-robot cooperative navigation problems. We have presented details of these different modes of planning and also talked about the modifications made in HATEB before including it in HATEB-2. These modifications remove some of the drawbacks of HATEB apart from the improvements. We have also presented the improvements made in human prediction and estimation. We performed several experiments in various intricate situations and then provided a detailed analysis of the results. Results show that HATEB-2 have an overall better performance. The framework was finally tested on a real robot platform, and the results were presented.

It is evident that HATEB-2 can handle intricate human-robot navigation scenarios better when compared to HATEB. However, it still carries some of the limitations of HATEB. One of the main drawbacks is the computational complexity with a growing hypergraph as more humans are added to the planning. Even though humans' predictions were improved in HATEB-2, path planning still needs improvements. In the next chapter, we discuss how these limitations were handled to scale our planning system to tens of humans. One of the limitations of the *TTC* constraint is that it works well when the robot and human are approximately on the same line of travel or with very small parallel gaps. If the line of travel has a large parallel gap, it is not very effective. The same limitations apply to *TTCplus* as well. Therefore, we introduce new social constraints to address these limitations and effectively handle parallel travel.

A Tunable Human-Aware Navigation Planner for Multi-Context Navigation

Contents

3.1	Introduction	57
3.2	Related Work	59
3.3	Co-operative Human-Aware Navigation Planner	60
3.3.1	Architecture	60
3.3.2	Types of Visible Humans and Costmap Layers	62
3.3.3	Human Path Prediction Mechanisms	63
3.3.4	An updated HATEB local planner	64
3.4	Testing CoHAN under Different Conditions	69
3.4.1	Door Crossing with Static and Moving Humans	69
3.4.2	A Very Narrow Corridor Scene	71
3.4.3	Co-operative Navigation Scenarios	73
3.4.4	Robot in a Crowd	74
3.4.5	Comparison with Another Human-Aware Planner	75
3.5	Real-world Experiments with CoHAN	76
3.6	Multi-context HAN and CoHAN	78
3.6.1	Modality-based HAN Planning	78
3.6.2	Multi-Context HAN Planning using Mode Shifting	79
3.6.3	CoHAN in Multiple Human-Robot Navigation Contexts	79
3.7	Conclusion	80

3.1 Introduction

In the previous chapter, we introduced situation assessment into human-aware navigation planning and showed how combining it with proactive planning can help some intricate human-robot navigation scenarios in HRI. In this chapter, we extend this planning scheme to address multi-context navigation in HAN. Depending on the shape of the local environment (large area, small rooms, narrow corridors or

passages) and the density and activity of the humans (individuals, crowds, domestic or public space motion activity) in that environment, HAN has to address various types of human-robot interaction contexts. These contexts can differ from situations where the current position of the human is enough to where a good estimate of the goal is necessary or even where path negotiation has to take place. For instance, if the robot is in the middle of a dense crowd, it should better be purely reactive and compliant to the overall motion flows than in a corridor where less reactive and more cooperative motion with path negotiation is preferable. Therefore, different navigation planners are developed for different types of environments with shared human spaces like malls [Foster 2019], streets [Ferrer 2013b], warehouses [Carmona 2019], offices [Truong 2014], labs, homes [Kollmitz 2015] etc. All these different planners emerged as there is no single algorithm that can cover all environments and situations. In order to address these issues, we propose a highly tunable HAN system with multiple modes of planning. It can be employed in a variety of human-robot contexts, with a small number of pertinent parameters that can be adjusted to the situation at hand. The main contributions of this chapter are threefold:

1. We propose a tunable HAN planner with different planning modes that can handle very complex indoor scenarios as well as crowded scenarios called the Cooperative Human-Aware Navigation (CoHAN) Planner.
2. We extend our planning framework, HATEB-2 [Singamaneni 2020], to effectively handle large numbers of people and to offer more legible and acceptable navigation.
3. We evaluate the proposed planner in several simulated human-robot scenarios and present both qualitative and quantitative analysis. Further, we also present the tests conducted on the real robot at our lab.

In the rest of the chapter, we simply use HATEB instead of HATEB-2 to refer to our HAN system combined with situation assessment. The rest of the chapter is organised as follows. Section 3.2 presents the related work. The proposed system’s architecture is presented in section 3.3 along with explanations of various modules and features. We also present in detail the updates in HATEB local planner and design choices made to handle multiple humans. Section 3.4 presents the evaluation of our planner in various simulated human contexts and comparisons with one of the existing human-aware navigation planners. Following this, in section 3.5, we talk about the tests conducted on the real robot in our lab. A discussion on HAN planning with multiple modalities and its extension to multi-context HAN is presented in section 3.6. This section also talks about the possible updates to CoHAN. Finally, section 3.7 concludes the chapter and presents the limitations and future direction.

3.2 Related Work

There are a variety of HAN planners designed for different human-robot contexts. In the context of a crowd or robot navigation in the street, Ferrer et al. [Ferrer 2013b] presents a potential field based navigation using the social force model. The authors of [Truong 2017b] extended this to human-object and human-group interactions by proposing the proactive social motion model. The work by Repiso et al. [Repiso 2017] shows the context of a robot accompanying a human. The authors of [Chen 2017a] address this crowd navigation problem by using reinforcement learning, and the works [Triebel 2016, Okal 2016] address the same with inverse reinforcement learning. Coming to other contexts, the works presented in [Sisbot 2007c, Truong 2014] and [Kollmitz 2015] show some interesting costmap based approaches for planning paths in complex indoor scenarios that can occur at homes or offices. In this work, we use a similar costmap based approach to handle static humans. Fernandez Carmona et al. [Carmona 2019] compare the performance of the existing navigation planners in a warehouse context and proposes an architecture to include humans in planning. The work of Guldenring et al. [Guldenring 2020] addresses the same context using reinforcement learning. Some other works like [Pérez-Higueras 2014, Pérez-Higueras 2018a] use inverse reinforcement learning for confined and public space navigation contexts. Khambhaita and Alami [Khambhaita 2017c] addressed the context of human-robot co-navigation using an optimization based approach. Note that none of the above planners was designed to handle multiple human contexts together. A multi-context human-aware navigation planning is a very new field, and not many works exist. Lu et al. [Lu 2014] proposed a layered costmap based approach for handling different navigation contexts. A more recent work by Banisetty et al. [Banisetty 2020] shows some promising results using a deep learning based context classification and multi-objective optimization for navigation planner [Banisetty 2019]. However, these results are validated only in indoor scenarios, and the authors do not present any results in a crowd, unlike the proposed system.

In order to handle the dynamic humans in our navigation planner and plan a socially acceptable trajectory for the robot, a human motion prediction system is required. One of the classic approaches of human motion prediction is based on the social force model [Helbing 1995]. Ferrer et al. [Ferrer 2015] used this social force model both to predict human motions and move the robot among the crowds. Kollmitz et al. [Kollmitz 2015] used a simple linear prediction based on current human velocity. Instead of predicting the trajectory, a possible human goal can also be predicted using reasoning over a probable set of goals [Bordallo 2015]. Our proposed navigation system uses one such goal prediction system [Ferrer 2014] as a part of the human path prediction module. Apart from this, our system offers three other human path prediction methods to handle different situations. In a recent work by Fisac et al. [Fisac 2018], the authors suggest a probabilistic human model with confidence to handle the uncertainties in a system.

One of the key elements of the proposed system is the context-based shift-

ing between different planning modes. This kind of modality shifting is discussed in the works of Qian et al. [Qian 2013] and Mehta et al. [Mehta 2016] based on Partially Observable Markov Decision Process (POMDP). Unlike these, our system uses situation assessment based modality shifting. In the previous chapter [Singamaneni 2020], we introduced this modality shifting with three different modes of planning. In the current chapter, we extend this to handle a large number of humans and also introduce some elements, including a new planning mode. This modified *HATEB local planner* is integrated into the proposed framework as the local planner.

3.3 Co-operative Human-Aware Navigation Planner

In this section, we present the overall architecture of the proposed HAN planning system and explain its features that allow us to deal with various kinds of human-robot contexts. The idea behind this architecture is to allow legible human-robot interactions and provide a scalable system that can easily be employed in various human environments. This system considers and differentiates between different types of visible humans in the environment to address the situations better. Further, some new human-aware constraints are defined that add more legibility to the robot’s navigation. We explain these things in detail, starting with the software architecture.

3.3.1 Architecture

The proposed HAN system, CoHAN, is developed over the ROS [Quigley 2009] Navigation Stack as previously, and its architecture is shown in Fig. 3.1. The red blocks shown in Fig. 3.1 are the modifications we introduced into the standard ROS Navigation Stack and can be considered as some of the major contributions of this thesis. As shown in the figure, we introduce *Human Safety* and *Human Visibility* costmap layers into both global and local costmaps. *Human Safety* layer takes care of the proxemics, whereas the *Human Visibility* layer is added to avoid any surprise appearances of the robot from behind the human [Sisbot 2007c]. The *Human Safety* layer is modelled as a 2D Gaussian around the human, and the *Human Visibility* layer as a 2D half Gaussian on the backside of the human. Both these layers have a cutoff radius of 2 m beyond which the cost is zero. However, the radius can be adjusted, and it is one of the tunable parameters of this system. These layers are implemented using a costmap plugin that we developed called the ***human_layers***, which is a part of the CoHAN system. These layers are added and updated based on the *Human States* determined by the situation assessment loop of the *HATEB local planner*.

The *Human Path Prediction* module predicts the possible paths for the requested humans using a selected prediction method. CoHAN provides some services for path prediction and updates the planning strategy to reduce the occurrence of ‘*entanglement*’, which are explained in detail in section 3.3.3.

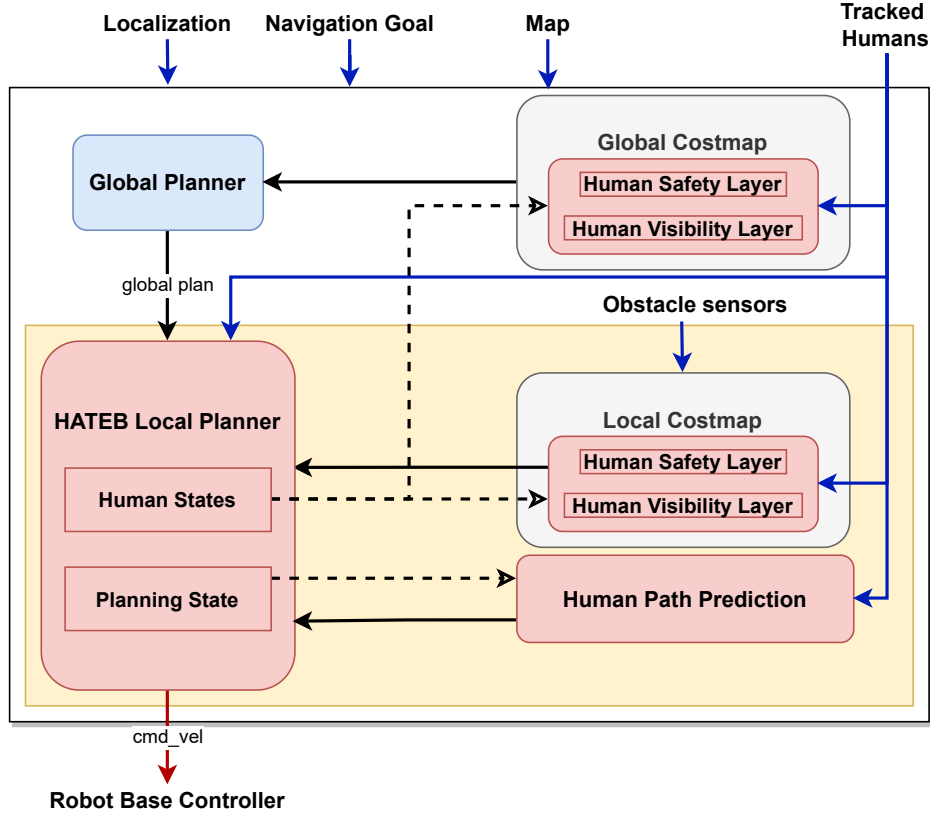


Figure 3.1: Architecture of the proposed HAN planner, CoHAN.

The *HATEB local planner* module in CoHAN accesses different human-robot scenarios and determines the *Human States* and the *Planning State* shown in the figure. Both these states together decide the planning mode of the system and also control the transition between different modes. Based on the *Planning State*, the appropriate path prediction method is selected for humans. After accepting a navigation goal, the system continuously accesses the human-robot scenario and appropriately chooses a planning mode that decides the command velocity sent to the robot's base controller as previously. Therefore, the planning mode need not be constant and can shift during the navigation depending on the context. However, the modes of planning and the situation assessment loop for mode shifting have been updated a little when compared to the previous version of HATEB. The parameter $Dist_{min}$ from the previous chapter is analogous to the planning radius in this version.

Further, our system is completely tunable, and the transition between different modes can be tuned (or changed) by changing the mode transition parameters [Singamaneni 2020] as per the requirement. This system is mainly designed to address most of the intricate indoor navigation scenarios, but it can be employed in semi-crowded scenarios with some proper tuning. We show an example of this in our results section 3.4. CoHAN is publicly available on GitHub at <https://github.com>.

com/sphanit/CoHAN_Planner and comes with inbuilt stage-ros¹ simulator to test some human-robot navigation scenarios quickly.

3.3.2 Types of Visible Humans and Costmap Layers

In CoHAN, we deal with different types of humans while navigating the robot to the goal. Fig. 3.2 shows all these types of humans along with the robot's visibility and the planning radius, R . While the robot is moving in the environment, the system considers only the humans within this planning radius that are in the visible region. Among these humans, it checks for the static and dynamic humans and updates the *Human States* for all the humans. In CoHAN, four states are defined for humans in the field of view of the robot, and they are presented below.

1. **STATIC**: This state defines the humans that do not have any velocity as long as they are in the robot's field of view.
2. **MOVING**: This state is attached to all the humans with non-zero velocity.
3. **STOPPED**: This state is for the humans who were moving previously but now have zero velocity.
4. **BLOCKED**: This is a conflict state of the human, which indicates that the human is either blocked by the robot or an *entanglement* is detected.

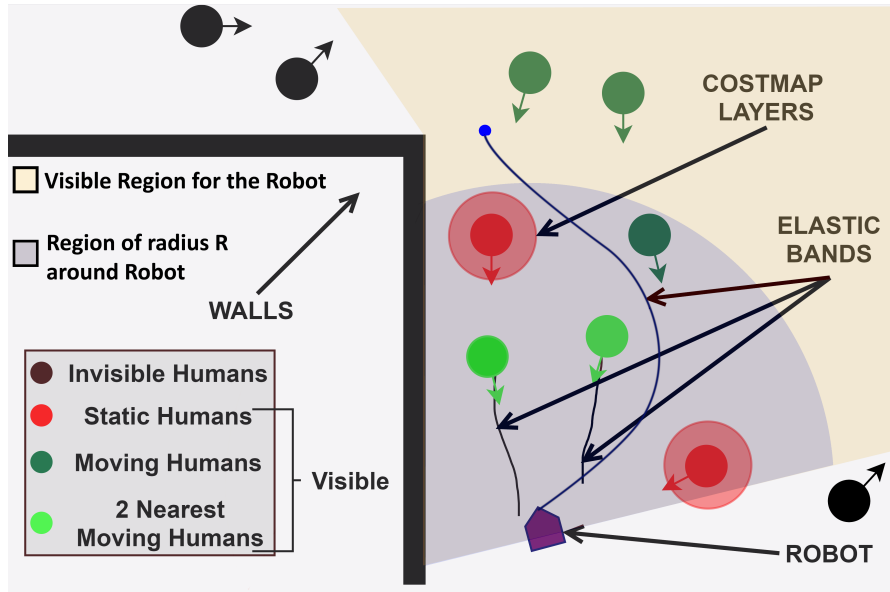


Figure 3.2: Different types of humans considered in our system. A sample trajectory of the robot is shown among different humans.

These states are updated in a separate state machine, which is different from the one that updates and analyses the situation in *HATEB*, except for the **BLOCKED**

¹http://wiki.ros.org/stage_ros

state. Since it is a conflict state, it is updated by *HATEB local planner* after analysing the situation. The *human_layers* plugin checks the *Human States* of all the observable humans and adds the corresponding costmap layers. For humans in **STATIC** state both *Human Safety* and *Human Visibility* layers are added around the humans. For the humans in all other states, only *Human Safety* layer is added, and that too with a reduced radius. This design choice is made based on the fact that *HATEB local planner* module specifically handles the conflicts with moving humans, and it has the necessary human-aware constraints that approximately imitate these costmap layers. The *Human Safety* with reduced radius in moving humans ensures better global plans with an additional safety margin. In the case of static humans, we try to reduce the computational complexity by using only global planning and running HATEB in **Single Band** mode, i.e., no elastic band is added to the static humans. Besides, static humans usually respond slowly compared to moving humans, and therefore the robot should maintain a larger safety distance as well as avoid surprise appearances from behind. Nevertheless, if a static human starts moving or in the presence of any other moving human, CoHAN quickly switches to **Dual Band** mode.

Another design choice that is made to reduce the computational complexity and scale the system for multiple humans is to restrict the addition of elastic bands to the two nearest dynamic humans within the planning radius. In addition, the proactive planning for these humans stops as soon as they go out of the field of view of the robot or the planning radius. Even though it seems intuitive, this restriction is at the planning level but not the perception level. If a motion capture system is employed to track humans, the system might receive data about humans who are not in the field of view, and the planning system should be aware of this. If a camera-based perception system with a limited field of view is employed, this will not be a problem, but our idea is to provide a planning system that is not affected by the type of perception.

3.3.3 Human Path Prediction Mechanisms

The *Human Path Prediction* module deals with different kinds of human goal predictions and building global plans for the required humans. Our system currently offers four types of human goal prediction and path planning methods.

1. *PredictBehind*: This method predicts that the human goal is behind the robot. The position of the robot when the human enters the visible planning radius is used for this. This goal is used to predict the path.
2. *PredictGoal*: This method predicts the most probable goal among the set of goals provided to the system using the approach described in [Ferrer 2014]. The predicted goal is then used for path prediction.
3. *PredictVelObs*: This method builds a path by extrapolating the current human velocity over a fixed duration and does not predict any goal. Currently, the duration is set to 5 s. This is the default prediction service in *VelObs* mode.

4. *PredictExternal*: This service accepts a goal from an external system and adds a global path prediction based on the provided goal.

The *PredictExternal* goal service allows CoHAN to communicate with any goal prediction system using ROS and extends its usability. All these services provide global plans (or paths) for the humans that are used by the *HATEB local planner* for planning the trajectories.

One of the major changes to the path prediction system in CoHAN is the path re-planning mechanism. The HAN presented in the previous chapter calculates the path of humans for proactive planning only once. Even though it is not the sole reason for *entanglement problem* that we discussed, it is one of the major ones. Hence, in this version of *Human Path Prediction* module, the human path is re-planned every time the human deviates from the previous plan above a certain threshold. The chosen threshold for this re-planning is 0.5 m, i.e., if the human deviates more than 0.5 m from the closest point in the previous path, a new path is planned. This ensures that human prediction is more consistent and reduces the *entanglement* issues. Even with this improvisation, the entanglement problem occurs from time to time, and *HATEB local planner* handles such occurrences and ensures acceptable HAN planning. The second advantage of re-planning is that it mitigates the homotopy class changes in proactive planning and eliminates the need for the changes in the elastic band introduced in the previous chapter. So, in CoHAN, we remove these changes to the elastic band and modify the situation assessment loop.

3.3.4 An updated HATEB local planner

It is the core module of the proposed HAN planning system. *HATEB local planner* in CoHAN is an extended version of the human-aware proactive planning system with situation assessment presented in the last chapter [Singamaneni 2020]. This module plans the robot's trajectory and the possible human trajectories for the two nearest humans in the visible planning radius based on the predicted human paths. It continuously assesses the current human-robot context and sets the *Planning State* and sometimes, the *Human States*. Depending on the value of these states, it shifts between different planning modes. As mentioned previously, mode shifting is needed in intricate human-robot contexts that cannot be solved using a single planning mode. CoHAN offers three different planning modes and one recovery mode that is selected based on the situation. Even though the first three modes of planning are similar to the one from chapter 2, the situation assessment loop has been modified significantly, as presented in Fig. 3.3. CoHAN also proposes two new human-aware constraints, which steer the robot to behave in a human-friendly manner. The details of these modifications and additions are presented below.

3.3.4.1 Modes of Planning

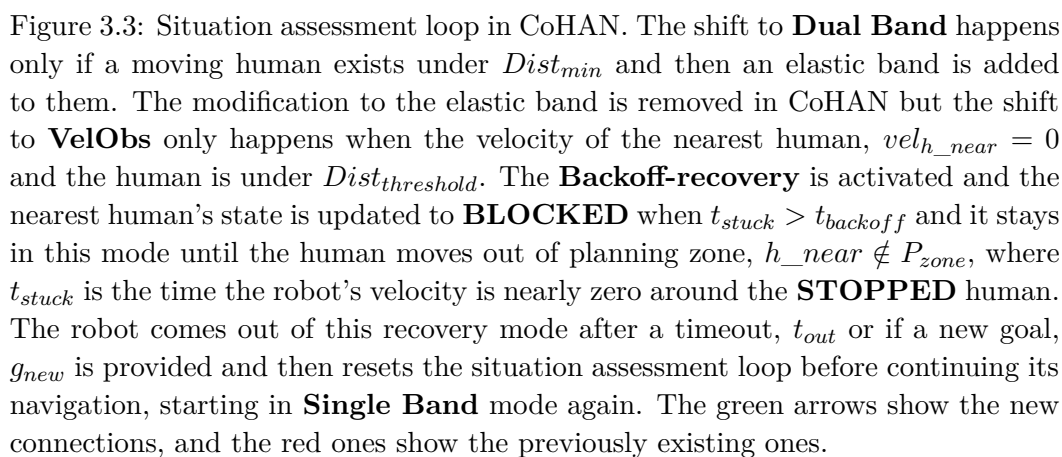
HATEB local planner provides four different modes of planning at the control level and intelligently shifts between them based on the situation. We briefly present these modes here.

1. **Single Band:** This is the mode in which the planning system starts and has an elastic band only for the robot. The system stays in this mode as long as there are no humans within the visible planning radius. The default planning radius or $Dist_{min}$ is 10 s as mentioned previously.
2. **Dual Band:** In this mode, elastic bands are added to the *two nearest* moving humans in visible planning radius and trajectories are planned for humans along with the robot.
3. **VelObs:** This mode uses all the human-aware criteria while planning but adds bands to humans only if they have some velocity.
4. **Backoff-recovery:** The **Backoff-recovery** mode is activated when there is no solution to the planning problem unless one of the agents completely clears the way for the other.

The **Dual Band** mode allows the robot to proactively plan its trajectory and adapt to the changing human plans. On top of providing human predictions, these trajectories of humans also offer a possible solution for the human-robot navigation context, which, if followed, will resolve any conflict that exists. **VelObs** mode is especially useful in crowded human scenarios or when the robot cannot move due to *entanglement* issues of the **Dual Band** mode. The new **Backoff-recovery** mode added in CoHAN is useful in situations that commonly occur in a very narrow corridor where only the person (or robot) can navigate at a time. If a human and a robot face each other in a very narrow corridor or another situation where one of them has to clear the way for the other, our system gives priority to the human and makes the robot clear the way for the human. Next, we move on to an explanation of the updated situation assessment loop and the implementation of **Backoff-recovery**.

3.3.4.2 Updated Situation Assessment and Backoff Recovery

The updates and modifications in the situation assessment loop of *HATEB* are shown in Fig. 3.3. In the figure, the updated or modified areas are zoomed using ellipses or rounded rectangular boxes, respectively. The part of the loop inside the ellipse shows small updates as compared to the previous version to make the transition decision from **Single Band** to **Dual Band** mode. In addition to $Dist_{min}$, this transition also depends on the *Human States*, which is denoted by the set, $\{H^{state}\}$. Even when there are humans within the planning radius, the robot stays in the **Single Band** mode as long as all the humans are static. If a moving human appears or one of the static humans starts moving, the *Human States* are updated,



and the planning mode switches to **Dual Band**. In **Dual Band** mode, elastic bands are added only to two of the nearest moving humans ($H^{state} \neq \text{STATIC}$).

The updated path planning mechanism for humans eliminated the need for elastic band change, but the shift to **VelObs** mode still happens under the threshold human-robot distance, $Dist_{threshold}$, when the nearest human stops moving, i.e., $vel_{H_{nearest}} = 0$. As this is a major modification, this is shown inside a rounded rectangular box. The rest of the situation analysis remains almost the same except for the major changes introduced around the **VelObs** mode and the **Backoff-recovery**. This is shown in another rounded rectangular box.

The **Backoff-recovery** is implemented by making the robot move back slowly until it can go either left or right to clear the way. This is done by querying the costmaps on all three sides (left, right and back) and setting a temporary goal at a small distance in the possible direction. The robot keeps moving back in steps of 0.5 m until a possible goal is found either on the left or the right at a distance of 1 m from the current position of the robot. Once the robot clears the way, it waits for the corresponding (nearest) human to complete the navigation or move out of the planning zone (visible region under planning radius), P_{zone} . If the system cannot determine this, the robot can get stuck in the waiting loop. To avoid this, a timeout, t_{out} , is implemented, and after this specified timeout, the robot automatically starts moving towards its previous goal in **Single Band** mode, resetting the situation assessment loop. The default timeout is set to 2 min. CoHAN can also accept a new goal in the waiting state discarding the existing goal. It then resets the loop and switches to **Single Band** (or **Dual Band**) mode to proceed with the navigation to the new goal. The **Backoff-recovery** mode is activated when the robot is in **VelObs** mode in the close vicinity of the human ($< (Dist_{threshold} = 2.5 \text{ m})$), and it is stuck without progressing towards the goal for more than a specific amount of time, $t_{backoff} = 5 \text{ s}$. The value of 5 s is selected based on empirical analysis in simulations, but it still needs tests in real-world settings to tune it better. This mode also updates the state of the nearest human to **BLOCKED**. In Fig. 3.3, new connections added to effectively handle the **Backoff-recovery** mode are shown using green arrows. The arrows shown in red represent the previous connections with existing blocks and are mostly unmodified.

3.3.4.3 New Social Constraints: Visibility and Relative Velocity

HATEB uses several human-aware constraints in its optimization scheme for proactive and legible planning around humans in the environment. Several of these constraints are listed in our previous works [Khambhaita 2017c, Singamaneni 2020]. In this chapter, we propose two new constraints, called the *Visibility* and *Relative Velocity*.

Visibility: It adds cost to the optimization when the robot is behind the human and it plans to cross or go in front of the human. This constraint tries to avoid the emergence of the robot suddenly from behind and makes the robot enter the human's field of view from a larger distance. It is implemented by adding a 2D

half-Gaussian behind the human.

$$cost_{visibility} = \eta \chi^{-(dx^2+dy^2)} \quad (3.1)$$

where χ and η determine the decay rate and amplitude, respectively, and dx and dy are the distances between the human and the robot in the x and y axes. In order to maintain a slow increase instead of sudden rises, we chose $\chi = 2$ with $\eta = 5$.

Relative Velocity: This constraint adds cost to the optimization based on the relative velocity between humans and the robot and their distance. The main effect of this constraint is low robot velocity in the close vicinity [Kruse 2014b] of the human if it cannot find a path to maintain a greater distance. If the robot can find a path with a greater distance from a human, it chooses that path with a normal velocity profile. Another effect of this constraint is early intention show of the robot similar to *TTC* or *TTCplus* constraint given in [Khambhaita 2017c, Singamaneni 2020]. The cost added is shown below.

$$cost_{rel_vel} = \frac{((\max(\vec{V}_{rel} \cdot \vec{P}_r \vec{P}_h), 0) + \|\vec{V}_r\| + 1)}{\|\vec{P}_r \vec{P}_h\|} \quad (3.2)$$

where \vec{P}_r, \vec{V}_r are the position and velocity of the robot, \vec{P}_h, \vec{V}_h are the position and velocity of the human and $\vec{V}_{rel} = \vec{V}_r - \vec{V}_h$. Since this constraint has similar effects as *TTC* or *TTCplus*, we activate this constraint alone and deactivate *TTC* and *TTCplus* constraints in all the experiments presented in this chapter. *HATEB* takes all the activated human-robot constraints and other necessary kinodynamic constraints and plans the trajectories of the robot and the humans. Since the local planner runs a computationally expensive optimization in each control loop, extending the planning beyond two humans does not yield real-time control of the robot. Hence we restricted our human planning to the two nearest humans.

3.3.4.4 Visibility in Planning

One final extension to *HATEB* is the addition of the field of view of the robot into the planning system. This means that while planning, CoHAN considers only the humans present in the field of view of the robot within the planning radius. This filtering is done in two steps. First, the humans outside the visibility cone are omitted from the list of humans to consider during planning. Then, ray-tracing is done from the robot's position to the humans' positions in the visibility region to check if any obstacles are obstructing the visibility. If the ray emerging from the robot successfully meets a human position without any overlap with obstacles in the environment map, then that human is added to the list of humans considered for planning. Since human tracking is provided by an external system, it is important to restrict the system to consider the humans present in its field of view. This is more natural and makes it easier to use our system with vision-based human tracking. Moreover, it also saves computational resources as the robot does not plan for humans which may not affect its navigation.

3.4 Testing CoHAN under Different Conditions

To validate our system, we applied it to various kinds of human-robot contexts that can occur in day-to-day life. These situations are generated in a simulated environment based on MORSE [Echeverria 2011]. The humans in these simulations are controlled in three different ways to test the robustness of the system: (1) Joystick based control by a human operator, (2) Using an improved human motion simulator we have developed in our lab, InHuS [Favier 2022] and (3) Using the human trajectories generated by *HATEB local planner* (an ideal situation where the human moves as expected by the robot). We present in detail some of these intricate scenarios in this section, along with some quantitative results. Further, we also present some details about the extension of our system to crowded scenarios using PedSim ROS². In all figures shown below, the trajectories of the robot and humans (if shown) are shown as coloured dots. These are the poses planned by *HATEB local planner*, and the colour visualizes the time. If the colour of the predicted human pose dot is the same as the colour of the robot pose dot, they are both estimated to be at those locations at the same time.

3.4.1 Door Crossing with Static and Moving Humans



Figure 3.4: Door crossing scenario in the simulated environment. The human moves towards the door. The robot sees the human and waits on the side of the door (right) until the human crosses.

Door crossing is a common situation in many human environments. If two humans try to pass through the same door, one of them has to compromise and clear the way for the other. We have placed the robot running our planning system in the door crossing situation shown in Fig. 3.4. The goal of the robot is beside the second human standing in the room, and the system uses *PredictBehind* human path prediction. The left part of the figure shows the simulated scenario and the corresponding trajectories planned for the human and the robot. The simulated human crossing the door was controlled using a joystick and, hence does not move as the planning system expects. The system quickly adapts to these changes and

²https://github.com/srl-freiburg/pedsim_ros

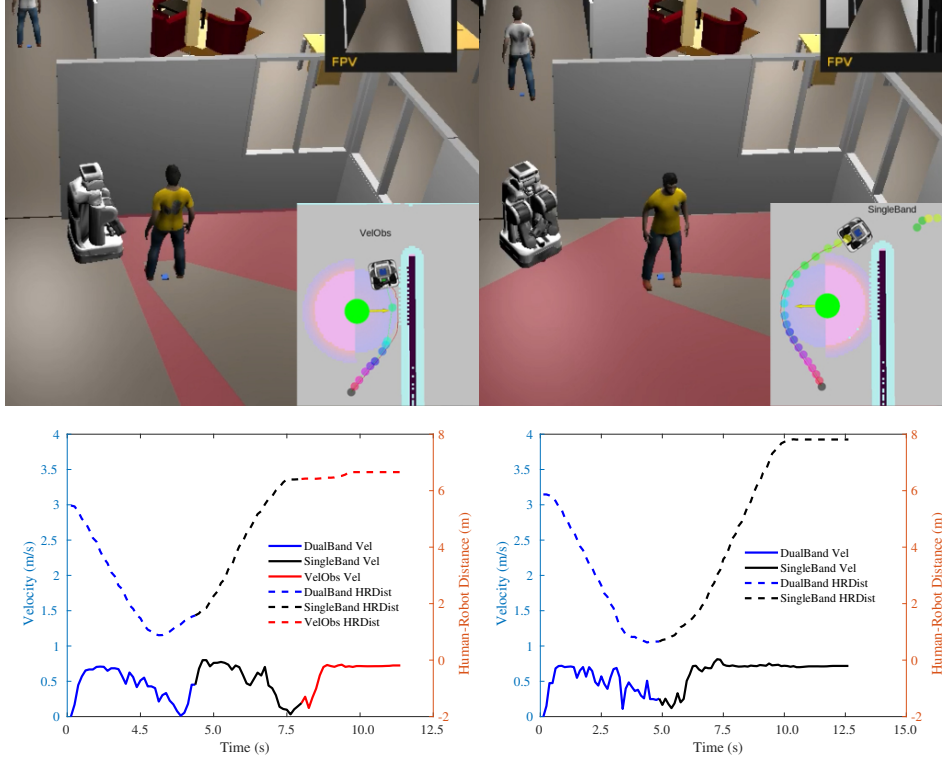


Figure 3.5: Door crossing scenario in the simulated environment with the static human in two different orientations. Top two pictures show the scenarios in simulation and the planned trajectory of the robot. The bottom two figures show the robot velocity and human-robot distance graphs over time.

makes the robot clear the way for the human by waiting on the side, as shown in the right part of Fig. 3.4. The robot continues to its goal after the human crosses the robot. The planning mode is **Dual Band** until the human crosses the robot, and then it switches to **Single Band**.

As soon as the robot crosses the door, it faces one more human, but this human is just standing in the same place and does not move. Since the human is static, our system adds the *human_layers* to the costmaps and re-plans its path. The same scenario is repeated with the second human placed in two different orientations and as shown in Fig. 3.5. In both scenarios, there is enough space between the wall and the human for the robot to reach its goal, maintaining a safe distance from the human. In the top left scenario of the figure, the human can see the robot, and so the planner makes the robot pass through this space. However, in the second scenario, the human cannot see the robot. Therefore, our planner completely re-plans the path as shown (top right) and makes the robot reach its goal by taking a longer and more visible path. It is due to the added *Human Visibility* layer

Fig. 3.5 also shows the plots corresponding to robot velocity (on the left y-axis) and the distance between the moving human and the robot (human-robot distance) (on the right y-axis) with respect to the time (on the x-axis). Different colours in

different portions of the plots correspond to different planning modes of the system, as indicated in the plots. The solid line represents the robot’s velocity (Vel), while the dashed line shows the human-robot distance (HRDist). The same conventions are followed across this chapter. From both the graphs (Fig. 3.5 bottom left and right), it can be observed that the Vel decreases as the HRdist decreases. It is a combined effect of several human-aware constraints of our system. However, the *Relative Velocity* constraint plays a major role here. Secondly, it can be seen in the graph of the first scenario (bottom left) that the Vel decreases one more time before the planner changes to **VelObs** mode. It is because the robot is trying to navigate a narrow space between the human and the wall. This causes the planner to slow down its velocity and check the state of the human. Since the human is static, it shifts to **VelObs** mode that has little reduced weights for the human-aware constraints and continues its navigation.

3.4.2 A Very Narrow Corridor Scene

This scenario occurs when a long corridor has to be traversed by two humans in opposite directions, and the corridor is wide enough only for a single human. In this case, one of them has to go back and wait for the other to cross. When one of the agents in this scenario is a robot, it becomes a little more complicated as the robot should back off giving priority to the human while taking legible actions. Most of the existing planners either re-plan a long deviation to reach the goal or fail in this complicated situation. A more natural way to handle this would be to clear the way for the human and wait until the human crosses the robot to resume its goal. The **Backoff-recovery** mode of our system does exactly this. To make the actions more legible, the robot moves back slowly without showing its back until it can go either left or right to clear the way.

The snapshots from the simulated version of this scenario are shown in Fig. 3.6. Each picture also shows the planned trajectory of the robot in each setting with the *Planning State* behind the robot. This scenario uses the *PredictGoal* human path prediction, and the goal of the robot is on the other side of the corridor. Fig. 3.6 (a) shows the initial situation when the two agents enter the narrow corridor. As the robot can see the human is moving, CoHAN operates in **Dual Band** mode until the human blocks its way completely. The human agent in this setting is controlled by InHuS [Favier 2021a]. As soon as the robot finds itself blocked, it switches to **VelObs** mode and checks for a possible solution. However, when it cannot find the solution after repeated checks, it switches to **Backoff-recovery** mode after few seconds (> 5 s) as shown in Fig. 3.6 (b). Fig. 3.6 (c) shows the robot waiting for the human to cross the corridor before it can resume its goal. CoHAN finally switches to **Single Band** mode and resumes the robot’s navigation to the goal as in Fig. 3.6 (d).

The plots of the Vel and HRDist with respect to time for this scenario are shown in Fig. 3.7. As the HRDist decreases after a certain threshold, Vel decreases, like in the door crossing scenario (blue part). When the robot switches its mode from

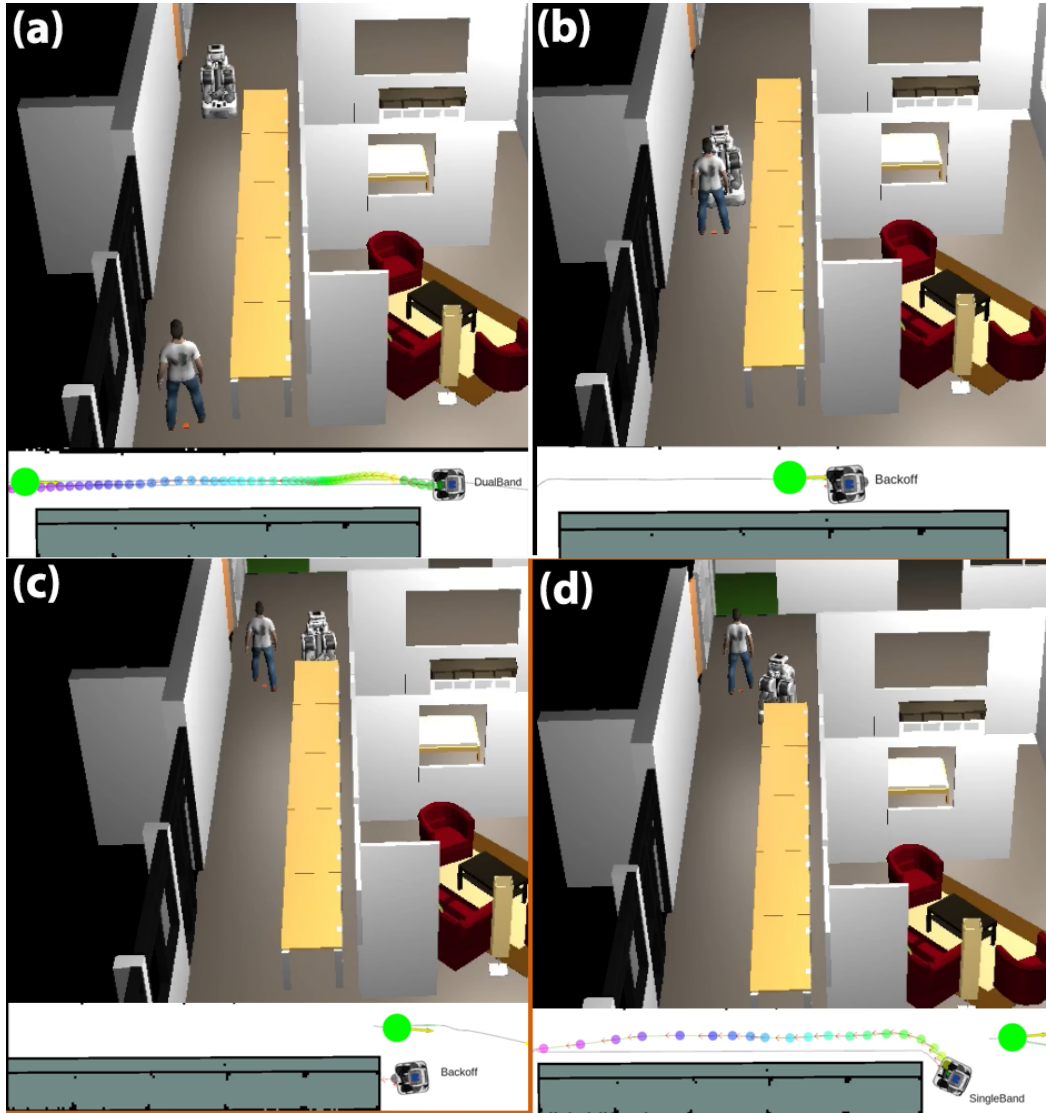


Figure 3.6: Narrow corridor scenario simulated in MORSE. (a) The initial planned trajectory of the robot in **Dual Band** mode. (b) The robot's way is blocked by the human and the system shifts to the **Backoff-recovery** mode. (c) The robot clears the way for the human and waits on the side until the human crosses the robot. (d) The robot continues to its goal in **Single Band** mode.

Dual Band to **VelObs**, the robot tries to move in different directions causing the oscillations seen in the plot (red part). In the **Backoff-recovery** mode, it maintains a constant velocity (green part) and halts, waiting for the human to cross. The human agent of the human simulator starts moving towards the robot as soon as it starts moving back. This explains a near-constant HRDist trend in green. Once the human passes the robot, it resumes its navigation in **SingleBand** mode (black part).

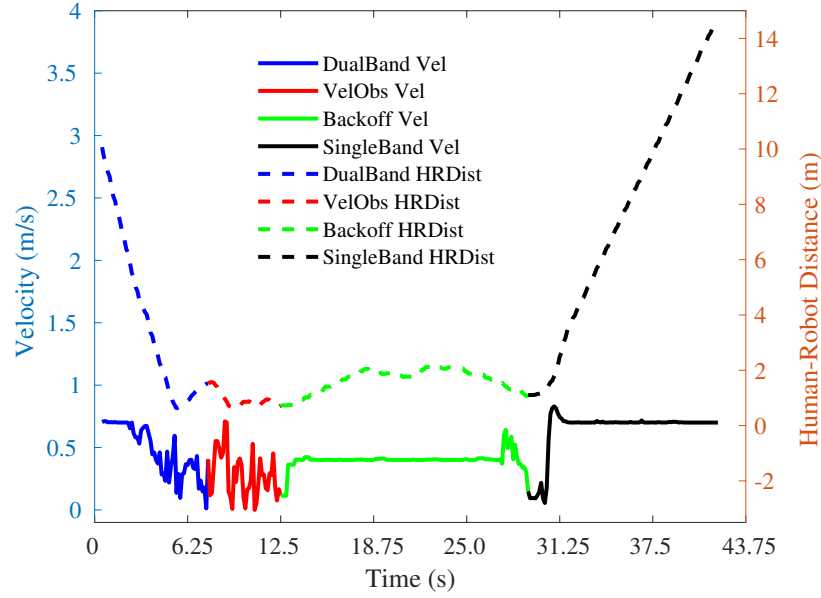


Figure 3.7: Plots of velocity and human-robot distance over time in the Narrow Corridor scenario.

3.4.3 Co-operative Navigation Scenarios

We have simulated three other scenarios in MORSE where cooperation is needed for successful navigation. These three scenarios are called ‘Pillar Corridor’, ‘Wide Corridor’ and ‘Open Space’. Fig. 3.8 shows the snapshots of these scenarios (top part) and part of the trajectory planned by CoHAN. In all three scenarios, human and robot goals are behind each other, and they start navigation at the same time. In the Wide Corridor scenario, CoHAN uses *PredictBehind* path prediction for the human, and hence the human’s goal is estimated to be at the back of the robot’s initial position. The path generated using this goal is then used to plan the human trajectory. In this case, we use this planned human trajectory to control the human, and thus it represents the ideal scenario for the planner. This scenario and its corresponding Vel and HRDist plots are shown in Fig. 3.8 (b). We can see from these plots that the Vel decreases as the HRDist decreases, but it never goes to zero as the human is behaving ideally. We can also see the shift from **Dual Band** to **Single Band** mode as soon as the human crosses the robot. This is true for all three cases.

In the other two scenarios, the human agent was controlled by InHuS, and the system uses *PredictGoal* human path prediction. The Vel and HRDist plots for these scenarios are shown in Fig. 3.8 (a) and (c). From the plots of the Pillar Corridor in Fig. 3.8 (a), we can see that the robot’s velocity decreases rapidly and momentarily goes to zero. This occurs as CoHAN plans proactively and makes the robot wait behind the pillar to let the human cross. This scenario highlights the priority given to humans over the robot in our planning system. The plots of the

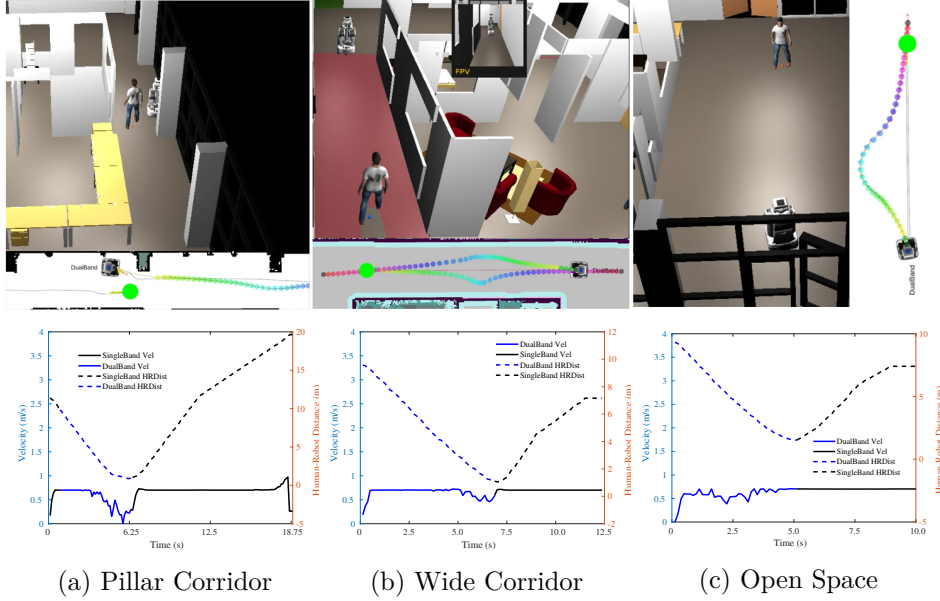


Figure 3.8: (a) A corridor with pillars, wide enough for only one agent at the side of the pillar. (b) A wide corridor where the two agents have enough space to cross each other maintaining safe distances. (c) An open space scenario where the robot has enough space to avoid and show its intention to the human well in advance. In (a), (b) and (c), the plots of robot's velocity and human-robot distance over time are shown below the scenario.

Open Space scenario in Fig. 3.8 (c) show a jerky velocity profile in the **Dual Band** mode and a smooth one in **Single Band** mode. However, the speed does not go down significantly. As the robot has enough space to move away from the human and continue its navigation at high speed, *Relative Velocity* constraint pushes the robot to a larger distance from the human. Hence, the jerky Vel profile is due to the directional changes in the velocity. Note that the HRDist, in this case, is always more than 1.5m, thanks to the *Relative Velocity* constraint. By looking at the initial trajectory generated by CoHAN in Fig. 3.8 (c), it is clear that the robot chooses to go to the left very early during its navigation, and we can argue that this intention show makes the robot's trajectory legible to the human.

3.4.4 Robot in a Crowd

For this experiment, we tuned the parameters of CoHAN and tested it in a simulated crowd. For the simulation of crowds, we used the PedSim ROS simulator, and CoHAN is run completely in the **VelObs** mode. Therefore, the human path prediction by default is set to *PredictVelObs*, which is a linear prediction methodology, as previously mentioned. Fig. 3.9 shows two snapshots from the tests. The robot adds elastic bands to two of the nearest humans in the environment and successfully navigates the crowd generated by the simulator (shown in video³). Further, it can

³https://youtu.be/DB_8HpjngJ4

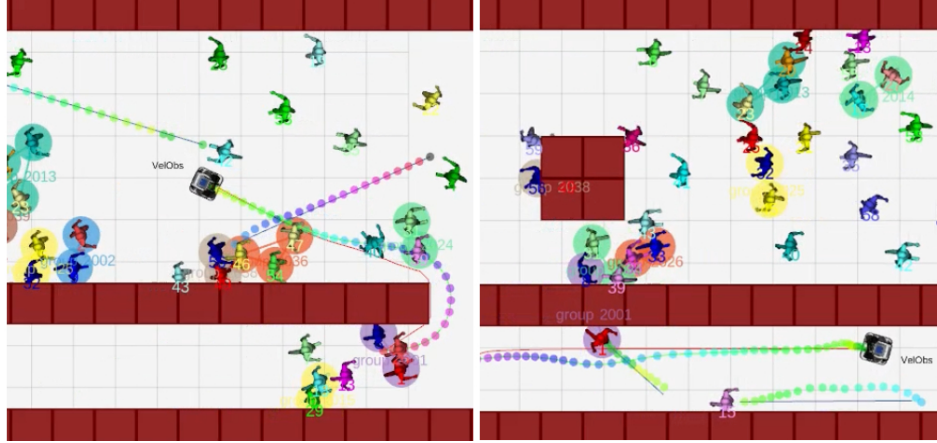


Figure 3.9: The robot running the CoHAN system in the PedSim ROS pedestrian simulator. The robot’s planned trajectory and the predicted trajectories of the two nearest humans in **VelObs** mode are shown.

be seen that the robot proactively clears the way for PedSim agents while navigating in the corridor, as shown in the snapshot to the right in Fig. 3.8. Even though CoHAN was mainly designed for indoor navigation scenarios, this example shows how it can be tuned and scaled to scenarios with a large number of humans. The choice of proactively planning only for two of the nearest humans did not affect the navigation performance as the system was able to quickly switch between different humans.

3.4.5 Comparison with Another Human-Aware Planner

Experiment	CoHAN		
	<i>Path Length (m)</i>	<i>Total Time (s)</i>	<i>Min HRDist (m)</i>
<i>Open Space</i>	9.23	16.01	1.29
<i>Narrow Passage</i>	9.73	17.80	0.71
<i>Pillar Corridor</i>	19.11	31.46	0.89
<i>Narrow Corridor</i>	23.54	48.38	0.66
TPF			
<i>Open Space</i>	8.79	24.77	0.95
<i>Narrow Passage</i>	9.40	27.16	0.93
<i>Pillar Corridor</i>	18.45	52.94	0.76
<i>Narrow Corridor</i>	-	-	-

Table 3.1: Mean values of the metrics over 10 repetitions in four different contexts. TPF failed in the Narrow Corridor case.

In order to check the repeatability of our system and evaluate its performance with respect to the existing human-aware navigation planners, we have selected four different simulated scenarios and repeated the same experiment 10 times in each of

the scenarios. The scenarios we considered here include Open Space, Narrow Passage (similar to the one in Chapter. 2), Pillar Corridor and Narrow Corridor. The human in all these scenarios is controlled by InHuS. In all these scenarios, our system produced consistent results over the repetitions with similar paths. Further, we compared it with the human-aware navigation planner presented in [Kollmitz 2015] that was designed for indoor home context. This system uses *Timed Path Follower* (TPF) as its local planner, and its trajectory is highly dependant on the path produced by its global planner, *Lattice Planner*. Note that our comparison was limited to only one planner as not many human-aware planners are available openly for indoor contexts. This planner was also made to run repeatedly 10 times in the above four scenarios. However, in the Narrow Corridor case, this planner failed to complete the navigation, and the robot got stuck in front of the human as the *Lattice Planner* could not find a path.

We used three different metrics to present the comparisons between these two planners, the total length of the path taken, the total time to complete the scenario and the minimum human-robot distance that the planner encountered while executing each scenario. The average over the 10 runs is taken and presented in Table 3.1. Note that the total time taken by the TPF is greater as its linear velocity was limited by the planner, even though the same velocity and acceleration limits of the robot were provided to both planners. In terms of the total path length, TPF always followed a lesser distance compared to *HATEB*. This is because *HATEB* took the larger deviations to either show intentions (in Open Space) or a clear path for the human (Pillar Corridor and Narrow Passage). However, in the Narrow Passage case, TPF produced better behaviour by waiting for the human to cross while *HATEB* blocked the way for a bit before clearing the way for the human⁴. This can also be seen by comparing the minimum human-robot distance in this case. Finally, in terms of the minimum human-robot distances, *HATEB* varies widely, as it handles each case differently. If space is available, it takes a greater distance than TPF, otherwise, it slows down the robot’s velocity and approaches a little closer to the human. In TPF, this metric produces similar results in two of the three scenarios. In the Pillar Corridor, this metric has a lesser value compared to *HATEB* as the robot goes towards the wall to the opposite side of the pillar and waits instead of going behind the pillar.

3.5 Real-world Experiments with CoHAN

We deployed CoHAN on a real robot platform, Pepper⁵ to conduct some real-world experiments in our lab. For human detection and tracking, we used the OptiTrack⁶ motion capture system and published the positions and velocities of the tracked humans at 10 Hz. The localization of the robot is done in the same manner as in

⁴https://youtu.be/DB_8HpjngJ4

⁵<https://www.ald.softbankrobotics.com/en/pepper>

⁶<http://www.optitrack.com/>

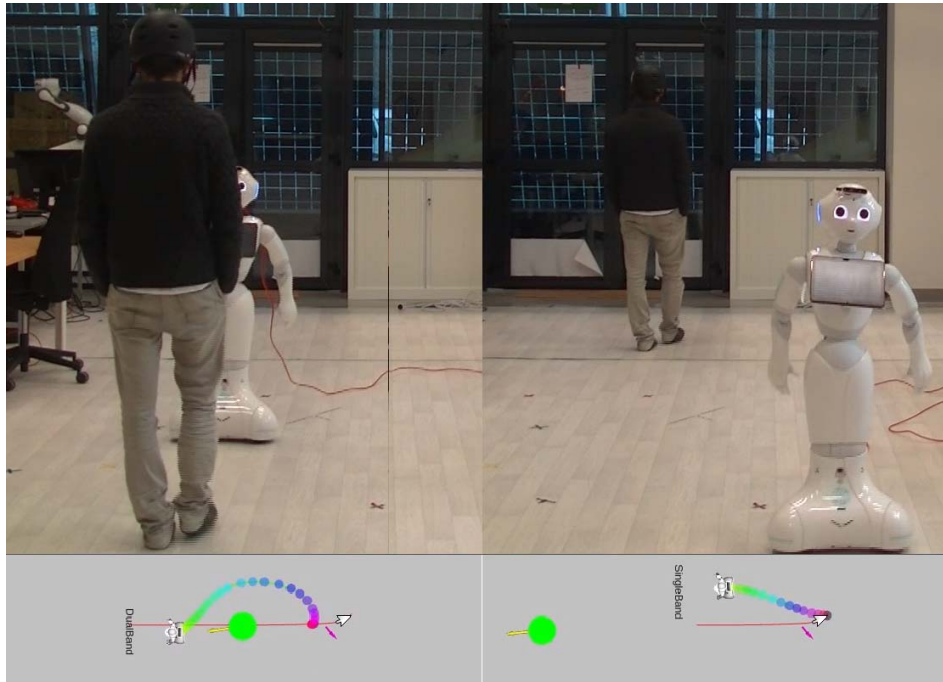


Figure 3.10: CoHAN running on Pepper in open space. The pictures show the planned trajectory of the robot around the moving human.




Figure 3.11: Testing **Backoff-recovery** on the real robot. As the human blocks the way, the robot cannot find any solution to move, and CoHAN makes the robot back off and wait for the human to pass.

Chapter. 2 using Aruco markers. However, during these experiments, the odometry of the robot shifted a lot, and even a small run drifted the robot far from its original position. Even with this bad odometry, we were able to capture two good runs, one in an open space and the other showing the **Backoff-recovery** in a narrow corridor.

Fig. 3.10 shows the instances from the run in open space. The figure also shows the trajectories of the robot at the bottom of each instance. In this scenario, as

the robot starts moving towards the goal, the human approaches faster than the robot anticipated. He comes close to the robot at the crossing point, but the robot immediately backs off and clears the way for the human. It slows down for a moment before re-planning its trajectory to the goal, as seen in the video⁷. The wrong estimate could be due to bad odometry or network delays. The second run shown in Fig. 3.11 shows instants from the test of **Backoff-recovery** mode in CoHAN. The human stands in the corridor blocking the robot's way. The planner starts in **Dual Band** mode and finally switches to **Backoff-recovery** mode as there is no solution. The robot slowly backs off and clears the way for the human by moving to the left side of the corridor, as seen in Fig. 3.11. The human then moves out of the corridor and clears the way for the robot.

3.6 Multi-context HAN and CoHAN

In this section, we present a discussion on HAN planning with multiple planning modalities and how this coupled with the situation assessment, can address multiple human-robot navigation contexts. We also present a brief discussion on how CoHAN addresses some of these contexts and talk about the possible extensions for CoHAN.

3.6.1 Modality-based HAN Planning

Some of the early works in HAN with multiple modes of planning like [Qian 2013, Mehta 2016] used POMDP based policies to choose different actions to avoid the freezing robot problem or the frequent re-planning. They show how having multiple policies benefits robot navigation and also increases the success rate. The increase in success rate does not necessarily mean that the navigation is acceptable and legible to the humans in the environment. We believe that if this multi-policy planning could be coupled with social norms of the environment, the resultant system could offer legible robot navigation among humans with high success rates. CoHAN is one such approach where situation assessment and decision-making are coupled with a complete HAN planning system and offers some interesting planning modalities. Unlike the previous works, we do not use a POMDP based approach but rather propose engineered ways to analyse situations and shift between modalities. A very recent work similar to our approach is presented by Banisetty [Banisetty 2020]. It uses a multi-objective optimization based local planner like ours and employs a deep learning based situation assessment module. Compared to the previous works with only multiple modes of planning, CoHAN and the HAN system by [Banisetty 2020] use situation assessment not only to successfully navigate the robot but also to add socially compliant behaviour to the navigation policies. These frameworks can hence be used to solve multiple human-robot navigation contexts.

⁷https://youtu.be/DB_8HpjngJ4

3.6.2 Multi-Context HAN Planning using Mode Shifting

Classical robot navigation can be generalized, and a planning scheme might work for different contexts without any modifications. On the other hand, HAN requires addressing different HRI contexts and hence, possibly requires different planning strategies. If these strategies can be employed as planning modalities, the frameworks can be used for solving multiple human-robot contexts. This approach is quite new, and limited research exists in this field. The work by [Banisetty 2020] presented above uses multi-objective optimization and geometric reasoning to address the multi-context navigation of a robot in human environments. It presents some very interesting contexts and scenarios that were addressed by combining situation assessment with HAN planning. Our approach, CoHAN, is a parallel work and addresses multi-context navigation in a similar fashion but uses proactive planning with an engineered situation assessment loop. Further, our system offers a variety of parameter settings that can choose prediction mode, the human-aware constraints to be used and tuning over these costs. Even the planning can be restricted to only one of the three planning modes (except **Backoff-recovery**). Hence, it can be further extended to many kinds of human-robot contexts by properly choosing the parameters and with simple tuning. With the addition of the costmap layers around the static humans, this framework can handle most of the scenarios presented in [Banisetty 2020]. CoHAN already uses some geometric reasoning to better understand the intricate HRI contexts with a human, and it can be easily extended to address group interactions.

3.6.3 CoHAN in Multiple Human-Robot Navigation Contexts

Currently, CoHAN is developed to address intricate human-robot scenarios that often occur in indoor environments like offices or labs. The added costmap layers and updated human-aware constraints provide promising results in situations like joining or leaving a group, over-taking a human from behind, various types of human-robot crossing scenarios at doors or in corridors and general HAN in wide spaces with few humans. With small modifications in the parameters, it can even be extended to crowded scenarios, as shown in the section 3.4. Still, there are HRI contexts CoHAN cannot address presently, like following or accompanying humans [Repiso 2017], approaching humans [Khambhaita 2016b] or taking an elevator with humans. We plan to extend CoHAN to address the scenarios mentioned above and modularise the design of CoHAN so that customized human-robot navigation contexts can be added easily. The approach modality is already studied with some implementations in HATEB, and so the next step would be completely integrating this modality into CoHAN. For the elevator scenario, a new custom modality needs to be developed such that the robot is allowed to violate the human proxemics but still acts in a human-friendly manner. Finally, the existing literature on the person-following robot can help design another modality for CoHAN.

3.7 Conclusion

In this chapter, we proposed a new HAN planner that can handle a variety of human-robot contexts. It was able to handle both outdoor crowd scenarios and indoor intricate scenarios, thanks to the different planning modes and tunable parameters. These planning modes are at the control level and, hence, differ from higher level modes as used in [Mehta 2016]. Consider the **Backoff-recovery** mode, for example, instead of going into the corridor, the robot could have stopped (or back off) as soon as it sees a person, or if the robot has progressed more than the human, the human can go back and let the robot go. By employing a higher level planner over our system, this case can be handled much more efficiently, but our focus is on providing this feature at the control level. We introduced *Human Safety* and *Human Visibility* layers into the system through costmaps to address the static human scenarios. For handling the dynamic humans, we have used a variety of human-aware constraints in *HATEB* along with visibility and planning radius. The proposed system also provided different types of human path prediction methods. We have proposed two new human-aware constraints in addition to the previous ones present in *HATEB* to offer a more legible trajectory. Further, our system was evaluated in a variety of simulated scenarios and presented both qualitative and quantitative results. Finally, real-world tests on a robotic platform were presented.

One of the major limitations of our system could be computational complexity as it performs optimization in each control loop. However, it does not affect the real-time performance of the robot in **Single Band** and **Backoff-recovery** modes (10Hz). In the other modes, it may lead to a little reduced control rate (8-9 Hz), however still in real-time. One of the immediate future works is to develop a higher-level planner on top of our system to handle the contexts more efficiently and to include more contexts like following or accompanying a human. Currently, the system is not designed for handling groups of people differently, and we plan to include it in the future version of the system.

The situations presented in this chapter and the possible extensions discussed above only deal with the humans that are currently visible and are within the planning environment. Designing a HAN system for handling all the scenarios that arise in such a setting is already complicated and requires further studies and new planning methodologies. However, there is no work addressing the humans that might emerge into the planning scene and can disturb the current motion of the robot. A complete HAN system should proactively anticipate such possibilities and be prepared instead of freezing in such scenarios. So, in the next chapter, we introduce the concept of ‘invisible humans’ to HAN to handle such scenarios.

Proactive Planning for Unseen Humans in the Environment

Contents

4.1	Introduction	81
4.2	Invisible Humans Detection	82
4.2.1	Corner Detection using Laser Contour	83
4.2.2	Estimation of Invisible Humans' Locations	84
4.3	Introducing Invisible Humans into CoHAN	87
4.3.1	The 'Invisible Humans Constraint'	87
4.3.2	Passage Detection and a New CoHAN Planning Mode	88
4.4	Results	89
4.4.1	The Effect of the Invisible Humans Constraint	89
4.4.2	Navigating in the Presence of Visible and Invisible Humans	93
4.4.3	Testing the Accuracy and Robustness	95
4.5	Real-World Tests	97
4.6	Discussion and Limitations	98
4.7	Conclusion	99

4.1 Introduction

Most of the HAN frameworks [Möller 2021, Kruse 2013] address only the visible humans and do not take into account the possible emergence of humans that are not visible currently. We believe that such '*invisible humans*' should be considered while developing a human-aware navigation framework to avoid any erratic behaviours of the robot planner when a human suddenly appears. Therefore in this chapter, we try to address these *invisible humans* in HAN settings.

There is no work that addresses this problem in the field of HAN. However, there are some existing works in classical robot navigation that address similar issues. Particularly, this work is inspired by the pioneering work of M. Krishna [Madhava Krishna 2006, Alami 2007, Madhava Krishna 2003] concerning the ability of a mobile robot, based on the model of its perception functions, to assess from where in the close environment of the robot a human can emerge and prepare

to react to ensure no-collision by adapting its path and velocity. Some recent works like [Chung 2009, Bouraine 2012] address the issues of robot navigation in occluded or unknown regions with a limited field of view. The work presented in [Miura 2006] talks about the adaptive speed control of the robot in unknown environments and also talks about the occluded regions. The authors of [Lambert 2008] propose a methodology to mitigate or avoid collisions while navigating. In our case, we are trying to mitigate possible future collisions with a human.

As it is evident that the unknown or occluded region could cause issues with classical navigation, the same applies to HAN. Hence, we propose the concept of ‘*invisible humans*’ to HAN planning in this chapter. As per our knowledge, there is no other work that addresses this problem in a human-aware navigation context. Firstly, we formulate the problem of *invisible humans* detection and propose an algorithm to determine the locations from which humans that are not currently visible can emerge suddenly. These *invisible humans* are then integrated into our human-aware navigation framework, CoHAN [Singamaneni 2021], by introducing a new human-aware constraint into our optimization scheme. The constraint modifies the path and speed of the robot, taking into account the anticipation of potential human appearances to avoid collisions and surprises. This kind of preventive planning is also a part of the proactive planning approach as it mitigates the occurrence of potential problems. We further show how the detected *invisible humans* can be exploited to identify some interesting places in a map, like doors or passages and address these by adding new modalities to CoHAN. The implementation and code can be found at https://github.com/sphanit/cohan_planner_multi/tree/model.

The organisation of this chapter is as follows. Section 4.2 presents the formulation and an algorithm to detect *invisible humans*. Section 4.3 shows how the *invisible humans* are integrated into CoHAN and talks about the issues that arise. It also presents a simple formulation to identify narrow passages. In section 4.4, various experiments to evaluate the proposed approach are presented, followed by the real-world experiments in section 4.5. A discussion on the limitations of this work is presented in section 4.6. Finally, section 4.7 concludes this chapter.

4.2 Invisible Humans Detection

The *invisible humans* are detected using an emulated laser scan on a 2D map in ROS. A custom laser scan is attached to the robot’s base, and it is continuously updated as the robot moves on a given map. The entire system is implemented in ROS [Quigley 2009] and requires the map that is published by the ROS Navigation Stack. In order to avoid too many detections, we limit the *invisible humans* detection to a radius of 5 m in front of the robot. The detection of *invisible humans* is a two-step process involving corner detection and locating invisible humans. Each step is explained in detail below.

4.2.1 Corner Detection using Laser Contour

A custom laser scan sensor attached to the robot's base scans the given 2D map to get the visible contour of the map. The laser data consists of a list of values showing the scan ranges in the field of vision of the sensor. The custom sensor data is used in place of the real laser scan data to ensure uniformity across different robots and sensors. An example laser contour built using this is shown in Fig. 4.1. Different parts of these contour lines are shown in different colours for ease of explanation. The red and the blue lines together constitute the regions on the real map where the laser has hit a wall or an obstacle. The black lines represent the laser data that did not hit anything and reached the end of their range (in our case, the range of the laser is 7 m). Lastly, the yellow lines are interpolated rays joining large separations between consecutive laser values and play a major role in our algorithm. The red circle and the arrow represent the robot's position and direction, while the green circles represent the estimated *invisible humans*. Corner detection is relatively easy once the laser contour is available. Firstly, all the pairs of consecutive laser range values separated by more than 0.5 m are determined and stored in a set $\{V\}$. The threshold of 0.5 m is chosen to filter out small gaps from where a human will not emerge. After this, the values closest to the robot's position in each of the above pairs are identified to be corners of interest and stored in a set $\{c\}$. These are shown as yellow circles in Fig. 4.1.

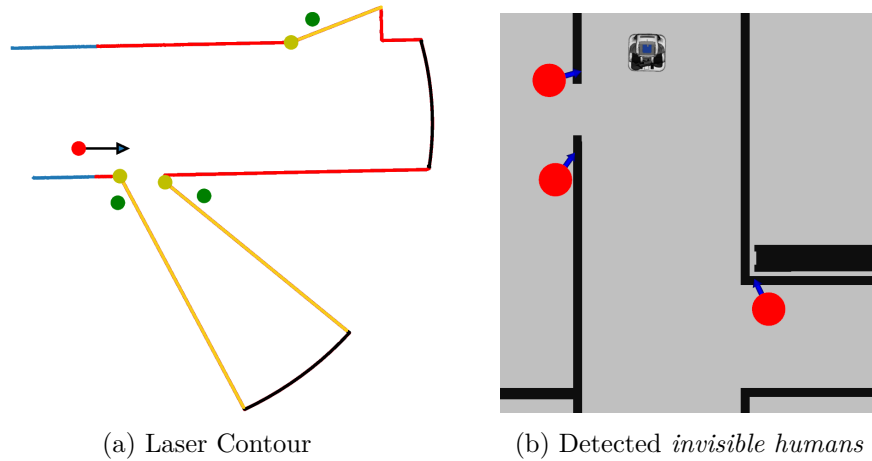


Figure 4.1: (a) Laser contour built by the custom laser sensor. The red and blue lines are the actual walls or obstacles in the front and back of the robot respectively. The black lines are the laser range boundaries and the yellow lines are interpolated lines between two gaps of laser data. The robot is shown as a red dot with an arrow. The detected corners are shown in yellow, while the detected *invisible humans* are shown in green. (b) The detected *invisible humans* on the map for the contour shown. The red circles show the location and the blue arrow shows the assumed direction, which is always oriented towards the robot.

4.2.2 Estimation of Invisible Humans' Locations

The estimation of possible locations for the *invisible humans* is not very straightforward. The laser contour forms a complex non-convex polygon, and we are searching for circles whose centres are outside this polygon and do not intersect the contour. We solve this problem using a combination of ray tracing and vector algebra. Consider a non-convex polygon as shown in Fig. 4.2. The vertices are numbered in the anti-clockwise direction. Consider a point P_1 that lies between the vertices V_1 and V_2 . If P_1 is outside the polygon, it should lie to the right of the vector $\overrightarrow{V_1V_2}$. Similarly, a point P_2 lying outside the polygon between V_2 and V_3 lies to the right of $\overrightarrow{V_2V_3}$ and so on. It holds true irrespective of the number of sides of the polygon. We exploit this property to determine the positions of the *invisible humans*. Note that a point lying outside the polygon will always be on the right side of the vectors, but not every point on the right of the vectors lies outside the polygon. This is because the methodology uses only a single side and does not consider the other sides. To handle this, we use the fact that the polygon, in our case, is determined by the laser contour, and any point outside this polygon is not visible to the laser.

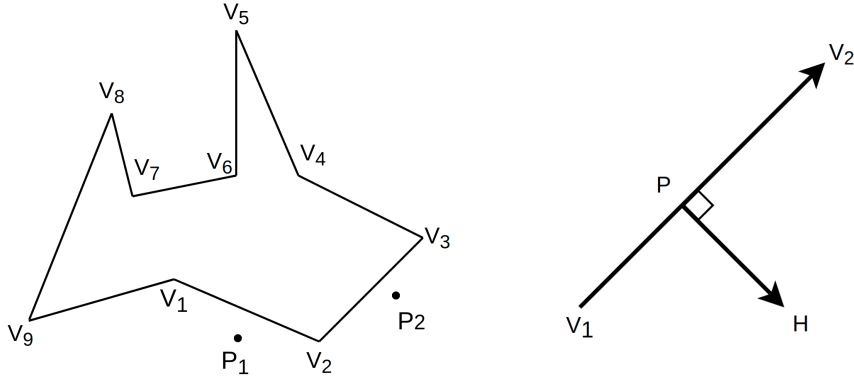


Figure 4.2: Non-convex polygon and vector formulation to determine *invisible humans*. We try to find a point H that lies between two vertices V_1 and V_2 and lies to the right of $\overrightarrow{V_1V_2}$. Note that the perpendicular distance of this point should be greater than the assumed human radius.

In Fig. 4.1, the yellow lines correspond to the edges of interest in the polygon. The numbering of the vertices is determined based on the indices of the laser scan data, which (the rays) move from right to left in the anti-clockwise direction. In order to determine a possible invisible human, we need to determine a point H that is to the right of $\overrightarrow{V_1V_2}$ and whose perpendicular distance is greater than an average human radius, h_{rad} as shown in Fig. 4.2. Consider a point P that lies on line segment $\overrightarrow{V_1V_2}$ such that \overrightarrow{PH} is perpendicular to $\overrightarrow{V_1V_2}$. If we know the point P , then H can be determined using the following equations:

$$\text{sign}(\overrightarrow{V_1V_2} \times \overrightarrow{PH}) = -1 \quad (4.1)$$

$$\overrightarrow{V_1V_2} \cdot \overrightarrow{PH} = 0 \quad (4.2)$$

$$\|\overrightarrow{PH}\| = h_{rad} + \varepsilon \quad (4.3)$$

where (\times) is the cross product, (\cdot) is the dot product and $(\|\cdot\|)$ is the euclidean norm, respectively. ε is the offset on the human radius to avoid unrealistic detections. In this work, we chose ε to be $\frac{h_{rad}}{2}$. Let $V_1 = (x_1, y_1)$, $V_2 = (x_2, y_2)$ and $P = (x_p, y_p)$, then by solving the above equations, $H = (x_h, y_h)$ is given by:

$$\begin{aligned} x_h &= x_p + \frac{d(y_2 - y_1)}{\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}} \\ y_h &= y_p - \frac{d(x_2 - x_1)}{\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}} \end{aligned} \quad (4.4)$$

where $d = (h_{rad} + \varepsilon)$. Similarly, the point on the left can be obtained by reversing the sign in Eq. (4.1), which yields:

$$\begin{aligned} x_h &= x_p - \frac{d(y_2 - y_1)}{\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}} \\ y_h &= y_p + \frac{d(x_2 - x_1)}{\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}} \end{aligned} \quad (4.5)$$

From the Eq. (4.4), we can see that (x_p, y_p) is required to determine (x_h, y_h) and it is still unknown as we cannot solve for four variables using only two equations (Eq. (4.1)-(4.2)). As it is already known that P lies on the line segment joining $\overline{V_1V_2}$, it can be determined by performing a search on this line segment, starting at one end and moving towards the other in small increments. The set of detected corners, $\{c\}$, are taken as the starting points of this search. In each iteration, a possible invisible human position is estimated using Eq. (4.4) and then projected onto the map to see if there is any overlap with an obstacle or wall. As mentioned before, we need another check to ensure that the point is outside the laser contour. Suppose the vector joining the robot and the point H is \overrightarrow{r} , and it subtends an angle β with the positive x -axis of the base frame of the robot. As the custom laser is also attached to the base frame of the robot, there should be laser scan data corresponding to this angle β . Hence, when the H is outside, the following condition is satisfied:

$$\|\overrightarrow{r}\| > \rho(\beta) \quad (4.6)$$

where $\beta = \text{atan2}(x_h - x_{rb}, y_h - y_{rb})$, (x_{rb}, y_{rb}) is the robot base frame's position and $\rho(\beta)$ is the laser scan reading at angle β . To refine this search further, two points, one on the left, P_l , and the other on the right, P_r , of the H are considered with incremental distances until $h_{rad} + \varepsilon$ and checked for overlap using the map.

The entire procedure is shown in Algorithm 1 where u_x and u_y (line 6) are the unit vectors along the direction of $\overline{V_1V_2}$ and α is a scalar determining (lines 32, 33) the step size or increment. The *invisible humans* detected using the above-mentioned algorithm are shown in Fig. 4.1 (b). The red circles are the detected location, while the blue arrows show the direction. We assume that the humans are always coming towards the robot, and hence, the direction is always oriented towards the robot. In the next section, we explain how this is integrated into the human-aware planning framework for social robot navigation.

Algorithm 1 Locate Invisible Humans

```

1: Determine the vertex pairs set  $\{V\}$  using laser contour
2: Determine the corners set  $\{c\}$  from  $\{V\}$ 
3: for each  $c$  do
4:    $V_1 = c = (x_1, y_1)$ 
5:    $V_2 = (x_2, y_2)$  ▷ Corresponding pair from  $\{V\}$ 
6:    $u_x = \frac{(x_2 - x_1)}{\|V_1 V_2\|}, u_y = \frac{(y_2 - y_1)}{\|V_1 V_2\|}$ 
7:   Set  $P = (x_p, y_p) = (x_1, y_1)$ 
8:   while True do
9:     Calculate  $H_{inv}$  using Eq. (4.4)
10:    Calculate  $\vec{r}$  and  $\beta$ 
11:    if  $\|\vec{r}\| < LaserData(\beta)$  then
12:       $x_p = x_p + \alpha u_x$ 
13:       $y_p = y_p + \alpha u_y$ 
14:      continue
15:    end if
16:    Check for overlap on the Map
17:    if no overlap then
18:       $advance = False$ 
19:      for  $i = 1$  to  $k$  do
20:         $d = \frac{i}{k} * (h_{rad} + \varepsilon)$ 
21:        Calculate  $P_r$  and  $P_l$  using Eqs. (4.4), (4.5)
22:        Check  $P_r, P_l$  for the overlap on Map
23:        if no overlap then
24:          continue
25:        else if overlap then
26:           $advance = True$ 
27:          break
28:        end if
29:      end for
30:    end if
31:    if  $advance == True$  then
32:       $x_p = x_p + \alpha u_x$ 
33:       $y_p = y_p + \alpha u_y$ 
34:    else if  $advance == False$  then
35:      break
36:    end if
37:  end while
38:  Add  $H_{inv}$  to the set of invisible humans,  $\{H_{inv}\}$ 
39: end for
40: return  $\{H_{inv}\}$ 

```

4.3 Introducing Invisible Humans into CoHAN

In the previous version of CoHAN, we address different types of visible humans by introducing new modalities and human-aware constraints [Singamaneni 2021]. In this chapter, we extend it further to address the *invisible humans*. The *invisible humans* are detected as explained above, and then they are published on a ROS topic. CoHAN subscribes to this topic and adds a new constraint to its optimization that is specifically designed for *invisible humans*. Note that we do not add any elastic band and only add an ‘invisible human-aware’ constraint to make the robot proactively plan its trajectory. Further, using these invisible human detections, we propose a methodology to identify doors and narrow passages.

4.3.1 The ‘Invisible Humans Constraint’

The *invisible humans* constraint takes into account the human reaction time, walking speed, and deceleration. It aims to make the robot cautious about sudden human emergence. The cost added by this constraint for the n^{th} pose of the robot’s trajectory is given as:

$$\begin{aligned} cost_{inv_human} &= \max\left(\frac{V - a\Delta t_n}{d}, 0\right) \quad \text{if } \Delta t_n > 0.5s \\ &= \frac{V}{d} \quad \text{otherwise} \end{aligned} \quad (4.7)$$

where d is the distance between the invisible human and the robot, V is the average human walking speed, 1.3 m/s [Teknomo 2002], a is the deceleration of the human, and Δt_n is the time difference between the n^{th} pose and the starting pose of the planned trajectory of the robot. The value of the deceleration, a , can vary and can be up to a maximum of 2.94 m/s² (0.3 g) [Lakoba 2005]. In this work, we take a reaction time of 0.5 s as discussed in [Lakoba 2005, Helbing 2000]. Hence the constraint adds the maximum possible cost until 0.5 s. Then we assume that the human will continuously decelerate to avoid collision with the robot over time and eventually stops, which is reflected in the upper part of Eq. (4.7). The time (Δt) and human detections are reset after every control cycle. The above cost makes the optimization to produce a trajectory that makes the robot take larger turns around the corners and other openings from where a human might emerge (if detected by the presented algorithm). Hence, our HAN system proactively mitigates potential collisions with unseen humans, and we believe that it is a more acceptable way than reacting after seeing a human.

4.3.1.1 Issue with the Constraint

The main objective of the constraint is to push the robot away from the opening, anticipating the emergence of *invisible humans*. However, when the robot needs to pass through this opening and if the passage is narrow (door or narrow corridor), the constraint pushes the robot away and makes it impossible to enter the passage.

To mitigate this, we devise a simple formulation that detects such scenarios. Once a narrow passage is detected, the *invisible humans* constraint is switched off, and the maximum robot's velocity is reduced until it passes through. The passage detection process is explained in detail in section 4.3.2.

4.3.2 Passage Detection and a New CoHAN Planning Mode

The detection of narrow passages or doorways not only allows us to overcome the issue of the *invisible humans* constraint but also to define a new modality of planning that needs to be handled separately. In this work, we try to address three different scenarios, as shown in Fig. 4.3. The first scenario, Fig. 4.3 (a), occurs in the case

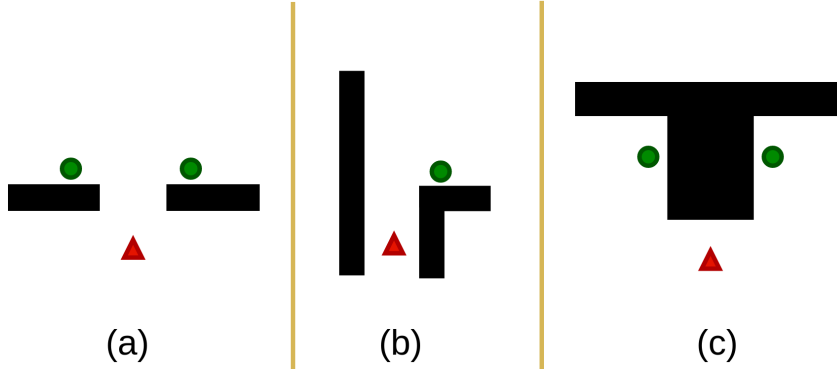


Figure 4.3: Different types of passages that are detected using *invisible humans*. (a) Doorways or openings and endings of the corridors (b) a narrow passage with an opening on one side and a wall on the other side (c) a large pillar or obstacle where the robot cannot see on either side. The green circles are the possible locations of the *invisible humans* and the red triangle shows the robot pointed towards the direction of its motion.

of a doorway or the openings and closings of a narrow passage. In such scenarios, the *invisible humans* exert equal forces from two different directions, which align the robot at the centre of the passage. If the opening is narrow, the resultant force is strong and does not let the robot pass through until the invisible constraint turns off. However, the threat of *invisible humans* still exists, so the robot should act cautiously. Therefore, we make the robot move slowly with a lower velocity ($\leq 0.3 \text{ m s}^{-1}$) until it passes through the passage. To detect this scenario, we use the positions of the *invisible humans* and the robot to check whether an *isosceles triangle* is formed with the three vertices. The robot lies on the vertex, which connects the approximately equal sides, and humans at the base vertices. In order to limit false detections, we set some numerical limits on the lengths of the equal sides and the base. Assuming that a human has 0.3 m radius and the robot has 0.5 m, the length of the base should be $\geq 1.6 \text{ m}$. When the clearance from obstacles or walls is taken into account, it increases further. In this work, we set the limit on base length as 3 m. Similarly, for the equal sides, there should be a minimum length of 0.8 m, and

we chose the limit to be 2 m. These values are chosen empirically based on the tests in several situations. If the above conditions are satisfied, a passage is detected, and CoHAN switches to a new modality called **Pass Through**, which sets the conditions mentioned above. This mode can be activated at any time and in any of the planning modes of CoHAN. Therefore, it can be seen as an asynchronous planning mode that can be triggered at any time, unlike the previous planning modes, and hence, it is independent of the situation assessment loop in CoHAN. The situation shown in Fig. 4.3 (c) is almost the same as the doorway. It is differentiated from the doorway case by reading the centre value of the laser scan data. If the value in the data is less than the length of the perpendicular bisector of the triangle's base, it is identified as a pillar or a large obstacle. CoHAN identifies this as a new case, but for now, we handle it the same we handle the doorway.

The situation shown in Fig. 4.3 (b) is different from the other two as the robot's passage is blocked on one side by an invisible human and an obstacle on the other. As the robot may or may not align in this case, it must be handled differently. We check the angle of the laser scan corresponding to the detected corner and read the value of the data that is symmetrical to this angle along the direction of the robot. If the difference between the distance of this laser scan data and the invisible human from the robot's position is ≥ 1 m, we identify it as a wall passage and set CoHAN to **Pass Through** mode once again. The threshold of 1 m is chosen empirically here.

4.4 Results

The proposed approach is tested in several settings after being completely integrated with CoHAN. In this section, we show four interesting scenarios and present a detailed analysis. In all these experiments, we assume $h_{rad} = 0.3$ m and set $k = 10$ and $\alpha = 0.2$. We use ROS-melodic with Ubuntu 18.04, and all the scenarios are simulated using MORSE [Echeverria 2011] simulator. The simulated human agents used in the experiments are controlled using InHuS [Favier 2021a], a human simulator developed in our lab.

4.4.1 The Effect of the Invisible Humans Constraint

We demonstrate the advantage of introducing proactive planning around the *invisible humans* with two different scenarios. A detailed analysis of these scenarios with and without the proposed constraint is presented. Subsequently, a comparison with some of the planners in *move_base* shows how the proposed approach can improve human comfort.

4.4.1.1 Door Crossing Scenario

To show the effect of introducing the *invisible humans* constraint into CoHAN, we present the robot with a door crossing scenario as shown in Fig. 4.4. We test the

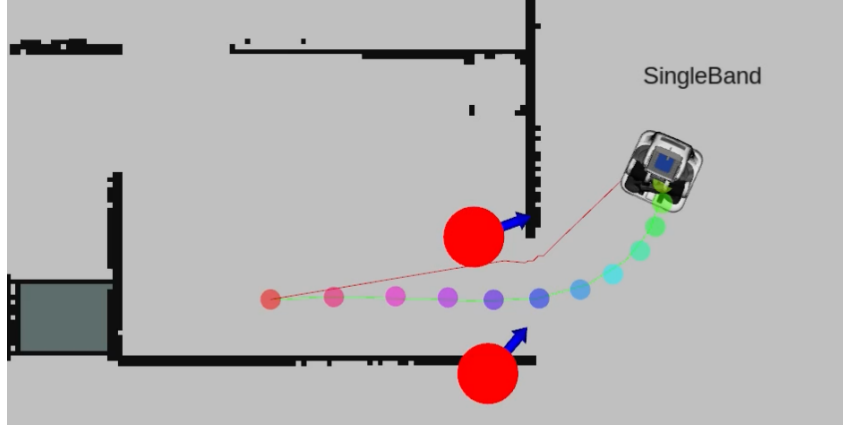


Figure 4.4: The robot passing through the door under the presence of *invisible humans*. The coloured circles represent the poses of the robot and the different colour corresponds to different time instance.

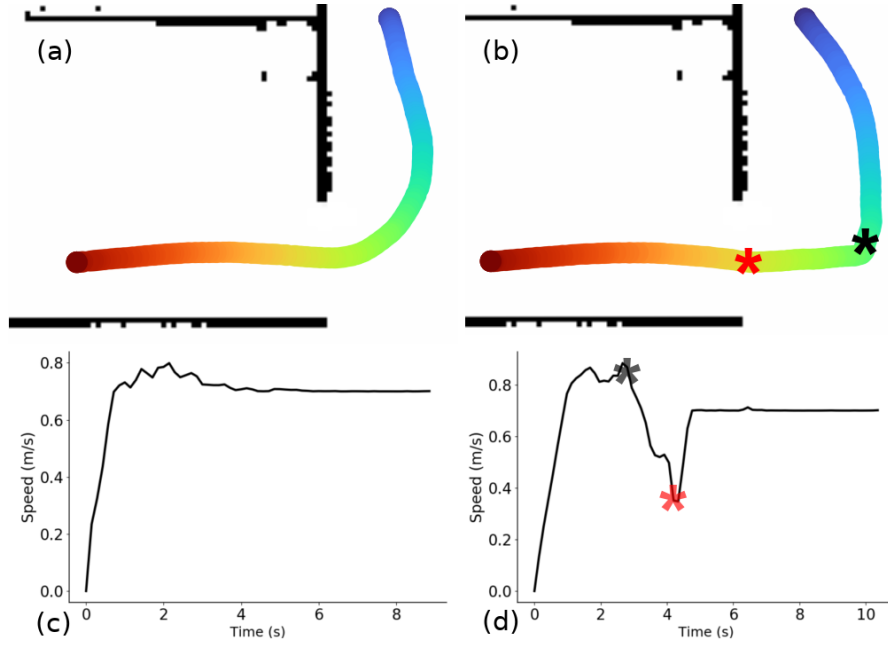


Figure 4.5: Paths and speed profiles of the robot passing the door without (a, c) and with (b, d) the *invisible humans* constraint. The colour of the paths indicates the time and progress of the robot, from blue to red (start to goal). In (a) the robot crosses the door “full speed”. In (b) it decelerates before entering the door (black star) and has the lowest speed at the entrance to the door (red star) around 4.2s corresponding to the shortest distance to *invisible humans*.

scenario without and with the *invisible humans* constraint and the corresponding paths of the robot are presented in Fig. 4.5 (a) and Fig. 4.5 (b) respectively. The paths are coloured, and the colour moves from blue to red as the robot moves from start to goal. It can be seen from these paths that the inclusion of the constraint

made the robot more cautious as it takes a larger turn and aligns its path before passing through the doorway.

The corresponding speed plots are shown in Fig. 4.5 (c) and (d). Comparing the plot in Fig. 4.5 (d) with the speed profile in Fig. 4.5 (c), it can be seen that the robot, after aligning itself with the entrance, starts decelerating slowly (black star) as it moves towards the door. At the door, there is a sharp deceleration again, and the robot passes this place (red star) at the lowest speed. This clearly reflects the cautious behaviour of the robot.

4.4.1.2 Sudden Emergence of a Human

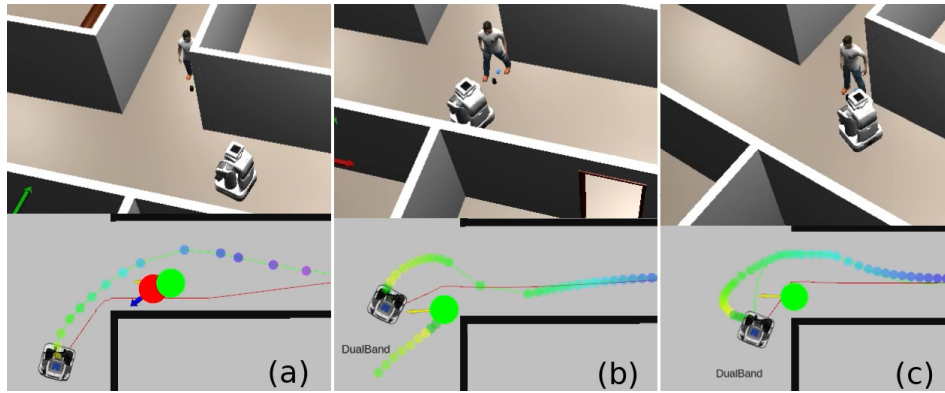


Figure 4.6: Sudden emergence scenario. The coloured paths with circles are the planned trajectory of the robot. (a) Shows the anticipated invisible human in red and the real human in green. The robot starts moving away from the corner. (b) The robot has seen the human and adjusted its trajectory to provide more space for the human. (c) The scenario without the *invisible humans* constraint. The robot moves close to the wall and blocks the human momentarily before adapting its path.

The next scenario we discuss in this section shows a situation where a human emerges suddenly from an occluded region. The snapshots of this scenario before and after the emergence are shown in Fig. 4.6. The added *invisible humans* detection predicts a possible position of the human as shown in Fig. 4.6 (a), which approximately overlaps with the real human. The robot starts moving away from the wall slowly because of this anticipation, and suddenly a real human appears in front of it (Fig. 4.6 (b)). The robot quickly adapts its trajectory and moves away from the human, slowing down a little before continuing to its goal. However, without this detection, as shown in Fig. 4.6 (c), the robot moves close to the wall and blocks the human's way for a moment before changing its path. The paths taken by the robot without and with the addition of *invisible humans* constraint to CoHAN are shown in Fig. 4.7 (a) and Fig. 4.7 (c) respectively. It is clear from these plots that the proposed constraint makes the robot move cautiously and lessens the surprise to humans. Further, the path of the robot is smoother in Fig. 4.7 (c) when compared to the path in Fig. 4.7 (a) as there are no sudden path changes.

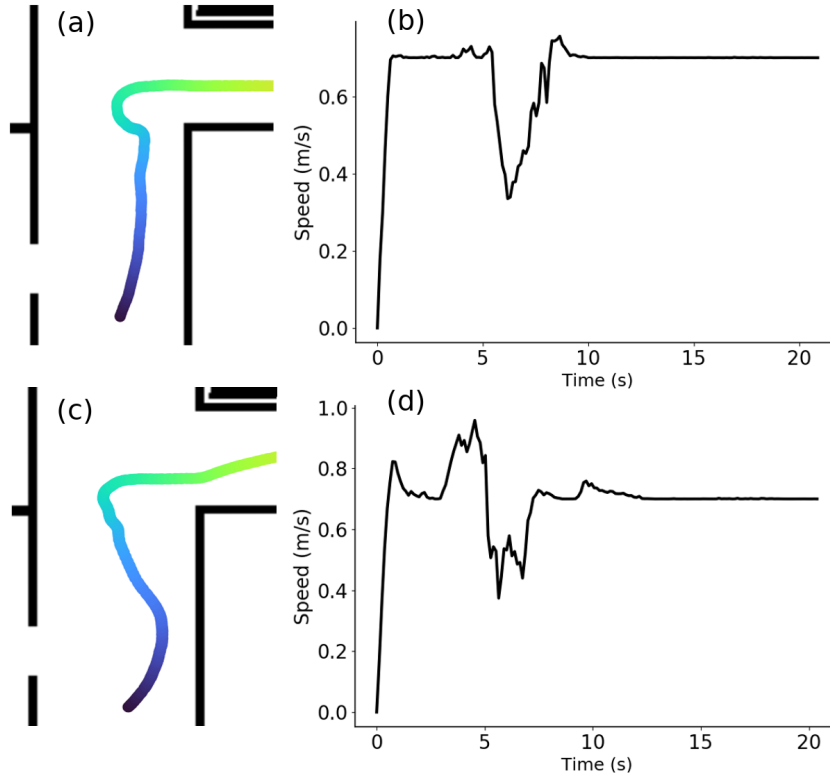


Figure 4.7: Path and speed profiles of the sudden emergence scenario without (a, b) and with (c, d) the *invisible humans* constraint. The colour of the path indicates the time and progress of the robot, from blue to green (start to goal).

The speed profile in Fig. 4.7 (b) shows a sudden drop in the velocity of the robot. This occurs because CoHAN adapts its speed to prevent a possible collision and shock to the human. Then, the robot slowly moves away and plans a new path to the goal. The speed profile in Fig. 4.7 (d) is completely different compared to the last one. The robot starts drifting away from the wall (both x and y velocities) before seeing the human, and this shows the increase in the velocity. When the human appears, it slows down and then changes its direction before continuing to the goal with almost a constant speed. This is the sharp change (slowdown) we see in Fig. 4.7 (d) between 5 and 10 seconds.

4.4.1.3 Comparison with Other Planners

To test the effectiveness of the proposed approach, we compare the sudden emergence scenario using three different planners. The first one is Simple Move Base (SMB), where humans are added using the laser scan. Then we use CoHAN with and without the proposed constraint as the other two planners. As the proposed work makes the robot cautious and tries to reduce the surprise to humans, we check the minimum human-robot distance while navigating using these planners. We have performed 5 runs of the same scenario with each planner, and the results

are presented in Table 4.1. From Table 4.1, we can see that adding the *invisible humans* constraint to CoHAN makes the robot maintain a larger distance from the human around a corner. Keeping distance from humans avoids surprise to humans and provides time for the robot planner to adapt slowly.

Planners	Min. Human-Robot Distance (m)
<i>SMB</i>	0.58
<i>CoHAN</i>	0.92
<i>CoHAN with constraint</i>	1.25

Table 4.1: Results of sudden emergence scenario using 3 different planners.

4.4.2 Navigating in the Presence of Visible and Invisible Humans

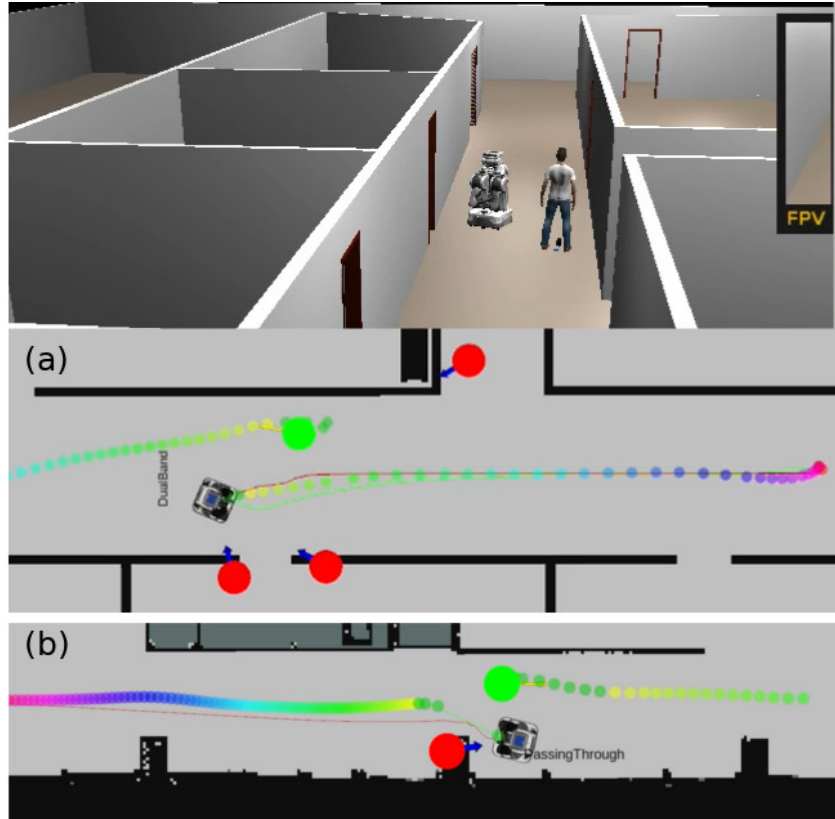


Figure 4.8: Corridor scenarios used for testing CoHAN. (a) Corridor with many openings where the robot continuously anticipates the emergence of humans. (b) Corridor with pillars between passage that creates complex navigation scenarios. In both these settings, the robot tries to find a balance between visible and *invisible humans*. The green circle is the visible human interacting with the robot, while the red circles are estimated *invisible humans*. The coloured path with circles is the planned trajectory of the robot.

The inclusion of the *invisible humans* into human-aware navigation planning should not cause discomfort to the visible humans that are moving around the robot. To show that CoHAN finds a balance between the invisible and visible humans, we present two corridor scenarios, one with many doors (or openings) as shown in Fig. 4.8 (a) and the other with pillars as shown in Fig. 4.8 (b). In both of these scenarios, the robot faces complex situations where it has to find a balance between the visible and the *invisible humans*.

In the case of the corridor with many openings, the robot anticipates that a human might emerge anytime and tries to move away from the openings. However, when it sees a human passing through the corridor, it tries to provide more space to the human by moving to one side. At the same instance, it faces the forces from the *invisible humans* and tries to find a balance between these and the visible human. By observing the path and speed profile of this scenario from Fig. 4.9 (a) and (c), we can see that the robot moves away as well as reduces its speed rapidly to accommodate the visible human. Nonetheless, it does not move very close to the door as it anticipates a human emergence.

In the corridor with pillars, the robot faces another complex situation where it has to pass through a very narrow opening and has to let the human pass through the same as shown in Fig. 4.8 (b). From the path and speed profiles of this scenario from Fig. 4.9 (b) and (d), we can see that the robot slows down rapidly while moving to a side and momentarily stops before moving forward again. Here, the robot stops and lets the visible human pass before it can continue its navigation. Further,

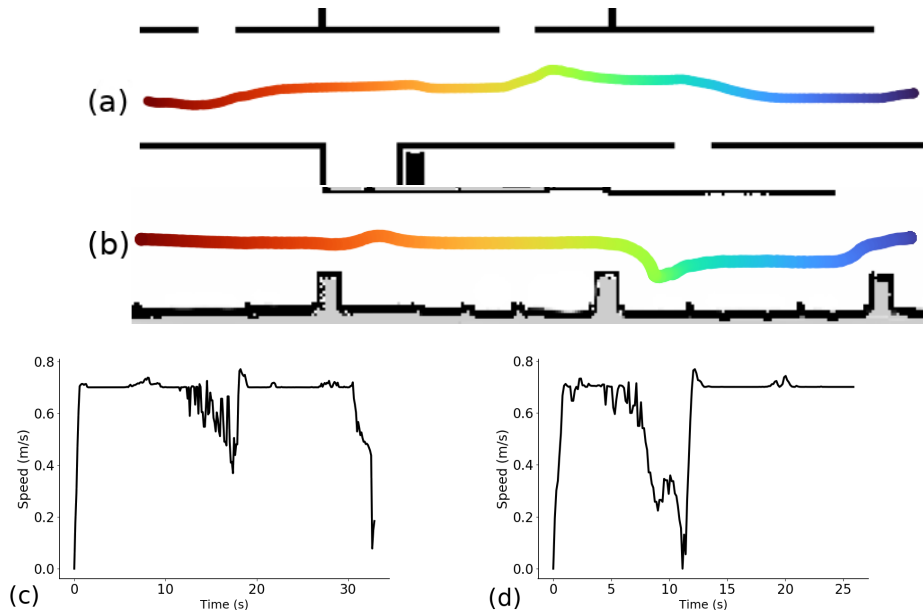
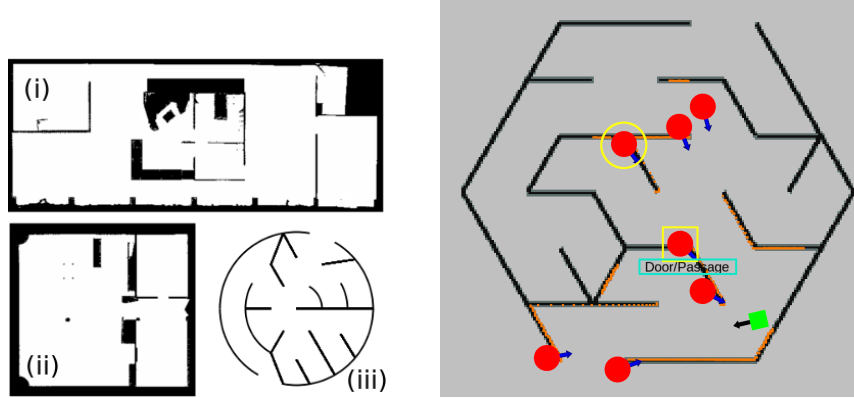


Figure 4.9: Paths and speeds profiles of the robot in corridor scenarios. (a), (c) correspond to the corridor with many openings. (b), (d) correspond to the pillar corridor. The colour of the paths indicates the time and progress of the robot, from blue (start) to red (goal).

it detects the narrow passage scenario discussed in section 4.3.2 and changes to **Pass Through** mode. We can, therefore, infer that CoHAN always tries to find a fine balance between visible and *invisible humans* and can mitigate very complex situations.

4.4.3 Testing the Accuracy and Robustness



(a) Maps used for testing (i) LAAS (ii) Bremen (iii) Random Maze (b) Failed detections

Figure 4.10: (a) Different maps used for testing the robustness and accuracy of the approach. The map in (iii) shows an example of a randomly generated maze. The other maps are standard ones. (b) Different types of failed detections. The detection in the yellow circle is a false positive whereas that in the yellow square is an overlap (possibly true). The passage detection is also wrong as the wall is detected as a passage. The green cube with the black arrow shows the robot and its direction. The red cylinders are the *invisible humans*.

For testing the proposed algorithm, we have designed randomised experiments. We either generate a random map using the Maze Generator¹ or randomise the position of the robot in the known map. The maps used for these experiments are shown in Fig. 4.10 (a). The LAAS and Bremen maps are collected from the models of real spaces. Fig. 4.10 (a) (iii) shows a random map generated using the Maze Generator. Fig. 4.10 (b) shows some failed detections using the proposed algorithm that are taken care of while calculating the accuracy. The red circles with blue arrows are the predicted *invisible humans*. The invisible human in the yellow circle is classified as a false positive, as no human could be located inside the wall. The detection in the yellow square is similar, but it is not completely inside the wall. We call this case an ‘overlap’ and classify this also as a false positive. The door/passage detected (cyan rectangle) in this picture is wrong, and we classify this as a false positive while calculating accuracy for passage detection. The green square with the black arrow is the robot in the figure.

¹<https://github.com/razimantv/mazegenerator>

4.4.3.1 Robustness and Accuracy

To test the robustness of the proposed algorithm, we perform several randomised experiments in different settings. In the LAAS and Bremen maps, we randomise the position of the robot and evaluate the detections manually. We did 50 such evaluations for each of the above two maps. In the next set of experiments, we generate a random map and place the robot in a random pose and then evaluate the detections. In this case, we have done 100 evaluations using 100 randomly generated maps. The calculated accuracy of our *invisible humans* detection algorithm based on these 200 experiments is 76.85%. However, if we include the overlaps as true detections, it increases by over 12% to 89.16%. These overlaps could be reduced by improving the filtering. Table 4.2 shows the list of experiments and the accuracy in each case.

Experiment	Invisible Humans Detection	
	<i>Accuracy (%)</i>	<i>Accuracy with overlap (%)</i>
<i>LAAS</i>	91.82	94.55
<i>Bremen</i>	65.28	81.94
<i>Random</i>	76.90	90.42
Overall	76.85	89.16

Table 4.2: Accuracy calculated based on 200 experiments in 3 different environments. By correcting the overlapping detections, the accuracy could be increased by over 12%.

For calculating the accuracy of passage detection, we have performed similar experiments and evaluated the detections in 200 experiments. Here, we classify the detection simply as either true or false. There are also cases where there are no detections. In such cases, no detection within limits is classified as false, and the out-of-the-limits is classified as a miss. Table 4.3 shows the accuracy of passage detection in different settings. The overall accuracy based on these experiments is around 62.50%. Note that the percentage of misses is around 23.50%. They can be eliminated by having higher or adaptive detection limits. The limits have to be set based on the requirement.

Experiment	Passage Detection	
	<i>Accuracy (%)</i>	<i>Misses due to limits (%)</i>
<i>LAAS</i>	66.00	30.00
<i>Bremen</i>	54.00	36.00
<i>Random</i>	65.00	14.0
Overall	62.50	23.50

Table 4.3: Accuracy calculated based on 200 experiments in 3 different environments. Increasing the limits of detection could increase the accuracy by over 23%. These limits could be decided based on the requirement.

4.5 Real-World Tests

The CoHAN system is installed on the PR2 robot in our lab and then tested in the doorway scenario discussed above. In these experiments, we do not use any human tracking as these are tests for *invisible humans*. The localisation of the robot is done using the *amcl* localization² package in ROS. The system runs entirely on the robot and only the goals were given from a remote computer.




Figure 4.11: Real world experiment setting (a) Human is inside the room and does not move. (b) The human starts coming out of the door as the robot approaches the door. In both scenarios, the robot tries to pass through the door.

We performed two kinds of experiments around the door, as shown in Fig. 4.11. In the first situation, shown in Fig. 4.11 (a), the human remains stationary, whereas in the second situation, shown in Fig. 4.11 (b), he comes out of the door as the robot approaches the door. The first situation is tested, with and without the *invisible humans* constraint, and the results are shown in Fig. 4.12. We can see from the figure that the real-world results match the results of the simulation approximately both in the paths and the speed profiles. The robot takes a larger turn and aligns itself with the door before moving towards the entrance. The video³ clearly shows the effect of the added constraint as well as the shift to the new modality **Pass Through**. Note that in Fig. 4.12 (c), the robot halts momentarily. This may be because of a sudden human appearance or moving very close to the wall.

The second situation is similar to the sudden emergence scenario, and the results of this run are shown in Fig. 4.13. The human emerges out of the door as the robot starts to move towards it. However, thanks to the *invisible humans* constraint, the robot is not very close to the door and leaves enough space for the human to pass through. In this scenario, the robot slows down twice, once when the human emerges and then to align itself with the door.

²<http://wiki.ros.org/amcl>

³<https://youtu.be/cbeFRkEdGgA>

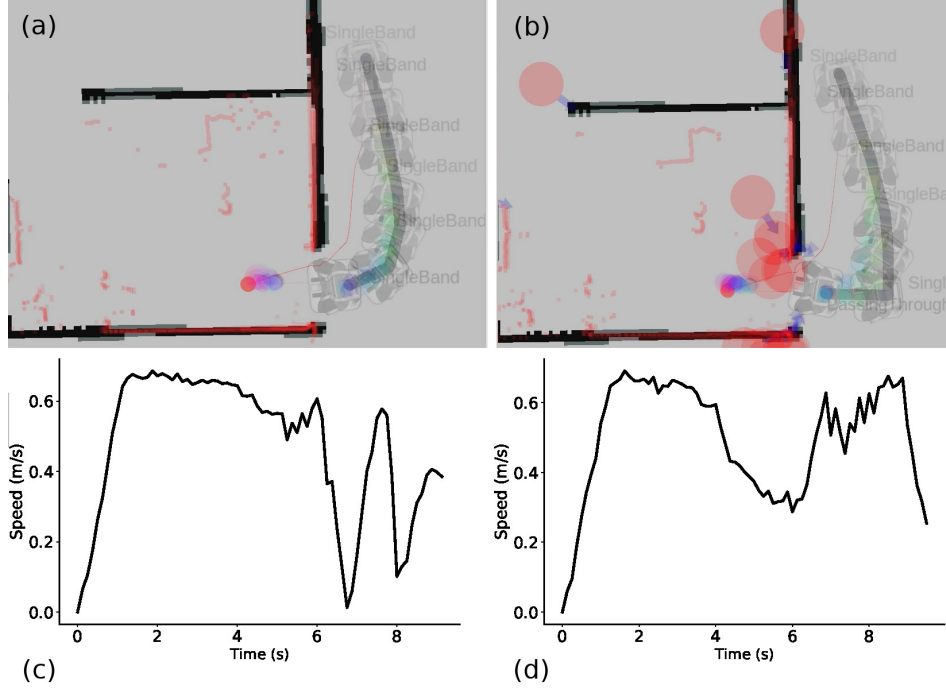


Figure 4.12: Paths and speed profiles without (a, c) and with (b, d) *invisible humans* constraint. The addition of the constraints makes robot takes larger turn (b).

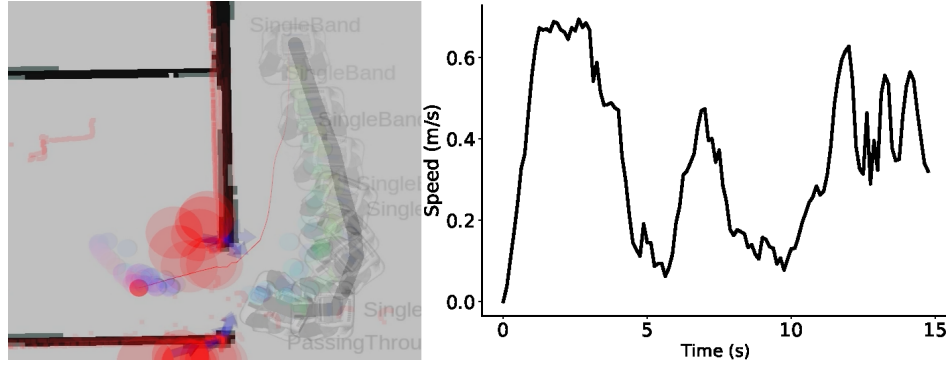


Figure 4.13: Path and speed profile in the sudden emergence scenario.

4.6 Discussion and Limitations

The introduction of *invisible humans* into human-aware navigation planning is relatively new and requires further research. We present one possible approach to address this issue. What is particularly interesting here is that our approach is modelled as a situation assessment and prediction ability to integrate into a mobile robot human-aware navigation. Having this, we can have a robot that can interact, using several modalities, with humans present in its field of view while making provisions to adapt to humans that are not yet seen. One difficulty we faced was the integration of all these features without being “too conservative” and avoiding an-

other case of the “freezing” robot. However, there are still some limitations to this approach. Since the approach is based on a 2D map, we can have false detections in the regions visible through the head of the robot but not through the base. There may also be some false detections in complex maps. Both these can be mitigated by augmenting the current approach with new sensor data and filtering further.

4.7 Conclusion

We have proposed an approach to estimate the locations of *invisible humans* on a 2D map. These *invisible humans* are then integrated into our human-aware navigation planner via a new constraint. We have also shown how narrow passages can be identified by exploiting the detected *invisible humans*. We have presented the qualitative analysis of several simulated scenarios, followed by the results of accuracy and comparisons with some planners. Finally, we have shown the real-world experiments and presented some discussion. In the future, we plan to refine this approach further and address the different modalities identified in a better manner. We also aim to build a complete human-aware navigation system that can address very intricate human-robot interactions.

Evaluation of Human-Aware Navigation

Contents

5.1	Introduction	101
5.2	Existing Evaluation methods	102
5.2.1	Navigation Metrics	102
5.2.2	Naturalness Metrics	103
5.2.3	Discomfort Metrics	103
5.2.4	Surveys and Studies	104
5.3	Proposed Discomfort Metrics	105
5.3.1	Velocity based Metrics	105
5.3.2	Vision based Metrics	108
5.4	Analysis	111
5.4.1	Scenario 1: Cross	111
5.4.2	Scenario 2: Follow and Overtake	112
5.4.3	Scenario 3: Approach	113
5.4.4	Scenario 4: Appear	114
5.5	Discussion	116

5.1 Introduction

Evaluating the efficiency and performance of a HAN planning system is one of the open questions in the field. Although there are already a set of evaluation metrics, we believe they may not do complete justice to evaluating a HAN system under intricate scenarios such as the ones addressed in this thesis. The majority of the works in HAN evaluate their systems based on the classical navigation metrics combined with proxemics violation. However, proxemics may not be applicable in all scenarios, and sometimes the robot has to intrude into the personal space. Given the circumstances, the human might not feel uncomfortable or even allow it to happen. For example, consider the case of a narrow corridor. Here, both agents have to pass close to each other to reach their goals, and proxemics violations always occur. Therefore, new evaluation metrics that are valid in such situations have

to be defined. Some of the existing studies [Takayama 2009, Lichtenthaler 2012a, Kruse 2014b] talk about how the velocity modulation is more legible and slower speeds are preferred in close vicinity by humans. Therefore, we propose a few new metrics for HAN based on distances, velocities and human visibility which can be applied to a wide range of settings.

Before proposing new evaluation metrics, we present some of the most commonly used evaluation metrics in section 5.2 and discuss some of their shortcomings. Then we present the proposed metrics in section 5.3 along with their mathematical formulation, if any. In section 5.4, we test the proposed metrics in different HAN settings and show a detailed analysis. Finally, some conclusions about HAN evaluation metrics are discussed in section 5.5.

5.2 Existing Evaluation methods

The existing evaluation methods [Gao 2021] and metrics of HAN can broadly be divided into four categories: 1) Navigation metrics, 2) Naturalness metrics, 3) Discomfort metrics and 4) Surveys and studies. We briefly describe each of these categories and present some of the commonly used metrics. Our main focus will be on the discomfort metrics as they deal with the evaluation of human-robot interaction.

5.2.1 Navigation Metrics

All the metrics of navigation can be used to evaluate the navigation performance of a HAN system. They can be used to test the robustness and stability of the designed system and are seldom useful to evaluate the interaction with a human. Some of these metrics are listed here:

- *Path Length*: Comparing the length of the path is one of the basic evaluations, and classically, the shorter the path, the better the system. In HAN, however, it may not be true.
- *Path Efficiency*: It is the ratio of the distance between two waypoints to the traversed distance between these two waypoints. This measure shows how well the controller is tracking the path.
- *Time to reach goal*: It is again one of the simplest metrics which computes the time taken by the robot to reach the final goal. A shorter time is usually preferred, but in HAN, it is not that straightforward.
- *Success Rate*: Success rate is used to quantify the navigation performance and measures how many times the robot navigation was successful over several trials. Irrespective of the type of system, this metric can be applied.
- *Number of collisions*: Similar to success rate, counting the number of collisions applies to all navigation systems. Some works in HAN use a modification to measure human safety, called collision index [Truong 2014] and the lower the index, the better.

5.2.2 Naturalness Metrics

The naturalness of the robot's trajectory is measured using similarity and smoothness metrics. Similarity metrics measure how similar the robot's trajectory is when it is compared to an expert human trajectory. The smoothness metrics, on the other hand, measure how smooth the robot's motion was during the navigation. Some of these metrics are presented below:

- *Path Deviation*: There are different ways in which path deviation is measured. These metrics measure how similar the robot's motion is when compared to an expert trajectory. Two of these ways are:
 - *Average Displacement Error (ADE)*: It is the average euclidean distance between the predicted robot's trajectory and the expert (or human) trajectory [Pellegrini 2009]. Some works used a non-linear version of this [Alahi 2016].
 - *Final Displacement Error (FDE)*: It is the distance between the final goal of the predicted trajectory and the given ground-truth data at the same time.
- *Cumulative heading changes*: The changes in the path and the amount of unnecessary turning over the whole path are good metrics to measure the smoothness of the robot's motion. It is also called path irregularity [Guzzi 2013].
- *Acceleration and Velocity*: The motion profiles of the robot can tell us a lot about the smoothness of the trajectory and are some of the basic metrics to use.
- *Topological Complexity*: Some works [Mavrogiannis 2018, Mavrogiannis 2019] measure the entanglement in paths using topological complexity index [Dynnikov 2007] and use it as a measure of legibility.

The smoothness and topological complexity metrics can be used to measure legibility in HAN systems. Legibility comes with expressiveness, and hence it is necessary to define some metrics for measuring navigation intention expressiveness apart from legibility.

5.2.3 Discomfort Metrics

HAN is essentially a human-robot interaction in the context of robot navigation, and it is necessary to measure how well or how bad the interaction was. Discomfort metrics aim to measure this and tell how well a HAN system is performing. Therefore, these metrics are naturally employed to compare different HAN systems. Most of the existing metrics are distance based and largely rely on the proxemics theory. The most commonly used metrics measure the intrusions into the spatial zones around humans, and some of these are:

- *Personal Space Intrusions*: The space surrounding a human is divided into different concentric circles (or other shapes) [Rios-Martinez 2015] with varying radii defining the proxemics zones. The intrusions into personal and intimate spaces are usually taken as a measure of discomfort, and many works [Okal 2016, Pérez-Higueras 2018a] count the number of these intrusions to quantify a HAN system. Truong et al. [Truong 2017a] defined a Social Individual Index (SII) based on proxemics (same as collision index in [Truong 2014]), and a good HAN system will always have this value below the given threshold.
- *Interaction Space Intrusions*: Similar to individual humans, a group of humans also have various interactions zones like *o-space*, *p-space* and *r-space* [Ferrer 2017, Rios-Martinez 2015], and their shape depends on the *f-formation* [Kendon 2010] maintained by the members of the group. In such settings, the intrusion into *p-space* and *r-space* are considered as violations and HAN systems try to minimize these intrusions. Sometimes, the human-object interactions are also considered [Truong 2017a] (use Social Group Index (SGI)) while measuring these intrusions.
- *Time spent in proxemics zones*: A better alternative to the number of intrusions is to measure the time spent in the areas associated with the human's personal zone or the group's interaction zone [Kostavelis 2017].
- *Distance to a human*: While the robot is navigating, there may be different distances from different humans in the environment. In general, HAN systems use either minimum or maximum distance to the closest human and the average distance over all the humans to quantify the system.

One of the very useful metrics combining the velocity and distance is called *Time-to-Collision (TTC)* [Biswas 2022], and it represents how the robot's motion is relative to the human's motion. The higher the value of *TTC*, the better the navigation. Truong et al. [Truong 2017a] define something similar called the Relative Motion Index (RMI), and the higher its value, the lower the social acceptance. After checking the formulation, one can notice that RMI is the inverse of *TTC*. One of the metrics proposed in this chapter is based on *TTC*, and for this, we provide a detailed description and mathematical formulation in the subsequent sections.

5.2.4 Surveys and Studies

Human interactions and their psychological impressions and states cannot be quantified using just numbers. It requires some well-designed studies and questionnaires, and experimental evaluations on a real system or through videos. HAN does the same to evaluate the psychological safety (discomfort and stress) and the sociability of the system under study. The perceived psychological safety is commonly measured using questionnaires [Butler 2001]. Some of the established questionnaires in social robotics like Godspeed [Bartneck 2008] already include perceived safety and

emotional states and the Robotic Social Attributes Scale (RoSAS) measures several psychological factors based on the Godspeed questionnaire.

The social intelligence of a robot, sometimes called sociability, is not very easy to quantify. A robot's motion may be perfectly natural, satisfies all the comfort-based metrics and can still be perceived as not completely social. In several intricate scenarios, the robot might not reach the high-level expectations of a human and fail to convey its intention or behaviour to the human. Barchard et al. [Barchard 2018] proposed Perceived Social Intelligence (PSI) scales to evaluate around 20 aspects of the social intelligence of a robot. This was used in some of the recent works [Banisetty 2021, Barchard 2020] to evaluate the HAN system. Some works [Vega 2018] have employed custom questions apart from these scales to study sociability.

5.3 Proposed Discomfort Metrics

We believe that social space violations by a robot can be permitted to some extent as long as the robot's intentions are conveyed clearly to the human. One of the recent studies by Joosse et al. [Joosse 2021] proposed that people are more lenient when a robot intrudes on their personal space than a human. Further, they showed that these intrusions could be mitigated by conveying the robot's intention to the humans. These studies and observations lay the basis for our metrics based on velocity. The second set of metrics is based on the human's visibility and recognition of the robot when it appears unexpectedly or when it is trying to approach a human for an interaction. We introduce each of these metrics and provide their mathematical formulation in this section.

5.3.1 Velocity based Metrics

Depending on the directions and the magnitudes of the velocities of humans and the robot, we define two metrics (or costs) to evaluate HAN systems. Fig. 5.1 shows two such scenarios where the human and robot can approach each other, or the robot can pass by the human. We define two costs to assess such situations namely, 1) $cost_{danger}$ and 2) $cost_{passby}$. The idea behind $cost_{danger}$ is to evaluate how fast a robot moves directly towards the human causing discomfort and, in some cases, fear. This is primarily associated with the danger of collision, hence the name $cost_{danger}$. If something or someone passes at a very high speed in close vicinity, it could lead to confusion or discomfort. The second metric, $cost_{passby}$, could be used to evaluate how the HAN system handles such scenarios. We now proceed to the details and the mathematical formulations of these costs.

5.3.1.1 Formulation

Suppose the circumscribed circle of the human has a radius of r_h and that of the robot has a radius of r_r . The human and the robot collide when the distance

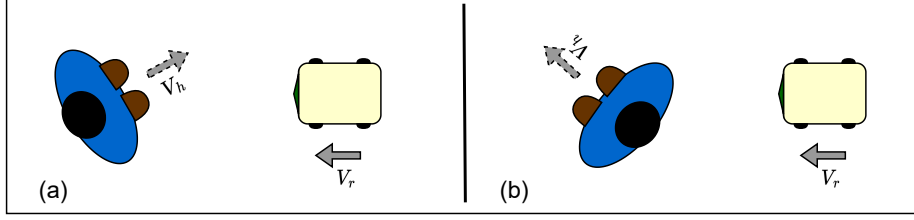


Figure 5.1: Human-Robot situations where velocity based metrics are important. (a) Cross: The human and the robot approach and cross each other at different velocities. (b) Overtake or Follow: The robot is behind the human and is about to overtake or follow the human. In both situations, the human can be either static or moving but the robot is always moving.

between their centres is less than or equal to the sum of these radii, $r_h + r_r = R$. If we assume the robot is a point and expand the radius of the human to R , the same conditions apply. This setting is shown in Fig. 5.2 along with the velocities of the human, \vec{V}_h and the robot, \vec{V}_r . The relative velocity of the robot with respect

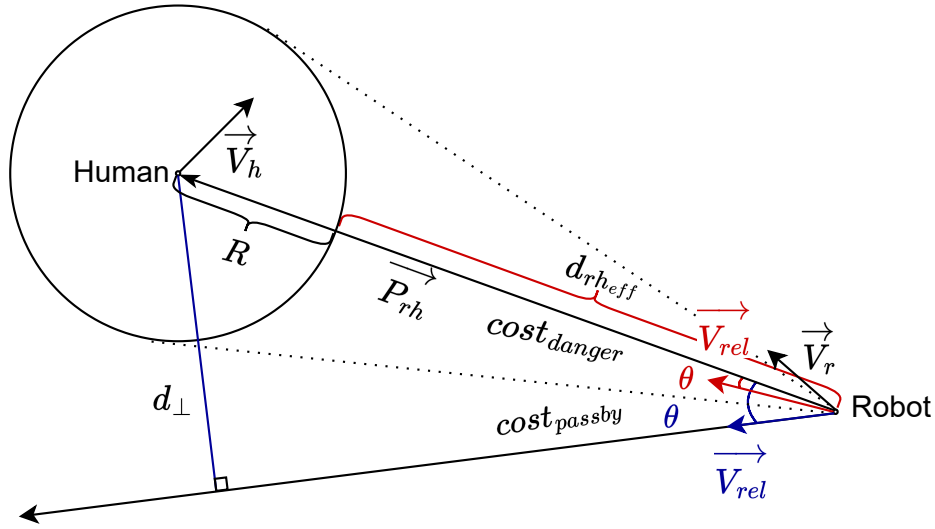


Figure 5.2: Mathematical Formulation of the costs. The figure shows the different vectors and the possible costs depending on the situation. If \vec{V}_{rel} falls within the zone indicated by the dotted lines then there is $cost_{danger}$ and if it falls outside this zone then we have the $cost_{passby}$. Human's velocity is represented by \vec{V}_h , the robot's velocity by \vec{V}_r , and the position vector from the robot to the human by \vec{P}_{rh} . θ is the angle between \vec{V}_{rel} and \vec{P}_{rh} . Note that only one of the \vec{V}_{rel} (blue or red one) exists at a time, but not both.

to the human is given by $\vec{V}_{rel} = \vec{V}_r - \vec{V}_h$ and depending on where it falls (see Fig. 5.2), we define two types of metrics or costs. If it falls within the collision zone, represented by dotted lines, then there is a danger of collision, and hence, we define

the $cost_{danger}$ in this setting. If it falls outside the collision zone, the danger of collision no longer exists, but the robot may pass by the human, and so in this setting, we define the $cost_{passby}$. Let the vector from the robot's position to the human's position be \vec{P}_{rh} and θ be the angle between \vec{V}_{rel} and \vec{P}_{rh} . While defining these costs, we use the effective distance between the human and the robot, d_{rheff} and the perpendicular component of \vec{P}_{rh} along \vec{V}_{rel} , d_{\perp} .

5.3.1.2 Metrics

The first of the two metrics is the cost of danger when there is a possibility of collision. This metric can be seen as a measure of the feeling of threat a robot can cause while it moves directly towards the human. For this, we use first calculate the TTC based on the above formulation,

$$TTC = \frac{d_{rheff}}{\|\vec{V}_{rel}\|} = \frac{\vec{P}_{rh} \cdot \vec{V}_{rel} - \sqrt{(\vec{P}_{rh} \cdot \vec{V}_{rel})^2 - \left(\|\vec{V}_{rel}\|^2 \left(\|\vec{P}_{rh}\|^2 - R^2\right)\right)}}{\|\vec{V}_{rel}\|^2}$$

when $\vec{P}_{rh} \cdot \vec{V}_{rel} > 0$ and $(\vec{P}_{rh} \cdot \vec{V}_{rel})^2 - \left(\|\vec{V}_{rel}\|^2 \left(\|\vec{P}_{rh}\|^2 - R^2\right)\right) > 0$

where ' $\|\cdot\|$ ' is the magnitude of the vector and ' \cdot ' is the dot product. Note that TTC is only defined when the \vec{V}_{rel} falls within the collision zone. Outside this zone, $TTC = \infty$. When $TTC > 0$, the cost of danger is defined as follows,

$$cost_{danger} = \frac{1}{TTC} = \frac{\|\vec{V}_{rel}\|}{d_{rheff}} \quad (5.1)$$

When the \vec{V}_{rel} falls outside the collision zone, the $cost_{danger} = 0$ and as $|\theta|$ or d_{rheff} decreases or $\|\vec{V}_{rel}\|$ increases, $cost_{danger}$ increases. Any HAN system can be defined to have a certain threshold for this cost beyond which some mitigating actions have to be taken. The second metric that we propose is valid when \vec{V}_{rel} is outside the collision zone. In such a setting, the robot may follow or pass by the human, and we define the $cost_{passby}$ as,

$$cost_{passby} = \frac{\|\vec{V}_{rel}\|}{d_{\perp} - R} |\sin(\theta)|$$

$$= \frac{\|\vec{V}_{rel}\| \left| \sqrt{\|\vec{V}_{rel}\|^2 \|\vec{P}_{rh}\|^2 - (\vec{P}_{rh} \cdot \vec{V}_{rel})^2} \right|}{\|\vec{P}_{rh}\| \left(\sqrt{\|\vec{V}_{rel}\|^2 \|\vec{P}_{rh}\|^2 - (\vec{P}_{rh} \cdot \vec{V}_{rel})^2} - \|\vec{V}_{rel}\| R \right)} \quad (5.2)$$

when $\vec{P}_{rh} \cdot \vec{V}_{rel} > 0$ and $d_{\perp} > R \implies \sqrt{\|\vec{V}_{rel}\|^2 \|\vec{P}_{rh}\|^2 - (\vec{P}_{rh} \cdot \vec{V}_{rel})^2} > \|\vec{V}_{rel}\| R$. The $cost_{passby}$ increases as the robot approaches the human. Specifically, the cost

increases as the perpendicular distance of the robot from the human decreases or the relative velocity increases. It also increases as θ increases, indicating that the robot is getting closer to the human. The main idea behind this metric is that a robot passing very close to a human at a very high speed could cause discomfort to the human.

5.3.2 Vision based Metrics

A human moving in an environment depends on different senses to identify and negotiate the path or the moving direction. In such conditions, if something or someone appears suddenly without any intimation, the human might be surprised, and sometimes he/she might not even feel comfortable with such appearances. This applies to static humans as well. We believe that in the case of static humans, the robot has to maintain a larger distance behind the human [Avrunin 2014] to avoid causing any discomfort or surprise. Moreover, when the robot is planning to enter a human's field of view (FoV), it should enter at a convenient distance and angle to avoid any compensatory actions by the human. An illustration of such conditions is provided in Fig. 5.3. The depiction on the left shows a robot trying to enter the human's FoV for interaction, and the one on the right shows the sudden appearance of a robot in front of the human without having any knowledge about the whereabouts of the human.

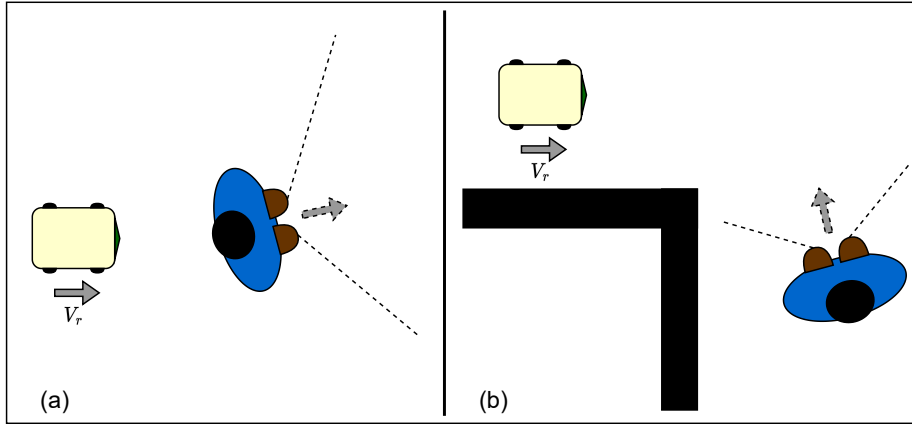


Figure 5.3: Human-Robot interaction scenarios where vision-based metrics are needed apart from velocity based ones. (a) Approach: The robot is approaching humans for an interaction. The robot needs to go from the back to the front to initiate the interaction. (b) Appear: The robot and human suddenly face each other without any prior knowledge about the other. The human can be static or dynamic in both these settings.

Considering all these, we propose $cost_{visibility}$, $cost_{surprise}$ and $cost_{react}$ to evaluate such situations. The first one, $cost_{visibility}$, measures how visible the robot is when it enters the FoV of a human below a certain threshold distance. The idea is to evaluate how well the HAN system adapts to the visibility of the robot.

$cost_{surprise}$, on the other hand, measures how surprised or uncomfortable a human is when a robot appears suddenly in a human's FoV, whereas $cost_{react}$ measures how safe the appearance is. These metrics are designed to evaluate how well a HAN system handles sudden or occluded emergences of humans.

5.3.2.1 Formulation

In this formulation as well, we consider the robot as a point, and the circumscribed radius of the human is taken as $R = r_h + r_r$. The vector from the human's position to the robot's position is represented by \vec{P}_{rh} , and the unit vector in the direction of orientation of the human is given by \hat{u}_h . The angle between \vec{P}_{rh} and \hat{u}_h is

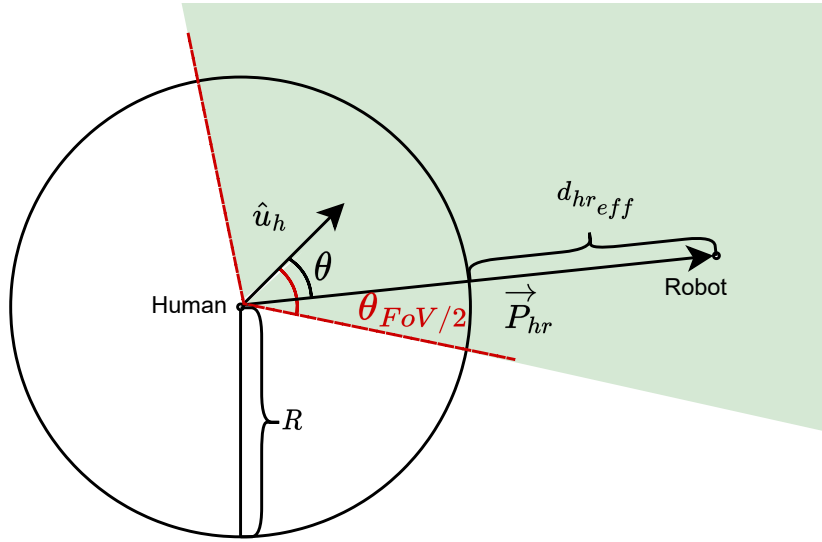


Figure 5.4: Mathematical formulation for vision-based metrics. θ is the angle between the unit vector of human's direction, \hat{u}_h and vector \vec{P}_{hr} . $\theta_{FoV/2}$ is the half angle of human's field of view and $d_{hr_{eff}} (= \|\vec{P}_{hr}\| - R)$ is the effective distance between the human and the robot.

represented by θ and the half angle of human's FoV by $\theta_{FoV/2}$. The effective distance between the human and the robot is given by $d_{hr_{eff}}$. When the robot comes into the FoV of the human (green area in the figure), we define three costs using $d_{hr_{eff}}$, θ and some studies based on human perception.

5.3.2.2 Metrics

We first define the visibility cost, $cost_{visibility}$ as follows,

$$cost_{visibility} = \frac{d_{proxemics}}{d_{hr_{eff}}} \left(\frac{\theta}{\theta_{FoV/2}} \right) = \alpha \left(\frac{\cos^{-1}(\hat{u}_h \cdot \vec{P}_{hr})}{\|\vec{P}_{hr}\| (\|\vec{P}_{hr}\| - R)} \right) \quad (5.3)$$

where $\alpha = \frac{d_{proxemics}}{\theta_{FoV/2}}$ and $d_{proxemics}$ is the defined proxemics-based distance that does not intrude the personal space (> 0.45 m). The lesser the cost of visibility,

the better the behaviour of the robot. This cost increases as θ increases indicating that more human effort is needed to look at the robot (head motion). The cost also increases as the distance, $d_{hr_{eff}}$, decreases to discourage close and sudden appearances.

When a human sees something, it takes a few milliseconds before it is registered by the brain. After the recognition, it takes more time to generate a response. The average recognition time for a human is around $t_{recognise} = 150$ milliseconds [Thorpe 1996, Rayner 2009] and the average time to react for visual stimuli is between $t_{react} = 400 - 600$ milliseconds [Eckner 2012, Wolfe 2020]. If something appears very close to a human before it can be recognised, it could result in a surprise or shock. Even after recognition, if there is no time to respond, it could result in a collision or may cause discomfort. Taking all of this into account, we propose two costs to measure surprise and response based on the robot's position when it appears. They follow similar formulations with some differences. We first define a linearly increasing function called the 'seen ratio' (SR) as given below,

$$SR = \begin{cases} \frac{t}{t_{react}}, & \text{if } 0 < t < t_{react} \\ 1, & \text{if } t \geq t_{react} \end{cases} \quad (5.4)$$

SR starts at zero the moment the robot enters the human's FoV, and as time, t , passes, this ratio slowly increases until it reaches one. When the value of SR is one, it means that the robot has been seen and identified by the human with enough time to react if needed. Hence, we can say that the SR becomes one at around $t = t_{react} = 600$ milliseconds and continues to stay at one until the robot moves out of FoV of the human. The costs are defined as follows,

$$cost_{surprise} = \max\left(\frac{d_{proxemics}}{d_{hr_{eff}}}(1 - \gamma SR), 0\right) \quad (5.5)$$

$$cost_{react} = \frac{d_{proxemics}}{d_{hr_{eff}}}(1 - SR) \quad (5.6)$$

where $\gamma = \frac{t_{react}}{t_{recognise}}$. Both $cost_{surprise}$ and $cost_{react}$ are high if the effective distance, $d_{hr_{eff}}$, is low when the robot enters the FoV of the human. $cost_{surprise}$ lasts only for a short period, and if the robot comes very close to a human during this period, the human can be marked as surprised with a corresponding cost. The $cost_{react}$ measures how the robot's appearance affects the human. If the $cost_{react}$ is high, that means the HAN system has done a poor job in handling sudden appearances, and it needs to be improved.

For all the costs (metrics) proposed, it is ideal to set a threshold below which the behaviour is acceptable, and these can be obtained based on real-world demonstrations and studies. However, such studies are yet to be performed, and therefore, we can only have a relative analysis with respect to another planner. In the next section, we do this comparative analysis using the standard ROS Navigation Stack and CoHAN.

5.4 Analysis

The proposed metrics are used to evaluate four human-robot interaction scenarios given in figures 5.1 and 5.3 i.e., 1) Cross (Fig. 5.1 a): the human and robot cross each other face-to-face, 2) Follow and Overtake (Fig. 5.1 b): the robot follows and then overtakes a slow-moving human, 3) Approach (Fig. 5.3 a): the robot moves from the back to the front of a static human for interaction and 4) Appear (Fig. 5.3 b): the human suddenly appears as the robot is taking an L-turn.

In all the experiments, $d_{proxemics} = 1.6$ m, $\theta_{FoV} = 120^\circ$ and $r_h = 0.3$ m. The interactions are simulated in the MORSE simulator, and the human agent is controlled using InHuS [Favier 2021a]. The robot is controlled using two different systems, CoHAN and Simple Move Base (SMB). SMB uses the standard ROS navigation stack and adds humans as obstacles. It neither uses human motion prediction nor social norms while navigating. Hence, we expect our metrics to differentiate clearly between SMB and CoHAN. Each scenario was run 5 times using both the planners and the mean values of the metrics are tabulated in Table 5.1. In this section, we analyse each of the above four scenarios in detail using these metrics.

Planner	Cross				
	$cost_{danger}$	$cost_{passby}$	$cost_{visibility}$	$cost_{surprise}$	$cost_{react}$
CoHAN	0.18	1.71	0.0	0.0	0.0
SMB	1.19	9.5	0.0	0.0	0.0
	Follow and Overtake				
	$cost_{danger}$	$cost_{passby}$	$cost_{visibility}$	$cost_{surprise}$	$cost_{react}$
CoHAN	0.31	0.90	1.21	0.98	1.19
SMB	0.27	0.92	2.66	2.15	2.61
	Approach				
	$cost_{danger}$	$cost_{passby}$	$cost_{visibility}$	$cost_{surprise}$	$cost_{react}$
CoHAN	0.28	1.05	2.27	1.91	2.32
SMB	0.31	1.95	3.53	3.01	3.66
	Appear				
	$cost_{danger}$	$cost_{passby}$	$cost_{visibility}$	$cost_{surprise}$	$cost_{react}$
CoHAN	1.46	7.22	0.49	0.70	0.85
SMB	3.15	2.27	0.77	0.91	1.11

Table 5.1: Proposed discomfort metrics in four different scenarios. SMB and CoHAN were run 5 times in each of the scenarios, and the mean values are presented.

5.4.1 Scenario 1: Cross

In this scenario, the robot is already in the FoV of the human as it starts moving. So, all the vision-based metrics will be zero, and only the velocity based metrics are relevant here. Comparing the values of $cost_{danger}$ and $cost_{passby}$ from Table 5.1, the costs corresponding to CoHAN are significantly lower than costs for SMB. As CoHAN is a human-aware planner, it tries to provide more ways for the human by moving away as shown in Fig. 5.5 (b) and thanks to the *Relative Velocity* constraint, it even slows down as it passes by the human. These behaviours result in low

$cost_{danger}$ and $cost_{passby}$. On the other hand, SMB does not modulate its velocity much and takes only a small deviation to avoid the collision, as shown in Fig. 5.5 (a). The slight decrease in human speed and the path change can be seen in Fig. 5.5 (a). Sometimes, the human took the full burden of avoiding the collision. These behaviours may not be acceptable, and it is reflected by the high values in $cost_{danger}$ and $cost_{passby}$.

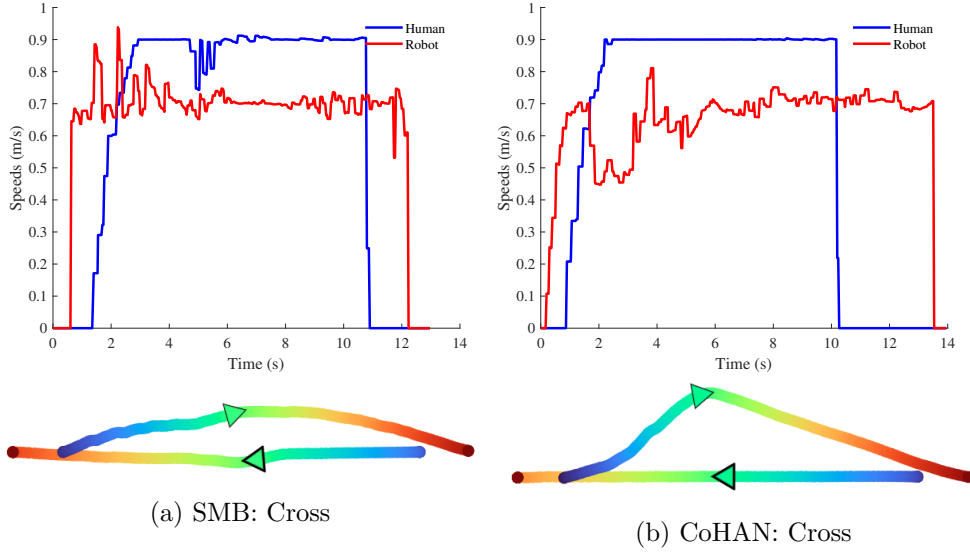


Figure 5.5: Speeds and paths of the human and robot in the Cross scenario. The triangle with the thick line is the human, and the other one is the robot. Both the human and robot move from the blue end to the red end of the paths. (a) In the case of SMB, the robot moves close to the human, and the human path is slightly modified. As they cross each other, there is a slight decrease in the velocity of the human. (b) The robot running CoHAN moves away showing the intention and also keeping its distance from the human. So the human path is almost a straight line, and the velocity remains constant.

5.4.2 Scenario 2: Follow and Overtake

As explained above, the robot starts following a slow-moving human and finally overtakes him before reaching its goal (Fig. 5.6). In this situation, a HAN system should have small values for all of the proposed metrics as the robot should pass by the human without colliding and not surprise the human as it overtakes him. From the values of the metrics in the table, it is evident that CoHAN performs comparatively better than SMB. Specifically, $cost_{surprise}$, $cost_{react}$ and $cost_{visibility}$ are significantly high for SMB, indicating that the robot might have entered a human's FoV suddenly and at a closer distance which is not ideal. This is true as one can see from Fig 5.6 (a) that the robot was very close as it overtook the human. The robot using CoHAN modifies its trajectory to accommodate the *Visibility* constraint

and enters the human's FoV in a better manner, as seen in Fig. 5.6 (b). However, we cannot claim that this is the best, as thresholds and benchmarks still need to be set.

In this setting, the robot and the human have some perpendicular offset distance, unlike in the previous case, and they move almost parallelly for the most part (see Fig. 5.6). As a result, both SMB and CoHAN have similar $cost_{danger}$ and $cost_{passby}$. The $cost_{passby}$ is slightly lower for CoHAN because of the velocity modulation, and $cost_{danger}$ is slightly higher because of the directional changes imposed by different constraints. The robot with SMB moves in a straight line, and when it is very close to the goal, it changes its path slightly.

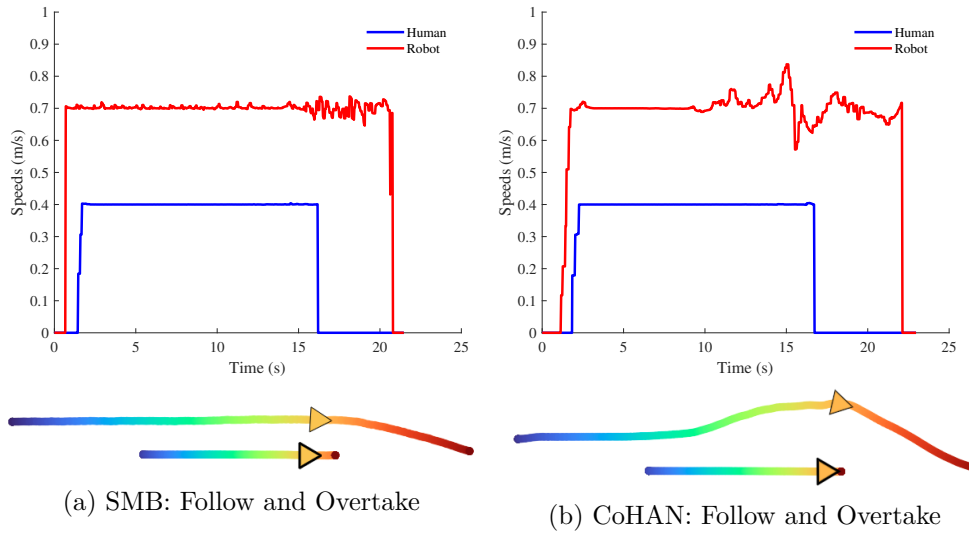


Figure 5.6: Speeds and paths of the human and robot in the Follow and Overtake scenario. The triangle with the thick line is the human, and the other one is the robot. Both the human and robot move from the blue end to the red end of the paths. (a) The human and the robot move parallelly until the very end without much change in their velocity profiles. (b) The robot with CoHAN takes a larger deviation as it plans to enter the FoV of the human. The velocity of the robot drops slightly when it overtakes the human and then it increases slowly.

5.4.3 Scenario 3: Approach

This is a fairly simple setting with a static human. The robot has to approach and face this static human for interaction. The robot starts somewhere behind the human and has to slowly enter his FoV, reducing the surprise as much as possible. As expected, CoHAN has comparatively lower values for all the metrics and in fact, $cost_{danger}$ and $cost_{passby}$ are similar to the previous case (Table 5.1). Although the vision-based metrics are significantly lower for CoHAN than SMB, their values are doubled compared to the *Follow and Overtake* case. This could mean that CoHAN has to improve its handling of static humans. The paths and the velocity profiles

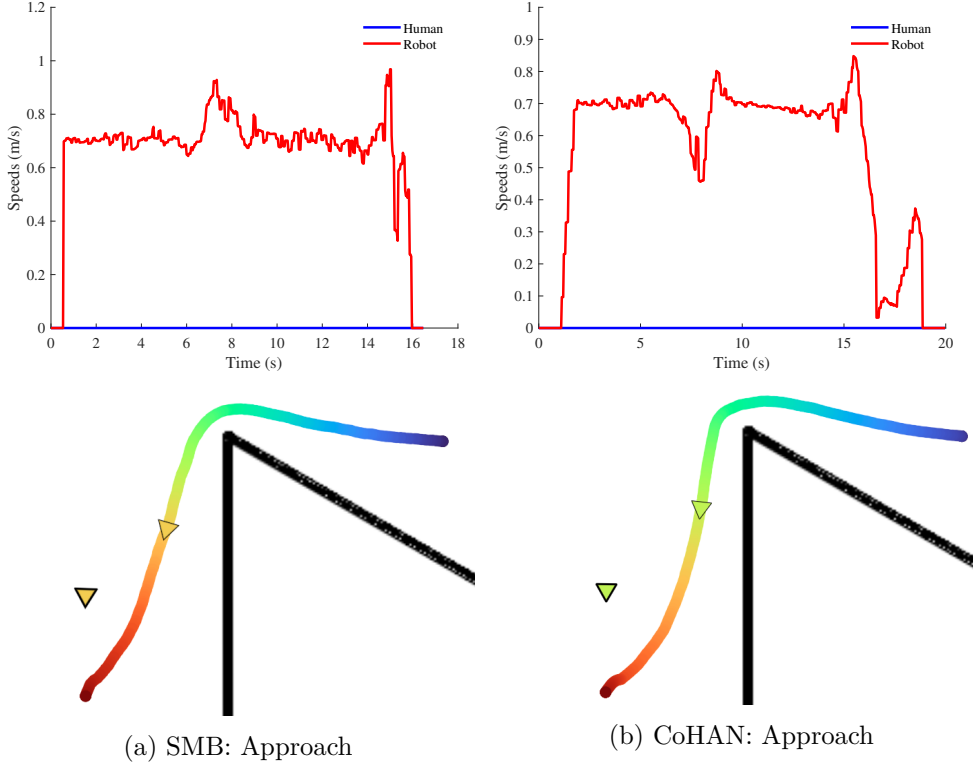


Figure 5.7: Speeds and paths of the human and robot in the Approach scenario. The triangle with the thick line is the human, and the other one is the robot. The robot moves from the blue end to the red end of the paths, and the human is static in the scenario. (a) The robot’s path does not deviate much except at the end. (b) The robot’s path slightly deviates as soon as the robot sees the human, and when the robot turns itself to align in the close vicinity of the human, CoHAN tries to keep the robot’s velocity low.

of the robot are shown in Fig. 5.7, and we can see that CoHAN’s path has slightly deviated, and the robot takes a larger turn. The speed profile in Fig. 5.7 (b) shows that the robot moves with lower velocities while aligning itself in front of the human.

5.4.4 Scenario 4: Appear

When the robot navigates in an environment, there could be several places from where a human can emerge, as discussed in Chapter 4. One such scenario is an L-turning, where a robot can find itself facing a previously unseen human suddenly. It is one of the most difficult cases to handle for a HAN system, as it has to prevent harm and shock to the human on top of avoiding the robot from freezing. As *invisible humans* are already a part of CoHAN, it is expected to perform better in this case.

From Table 5.1, we can see that CoHAN reduces the surprise and has a lesser reaction cost ($cost_{react}$) as well, indicating slightly better performance compared

to SMB. The $cost_{danger}$ also indicates the same in comparison with SMB, but the numerical values are high compared to all the previous cases. As the robot was modulating only its path to handle a sudden emergence, the high velocity explains the higher value. With SMB, it is even higher as it does not even modulate its path. The most important metric to discuss in this scenario is the $cost_{passby}$. Looking at the numerical values, it seems like CoHAN is performing very badly, and SMB has a better performance. It means that the relative velocity between the human and the robot is high in the case of CoHAN.

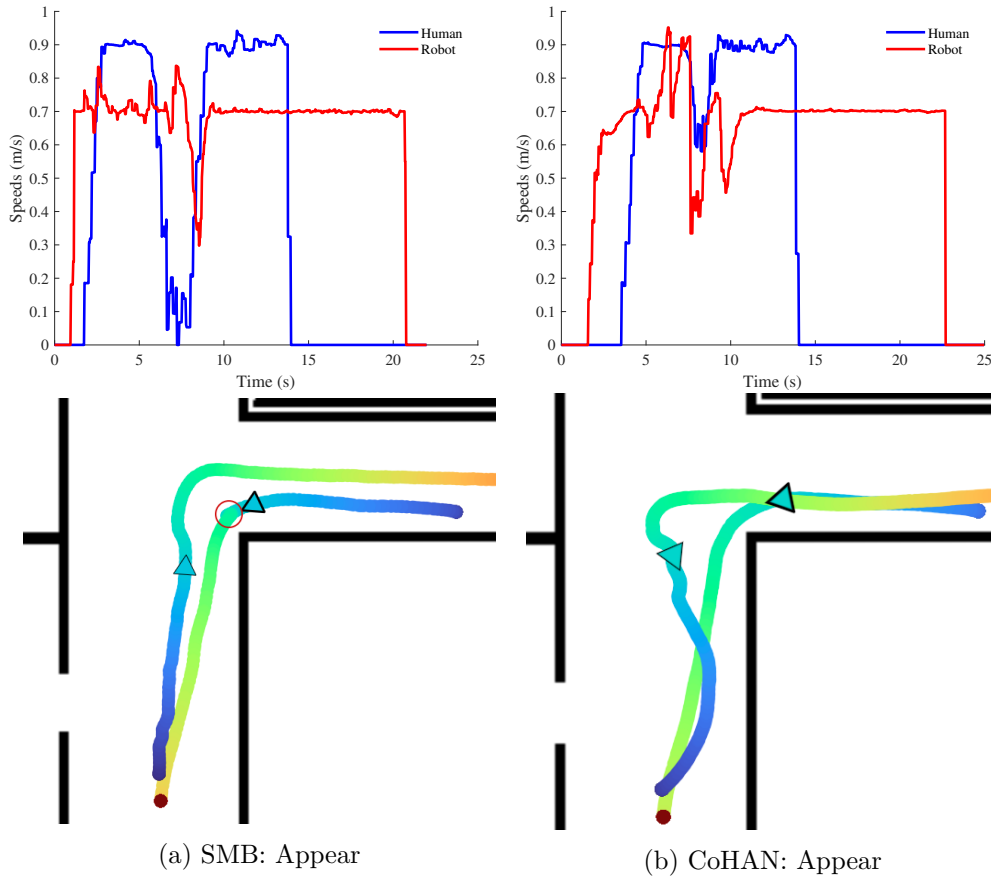


Figure 5.8: Speeds and paths of the human and robot in the Appear scenario. The triangle with the thick line is the human, and the other one is the robot. Both the human and robot move from the blue end to the red end of the paths. (a) The speed of the human drastically changes and momentarily the human even halts as the robot faces him (red circle on the path). (b) The robot moves away from the corner and takes the turn with a larger radius so that the human does not have to slow down much. The robot speeds up a little to move further away as it sees the human and then slows down while it passes by him.

The speeds and paths of the human and the robot for SMB and CoHAN are shown in Fig. 5.8. From the speed plots of SMB in Fig. 5.8 (a), we can see that the human's speed drastically falls between 5 and 10 seconds, and momentarily, it

is even zero. At the same time, the robot's speed seems to be dropping too. This happens as the robot and human face each other, and the robot blocks the human's way (red circle on the path in Fig. 5.8 (a)) before changing its trajectory. This results in a lower relative velocity, and hence, $cost_{passby}$ is small. However, this is not acceptable behaviour and needs to be avoided. The speed plots of CoHAN in Fig. 5.8 (b) also show the speed drops, but the human speed drops only slightly and for a short duration. The speed of the robot, on the other hand, initially increased to drift away from the human and then dropped because of the *Relative Velocity* constraint as it passed by the human. Based on the path and speed profiles from Fig. 5.8 (b), we can say that the human is disturbed less by CoHAN and may call this behaviour somewhat acceptable. In an ideal situation, the human should not be disturbed at all, and the human's speed should be almost constant. We need more studies to identify practical behaviours in such situations.

5.5 Discussion

Evaluation of HAN is not really straightforward, and researchers follow different methodologies to do this. For numerical evaluation, several metrics were proposed based on proxemic zone violations. As most of these discomfort metrics are based only on distance, we believe that they might not do the complete justification in evaluating a complex situation. Therefore, we have proposed a set of discomfort metrics based on velocity and visibility along with the distance. These metrics were applied to evaluate some commonly-occurring human-robot navigation scenarios with the robot running different planners. The evaluations in these settings revealed that the proposed metrics could differentiate a HAN planning system and a classical navigation planner.

Although the idea behind the proposed metrics was to evaluate a HAN system, they might not be sufficient to completely assess a situation. During our analysis, we used the paths and the velocity profiles to provide a complete picture of what was happening. Hence, the proposed metrics should be used together with existing metrics to evaluate a situation better. Furthermore, user studies are required to set thresholds and benchmark acceptable navigation behaviours. It should also be noted that assessing a scenario using only one of the costs may be erroneous. The proposed costs, when studied together, gives better judgement. Hence, one can combine one or more of these costs to formulate a better metric. For example, the $cost_{surprise}$ can be combined with $cost_{passby}$ or $cost_{danger}$ to check how admissible or undesirable the robot's behaviour is in a situation. The exact way to combine these metrics could be tricky, but one can even come up with a single metric by combining all the costs to evaluate HAN. As a part of future work, we plan to explore how the proposed metrics can be combined to build better metrics. We also plan to devise new metrics that are more generic.

Conclusions

In this thesis, we have explored the problem of mobile robot navigation in human environments, generally called human-aware navigation (HAN). We have presented detailed literature on how the field evolved and some existing challenges. Numerous solutions were proposed for HAN in motion planning literature by modelling humans as special obstacles. However, we believe that HAN is essentially an interaction between humans and robots. Therefore, it should satisfy the principles of HRI. In this thesis, we have put together these ideas and proposed a HAN that assesses a situation to take appropriate action. Although situation assessment and behaviour shifting have already been explored in HAN, they were mostly used on top of motion planning systems. Unlike the previous approaches, we have introduced situation assessment at the level of trajectory planning to shift between different modes of planning while the robot navigates to the goal. As this low-level mode shifting was combined with proactive planning, the robot can deal with complex situations in an efficient manner before it is too late. Proactive planning itself has some very interesting features, like quick plan adaptations and early intention shows, but there could be some very complicated situations where proactive planning may not be sufficient. The situation assessment based modality shifting is useful in such places.

We have presented three versions of our HAN system, starting with the idea of introducing modality switching inside the local trajectory planner. There were several improvements over the previous version of HATEB, and these improvements, combined with the mode-switching scheme, are some of the major contributions of this thesis. Qualitative and quantitative results have shown the advantages of such an approach in various settings. We then moved on to propose a Human-Aware Navigation Stack called CoHAN with many changes to scale the system to multiple humans and to address different types of humans. CoHAN can be seen as a complete navigation stack for HAN with different costmap layers, human path prediction mechanisms, and several modes of planning that can solve most of the intricate human-robot navigation scenarios. The early intention show and the Backoff-recovery act as implicit communication mechanisms to tell the human what the robot is going to do. In the future, we plan to integrate more explicit communication through head orientation and probably voice. Various kinds of human states combined with the situation assessment can address more complicated scenarios, and to ease this, we plan to modularize the future version of CoHAN with detailed documentation. CoHAN has been tested on simulated crowds and has already shown some promising results. However, we believe that crowd navigation

could be more complex, and we need more modalities and mechanisms to handle crowds efficiently. We have already presented some ideas we plan to integrate into the future version, and we expect to add more.

Apart from addressing multi-context navigation, we have also focused on one key aspect throughout the development of our HAN system, legibility. We have introduced some new human-aware constraints to make the robot's motion more legible to the human. To show the intention of the robot early, we have proposed TTC, TTCplus and Relative Velocity constraints. The Relative Velocity constraint also addresses the issue of directionality in the crossing or close-to-human navigation scenarios. We have introduced the Visibility constraint to avoid any surprises or shocks to a human when the robot overtakes the human. One more attempt to improve legibility was done through 'invisible humans'. We have introduced this concept into HAN to make sure that the robot is aware not only of the humans it is seeing but also of the environment and the places humans might emerge. We believe that this makes the system more complete and ready to face any kind of environment without behaving erratically or freezing. The 'invisible humans' concept has also made it possible to identify different places on the map through geometric reasoning and introduce different modalities of planning depending on the situation. Further, the algorithm was rigorously tested and showed satisfying performance in some very complex maps. The final version of CoHAN integrated with the 'invisible humans' has been shown to perform better and move smoothly without having any freezing robot problems.

One of the open challenges in HAN is its evaluation, and the currently existing metrics do not do justice while the robot is navigating very dense crowds or confined spaces. As most of them are based on proxemic zones that are variable from place to place and the experiences of the person, it makes it more difficult to generalise the metrics to all cases of HAN. Therefore, we proposed some new metrics of evaluation using velocities and visibility along with distance. The velocity-based metrics, $cost_{danger}$ and $cost_{passby}$, aim to measure the feeling of threat and discomfort caused to the human depending on the direction and speed of the robot's heading. Since these metrics combine distance and velocity, they can explain intricate settings better than proxemic zone violations. The first one of the vision-based metrics, $cost_{visibility}$, was designed to measure how well a robot can approach a human from behind. The other metrics, $cost_{surprise}$ and $cost_{react}$, aim to measure the surprise(or shock) and risk caused when a robot appears in front of a human suddenly. The comparison between a standard navigation planner and a HAN planner showed how these metrics evaluate the 'human awareness' of the system.

Human-aware navigation is not a very simple problem to address, and it taught us some valuable lessons. The developed systems are hard to validate in real-world settings. The experiments could be tedious, not easily reproducible and can be limited. Not every robot is the same. The structure of the robot matters to humans in the environment. Their behaviour towards the robot changes as the shapes and designs of the robot change. Humanoid structures are accepted better than simple mobile bases, even with the same algorithm. One of the important aspects

while developing a HAN system is to have a good simulator that can generate realistic human behaviours. To this date, there is no perfect human simulator. The HAN research community has been using crowd simulators or directly testing with real humans. Crowd simulators are mostly reactive and are not realistic. On the other hand, tests with real humans are tedious and require a lot of resources and time. Although there is rapid growth in this aspect, the community is still waiting for a reliable robot simulator that can simulate rational human agents. We require more user studies to understand the intricacies of human-robot navigation. This is the right moment to invest more time into these studies, as drones and autonomous vehicles have also entered the field. The robotics community seems to have realised this, and now, there are special sessions dedicated to human-aware motion planning and human-aware navigation. The number of studies has already increased, and many researchers in motion planning are coming together to address this complicated yet interesting problem of robotics that can eventually lead to a society where humans and robots coexist.

There are already many immediate future perspectives for HAN. As autonomous vehicles have already hit the roads, it is time to make them behave socially by making them aware of humans and their intentions. An autonomous vehicle not only needs to be safe but also needs to obey certain rules and untold norms followed in society. HAN studies this problem exactly. Drones have become very affordable, and some companies are planning on using them for deliveries, while others are planning to use them as helpers in construction. All these applications require drones to navigate around humans and communicate with them. HAN in drones needs to study these in more detail and come up with better navigation systems for drones. Apart from guiding people in public places, some other applications of HAN lie within warehouses where they need to work or deliver goods to different locations. These environments are more structured, and solving HAN in such places is much easier. HAN has been and will always be used in the sector of service robotics. Be it a robotic companion, a robotic wheelchair or robots in hospitals delivering medicines or equipment the HAN system faces dynamically changing environments where safety is one of the main concerns. A multi-context tunable navigation system could address most of these scenarios by choosing a set of parameters suitable for the setting.

Simulating Human Agents for Testing HAN

One of the challenges during the development of a HAN system is to test the system before its final deployment and real-world runs. Robotic simulators are of great use during this period as we can test the system under various conditions and in several environments. Unlike the classical setting, testing HAN requires the simulation of humans, which is still research under development. Until recently, the HAN community used the crowd simulators like Pedsim or MengeROS [Aroor 2017] to simulate humans in semi-crowded or crowded scenarios. However, the motion generated by these simulators uses reactive schemes like SFM or ORCA, which are good for generating crowds but lack intelligence at the level of an individual human. Recent simulators like SEAN [Tsoi 2020, Tsoi 2022] and SocNavBench [Biswas 2022] tried to generate somewhat intelligent behaviours using new approaches and real-world data. However, these human agents are either not reactive (as they replay real-world trajectories without considering the robot) or use schemes similar to ORCA. Although they have better human agents and environments for testing HAN, they still lack intelligent agents that can be used to test intricate scenarios in spaces like offices, labs or homes. Hence, we have used different ways to control the human(s) while testing the HAN proposed in this thesis. These ways include manual control using a joystick, a simple human controller that follows the generated trajectories, and finally, an intelligent human with rational decision-making capabilities.

A.1 Planning and Control for Human Agents

In this section, we talk about the simple modes and planners employed to control the human avatar in the robot simulator. A human is assumed to be a robot with special requirements.

A.1.1 Manual Control

One of the simplest ways to control a human is to move the human manually using a keyboard or joystick. This is efficient in testing some very complicated scenarios involving intelligent decisions. Since a real human is already controlling the human avatar in the simulator, all the decision-making process is handled by the human operator. To integrate such a human avatar into our system, we used the *joy*¹ ROS

¹<http://wiki.ros.org/joy>

package and then mapped the inputs of the joystick to the avatar’s velocity with a cap at 2 m/s.

Manual control is good for testing some interesting and particular cases, but it becomes tiresome to run several scenarios to benchmark or quantify the results. Moreover, the simulation runs cannot be completely automated as the human operator is always involved in the loop. So, the next solution was to plan and control the human, like the robot. It is different from collision avoidance algorithms as the human has a global path to trace and a separate local planner to move the human.

A.1.2 Planning based Control

To automate and replicate the tasks easily, we have developed a simple *humans navigation*² package using the *global planner*³ in ROS and a simple controller. The developed system has two modes of operation:

1. *Trajectory following*: In this mode, the human follows the trajectory that is provided externally through a ROS topic. We used this mode to test the ideal situations where the human follows the trajectory planned by CoHAN.
2. *Goal-based Control*: This mode is more autonomous as we only provide a goal via a ROS topic, and the system plans and moves the human to the goal. The planning module updates the plan as the human moves, and the simple controller traces the path. This mode was used to run multiple tests to check how well the robot adapts to the human.

Both these modes can control more than one human simultaneously. The *Trajectory Following* mode used the trajectories planned by CoHAN to move the humans. The trajectory provided the desired velocities, which were sent directly to the human controller. However, in the *Goal-based Control* mode, the velocities were calculated based on the current position and planned paths of the humans. To accept multiple goals and plan for all the humans together, a *multi-goal planner* was developed, and it is used internally by the *humans navigation* package to get plans based on the provided goals.

Even though this kind of system solves the issues of automation and is less tiresome to the developer, the human agent is still not intelligent and simply follows the given trajectory. Although the trajectory provided by CoHAN takes care of many human-robot social constraints, this ideal behaviour may not be expected from humans. In the second mode of control, the human agent might have better behaviour, but the agent is still not intelligent and somewhat reactive, like in collision avoidance schemes.

²https://github.com/sphanit/humans_nav

³http://wiki.ros.org/global_planner

A.2 InHuS

InHuS⁴ [Favier 2022] stands for **I**ntelligent **H**uman **S**imulator, and it was developed to specifically simulate a human agent that is rational and persistent about its goal, unlike the reactive schemes.

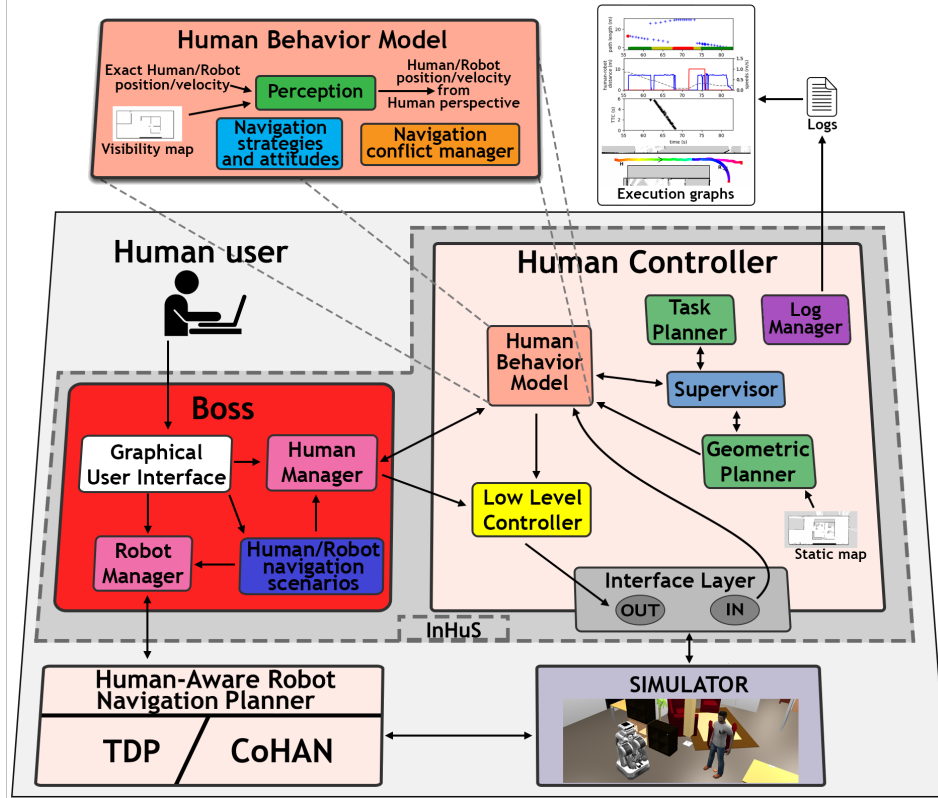


Figure A.1: InHuS Architecture. The system consists of the Boss and the InHuS Human Controller macro components. The human operator interacts with the system through the Boss which in turn interacts with the Human Controller and the HAN Planner. Both the human and the robot controllers are connected to the same simulator where they control different components but share the same environment.

A.2.1 Architecture

The architecture of the proposed system is shown in Fig. A.1. InHuS is composed of two major components namely, 1) *Boss* and 2) *Human Controller*. These components are explained in detail below:

Boss: The *Boss* component is responsible for taking input from the user and sending the appropriate instructions to the human and the robot planners.

⁴https://github.com/AnthonyFavier/InHuS_Social_Navigation

Hence it is provided with a graphical user interface through which a user can give individual goals to each agent, run or re-run predefined scenarios or initiate an endless loop of the human and robot navigating continuously from one goal to another. The endless loop can be used to identify the limits of the HAN system under test. After taking the input from the user, this component communicates the navigation goals to the *Robot Manager* and *Human Manager* at the appropriate times. Once the goals are communicated to these components, the *Boss* does not interfere with the execution unless the human operator selects a new goal or scenario.

Human Controller: This is the main component of InHuS that controls human motion and decides what to do in case of a conflict with the robot. It has an internal '*Human Behaviour Model*' that makes these decisions and sets some attitudes to humans. The other major sub-component is the '*Supervisor*', which supervises the execution of the goal and the progress, and activates the respective components as needed. While the '*Geometric Planner*' provides the geometric path and trajectory to the goal, the '*Low Level Controller*' sub-component sends the command velocity to the human avatar. The '*Task Planner*' was used to define different kinds of navigation tasks like *go to goal*, *wait* etc. Finally, the *Log Manager* of this component logs the data and sends it to a GUI for visualisation of the calculated metrics.

A.2.2 Supervisor and Geometric Planner

The navigation goal of the human avatar is sent to the *Supervisor*, which in turn asks the *Task Planner* for a plan to achieve this goal. The *Supervisor* then supervises and coordinates the execution of all the actions in the plan while managing the conflicts. The navigation plan generally consists of 'moving' and 'waiting' actions. This kind of architecture allows us to define complex navigation goals with multiple steps. The *Supervisor* queries the *Human Behaviour Model* from time to time to detect any potential conflicts. It has the power to suspend the execution of a plan in case of a conflict and resume it whenever necessary. This is especially useful to execute other actions in case of conflict to show the navigation intention and goal persistence of the human agent, unlike the reactive or simple planning approaches.

When the *Supervisor* has to execute a 'moving' action, it sends the navigation goal to the *Geometric Planner*, which generates the path and then calculates the velocity commands to make the human move towards the goal. Depending on the type of trajectory planner selected, human motion can be different. In the current version, the standard ROS Navigation Stack or CoHAN can be selected as the *Geometric Planner*. The velocity command given by this component is not directly sent to the human. The *Low Level Controller* receives this velocity and, if necessary, modifies or perturbs this velocity before sending it to the human avatar. This is used to emulate some reactions while navigation called '*Attitudes*', which is presented in the next subsection.

A.2.3 Human Behaviour Model

The *Human Behaviour Model* is the most important component of the proposed architecture. It controls the human avatar and is responsible for the behaviours exhibited by the human agent. As mentioned previously, this component manages the navigation conflicts, and in this version, only the blockage of the path by the robot is handled. Whenever the *Geometric Planner* is called for the first time, the shortest path to the goal without any moving agents is calculated and sent to the ‘*Navigation Conflict Manager*’ of this component. If any blockage of this path by the robot is detected by this component, it changes the human state, and then the *Supervisor* suspends the goal. The human avatar then performs an ‘approach’ action where it moves towards the goal until it reaches a limiting distance from the robot. After this, it stops and waits for the robot to clear the way. Note that any collision avoidance or simple planning strategies will fail to show such behaviour as they will completely change the path of the human agent instead of showing persistence towards the goal.

This component can also set the goals for the human agent apart from the *Boss*. The internal goal selection mechanism is responsible for different *Attitudes* of the human avatar. Three kinds of attitudes are provided in InHuS:

1. **Stop and Look:** It emulates a curious behaviour, where the human avatar navigating to the goal stops and looks at the robot shortly if the robot is in close vicinity of the human avatar. After this action, the navigation to the original goal is resumed.
2. **Harass:** This attitude emulates a behaviour where the human avatar continuously disturbs the robot by blocking the robot’s path. The idea here is to generate a child-like behaviour for the human avatar.
3. **Random Goal:** In this attitude, a new random goal is set to the human avatar while it is already moving towards a goal, emulating something like a change of mind.

To make the system more realistic, this module builds the perception of the human avatar using the map of the environment and the location of the other agents. The information about the other agents is taken directly from the simulator rather than using simulated cameras or lasers. Therefore, the human avatar does not consider the robot that it cannot see, even if it is below the threshold distance geometrically.

A.2.4 Logs and Metrics

The *Log Manager* logs the data of the human-robot navigation interactions and sends it to GUI based data visualiser. This visualiser shows the different states of the human, the paths of the human and the robot, and some metrics to evaluate the robot’s navigation. The logged data can also be used to calculate new metrics or methodologies for evaluation. A screenshot of this visualisation is shown in Fig.

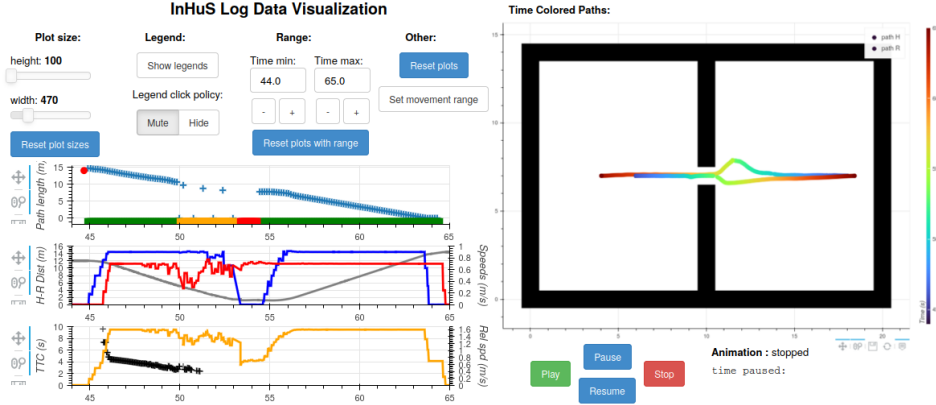


Figure A.2: The data visualisation in GUI. On the right, the paths taken by the agents are shown, while on the left, the human states and the calculated metrics are shown.

A.2. The paths shown on the right in Fig. A.2 are coloured over time. It means that the same colour on the paths represents the same time instant, and using this, one can interpret the behaviour of the agents better. On the left, the plot on the top shows the human avatar's distance to the goal and its estimated state over time. If no conflict occurs, the human stays in a single state, and the distance to the goal decreases linearly. The plots below the first one show some of the calculated metrics and the agents' velocities over time. One can calculate and add more metrics as needed using the logged data.

A.2.5 Generating Different Behaviours

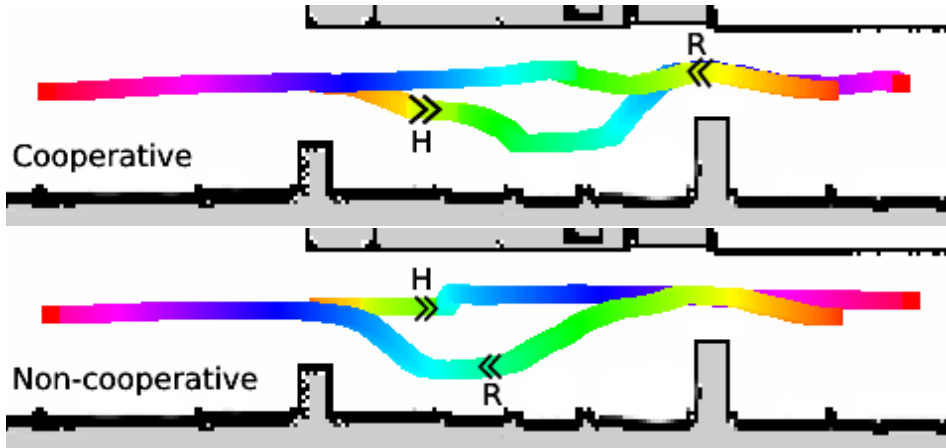


Figure A.3: Traversed paths generated by InHuS during the Pillar corridor scenario. The top part is with cooperative settings and the bottom part with non-cooperative settings along with the *Stop and Look* attitude.

Depending on the *Geometric Planner* and the *Attitude*, different navigation

behaviours can be emulated for the human avatar. For example, using the standard ROS Navigation stack and **Stop and Look** attitude, we can simulate a non-cooperative human who contributes nothing in a setting like a corridor. If CoHAN is used, a cooperative yet curious human can be simulated. Moreover, CoHAN can be tuned to set a degree of cooperative behaviour. The comparisons of different combinations and behaviours generated are presented in more detail in [Favier 2021b]. Fig. A.3 shows the paths of the robot and human in two situations, one where the human is non-cooperative and curious and the other in which the human is completely cooperative.

Effects of Social Constraints

Each of the proposed social constraints in this thesis has some particular effect on the behaviour of the robot. Some of the predominant effects of these are briefly presented here. Each case presents a scenario without and with the social constraint activated. The figures of each scenario show the paths taken by the human and the robot (starting at blue and moving towards red) and their velocities (robot's velocity in red and human's velocity in blue) below. The velocity plot also includes the distance between the human and the robot during the execution of the scenario.

B.1 TTCplus Constraint

B.1.1 Approach

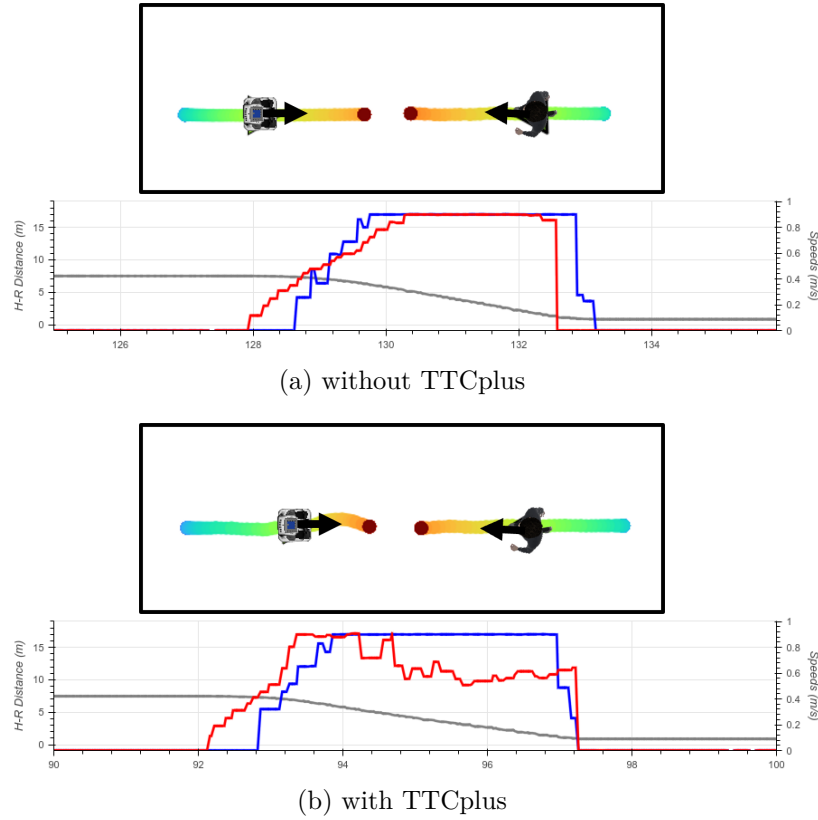


Figure B.1: The robot approaches a human head-on. The addition of the TTCplus constraint makes the robot deviate a little and slow down as it nears the human.

In this scenario, the robot and human move towards each other and stop at a very close distance from each other. From Fig. B.1 (a), it can be seen that the robot and human move at their full speeds towards each. However, with the addition of the TTCplus constraint, the robot has a decreasing velocity profile as the human-robot distance decreases, showing the robot's intention to stop (shown in Fig. B.1 (b)).

B.1.2 Open Space Crossing

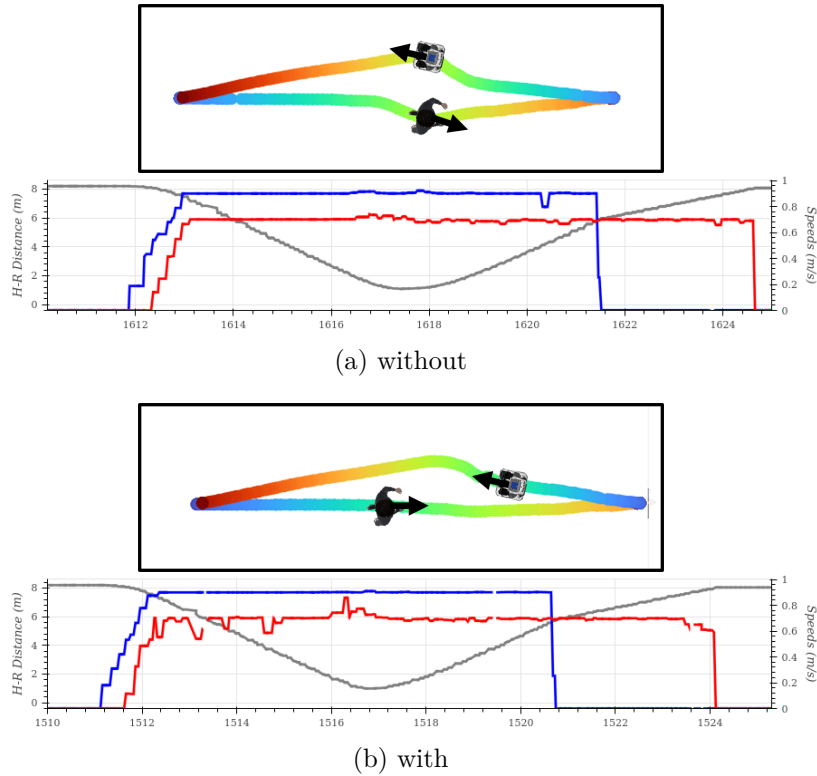


Figure B.2: The human and the robot cross each other in an open space. The addition of the TTCplus constraint makes the robot move aside quickly, showing its intention to give way and the choice of its side to move.

This scenario simulates a robot crossing a human in an open area where there is enough space to move away and not disturb the human. In Fig. B.2 (a), the robot and human move directly towards each other and only avoid each other at the last minute before the collision. This puts on an additional burden on the human to deviate from his path to avoid a collision with the robot. A more human-aware robot should avoid the occurrence of such path deviation, which is similar to what is seen in Fig. B.2 (b). Therefore, the TTCplus constraint not only shows its intention to move away early but also reduces the additional navigational burden that might be imposed on the human.

B.1.3 Corridor Crossing

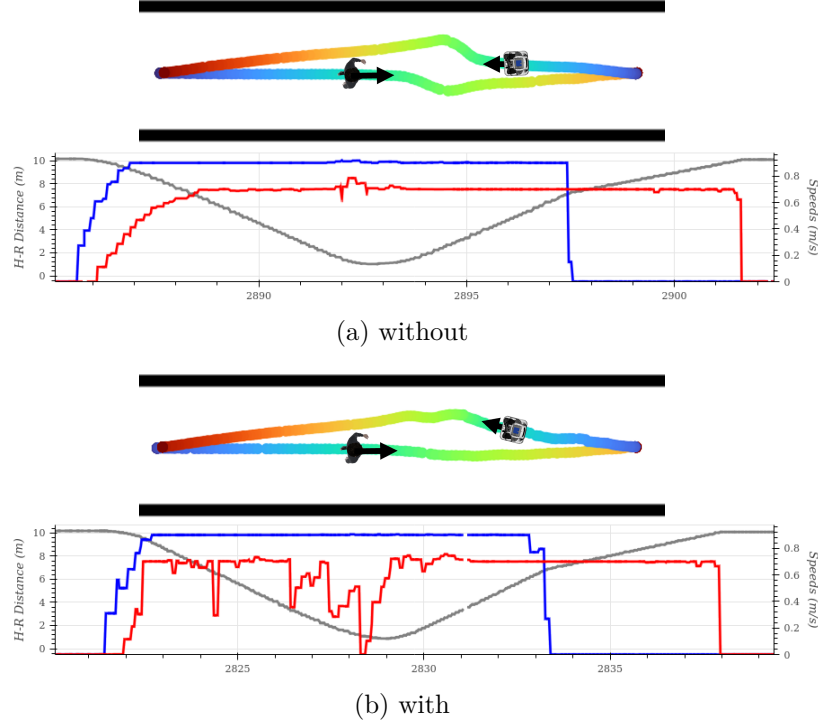


Figure B.3: The human and the robot cross each other in a corridor. TTCplus constraint not only makes the robot take a side early but also slows down the robot as it crosses the human.

This scenario is similar to the previous one, but the space is narrower. As there is not enough space to move away, the robot with the TTCplus constraint moves to a side as well as slows down, as shown in Fig. B.3 (b). Without this constraint, it behaves exactly like in the previous case (Fig. B.3 (a)).

B.2 Relative Velocity Constraint

B.2.1 Corridor Crossing

The corridor crossing scenario is the same as the one shown in the previous section. In the case of the Relative Velocity constraint, the robot should try to move away as quickly as possible and provide more space for the human even when the line of travel is not the same. This can be clearly seen from the path and the velocity profile of the robot in Fig. B.4. The robot starts to move to one side very quickly and slows down as it crosses the human. If there was enough space, the robot could have moved with a larger velocity while crossing. This constraint addresses parallel travel better when compared to TTCplus.

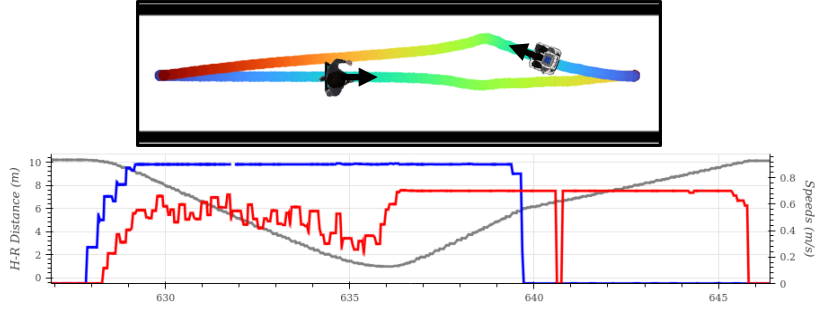


Figure B.4: The human and the robot cross each other in a corridor. Relative Velocity constraint makes the robot clear the way quickly and move with a slower speed robot as it crosses the human at a small parallel distance.

B.2.2 Open Space Crossing

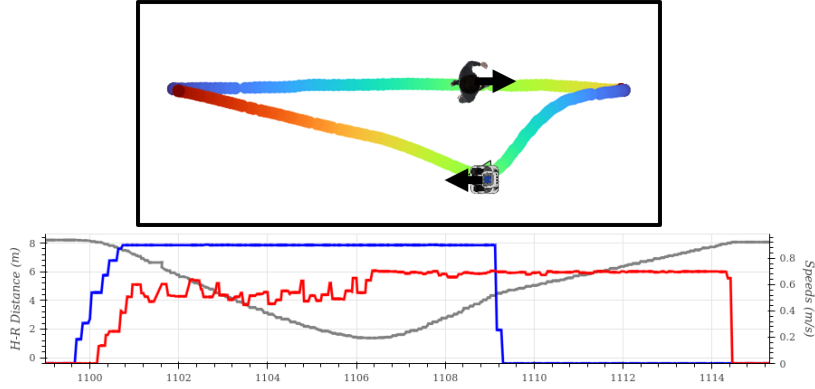


Figure B.5: The human and the robot cross each other in an open space. The addition of the Relative Velocity constraint makes the robot take a large deviation by exploiting the available space while showing its intention to give way. It also facilitates the robot to move at a larger velocity towards the goal.

In this setting, the robot has enough space to move away and then travel with a larger velocity. Without the addition of the Relative Velocity constraint, the robot and human avoid each other moments before the collision, similar to Fig. B.2 (a). This constraint makes the robot move away very quickly and exploit the available space to move at full speed towards its goal. It can be observed from the path and the speed profile of the robot in Fig. B.5. This clearly shows how the Relative Velocity constraint addresses the parallel travel better compared to the TTCplus constraint (see Fig. B.2 (b)).

B.3 Visibility Constraint

To show the advantage of adding the visibility constraint, an overtaking scenario is simulated. In this setting, the robot encounters a human moving very slowly and

partially blocking its way. The robot has to overtake the human in order to move to its goal faster. In the first case, shown in Fig. B.6 (a), without the addition of Visibility constraint, the robot overtakes the human very closely and also disturbs his navigation. With the addition of this constraint, however, the robot takes a large deviation and tries to enter the human's field of view as far as possible without disturbing him. This can be observed from the plots in Fig. B.6 (b).

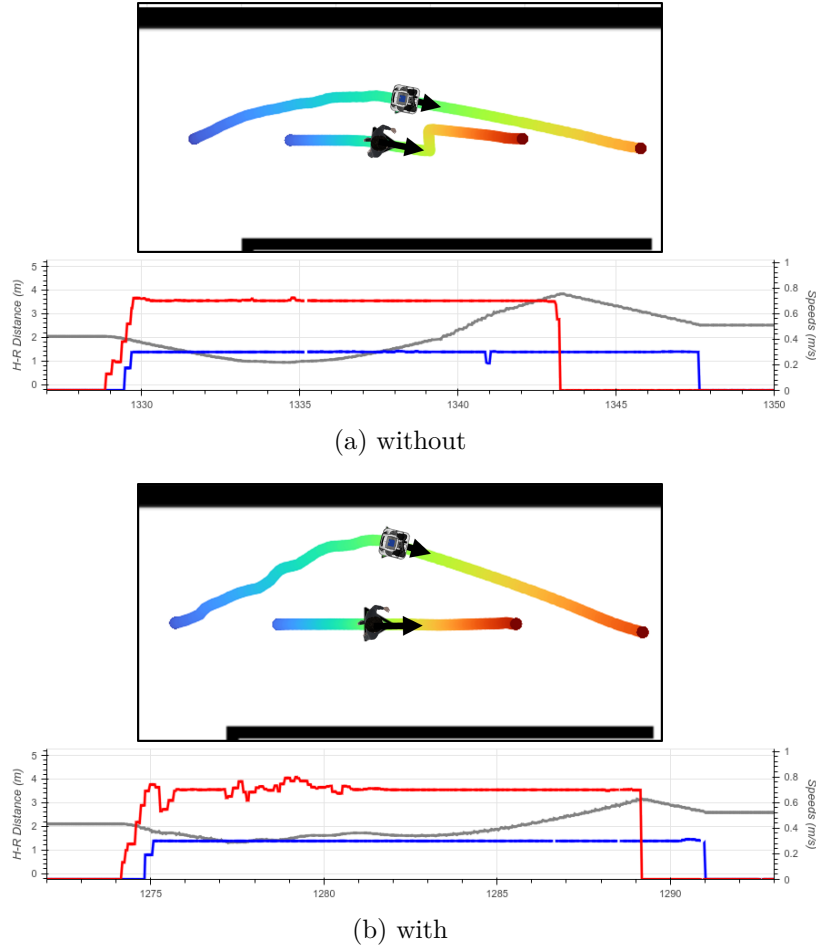


Figure B.6: The robot overtakes a human who is moving very slowly. The addition of the Visibility constraint makes the robot enter the human's field of view slowly without surprising or disturbing the human.

B.4 Updated Invisible Humans Constraint

As shown in chapter 4, the current version of the 'Invisible Humans' constraint already addresses a lot of scenarios to proactively accommodate sudden human appearances. However, the defined formulation had some issues which needed to be addressed using passage detection and mode shifting. After testing the robot navigation in more complicated scenarios, we observed that the current formulation could lead to some deadlocks even after the passage detection. One such deadlock

situation is shown in Fig. B.7. Here, the robot faces opposing forces from the obstacles and the invisible humans and freezes before it can even detect an opening. Therefore, we update our formulation by taking inspiration from the Relative Ve-

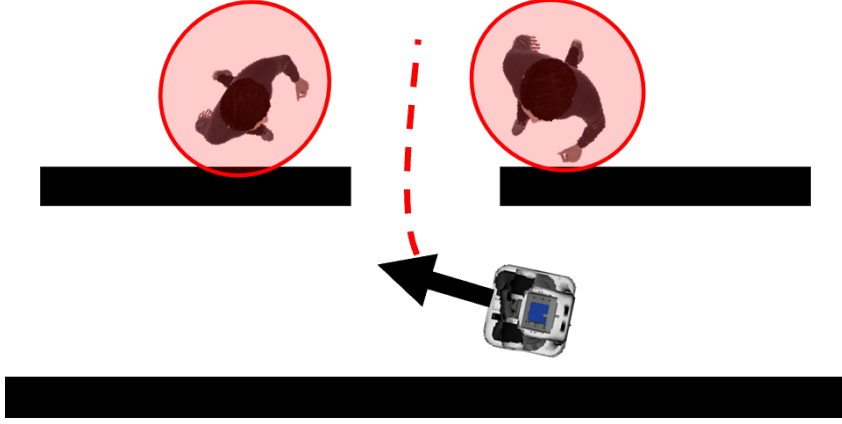


Figure B.7: A situation where the current formulation of the Invisible Humans constraint could fail. The opposing forces from the obstacles and the invisible humans make the robot freeze without moving.

locity constraint. Instead of using the velocity of the visible humans, we use the defined invisible humans' velocity in the formulation and update the formulation as follows:

$$\begin{aligned} cost_{inv_human} &= \max \left(\frac{V - a\Delta t_n + \|\vec{V}_r\| + 1}{d}, 0 \right) \quad \text{if } \Delta t_n > 0.5s \\ &= \frac{V}{d} \quad \text{otherwise} \end{aligned} \quad (\text{B.1})$$

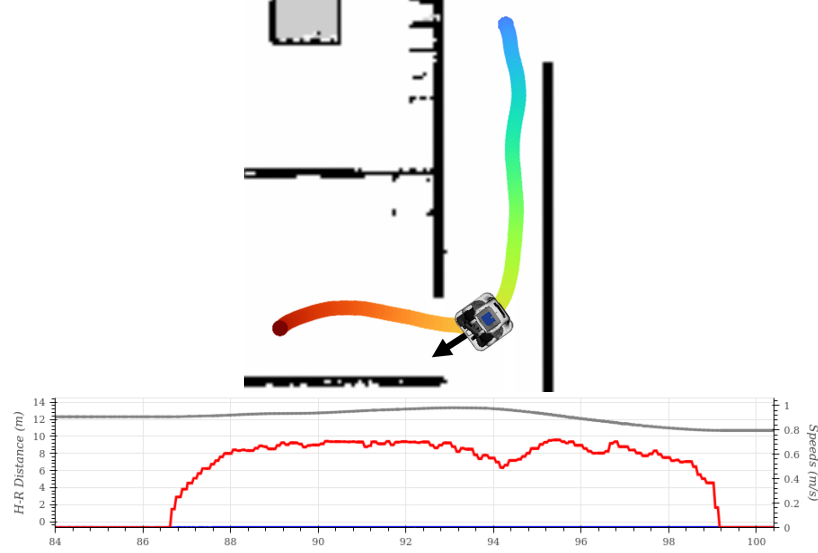
In the latest version of CoHAN, the above formulation is used instead of the previous one. The rigorous testing of the updated formulation is still pending, but we already see some improvements over the previous one. An example of the constrained door crossing is presented below.

B.4.1 Testing the Updated Constraint

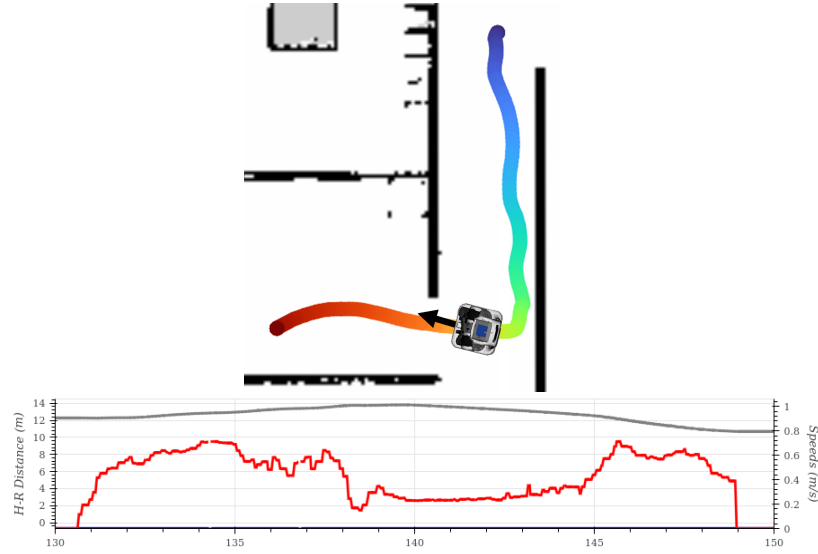
The above formulation acts on the robot's velocity during the possible freezing scenarios and makes the robot move with lower velocities, and reduces the cost. Since it is an updated formulation of the Invisible Humans constraint, it should still hold the properties of the previous formulation. To show this, we have simulated the door-crossing scenario again with a wall on the side that limits the space.

In Fig. B.8 (a), the robot moves without considering and accounting for the invisible humans in the environment. Therefore, it moves at almost full speed and takes the shortest path to reach the goal. With the formulation in chapter 4, the robot froze between the wall and the entry to the door and did not move. However, with the updated formulation, the robot moves away from the door as

much as possible without colliding with the wall and also aligns itself to properly pass through the door. Further, while crossing the door, the robot moves very cautiously with a slower velocity, as seen in Fig. B.8 (b) between 140 – 145s.



(a) Without Invisible Humans constraint



(b) with the updated Invisible Humans constraint

Figure B.8: Constrained door crossing scenario. The robot has to enter a door in a narrow space and try to accommodate humans as much as possible. The updated Invisible Humans constraint makes the robot move close to the wall before aligning itself towards the door and carefully entering it.

Bibliography

- [Adorni 2001] Giovanni Adorni, Stefano Cagnoni, Stefan Enderle, Gerhard K Kraetzschmar, Monica Mordonini, Michael Plagge, Marcus Ritter, Stefan Sablatnög and Andreas Zell. *Vision-based localization for mobile robots*. Robotics and Autonomous Systems, vol. 36, no. 2-3, pages 103–119, 2001. (Cited in page 11.)
- [Alahi 2016] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei and Silvio Savarese. *Social lstm: Human trajectory prediction in crowded spaces*. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 961–971, 2016. (Cited in pages 27 and 103.)
- [Alami 2000] R. Alami, I. Belousov, S. Fleury, M. Herrb, F. Ingrand, J. Minguez and B. Morisset. *Diligent: towards a human-friendly navigation system*. In Proceedings. 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000) (Cat. No.00CH37113), volume 1, pages 21–26, Takamatsu, Japan, 2000. IEEE. (Cited in pages 22 and 24.)
- [Alami 2007] Rachid Alami, K. Madhava Krishna and Thierry Simeon. *Provably Safe Motions Strategies for Mobile Robots in Dynamic Domains*. In Autonomous Navigation in Dynamic Environments, volume 35 of *Springer Tracts in Advanced Robotics*, pages 85–106. Springer Berlin Heidelberg, 2007. (Cited in page 81.)
- [Althaus 2004] Philipp Althaus, Hiroshi Ishiguro, Takayuki Kanda, Takahiro Miyashita and Henrik I Christensen. *Navigation for human-robot interaction tasks*. In IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004, volume 2, pages 1894–1900. IEEE, 2004. (Cited in page 15.)
- [Araujo 2015] Arthur R Araujo, Daniel D Caminhas and Guilherme AS Pereira. *An architecture for navigation of service robots in human-populated office-like environments*. IFAC-PapersOnLine, vol. 48, no. 19, pages 189–194, 2015. (Cited in page 20.)
- [Ardakani 2015] M Mahdi Ghazaei Ardakani, Björn Olofsson, Anders Robertsson and Rolf Johansson. *Real-time trajectory generation using model predictive control*. In 2015 IEEE International Conference on Automation Science and Engineering (CASE), pages 942–948. IEEE, 2015. (Cited in page 10.)
- [Aroor 2017] Anoop Aroor, Susan L Esptein and Raj Korpan. *Mengeros: A crowd simulation tool for autonomous robot navigation*. In 2017 AAAI Fall Symposium Series, 2017. (Cited in page 121.)

- [Arras 2003] Kai O Arras, Nicola Tomatis and Roland Siegwart. *Robox, a remarkable mobile robot for the real world*. In Experimental robotics VIII, pages 178–187. Springer, 2003. (Cited in page 15.)
- [Avrunin 2014] Eleanor Avrunin and Reid Simmons. *Socially-appropriate approach paths using human data*. In The 23rd IEEE International Symposium on Robot and Human Interactive Communication, pages 1037–1042. IEEE, 2014. (Cited in page 108.)
- [Banisetty 2018] Santosh Balajee Banisetty and David Feil-Seifer. *Towards a unified planner for socially-aware navigation*. arXiv preprint arXiv:1810.00966, 2018. (Cited in page 20.)
- [Banisetty 2019] Santosh Balajee Banisetty, Scott Forer et al. *Socially-aware navigation: A non-linear multi-objective optimization approach*. arXiv preprint arXiv:1911.04037, 2019. (Cited in page 59.)
- [Banisetty 2020] Santosh Balajee Banisetty, Vineeth Rajamohan et al. *A deep learning approach to multi-context socially-aware navigation*. In IEEE/RSJ IROS, 2020. (Cited in pages 20, 28, 59, 78, and 79.)
- [Banisetty 2021] Santosh Balajee Banisetty and Tom Williams. *Implicit communication through social distancing: Can social navigation communicate social norms?* In Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction, pages 499–504, 2021. (Cited in page 105.)
- [Barchard 2018] Kimberly A Barchard, Leiszle Lapping-Carr, R Shane Westfall, Santosh Balajee Banisetty and David Feil-Seifer. *Perceived Social Intelligence (PSI) Scales Test Manual (August, 2018)*. 2018. (Cited in page 105.)
- [Barchard 2020] Kimberly A Barchard, Leiszle Lapping-Carr, R Shane Westfall, Andrea Fink-Armold, Santosh Balajee Banisetty and David Feil-Seifer. *Measuring the perceived social intelligence of robots*. ACM Transactions on Human-Robot Interaction (THRI), vol. 9, no. 4, pages 1–29, 2020. (Cited in page 105.)
- [Bartneck 2008] Christoph Bartneck, Elizabeth Croft and Dana Kulic. *Measuring the anthropomorphism, animacy, likeability, perceived intelligence and perceived safety of robots*. 2008. (Cited in page 104.)
- [Becerra 2020] Israel Becerra, Markku Suomalainen, Eliezer Lozano, Katherine J Mimnaugh, Rafael Murrieta-Cid and Steven M LaValle. *Human perception-optimized planning for comfortable vr-based telepresence*. IEEE Robotics and Automation Letters, vol. 5, no. 4, pages 6489–6496, 2020. (Cited in page 20.)

- [Bennewitz 2004] Maren Bennewitz. *Mobile robot navigation in dynamic environments*. PhD thesis, Freiburg (Breisgau), Univ., Diss., 2004, 2004. (Cited in page 15.)
- [Bennewitz 2005] Maren Bennewitz, Wolfram Burgard, Grzegorz Cielniak and Sebastian Thrun. *Learning motion patterns of people for compliant robot motion*. The International Journal of Robotics Research, vol. 24, no. 1, pages 31–48, 2005. (Cited in page 15.)
- [Berg 2011] Jur van den Berg, Stephen J Guy, Ming Lin and Dinesh Manocha. *Reciprocal n-body collision avoidance*. In Robotics research, pages 3–19. Springer, 2011. (Cited in pages 13 and 38.)
- [Biswas 2022] Abhijat Biswas, Allan Wang, Gustavo Silvera, Aaron Steinfeld and Henny Admoni. *Socnavbench: A grounded simulation testing framework for evaluating social navigation*. ACM Transactions on Human-Robot Interaction (THRI), vol. 11, no. 3, pages 1–24, 2022. (Cited in pages 22, 104, and 121.)
- [Bogue 2016] Robert Bogue. *Growth in e-commerce boosts innovation in the warehouse robot market*. Industrial Robot: An International Journal, 2016. (Cited in page 1.)
- [Bordallo 2015] Alejandro Bordallo, Fabio Previtali, Nantas Nardelli and Subramanian Ramamoorthy. *Counterfactual Reasoning about Intent for Interactive Navigation in Dynamic Environments*. In Int. Conf. on Intelligent Robots and Systems, pages 2943–2950, 2015. (Cited in pages 39 and 59.)
- [Bouraine 2012] Sara Bouraine, Thierry Fraichard and Hassen Salhi. *Provably safe navigation for mobile robots with limited field-of-views in unknown dynamic environments*. In 2012 IEEE International Conference on Robotics and Automation, pages 174–179, 2012. (Cited in page 82.)
- [Boy 2002] Eng Seng Boy, Chee Leong Teo and Etienne Burdet. *Collaborative wheelchair assistant*. In IEEE/RSJ International Conference on Intelligent Robots and Systems, volume 2, pages 1511–1516. IEEE, 2002. (Cited in page 15.)
- [Bretto 2013] Alain Bretto. *Hypergraph theory. An introduction*. Mathematical Engineering. Cham: Springer, 2013. (Cited in page 30.)
- [Bruckschen 2020] Lilli Bruckschen, Kira Bungert, Nils Dengler and Maren Bennewitz. *Human-aware robot navigation by long-term movement prediction*. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 11032–11037. IEEE, 2020. (Cited in page 20.)
- [Buechegger 2018] Klaus Buechegger, George Todoran and Markus Bader. *Safe and Efficient Autonomous Navigation in the Presence of Humans at Control*

- Level*. In International Conference on Robotics in Alpe-Adria Danube Region, pages 504–511. Springer, 2018. (Cited in page 19.)
- [Burgard 1999a] Wolfram Burgard, Armin B Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner and Sebastian Thrun. *Experiences with an interactive museum tour-guide robot*. Artificial intelligence, vol. 114, no. 1-2, pages 3–55, 1999. (Cited in page 14.)
- [Burgard 1999b] Wolfram Burgard, Armin B Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner and Sebastian Thrun. *The museum tour-guide robot RHINO*. In Autonome Mobile Systeme 1998, pages 245–254. Springer, 1999. (Cited in page 14.)
- [Butler 2001] John Travis Butler and Arvin Agah. *Psychological effects of behavior patterns of a mobile personal robot*. Autonomous Robots, vol. 10, no. 2, pages 185–202, 2001. (Cited in pages 16 and 104.)
- [Cai 2020] Kuanqi Cai, Chaoqun Wang, Jiyu Cheng, Clarence W De Silva and Max Q-H Meng. *Mobile robot path planning in dynamic environments: A survey*. arXiv preprint arXiv:2006.14195, 2020. (Cited in page 12.)
- [Carmona 2019] Manuel Fernandez Carmona, Tejas Parekhet *al.* *Making the case for human-aware navigation in warehouses*. In Annual Conference Towards Autonomous Robotic Systems. Springer, 2019. (Cited in pages 58 and 59.)
- [Chen 2017a] Yu Fan Chen, Michael Everett, Miao Liu and Jonathan P How. *Socially aware motion planning with deep reinforcement learning*. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1343–1350. IEEE, 2017. (Cited in pages 13, 18, and 59.)
- [Chen 2017b] Yu Fan Chen, Miao Liu, Michael Everett and Jonathan P How. *Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning*. In 2017 IEEE international conference on robotics and automation (ICRA), pages 285–292. IEEE, 2017. (Cited in page 13.)
- [Chen 2019] Changan Chen, Yuejiang Liu, Sven Kreiss and Alexandre Alahi. *Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning*. In 2019 International Conference on Robotics and Automation (ICRA), pages 6015–6022. IEEE, 2019. (Cited in page 18.)
- [Chen 2020] Changan Chen, Sha Hu, Payam Nikdel, Greg Mori and Manolis Savva. *Relational graph learning for crowd navigation*. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 10007–10013. IEEE, 2020. (Cited in pages 18 and 19.)
- [Chung 2009] Woojin Chung, Seokgyu Kim, Minki Choi, Jaesik Choi, Hoyeon Kim, Chang-bae Moon and Jae-Bok Song. *Safe navigation of a mobile robot considering visibility of environment*. IEEE Transactions on Industrial Electronics, vol. 56, no. 10, pages 3941–3950, 2009. (Cited in page 82.)

- [Chung 2010] Shu-Yun Chung and Han-Pang Huang. *A mobile robot that understands pedestrian spatial behaviors*. In 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 5861–5866. IEEE, 2010. (Cited in page 18.)
- [Ciou 2018] Pei-Huai Ciou, Yu-Ting Hsiao, Zong-Ze Wu, Shih-Huan Tseng and Li-Chen Fu. *Composite reinforcement learning for social robot navigation*. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 2553–2558. IEEE, 2018. (Cited in page 13.)
- [Clodic 2006] Aurelie Clodic, Sara Fleury, Rachid Alami, Raja Chatila, Gerard Bailly, Ludovic Brethes, Maxime Cottret, Patrick Danes, Xavier Dollat, Frederic Elisei, Isabelle Ferrane, Matthieu Herrb, Guillaume Infantes, Christian Lemaire, Frederic Lerasle, Jerome Manhes, Patrick Marcoul, Paulo Menezes and Vincent Montreuil. *Rackham: An Interactive Robot-Guide*. In ROMAN 2006 - The 15th IEEE International Symposium on Robot and Human Interactive Communication, pages 502–509, Univ. of Hertfordshire, Hatfield, UK, September 2006. IEEE. (Cited in page 23.)
- [Cunningham 2015] Alexander G Cunningham, Enric Galceran, Ryan M Eustice and Edwin Olson. *MPDM: Multipolicy decision-making in dynamic, uncertain environments for autonomous driving*. In 2015 IEEE International Conference on Robotics and Automation (ICRA), pages 1670–1677. IEEE, 2015. (Cited in page 19.)
- [Curioni 2019] Arianna Curioni, Gunther Knoblich, Natalie Sebanz, A Goswami and P Vadakkepat. *Joint action in humans: A model for human-robot interactions*. Humanoid Robotics: A Reference, pages 2149–2167, 2019. (Cited in pages 1, 25, and 33.)
- [Dautenhahn 2006] K. Dautenhahn, M. Walters, S. Woods, K. L. Koay, C. L. Nehaniv, A. Sisbot, R. Alami and T. Siméon. *How may I serve you?: a robot companion approaching a seated person in a helping context*. In Proceeding of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction - HRI '06, page 172, Salt Lake City, Utah, USA, 2006. ACM Press. (Cited in pages 22 and 23.)
- [Dehais 2011] Frédéric Dehais, Emrah Akin Sisbot, Rachid Alami and Mickaël Causse. *Physiological and subjective evaluation of a human-robot object hand-over task*. Applied Ergonomics, vol. 42, no. 6, pages 785–791, November 2011. (Cited in page 23.)
- [Demeester 2008] Eric Demeester, Alexander Hüntemann, Dirk Vanhooydonck, Gerolf Vanacker, Hendrik Van Brussel and Marnix Nuttin. *User-adapted plan recognition and user-adapted shared control: A Bayesian approach to semi-autonomous wheelchair driving*. Autonomous Robots, vol. 24, no. 2, pages 193–211, 2008. (Cited in page 17.)

- [Dondrup 2016] Christian Dondrup and Marc Hanheide. *Qualitative constraints for human-aware robot navigation using velocity costmaps*. In 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pages 586–592. IEEE, 2016. (Cited in page 20.)
- [Dugas 2020] Daniel Dugas, Juan Nieto, Roland Siegwart and Jen Jen Chung. *Ian: Multi-behavior navigation planning for robots in real, crowded environments*. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 11368–11375. IEEE, 2020. (Cited in pages 19 and 28.)
- [Dyannikov 2007] Ivan Dynnikov and Bert Wiest. *On the complexity of braids*. Journal of the European Mathematical Society, vol. 9, no. 4, pages 801–840, 2007. (Cited in page 103.)
- [Echeverria 2011] Gilberto Echeverria, Nicolas Lassabe, Arnaud Degroote and Séverin Lemaignan. *Modular open robots simulation engine: Morse*. In 2011 IEEE International Conference on Robotics and Automation, pages 46–51. IEEE, 2011. (Cited in pages 48, 69, and 89.)
- [Eckner 2012] James T Eckner, James K Richardson, Hogene Kim, David B Lipps and James A Ashton-Miller. *A novel clinical test of recognition reaction time in healthy adults*. Psychological assessment, vol. 24, no. 1, page 249, 2012. (Cited in page 110.)
- [Escobedo 2012] Arturo Escobedo, Jorge Rios-Martinez, Anne Spalanzani and Christian Laugier. *Context-based face control of a robotic wheelchair*. 2012. (Cited in page 20.)
- [Evens 2022] Brecht Evens, Mathijs Schuurmans and Panagiotis Patrinos. *Learning MPC for Interaction-Aware Autonomous Driving: A Game-Theoretic Approach*. In 2022 European Control Conference (ECC), pages 34–39. IEEE, 2022. (Cited in page 22.)
- [Fahad 2020] Muhammad Fahad, Guang Yang and Yi Guo. *Learning Human Navigation Behavior Using Measured Human Trajectories in Crowded Spaces*. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 11154–11160. IEEE, 2020. (Cited in page 20.)
- [Favier 2021a] Anthony Favier, Phani Singamaneni and Rachid Alami. *An Intelligent Human Simulation (InHuS) for developing and experimenting human-aware and interactive robot abilities*. 2021. (Cited in pages 71, 89, and 111.)
- [Favier 2021b] Anthony Favier, Phani-Teja Singamaneni and Rachid Alami. *Simulating Intelligent Human Agents for Intricate Social Robot Navigation*. 2021. (Cited in page 127.)

- [Favier 2022] Anthony Favier, Phani Teja Singamaneni and Rachid Alami. *An Intelligent Human Avatar to Debug and Challenge Human-Aware Robot Navigation Systems*. In Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction. IEEE Press, 2022. (Cited in pages 69 and 123.)
- [Feil-Seifer 2011] David Feil-Seifer and Maja Matarić. *People-aware navigation for goal-oriented behavior involving a human partner*. In 2011 IEEE International Conference on Development and Learning (ICDL), volume 2, pages 1–6. IEEE, 2011. (Cited in pages 16 and 28.)
- [Fernandez Carmona 2019] Manuel Fernandez Carmona, Tejas Parekh and Marc Hanheide. *Making the case for human-aware navigation in warehouses*. In Annual Conference Towards Autonomous Robotic Systems, pages 449–453. Springer, 2019. (Cited in page 20.)
- [Ferrer 2013a] Gonzalo Ferrer, Anais Garrell and Alberto Sanfeliu. *Robot companion: A social-force based approach with human awareness-navigation in crowded environments*. In 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1688–1694. IEEE, 2013. (Cited in pages 18 and 38.)
- [Ferrer 2013b] Gonzalo Ferrer, Anais Garrell and Alberto Sanfeliu. *Social-aware robot navigation in urban environments*. In 2013 European Conference on Mobile Robots, pages 331–336. IEEE, 2013. (Cited in pages 1, 2, 58, and 59.)
- [Ferrer 2014] Gonzalo Ferrer and Alberto Sanfeliu. *Bayesian human motion intentionality prediction in urban environments*. Pattern Recognition Letters, vol. 44, pages 134–140, 2014. (Cited in pages 48, 59, and 63.)
- [Ferrer 2015] Gonzalo Ferrer and Alberto Sanfeliu. *Multi-objective cost-to-go functions on robot navigation in dynamic environments*. In IEEE/RSJ IROS, 2015. (Cited in pages 39 and 59.)
- [Ferrer 2017] Gonzalo Ferrer, Anaís Garrell Zulueta, Fernando Herrero Cotarelo and Alberto Sanfeliu. *Robot social-aware navigation framework to accompany people walking side-by-side*. Autonomous robots, vol. 41, no. 4, pages 775–793, 2017. (Cited in pages 18 and 104.)
- [Fisac 2018] Jaime F Fisac, Andrea Bajcsy, Sylvia L Herbert, David Fridovich-Keil, Steven Wang, Claire J Tomlin and Anca D Dragan. *Probabilistically safe robot planning with confidence-based human predictions*. arXiv preprint arXiv:1806.00109, 2018. (Cited in pages 39 and 59.)
- [Forer 2018] Scott Forer, Santosh Balajee Banisetty, Logan Yliniemi, Monica Nicolescu and David Feil-Seifer. *Socially-aware navigation using non-linear multi-objective optimization*. In 2018 IEEE/RSJ International Conference

- on Intelligent Robots and Systems (IROS), pages 1–9. IEEE, 2018. (Cited in page 28.)
- [Foster 2016] Mary Ellen Foster, Rachid Alami, Olli Gestranus, Oliver Lemon, Marketta Niemelä, Jean-Marc Odobez and Amit Kumar Pandey. *The MuM-MER project: Engaging human-robot interaction in real-world public spaces*. In International Conference on Social Robotics, pages 753–763. Springer, 2016. (Cited in page 24.)
- [Foster 2019] Mary Ellen Foster, Bart Craenen *et al.* *Mummer: Socially intelligent human-robot interaction in public spaces*. arXiv preprint arXiv:1909.06749, 2019. (Cited in page 58.)
- [Fox 1997] Dieter Fox, Wolfram Burgard and Sebastian Thrun. *The dynamic window approach to collision avoidance*. IEEE Robotics & Automation Magazine, vol. 4, no. 1, pages 23–33, 1997. (Cited in page 12.)
- [Fox 2001] Dieter Fox, Sebastian Thrun, Wolfram Burgard and Frank Dellaert. *Particle filters for mobile robot localization*. In Sequential Monte Carlo methods in practice, pages 401–428. Springer, 2001. (Cited in page 11.)
- [Fox 2003] Dieter Fox. *Adapting the sample size in particle filters through KLD-sampling*. The international Journal of robotics research, vol. 22, no. 12, pages 985–1003, 2003. (Cited in page 11.)
- [Galindo 2006] Cipriano Galindo, Javier Gonzalez and J-A Fernandez-Madrigal. *Control architecture for human-robot integration: application to a robotic wheelchair*. IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), vol. 36, no. 5, pages 1053–1067, 2006. (Cited in page 15.)
- [Gao 2021] Yuxiang Gao and Chien-Ming Huang. *Evaluation of socially-aware robot navigation*. Frontiers in Robotics and AI, page 420, 2021. (Cited in page 102.)
- [Garrell 2010] Anaís Garrell and Alberto Sanfeliu. *Local optimization of cooperative robot movements for guiding and regrouping people in a guiding mission*. In 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 3294–3299. IEEE, 2010. (Cited in page 18.)
- [Garrell 2019] Anaís Garrell, Carles Coll, René Alquézar and Alberto Sanfeliu. *Teaching a drone to accompany a person from demonstrations using non-linear asfm*. In 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1985–1991. IEEE, 2019. (Cited in page 22.)
- [Garrido 2006] Santiago Garrido, Luis Moreno, Mohamed Abderrahim and Fernando Martin. *Path planning for mobile robot navigation using voronoi diagram and fast marching*. In 2006 IEEE/RSJ International Conference on

- Intelligent Robots and Systems, pages 2376–2381. IEEE, 2006. (Cited in page 9.)
- [Garza Elizondo 2016] Luis Alberto Garza Elizondo. Social-aware drone navigation using social force model. Master’s thesis, Universitat Politècnica de Catalunya, 2016. (Cited in page 22.)
- [Gasparetto 2015] Alessandro Gasparetto, Paolo Boscariol, Albano Lanzutti and Renato Vidoni. *Path planning and trajectory planning algorithms: A general overview*. Motion and operation planning of robotic systems, pages 3–27, 2015. (Cited in page 10.)
- [Gharbi 2013] Mamoun Gharbi, Severin Lemaignan, Jim Mainprice and Rachid Alami. *Natural interaction for object hand-over*. In 2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pages 401–401, Tokyo, Japan, March 2013. IEEE. (Cited in page 23.)
- [Gharbi 2015] Mamoun Gharbi, Pierre-Vincent Paubel, Aurelie Clodic, Ophelie Carreras, Rachid Alami and Jean-Marie Cellier. *Toward a better understanding of the communication cues involved in a human-robot object transfer*. In 2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pages 319–324, Kobe, Japan, August 2015. IEEE. (Cited in page 24.)
- [Gockley 2007] Rachel Gockley, Jodi Forlizzi and Reid Simmons. *Natural person-following behavior for social robots*. In Proceedings of the ACM/IEEE international conference on Human-robot interaction, pages 17–24, 2007. (Cited in page 16.)
- [Gonzalez 2006a] J Gonzalez, AJ Muaeoz, C Galindo, JA Fernandez-Madrigal and JL Blanco. *A description of the SENA robotic wheelchair*. In MELECON 2006-2006 IEEE Mediterranean Electrotechnical Conference, pages 437–440. IEEE, 2006. (Cited in page 15.)
- [Gonzalez 2006b] Javier Gonzalez, Cipriano Galindo, Jose-Luis Blanco, AJ Muñoz, Vicente Arevalo and Juan-Antonio Fernandez-Madrigal. *The SENA robotic wheelchair project*. DAAAM International Scientific Book, pages 235–251, 2006. (Cited in page 15.)
- [Govea 2010] Alejandro Dizan Vasquez Govea. *Growing hidden markov models*. In Incremental Learning for Motion Prediction of Pedestrians and Vehicles, pages 71–82. Springer, 2010. (Cited in page 20.)
- [Granata 2012] Consuelo Granata and Philippe Bidaud. *A framework for the design of person following behaviors for social mobile robots*. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 4652–4659. IEEE, 2012. (Cited in pages 16 and 28.)

- [Grisetti 2011] Giorgio Grisetti, Rainer Kümmerle, Hauke Strasdat and Kurt Konolige. *g2o: A general framework for (hyper) graph optimization*. In Proceedings of the IEEE international conference on robotics and automation (ICRA), Shanghai, China, pages 9–13, 2011. (Cited in page 31.)
- [Gross 2009] H-M Gross, H Boehme, Ch Schroeter, Steffen Müller, Alexander König, Erik Einhorn, Ch Martin, Matthias Merten and Andreas Bley. *TOOMAS: interactive shopping guide robots in everyday use-final implementation and experiences from long-term field trials*. In 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 2005–2012. IEEE, 2009. (Cited in page 17.)
- [Guldenring 2020] Ronja Guldenring, Michael Görner, Norman Hendrich, Niels Jul Jacobsen and Jianwei Zhang. *Learning local planners for human-aware navigation in indoor environments*. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 6053–6060. IEEE, 2020. (Cited in pages 1, 20, and 59.)
- [Guzzi 2013] Jérôme Guzzi, Alessandro Giusti, Luca M Gambardella, Guy Theraulaz and Gianni A Di Caro. *Human-friendly robot navigation in dynamic environments*. In 2013 IEEE international conference on robotics and automation, pages 423–430. IEEE, 2013. (Cited in page 103.)
- [Hall 1966] E. T. Hall. The hidden dimension: Man’s use of space in public and private. 1966. (Cited in page 26.)
- [Han 2010] Byung-Ok Han, Young-Ho Kim, Kyusung Cho and Hyun S Yang. *Museum tour guide robot with augmented reality*. In 2010 16th International Conference on Virtual Systems and Multimedia, pages 223–229. IEEE, 2010. (Cited in page 17.)
- [Hansen 2009] Soren Tranberg Hansen, Mikael Svenstrup, Hans Jorgen Andersen and Thomas Bak. *Adaptive human aware navigation based on motion pattern analysis*. In RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication, pages 927–932. IEEE, 2009. (Cited in pages 17 and 28.)
- [Hart 1968] Peter E Hart, Nils J Nilsson and Bertram Raphael. *A formal basis for the heuristic determination of minimum cost paths*. IEEE transactions on Systems Science and Cybernetics, vol. 4, no. 2, pages 100–107, 1968. (Cited in page 9.)
- [Hart 2020] Justin Hart, Reuth Mirsky, Xuesu Xiao, Stone Tejeda, Bonny Mahajan, Jamin Goo, Kathryn Baldauf, Sydney Owen and Peter Stone. *Using human-inspired signals to disambiguate navigational intentions*. In International Conference on Social Robotics, pages 320–331. Springer, 2020. (Cited in page 21.)

- [Hayashi 2012] Kotaro Hayashi, Masahiro Shiomi, Takayuki Kanda, Norihiro Hagita and AI Robotics. *Friendly patrolling: A model of natural encounters*. In Proc. RSS, page 121, 2012. (Cited in pages 17 and 28.)
- [Helbing 1995] Dirk Helbing and Peter Molnar. *Social force model for pedestrian dynamics*. Physical review E, vol. 51, no. 5, page 4282, 1995. (Cited in pages 15, 38, and 59.)
- [Helbing 2000] Dirk Helbing, Illés Farkas and Tamas Vicsek. *Simulating dynamical features of escape panic*. Nature, vol. 407, no. 6803, pages 487–490, 2000. (Cited in page 87.)
- [Hennes 2012] Daniel Hennes, Daniel Claes, Wim Meeussen and Karl Tuyls. *Multi-robot collision avoidance with localization uncertainty*. In AAMAS, pages 147–154, 2012. (Cited in page 13.)
- [Henry 2010] Peter Henry, Christian Vollmer, Brian Ferris and Dieter Fox. *Learning to navigate through crowded environments*. In 2010 IEEE international conference on robotics and automation, pages 981–986. IEEE, 2010. (Cited in page 17.)
- [Hetherington 2021] Nicholas J Hetherington, Ryan Lee, Marlene Haase, Elizabeth A Croft and HF Machiel Van der Loos. *Mobile Robot Yielding Cues for Human-Robot Spatial Interaction*. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3028–3033. IEEE, 2021. (Cited in page 21.)
- [Hoeller 2007] Frank Hoeller, Dirk Schulz, Mark Moors and Frank E Schneider. *Accompanying persons with a mobile robot using motion prediction and probabilistic roadmaps*. In 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1260–1265. IEEE, 2007. (Cited in page 16.)
- [Holtz 2021] Jarrett Holtz and Joydeep Biswas. *SOCIALGYM: A Framework for Benchmarking Social Robot Navigation*. arXiv preprint arXiv:2109.11011, 2021. (Cited in page 22.)
- [Huang 2010] Kuo-Chen Huang, Jiun-Yi Li and Li-Chen Fu. *Human-oriented navigation for service providing in home environment*. In Proceedings of SICE Annual Conference 2010, pages 1892–1897. IEEE, 2010. (Cited in page 17.)
- [IDM 2011] IDMIND-ENGENHARIA DE SISTEMAS LDA IDM. *FROG-Fun Robotic Outdoor Guide*. 2011. (Cited in page 18.)
- [Imai 1999] Michita Imai, Tetsuo Ono and Tameyuki Etani. *Agent migration: communications between a human and robot*. In IEEE SMC’99 Conference Proceedings. 1999 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No. 99CH37028), volume 4, pages 1044–1048. IEEE, 1999. (Cited in page 14.)

- [Jensen 2005] Björn Jensen, Nicola Tomatis, Laetitia Mayor, Andrzej Drygajlo and Roland Siegwart. *Robots meet humans-interaction in public spaces*. IEEE Transactions on Industrial Electronics, vol. 52, no. 6, pages 1530–1546, 2005. (Cited in page 15.)
- [Joosse 2021] Michiel Joosse, Manja Lohse, Niels Van Berkel, Aziez Sardar and Vanessa Evers. *Making Appearances: How Robots Should Approach People*. ACM Transactions on Human-Robot Interaction (THRI), vol. 10, no. 1, pages 1–24, 2021. (Cited in page 105.)
- [Kaiser 1997] Michael Kaiser. *Transfer of elementary skills via human-robot interaction*. Adaptive Behavior, vol. 5, no. 3-4, pages 249–280, 1997. (Cited in page 24.)
- [Kanda 2001] Takayuki Kanda, Hiroshi Ishiguro and Toru Ishida. *Psychological analysis on human-robot interaction*. In Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No. 01CH37164), volume 4, pages 4166–4173. IEEE, 2001. (Cited in page 16.)
- [Karnan 2022] Haresh Karnan, Anirudh Nair, Xuesu Xiao, Garrett Warnell, Sören Pirk, Alexander Toshev, Justin Hart, Joydeep Biswas and Peter Stone. *Socially Compliant Navigation Dataset (SCAND): A Large-Scale Dataset of Demonstrations for Social Navigation*. arXiv preprint arXiv:2203.15041, 2022. (Cited in page 22.)
- [Kavraki 1996] Lydia E Kavraki, Petr Svestka, J-C Latombe and Mark H Overmars. *Probabilistic roadmaps for path planning in high-dimensional configuration spaces*. IEEE transactions on Robotics and Automation, vol. 12, no. 4, pages 566–580, 1996. (Cited in pages 9 and 12.)
- [Kawamura 1996] Kazuhiko Kawamura, Robert T Pack, Magued Bishay and Moenes Iskarous. *Design philosophy for service robots*. Robotics and Autonomous Systems, vol. 18, no. 1-2, pages 109–116, 1996. (Cited in page 14.)
- [Kendon 2010] Adam Kendon. *Spacing and orientation in co-present interaction*. In Development of multimodal interfaces: Active listening and synchrony, pages 1–15. Springer, 2010. (Cited in pages 27 and 104.)
- [Kenk 2019] Mourad A Kenk, M Hassaballah and Jean-François Brethé. *Human-aware Robot Navigation in Logistics Warehouses*. In ICINCO (2), pages 371–378, 2019. (Cited in page 20.)
- [Kessler 2011] Jens Kessler, Christof Schroeter and Horst-Michael Gross. *Approaching a person in a socially acceptable manner using a fast marching planner*. In International Conference on Intelligent Robotics and Applications, pages 368–377. Springer, 2011. (Cited in page 16.)

- [Khambhaita 2016a] Harmish Khambhaita, Jorge Rios-Martinez and Rachid Alami. *Head-body motion coordination for human aware robot navigation*. In 9th International workshop on Human-Friendly Robotics (HFR 2016), page 8p, 2016. (Cited in page 24.)
- [Khambhaita 2016b] Harmish Khambhaita, Jorge Rios-Martinez and Rachid Alami. *Head-Body Motion Coordination for Human Aware Robot Navigation*. In Int. Ws. on Human-Friendly Robotics, 2016. (Cited in page 79.)
- [Khambhaita 2017a] Harmish Khambhaita and Rachid Alami. *Assessing the social criteria for human-robot collaborative navigation: A comparison of human-aware navigation planners*. In 2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pages 1140–1145, Lisbon, August 2017. IEEE. (Cited in page 24.)
- [Khambhaita 2017b] Harmish Khambhaita and Rachid Alami. *A Human-Robot Cooperative Navigation Planner*. In Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, pages 161–162, Vienna Austria, March 2017. ACM. (Cited in page 24.)
- [Khambhaita 2017c] Harmish Khambhaita and Rachid Alami. *Viewing Robot Navigation in Human Environment as a Cooperative Activity*. In International Symposium on Robotics Research (ISSR 2017), page 18p, 2017. (Cited in pages 24, 31, 32, 37, 39, 46, 59, 67, and 68.)
- [Kirby 2009] Rachel Kirby, Reid Simmons and Jodi Forlizzi. *Companion: A constraint-optimizing method for person-acceptable navigation*. In RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication, pages 607–612. IEEE, 2009. (Cited in page 18.)
- [Kitazawa 2010] Kay Kitazawa and Taku Fujiyama. *Pedestrian vision and collision avoidance behavior: Investigation of the information process space of pedestrians using an eye tracker*. In Pedestrian and evacuation dynamics 2008, pages 95–108. Springer, 2010. (Cited in page 26.)
- [Kluge 2004] Boris Kluge and Erwin Prassler. *Reflective navigation: Individual behaviors and group behaviors*. In IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004, volume 4, pages 4172–4177. IEEE, 2004. (Cited in page 16.)
- [Koay 2007] Kheng Lee Koay, Emrah Akin Sisbot, Dag Sverre Syrdal, Mick L Walters, Kerstin Dautenhahn and Rachid 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, pages 18–24. Stanford, CA, 2007. (Cited in page 23.)

- [Köhler 2012] Sebastian Köhler, Michael Goldhammer, Sebastian Bauer, Konrad Doll, Ulrich Brunsmann and Klaus Dietmayer. *Early detection of the pedestrian's intention to cross the street*. In 2012 15th International IEEE Conference on Intelligent Transportation Systems, pages 1759–1764. IEEE, 2012. (Cited in page 27.)
- [Kollmitz 2015] Marina Kollmitz, Kaijen Hsiao, Johannes Gaa and Wolfram Burgard. *Time dependent planning on a layered social cost map for human-aware robot navigation*. In 2015 European Conference on Mobile Robots (ECMR), pages 1–6. IEEE, 2015. (Cited in pages 20, 27, 58, 59, and 76.)
- [Kollmitz 2020] Marina Kollmitz, Torsten Koller, Joschka Boedecker and Wolfram Burgard. *Learning human-aware robot navigation from physical interaction via inverse reinforcement learning*. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 11025–11031. IEEE, 2020. (Cited in page 20.)
- [Kostavelis 2016] Ioannis Kostavelis, Dimitrios Giakoumis, Sotiris Malassiotis and Dimitrios Tzovaras. *Human aware robot navigation in semantically annotated domestic environments*. In International Conference on Universal Access in Human-Computer Interaction, pages 414–423. Springer, 2016. (Cited in page 26.)
- [Kostavelis 2017] Ioannis Kostavelis, Andreas Kargakos, Dimitrios Giakoumis and Dimitrios Tzovaras. *Robot's workspace enhancement with dynamic human presence for socially-aware navigation*. In international conference on computer vision systems, pages 279–288. Springer, 2017. (Cited in page 104.)
- [Kruse 2010a] Thibault Kruse, Alexandra Kirsch, E Akin Sisbot and Rachid Alami. *Exploiting human cooperation in human-centered robot navigation*. In 19th International Symposium in Robot and Human Interactive Communication, pages 192–197. IEEE, 2010. (Cited in page 23.)
- [Kruse 2010b] Thibault Kruse, Alexandra Kirsch, Emrah Akin Sisbot and Rachid Alami. *Dynamic generation and execution of human aware navigation plans*. In The 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2010. (Cited in page 23.)
- [Kruse 2012a] Thibault Kruse, Patrizia Basili, Stefan Glasauer and Alexandra Kirsch. *Legible robot navigation in the proximity of moving humans*. In 2012 IEEE workshop on advanced robotics and its social impacts (ARSO), pages 83–88. IEEE, 2012. (Cited in pages 17, 19, and 24.)
- [Kruse 2012b] Thibault Kruse, Patrizia Basili, Stefan Glasauer and Alexandra Kirsch. *Legible Robot Navigation in the Proximity of Moving Humans*. In Ws. on Advanced Robotics and Its Social Impacts, pages 83–88, 2012. (Cited in page 38.)

- [Kruse 2013] Thibault Kruse, Amit Kumar Pandey, Rachid Alami and Alexandra Kirsch. *Human-Aware Robot Navigation: A Survey*. Robotics and Autonomous Systems, vol. 61, no. 12, pages 1726–1743, 2013. (Cited in pages 38 and 81.)
- [Kruse 2014a] Thibault Kruse, Alexandra Kirsch, Harmish Khambhaita and Rachid Alami. *Evaluating directional cost models in navigation*. In Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction, pages 350–357, Bielefeld Germany, March 2014. ACM. (Cited in page 24.)
- [Kruse 2014b] Thibault Kruse, Alexandra Kirsch, Harmish Khambhaita and Rachid Alami. *Evaluating directional cost models in navigation*. In Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction, pages 350–357, 2014. (Cited in pages 28, 38, 45, 46, 68, and 102.)
- [Kuderer 2012] Markus Kuderer, Henrik Kretzschmar, Christoph Sprunk and Wolfram Burgard. *Feature-Based Prediction of Trajectories for Socially Compliant Navigation*. In Robotics: Science and Systems, 2012. (Cited in page 39.)
- [Kuffner 2000] James J Kuffner and Steven M LaValle. *RRT-connect: An efficient approach to single-query path planning*. In Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065), volume 2, pages 995–1001. IEEE, 2000. (Cited in page 12.)
- [Lakoba 2005] Taras I Lakoba, David J Kaup and Neal M Finkelstein. *Modifications of the Helbing-Molnar-Farkas-Vicsek social force model for pedestrian evolution*. Simulation, vol. 81, no. 5, pages 339–352, 2005. (Cited in page 87.)
- [Lam 2010] Chi-Pang Lam, Chen-Tun Chou, Kuo-Hung Chiang and Li-Chen Fu. *Human-centered robot navigation—towards a harmoniously human-robot co-existing environment*. IEEE Transactions on Robotics, vol. 27, no. 1, pages 99–112, 2010. (Cited in page 17.)
- [Lambert 2008] Alain Lambert, Dominique Gruyer, Guillaume Saint Pierre and Alexandre Ndjeng Ndjeng. *Collision probability assessment for speed control*. In 2008 11th International IEEE Conference on Intelligent Transportation Systems, pages 1043–1048. IEEE, 2008. (Cited in page 82.)
- [Latombe 1999] Jean-Claude Latombe. *Motion planning: A journey of robots, molecules, digital actors, and other artifacts*. The International Journal of Robotics Research, vol. 18, no. 11, pages 1119–1128, 1999. (Cited in page 8.)
- [LaValle 1998] Steven M LaValle et al. *Rapidly-exploring random trees: A new tool for path planning*. 1998. (Cited in page 9.)

- [Leibo 2017] Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki and Thore Graepel. *Multi-agent reinforcement learning in sequential social dilemmas*. arXiv preprint arXiv:1702.03037, 2017. (Cited in page 19.)
- [Lichtenthaler 2012a] Christina Lichtenthaler, Tamara Lorenz, Michael Karg and Alexandra Kirsch. *Increasing perceived value between human and robots—measuring legibility in human aware navigation*. In 2012 IEEE workshop on advanced robotics and its social impacts (ARSO), pages 89–94. IEEE, 2012. (Cited in pages 21 and 102.)
- [Lichtenthaler 2012b] Christina Lichtenthaler, Tamara Lorenzy and Alexandra Kirsch. *Influence of legibility on perceived safety in a virtual human-robot path crossing task*. In 2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication, pages 676–681. IEEE, 2012. (Cited in page 17.)
- [Lichtenthaler 2013a] Christina Lichtenthaler and Alexandra Kirsch. *Towards legible robot navigation-how to increase the intend expressiveness of robot navigation behavior*. In International Conference on Social Robotics-Workshop Embodied Communication of Goals and Intentions, 2013. (Cited in page 28.)
- [Lichtenthaler 2013b] Christina Lichtenthaler, Annika Peters, Sascha Griffiths and Alexandra Kirsch. *Social navigation-identifying robot navigation patterns in a path crossing scenario*. In International Conference on Social Robotics, pages 84–93. Springer, 2013. (Cited in page 21.)
- [Liu 2020] Lucia Liu, Daniel Dugas, Gianluca Cesari, Roland Siegwart and Renaud Dub  . *Robot navigation in crowded environments using deep reinforcement learning*. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 5671–5677. IEEE, 2020. (Cited in pages 18 and 19.)
- [Long 2018] Pinxin Long, Tingxiang Fan, Xinyi Liao, Wenxi Liu, Hao Zhang and Jia Pan. *Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning*. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 6252–6259. IEEE, 2018. (Cited in page 19.)
- [Lu 2013] David V Lu and William D Smart. *Towards more efficient navigation for robots and humans*. In 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1707–1713. IEEE, 2013. (Cited in pages 19 and 28.)
- [Lu 2014] David V. Lu, Dave Hershberger and William D. Smart. *Layered Costmaps for Context-Sensitive Navigation*. In Int. Conf. on Intelligent Robots and Systems, pages 709–715, 2014. (Cited in page 59.)

- [Luber 2012] Matthias Luber, Luciano Spinello, Jens Silva and Kai O Arras. *Socially-aware robot navigation: A learning approach*. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 902–907. IEEE, 2012. (Cited in page 18.)
- [Luo 2018] Yuanfu Luo, Panpan Cai, Aniket Bera, David Hsu, Wee Sun Lee and Dinesh Manocha. *Porca: Modeling and planning for autonomous driving among many pedestrians*. IEEE Robotics and Automation Letters, vol. 3, no. 4, pages 3418–3425, 2018. (Cited in page 22.)
- [Macaluso 2005] Irene Macaluso, Edoardo Ardizzone, Antonio Chella, Massimo Cossentino, Antonio Gentile, R Gradino, Ignazio Infantino, Marilia Liotta, Riccardo Rizzo and Giuseppe Scardino. *Experiences with CiceRobot, a museum guide cognitive robot*. In Congress of the Italian Association for Artificial Intelligence, pages 474–482. Springer, 2005. (Cited in page 15.)
- [Madhava Krishna 2003] K. Madhava Krishna, Rachid Alami and Thierry Simeon. *Moving Safely but not Slowly - reactively Adapting Paths for Better Trajectory Times*. In 1th International Conference on Advanced Robotics, Coimbra, Portugal, June 2003. (Cited in page 81.)
- [Madhava Krishna 2006] K. Madhava Krishna, Rachid Alami and Thierry Simeon. *Safe proactive plans and their execution*. Robotics and Autonomous Systems, vol. 54, no. 3, pages 244–255, March 2006. (Cited in page 81.)
- [Mainprice 2010] Jim Mainprice, Emrah Akin Sisbot, Thierry Siméon and Rachid Alami. *Planning safe and legible hand-over motions for human-robot interaction*. In IARP/IEEE-RAS/EURON workshop on technical challenges for dependable robots in human environments, 2010. (Cited in page 23.)
- [Mainprice 2012a] Jim Mainprice, Mamoun Gharbi, Thierry Simeon and Rachid Alami. *Sharing effort in planning human-robot handover tasks*. In 2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication, pages 764–770, Paris, France, September 2012. IEEE. (Cited in page 23.)
- [Mainprice 2012b] Jim Mainprice, Mamoun Gharbi, Thierry Siméon and Rachid Alami. *Sharing Effort in Planning Human-Robot Handover Tasks*. In Int. Symposium on Robot and Human Interactive Communication, pages 764–770, 2012. (Cited in page 39.)
- [Majd 2021] Keyvan Majd, Shakiba Yaghoubi, Tomoya Yamaguchi, Bardh Hoxha, Danil Prokhorov and Georgios Fainekos. *Safe navigation in human occupied environments using sampling and control barrier functions*. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 5794–5800. IEEE, 2021. (Cited in page 20.)

- [Manso 2020a] Luis J Manso, Ronit R Jorvekar, Diego R Faria, Pablo Bustos and Pilar Bachiller. *Graph neural networks for human-aware social navigation*. In Workshop of Physical Agents, pages 167–179. Springer, 2020. (Cited in page 20.)
- [Manso 2020b] Luis J Manso, Pedro Nuñez, Luis V Calderita, Diego R Faria and Pilar Bachiller. *Socnav1: A dataset to benchmark and learn social navigation conventions*. Data, vol. 5, no. 1, page 7, 2020. (Cited in page 22.)
- [Martinez-Garcia 2005] Edgar A Martinez-Garcia, Ohya Akihisa et al. *Crowding and guiding groups of humans by teams of mobile robots*. In IEEE Workshop on Advanced Robotics and its Social Impacts, 2005., pages 91–96. IEEE, 2005. (Cited in page 15.)
- [Mavrogiannis 2018] Christoforos I Mavrogiannis, Wil B Thomason and Ross 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, pages 361–369, 2018. (Cited in pages 19 and 103.)
- [Mavrogiannis 2019] Christoforos Mavrogiannis, Alena M Hutchinson, John Macdonald, Patrícia Alves-Oliveira and Ross A Knepper. *Effects of distinct robot navigation strategies on human behavior in a crowded environment*. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pages 421–430. IEEE, 2019. (Cited in pages 22 and 103.)
- [May 2015] Alyxander David May, Christian Dondrup and Marc Hanheide. *Show me your moves! Conveying navigation intention of a mobile robot to humans*. In 2015 European Conference on Mobile Robots (ECMR), pages 1–6. IEEE, 2015. (Cited in pages 21 and 28.)
- [Mehta 2016] Dhanvin Mehta, Gonzalo Ferrer and Edwin Olson. *Autonomous navigation in dynamic social environments using multi-policy decision making*. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1190–1197. IEEE, 2016. (Cited in pages 38, 39, 45, 60, 78, and 80.)
- [Minguez 2000] Javier Minguez and Luis Montano. *Nearness diagram navigation (nd): A new real time collision avoidance approach*. In Proceedings. 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000)(Cat. No. 00CH37113), volume 3, pages 2094–2100. IEEE, 2000. (Cited in page 22.)
- [Mishra 1983] Pravash K Mishra. *Proxemics: Theory and research*. Perspectives in Psychological Researches, 1983. (Cited in page 26.)

- [Miura 2006] Jun Miura, Yoshiro Negishi and Yoshiaki Shirai. *Adaptive robot speed control by considering map and motion uncertainty*. Robotics and Autonomous Systems, vol. 54, no. 2, pages 110–117, 2006. (Cited in page 82.)
- [Mizuchi 2017] Yoshiaki Mizuchi and Tetsunari Inamura. *Cloud-based multimodal human-robot interaction simulator utilizing ros and unity frameworks*. In 2017 IEEE/SICE International Symposium on System Integration (SII), pages 948–955. IEEE, 2017. (Cited in page 22.)
- [Möller 2021] Ronja Möller, Antonino Furnari, Sebastiano Battiato, Aki Härmä and Giovanni Maria Farinella. *A survey on human-aware robot navigation*. Robotics and Autonomous Systems, vol. 145, page 103837, 2021. (Cited in page 81.)
- [Morales 2017] Yoichi Morales, Takahiro Miyashita and Norihiro Hagita. *Social robotic wheelchair centered on passenger and pedestrian comfort*. Robotics and Autonomous Systems, vol. 87, pages 355–362, 2017. (Cited in page 21.)
- [Morris 2003] Aaron Morris, Raghavendra Donamukkala, Anuj Kapuria, Aaron Steinfeld, Judith T Matthews, Jacqueline Dunbar-Jacob and Sebastian Thrun. *A robotic walker that provides guidance*. In 2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422), volume 1, pages 25–30. IEEE, 2003. (Cited in page 15.)
- [Müller 2008] Jörg Müller, Cyrill Stachniss, Kai O Arras and Wolfram Burgard. *Socially inspired motion planning for mobile robots in populated environments*. In Proc. of International Conference on Cognitive Systems, 2008. (Cited in page 17.)
- [Nagariya 2015] Akhil Nagariya, Bharath Gopalakrishnan, Arun Kumar Singh, Krishna Gupta and K. Madhava Krishna. *Mobile Robot Navigation amidst Humans with Intents and Uncertainties: A Time Scaled Collision Cone Approach*. In Conf. on Decision and Control, pages 2773–2779, 2015. (Cited in page 39.)
- [Narayanan 2016] Vishnu K Narayanan, Anne Spalanzani and Marie Babel. *A semi-autonomous framework for human-aware and user intention driven wheelchair mobility assistance*. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 4700–4707. IEEE, 2016. (Cited in page 20.)
- [Narayanan 2018a] Vishnu K Narayanan, Takahiro Miyashita and Norihiro Hagita. *Formalizing a transient-goal driven approach for pedestrian-aware robot navigation*. In 2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pages 862–867. IEEE, 2018. (Cited in page 21.)

- [Narayanan 2018b] Vishnu K Narayanan, Takahiro Miyashita, Yukiko Horikawa and Norihiro Hagita. *A transient-goal driven communication-aware navigation strategy for large human-populated environments*. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1–9. IEEE, 2018. (Cited in page 21.)
- [Nishimura 2020] Mai Nishimura and Ryo Yonetani. *L2B: Learning to balance the safety-efficiency trade-off in interactive crowd-aware robot navigation*. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 11004–11010. IEEE, 2020. (Cited in page 19.)
- [Okal 2016] Billy Okal and Kai O Arras. *Learning socially normative robot navigation behaviors with bayesian inverse reinforcement learning*. In IEEE ICRA, 2016. (Cited in pages 59 and 104.)
- [Oleiwi 2014] Bashra Kadhim Oleiwi, Rami Al-Jarrah, Hubert Roth and Bahaa I Kazem. *Multi objective optimization of trajectory planning of non-holonomic mobile robot in dynamic environment using enhanced GA by fuzzy motion control and A*. In International Conference on Neural Networks and Artificial Intelligence, pages 34–49. Springer, 2014. (Cited in page 10.)
- [Othman 2020] KM Othman and AB Rad. *SRIN: A New Dataset for Social Robot Indoor Navigation*. Glob. J. Eng. Sci, vol. 4, 2020. (Cited in page 22.)
- [Pacchierotti 2005] Elena Pacchierotti, Henrik I Christensen and Patric Jensfelt. *Human-robot embodied interaction in hallway settings: a pilot user study*. In ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005., pages 164–171. IEEE, 2005. (Cited in page 16.)
- [Pacchierotti 2006] Elena Pacchierotti, Henrik I Christensen and Patric Jensfelt. *Evaluation of passing distance for social robots*. In Roman 2006-the 15th iee international symposium on robot and human interactive communication, pages 315–320. IEEE, 2006. (Cited in page 16.)
- [Paez-Granados 2022] Diego Paez-Granados, Vaibhav Gupta and Aude Billard. *Un-freezing Social Navigation: Dynamical Systems based Compliance for Contact Control in Robot Navigation*. arXiv preprint arXiv:2203.01053, 2022. (Cited in page 21.)
- [Palinko 2020] Oskar Palinko, Eduardo Ruiz Ramírez, William Kristian Juel, Norbert Krüger and Leon Bodenhagen. *Intention Indication for Human Aware Robot Navigation*. In VISIGRAPP (2: HUCAPP), pages 64–74, 2020. (Cited in page 21.)
- [Pandey 2009a] Amit Kumar Pandey and Rachid Alami. *A framework for adapting social conventions in a mobile robot motion in human-centered environment*. In 2009 International Conference on Advanced Robotics, pages 1–8. IEEE, 2009. (Cited in page 23.)

- [Pandey 2009b] Amit Kumar Pandey and Rachid Alami. *A step towards a sociable robot guide which monitors and adapts to the person's activities*. In 2009 International Conference on Advanced Robotics, pages 1–8. IEEE, 2009. (Cited in page 23.)
- [Pandey 2010] A K Pandey and R Alami. *A framework towards a socially aware Mobile Robot motion in Human-Centered dynamic environment*. In 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 5855–5860, Taipei, October 2010. IEEE. (Cited in page 23.)
- [Parikh 2003] Sarangi P Parikh, Rahul Rao, Sang-Hack Jung, Vijay Kumar, James P Ostrowski and Camillo J Taylor. *Human robot interaction and usability studies for a smart wheelchair*. In Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453), volume 4, pages 3206–3211. IEEE, 2003. (Cited in page 15.)
- [Park 2016] Chonhyon Park, Jan Ondřej, Max Gilbert, Kyle Freeman and Carol O’Sullivan. *HI Robot: Human intention-aware robot planning for safe and efficient navigation in crowds*. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3320–3326. IEEE, 2016. (Cited in page 20.)
- [Pasteau 2013] François Pasteau, Marie Babel and Rafiq Sekkal. *Corridor following wheelchair by visual servoing*. In 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 590–595. IEEE, 2013. (Cited in page 21.)
- [Pasteau 2014] François Pasteau, Alexandre Krupa and Marie Babel. *Vision-based assistance for wheelchair navigation along corridors*. In 2014 IEEE International Conference on Robotics and Automation (ICRA), pages 4430–4435. IEEE, 2014. (Cited in page 21.)
- [Pasteau 2016] François Pasteau, Vishnu K Narayanan, Marie Babel and François Chaumette. *A visual servoing approach for autonomous corridor following and doorway passing in a wheelchair*. Robotics and Autonomous Systems, vol. 75, pages 28–40, 2016. (Cited in page 21.)
- [Patompak 2016] Pakpoom Patompak, Sungmoon Jeong, Nak Young Chong and Itthisek Nilkhamhang. *Mobile robot navigation for human-robot social interaction*. In 2016 16th International Conference on Control, Automation and Systems (ICCAS), pages 1298–1303. IEEE, 2016. (Cited in page 18.)
- [Peddi 2020] Rahul Peddi, Carmelo Di Franco, Shijie Gao and Nicola Bezzo. *A data-driven framework for proactive intention-aware motion planning of a robot in a human environment*. In 2020 IEEE/RSJ International Conference

- on Intelligent Robots and Systems (IROS), pages 5738–5744. IEEE, 2020. (Cited in pages 20 and 27.)
- [Pellegrini 2009] Stefano Pellegrini, Andreas Ess, Konrad Schindler and Luc Van Gool. *You’ll never walk alone: Modeling social behavior for multi-target tracking*. In 2009 IEEE 12th international conference on computer vision, pages 261–268. IEEE, 2009. (Cited in page 103.)
- [Pérez-Higueras 2014] Noé Pérez-Higueras, Rafael Ramón-Vigo, Fernando Caballero and Luis Merino. *Robot local navigation with learned social cost functions*. In 2014 11th International Conference on Informatics in Control, Automation and Robotics (ICINCO), volume 2, pages 618–625. IEEE, 2014. (Cited in page 18.)
- [Pérez-Higueras 2018a] Noé Pérez-Higueras, Fernando Caballero *et al.* *Teaching robot navigation behaviors to optimal RRT planners*. International Journal of Social Robotics, 2018. (Cited in pages 59 and 104.)
- [Pérez-Higueras 2018b] Noé Pérez-Higueras, Fernando Caballero and Luis Merino. *Learning human-aware path planning with fully convolutional networks*. In 2018 IEEE international conference on robotics and automation (ICRA), pages 5897–5902. IEEE, 2018. (Cited in page 20.)
- [Philippsen 2003] Roland Philippsen and Roland Siegwart. *Smooth and efficient obstacle avoidance for a tour guide robot*. In Proceedings. 2003 IEEE International Conference on Robotics and Automation, September 14–19, 2003, The Grand Hotel, Taipei, Taiwan, volume 1, pages 446–451. IEEE Operations Center, 2003. (Cited in page 15.)
- [Prassler 1998] Erwin Prassler, Jens Scholz and Matthias Strobel. *Maid: Mobility assistance for elderly and disabled people*. In IECON’98. Proceedings of the 24th Annual Conference of the IEEE Industrial Electronics Society (Cat. No. 98CH36200), volume 4, pages 2493–2498. IEEE, 1998. (Cited in page 14.)
- [Prassler 1999] Erwin Prassler, Jens Scholz and Paolo Fiorini. *Navigating a robotic wheelchair in a railway station during rush hour*. The international journal of robotics research, vol. 18, no. 7, pages 711–727, 1999. (Cited in pages 13 and 14.)
- [Prassler 2002] Erwin Prassler, Dirk Bank and Boris Kluge. *Key technologies in robot assistants: Motion coordination between a human and a mobile robot*. Transactions on Control, Automation and Systems Engineering, vol. 4, no. 1, pages 56–61, 2002. (Cited in page 15.)
- [Pérez-Higueras 2014] Noé Pérez-Higueras, Rafael Ramón-Vigo *et al.* *Robot local navigation with learned social cost functions*. In 11th International Conference on Informatics in Control, Automation and Robotics, 2014. (Cited in page 59.)

- [Qian 2013] Kun Qian, Xudong Ma, Xianzhong Dai, Fang Fang and Bo Zhou. *Decision-theoretical navigation of service robots using POMDPs with human-robot co-occurrence prediction*. International Journal of Advanced Robotic Systems, vol. 10, no. 2, page 143, 2013. (Cited in pages 19, 28, 38, 39, 45, 60, and 78.)
- [Qiu 2022] Quecheng Qiu, Shunyi Yao, Jing Wang, Jun Ma, Guangda Chen and Jianmin Ji. *Learning to Socially Navigate in Pedestrian-rich Environments with Interaction Capacity*. arXiv preprint arXiv:2203.16154, 2022. (Cited in page 20.)
- [Quigley 2009] Morgan Quigley, Ken Conley *et al.* *ROS: an open-source Robot Operating System*. In ICRA workshop on open source software, 2009. (Cited in pages 13, 60, and 82.)
- [Quinlan 1993] Sean Quinlan and Oussama Khatib. *Elastic bands: Connecting path planning and control*. In [1993] Proceedings IEEE International Conference on Robotics and Automation, pages 802–807. IEEE, 1993. (Cited in pages 12 and 29.)
- [Ramírez 2016] Omar A Islas Ramírez, Harmish Khambhaita, Raja Chatila, Mohamed Chetouani and Rachid Alami. *Robots learning how and where to approach people*. In 2016 25th IEEE international symposium on robot and human interactive communication (RO-MAN), pages 347–353. IEEE, 2016. (Cited in pages 23 and 28.)
- [Rao 2002] RS Rao, K Conn, Sang-Hack Jung, Jayantha Katupitiya, Terry Kientz, Vijay Kumar, J Ostrowski, Sarangi Patel and Camillo J Taylor. *Human robot interaction: application to smart wheelchairs*. In Proceedings 2002 IEEE international conference on robotics and automation (Cat. No. 02CH37292), volume 4, pages 3583–3588. IEEE, 2002. (Cited in page 15.)
- [Rasouli 2019] Amir Rasouli and John K Tsotsos. *Autonomous vehicles that interact with pedestrians: A survey of theory and practice*. IEEE transactions on intelligent transportation systems, vol. 21, no. 3, pages 900–918, 2019. (Cited in page 1.)
- [Ratsamee 2013] Photchara Ratsamee, Yasushi Mae, Kenichi Ohara, Masaru Kojima and Tatsuo Arai. *Social navigation model based on human intention analysis using face orientation*. In 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1682–1687. IEEE, 2013. (Cited in pages 20 and 27.)
- [Rayner 2009] Keith Rayner, Tim J Smith, George L Malcolm and John M Henderson. *Eye movements and visual encoding during scene perception*. Psychological science, vol. 20, no. 1, pages 6–10, 2009. (Cited in page 110.)

- [Repiso 2017] Ely Repiso, Gonzalo Ferrer and Alberto Sanfeliu. *On-line adaptive side-by-side human robot companion in dynamic urban environments*. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 872–877. IEEE, 2017. (Cited in pages 1, 18, 59, and 79.)
- [Rios-Martinez 2011] Jorge Rios-Martinez, Anne Spalanzani and Christian Laugier. *Understanding human interaction for probabilistic autonomous navigation using Risk-RRT approach*. In 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 2014–2019. IEEE, 2011. (Cited in page 20.)
- [Rios-Martinez 2012a] Jorge Rios-Martinez, Arturo Escobedo, Anne Spalanzani and Christian Laugier. *Intention driven human aware navigation for assisted mobility*. In Workshop on Assistance and Service robotics in a human environment at IROS, 2012. (Cited in page 20.)
- [Rios-Martinez 2012b] Jorge Rios-Martinez, Alessandro Renzaglia, Anne Spalanzani, Agostino Martinelli and Christian Laugier. *Navigating between people: A stochastic optimization approach*. In 2012 IEEE International Conference on Robotics and Automation, pages 2880–2885. IEEE, 2012. (Cited in page 17.)
- [Rios-Martinez 2015] Jorge Rios-Martinez, Anne Spalanzani and Christian Laugier. *From proxemics theory to socially-aware navigation: A survey*. International Journal of Social Robotics, vol. 7, no. 2, pages 137–153, 2015. (Cited in pages 26, 27, and 104.)
- [Rodriguez-Losada 2008] Diego Rodriguez-Losada, Fernando Matia, Juan Manuel Lucas, Juan Manuel Montero, Miguel Hernando and Ramon Galan. Urbano, an interactive mobile tour-guide robot. INTECH Open Access Publisher London, UK., 2008. (Cited in page 17.)
- [Rösmann 2013] Christoph Rösmann, Wendelin Feiten, Thomas Wösch, Frank Hoffmann and Torsten Bertram. *Efficient trajectory optimization using a sparse model*. In 2013 European Conference on Mobile Robots, pages 138–143. IEEE, 2013. (Cited in pages 10, 13, 28, 29, 31, and 41.)
- [Rösmann 2021] Christoph Rösmann, Artemi Makarow and Torsten Bertram. *On-line motion planning based on nonlinear model predictive control with non-Euclidean rotation groups*. In 2021 European Control Conference (ECC), pages 1583–1590. IEEE, 2021. (Cited in page 28.)
- [Rudenko 2020] Andrey Rudenko, Tomasz P Kucner, Chittaranjan S Swaminathan, Ravi T Chadalavada, Kai O Arras and Achim J Lilienthal. *Thör: Human-robot navigation data collection and accurate motion trajectories dataset*. IEEE Robotics and Automation Letters, vol. 5, no. 2, pages 676–682, 2020. (Cited in page 22.)

- [Salvini 2022] Pericle Salvini, Diego Paez-Granados and Aude Billard. *Safety Concerns Emerging from Robots Navigating in Crowded Pedestrian Areas*. In *International Journal of Social Robotics*, vol. 14, no. 2, pages 441–462, 2022. (Cited in page 22.)
- [Satake 2009] Satoru Satake, Takayuki Kanda, Dylan F Glas, Michita Imai, Hiroshi Ishiguro and Norihiro Hagita. *How to approach humans? Strategies for social robots to initiate interaction*. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 109–116, 2009. (Cited in page 16.)
- [Sathiamoorthy 2020] Adarsh Jagan Sathiamoorthy, Utsav Patel, Tianrui Guan and Dinesh Manocha. *Frozone: Freezing-free, pedestrian-friendly navigation in human crowds*. *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pages 4352–4359, 2020. (Cited in page 19.)
- [Schulte 1999] Jamieson Schulte, Charles Rosenberg and Sebastian Thrun. *Spontaneous, short-term interaction with mobile robots*. In *Proceedings 1999 IEEE International Conference on Robotics and Automation* (Cat. No. 99CH36288C), volume 1, pages 658–663. IEEE, 1999. (Cited in page 14.)
- [Senft 2020] Emmanuel Senft, Satoru Satake and Takayuki Kanda. *Would You Mind Me if I Pass by You?: Socially-Appropriate Behaviour for an Omnibased Social Robot in Narrow Environment*. In *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 539–547. IEEE, 2020. (Cited in page 21.)
- [Shi 2008] Dongqing Shi, Emmanuel G Collins Jr, Brian Goldiez, Arturo Donate, Xiuwen Liu and Damion Dunlap. *Human-aware robot motion planning with velocity constraints*. In *2008 International Symposium on Collaborative Technologies and Systems*, pages 490–497. IEEE, 2008. (Cited in page 17.)
- [Shin 2020] Heechan Shin and Sung-Eui Yoon. *Optimization-based path planning for person following using following field*. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11352–11359. IEEE, 2020. (Cited in page 20.)
- [Shrestha 2018] Moondeep C Shrestha, Tomoya Onishi, Ayano Kobayashi, Mitsuhiro Kamezaki and Shigeki Sugano. *Communicating directional intent in robot navigation using projection indicators*. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 746–751. IEEE, 2018. (Cited in page 28.)
- [Siegwart 2003] Roland Siegwart, KO Arras, Björn Jensen, Roland Philippsen and Nicola Tomatis. *Design, implementation and exploitation of a new fully autonomous tour guide robot*. In *Proceedings. ASER’03, 1st International*

- Workshop on Advances in Service Robotics: March 13-15, 2003-Bardolino, Italy, pages 146–151. Fraunhofer IRB-Verl., 2003. (Cited in page 15.)
- [Singamaneni 2020] Phani-Teja Singamaneni and Rachid Alami. *HATEB-2: Reactive Planning and Decision making in Human-Robot Co-navigation*. In International Conference on Robot & Human Interactive Communication, 2020. (Cited in pages 58, 60, 61, 64, 67, and 68.)
- [Singamaneni 2021] Phani Teja Singamaneni, Anthony Favier and Rachid Alami. *Human-Aware Navigation Planner for Diverse Human-Robot Interaction Contexts*. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 5817–5824. IEEE, 2021. (Cited in pages 82 and 87.)
- [Sisbot 2005] E.A. Sisbot, R. Alami, T. Simeon, K. Dautenhahn, M. Walters and S. Woods. *Navigation in the presence of humans*. In 5th IEEE-RAS International Conference on Humanoid Robots, 2005., pages 181–188, San Diego, Cali, USA, 2005. IEEE. (Cited in pages 22 and 23.)
- [Sisbot 2006] Emrah Sisbot, Luis Marin, Rachid Alami and Thierry Simeon. *A mobile robot that performs human acceptable motions*. In 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1811–1816, Beijing, China, October 2006. IEEE. (Cited in page 22.)
- [Sisbot 2007a] E.A. Sisbot, L.F. Marin-Urias, R. Alami and T. Simeon. *A Human Aware Mobile Robot Motion Planner*. IEEE Trans. Robot., vol. 23, no. 5, pages 874–883, October 2007. (Cited in page 22.)
- [Sisbot 2007b] Emrah Akin Sisbot, Luis F. Marin and Rachid Alami. *Spatial reasoning for human robot interaction*. In 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 2281–2287, San Diego, CA, USA, October 2007. IEEE. (Cited in pages 22 and 23.)
- [Sisbot 2007c] Emrah Akin Sisbot, Luis F. Marin-Urias, Rachid Alami and Thierry Siméon. *A Human Aware Mobile Robot Motion Planner*. Transactions on Robotics, vol. 23, no. 5, pages 874–883, 2007. (Cited in pages 38, 59, and 60.)
- [Sisbot 2010] Emrah Akin Sisbot, Luis F. Marin-Urias, Xavier Broquère, Daniel Sidobre and Rachid Alami. *Synthesizing Robot Motions Adapted to Human Presence: A Planning and Control Framework for Safe and Socially Acceptable Robot Motions*. Int J of Soc Robotics, vol. 2, no. 3, pages 329–343, September 2010. (Cited in page 23.)
- [Snape 2011] Jamie Snape, Jur Van Den Berg, Stephen J Guy and Dinesh Manocha. *The hybrid reciprocal velocity obstacle*. IEEE Transactions on Robotics, vol. 27, no. 4, pages 696–706, 2011. (Cited in pages 18 and 38.)

- [Sorrentino 2021] Alessandra Sorrentino, O Khalid, Luigi Coviello, Filippo Cavallo and Laura Fiorini. *Modeling human-like robot personalities as a key to foster socially aware navigation*. In 2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN), pages 95–101. IEEE, 2021. (Cited in page 22.)
- [Stentz 1997] Anthony Stentz. *Optimal and efficient path planning for partially known environments*. In Intelligent unmanned ground vehicles, pages 203–220. Springer, 1997. (Cited in page 9.)
- [Svenstrup 2010] Mikael Svenstrup, Thomas Bak and Hans Jørgen Andersen. *Trajectory planning for robots in dynamic human environments*. In 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 4293–4298. IEEE, 2010. (Cited in page 17.)
- [Tadokoro 1995] Satoshi Tadokoro, Masaki Hayashi, Yasuhiro Manabe, Yoshihiro Nakami and Toshi Takamori. *On motion planning of mobile robots which coexist and cooperate with human*. In Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots, volume 2, pages 518–523. IEEE, 1995. (Cited in page 14.)
- [Taha 2008] Tarek Taha, J Valls Miro and Gamini Dissanayake. *Intention driven assistive wheelchair navigation*. ACM/IEEE Human Robot Interaction, 2008. (Cited in page 17.)
- [Takayama 2009] Leila Takayama and Caroline Pantofaru. *Influences on proxemic behaviors in human-robot interaction*. In 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 5495–5502. IEEE, 2009. (Cited in pages 17 and 102.)
- [Teknomo 2002] Kardi Teknomo. *Microscopic Pedestrian Flow Characteristics: Development of an Image Processing Data Collection and Simulation Model*. PhD thesis, 03 2002. (Cited in page 87.)
- [Thompson 2009] Simon Thompson, Takehiro Horiuchi and Satoshi Kagami. *A probabilistic model of human motion and navigation intent for mobile robot path planning*. In 2009 4th International Conference on Autonomous Robots and Agents, pages 663–668. IEEE, 2009. (Cited in page 18.)
- [Thorpe 1996] Simon Thorpe, Denis Fize and Catherine Marlot. *Speed of processing in the human visual system*. nature, vol. 381, no. 6582, pages 520–522, 1996. (Cited in page 110.)
- [Thrun 1999] Sebastian Thrun, Maren Bennewitz, Wolfram Burgard, Armin B Cremers, Frank Dellaert, Dieter Fox, Dirk Hahnel, Charles Rosenberg, Nicholas Roy, Jamieson Schulteet al. *MINERVA: A second-generation museum tour-guide robot*. In Proceedings 1999 IEEE International Conference on Robotics

- and Automation (Cat. No. 99CH36288C), volume 3. IEEE, 1999. (Cited in page 14.)
- [Thrun 2000] Sebastian Thrun, Michael Beetz, Maren Bennewitz, Wolfram Burgard, Armin B Cremers, Frank Dellaert, Dieter Fox, Dirk Haehnel, Chuck Rosenberg, Nicholas Roy *et al.* *Probabilistic algorithms and the interactive museum tour-guide robot minerva*. The International Journal of Robotics Research, vol. 19, no. 11, pages 972–999, 2000. (Cited in page 14.)
- [Thrun 2006] Sebastian Thrun, Wolfram Burgard and Dieter Fox. *Probalistic robotics*. Kybernetes, 2006. (Cited in pages 11 and 12.)
- [Thrun 2007] Sebastian Thrun. *Simultaneous localization and mapping*. In Robotics and cognitive approaches to spatial mapping, pages 13–41. Springer, 2007. (Cited in page 11.)
- [Tian 2004] Lianfang Tian and Curtis Collins. *An effective robot trajectory planning method using a genetic algorithm*. Mechatronics, vol. 14, no. 5, pages 455–470, 2004. (Cited in page 10.)
- [Tomari 2012] Razali Tomari, Yoshinori Kobayashi and Yoshinori Kuno. *Empirical framework for autonomous wheelchair systems in human-shared environments*. In 2012 IEEE International Conference on Mechatronics and Automation, pages 493–498. IEEE, 2012. (Cited in page 17.)
- [Trautman 2010] Peter Trautman and Andreas Krause. *Unfreezing the robot: Navigation in dense, interacting crowds*. In 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 797–803. IEEE, 2010. (Cited in pages 1 and 17.)
- [Trautman 2015] Pete Trautman, Jeremy Ma, Richard M Murray and Andreas Krause. *Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation*. The International Journal of Robotics Research, vol. 34, no. 3, pages 335–356, 2015. (Cited in page 19.)
- [Triebel 2016] Rudolph Triebel, Kai Arras, Rachid Alami, Lucas Beyer, Stefan Breuers, Raja Chatila, Mohamed Chetouani, Daniel Cremers, Vanessa Evers, Michelangelo Fiore *et al.* *Spencer: A socially aware service robot for passenger guidance and help in busy airports*. In Field and service robotics, pages 607–622. Springer, 2016. (Cited in pages 18 and 59.)
- [Truc 2022] Jérôme Truc, Phani-Teja Singamaneni, Daniel Sidobre, Serena Ivaldi and Rachid Alami. *KHAOS: a Kinematic Human Aware Optimization-based System for Reactive Planning of Flying-Coworker*. In ICRA 2022, 2022. (Cited in page 24.)
- [Truong 2014] Xuan-Tung Truong, Voo Nyuk Yoong and Trung-Dung Ngo. *Dynamic social zone for human safety in human-robot shared workspaces*. In

- 2014 11th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI), pages 391–396. IEEE, 2014. (Cited in pages 19, 27, 58, 59, 102, and 104.)
- [Truong 2016] Xuan Tung Truong, Yong Sheng Ou and Trung-Dung Ngo. *Towards culturally aware robot navigation*. In 2016 IEEE International Conference on Real-time Computing and Robotics (RCAR), pages 63–69. IEEE, 2016. (Cited in page 20.)
- [Truong 2017a] Xuan-Tung Truong and Trung Dung Ngo. *Toward socially aware robot navigation in dynamic and crowded environments: A proactive social motion model*. IEEE Transactions on Automation Science and Engineering, no. 4, 2017. (Cited in pages 18 and 104.)
- [Truong 2017b] Xuan-Tung Truong and Trung Dung Ngo. *Toward Socially Aware Robot Navigation in Dynamic and Crowded Environments: A Proactive Social Motion Model*. IEEE Transactions on Automation Science and Engineering, 2017. (Cited in page 59.)
- [Truong 2017c] Xuan-Tung Truong and Trung-Dung Ngo. *“To approach humans?”: A unified framework for approaching pose prediction and socially aware robot navigation*. IEEE Transactions on Cognitive and Developmental Systems, vol. 10, no. 3, pages 557–572, 2017. (Cited in page 20.)
- [Truong 2017d] Xuan-Tung Truong, Voo Nyuk Yoong and Trung-Dung Ngo. *Socially aware robot navigation system in human interactive environments*. Intelligent Service Robotics, vol. 10, no. 4, pages 287–295, 2017. (Cited in page 20.)
- [Tsoi 2020] Nathan Tsoi, Mohamed Hussein, Jeacy Espinoza, Xavier Ruiz and Marynel Vázquez. *Sean: Social environment for autonomous navigation*. In Proceedings of the 8th International Conference on Human-Agent Interaction, pages 281–283, 2020. (Cited in pages 22 and 121.)
- [Tsoi 2022] Nathan Tsoi, Alec Xiang, Peter Yu, Samuel S Sohn, Greg Schwartz, Subashri Ramesh, Mohamed Hussein, Anjali W Gupta, Mubbasir Kapadia and Marynel Vázquez. *Sean 2.0: Formalizing and generating social situations for robot navigation*. IEEE Robotics and Automation Letters, 2022. (Cited in page 121.)
- [Vadakkepat 2000] Prahlad Vadakkepat, Kay Chen Tan and Wang Ming-Liang. *Evolutionary artificial potential fields and their application in real time robot path planning*. In Proceedings of the 2000 congress on evolutionary computation. CEC00 (Cat. No. 00TH8512), volume 1, pages 256–263. IEEE, 2000. (Cited in page 9.)
- [Van den Berg 2008] Jur Van den Berg, Ming Lin and Dinesh Manocha. *Reciprocal velocity obstacles for real-time multi-agent navigation*. In 2008 IEEE

- international conference on robotics and automation, pages 1928–1935. Ieee, 2008. (Cited in page 13.)
- [Vanhooydonck 2010] Dirk Vanhooydonck, Eric Demeester, Alexander Hüntemann, Johan Philips, Gerolf Vanacker, Hendrik Van Brussel and Marnix Nuttin. *Adaptable navigational assistance for intelligent wheelchairs by means of an implicit personalized user model*. Robotics and Autonomous Systems, vol. 58, no. 8, pages 963–977, 2010. (Cited in page 17.)
- [Vasconcelos 2015] Phelipe AA Vasconcelos, Henrique NS Pereira, Douglas G Macharet and Erickson R Nascimento. *Socially acceptable robot navigation in the presence of humans*. In 2015 12th Latin American Robotics Symposium and 2015 3rd Brazilian Symposium on Robotics (LARS-SBR), pages 222–227. IEEE, 2015. (Cited in page 19.)
- [Vasquez 2013] Dizan Vasquez, Procópio Stein, Jorge Rios-Martinez, Arturo Escobedo, Anne Spalanzani and Christian Laugier. *Human aware navigation for assistive robotics*. In experimental robotics, pages 449–462. Springer, 2013. (Cited in page 20.)
- [Vasquez 2014] Dizan Vasquez, Billy Okal and Kai O Arras. *Inverse reinforcement learning algorithms and features for robot navigation in crowds: an experimental comparison*. In 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1341–1346. IEEE, 2014. (Cited in page 18.)
- [Vathsal 1977] S Vathsal, KAP Menon and R Swaminathan. *Minimax approach to trajectory optimization of multistage launch vehicles*. IEEE Transactions on Aerospace and Electronic Systems, no. 2, pages 179–187, 1977. (Cited in page 10.)
- [Vega-Magro 2018] Araceli Vega-Magro, Luis J Manso, Pablo Bustos and Pedro Núñez. *A flexible and adaptive spatial density model for context-aware social mapping: towards a more realistic social navigation*. In 2018 15th International Conference on Control, Automation, Robotics and Vision (ICARCV), pages 1727–1732. IEEE, 2018. (Cited in pages 20 and 24.)
- [Vega 2017] Araceli Vega, LJ Manso, Douglas G Macharet, Pablo Bustos and Pedro Núñez. *A new strategy based on an adaptive spatial density function for social robot navigation in human-populated environments*. In Proceedings of REACTS workshop at the International Conference on Computer Analysis and Patterns, 2017. (Cited in page 20.)
- [Vega 2018] Araceli Vega, Luis J Manso, Ramón Cintas and Pedro Núñez. *Planning human-robot interaction for social navigation in crowded environments*. In Workshop of Physical Agents, pages 195–208. Springer, 2018. (Cited in pages 24 and 105.)

- [Waldhart 2015a] Jules Waldhart, Mamoun Gharbi and Rachid Alami. *Planning handovers involving humans and robots in constrained environment*. In 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 6473–6478, Hamburg, Germany, September 2015. IEEE. (Cited in page 23.)
- [Waldhart 2015b] Jules Waldhart, Mamoun Gharbi and Rachid Alami. *Planning Handovers Involving Humans and Robots in Constrained Environment*. In Int. Conf. on Intelligent Robots and Systems, 2015. (Cited in page 39.)
- [Wang 2020] Allan Wang and Aaron Steinfeld. *Group split and merge prediction with 3D convolutional networks*. IEEE Robotics and Automation Letters, vol. 5, no. 2, pages 1923–1930, 2020. (Cited in page 19.)
- [Wang 2022] Allan Wang, Christoforos Mavrogiannis and Aaron Steinfeld. *Group-based motion prediction for navigation in crowded environments*. In Conference on Robot Learning, pages 871–882. PMLR, 2022. (Cited in page 19.)
- [Warren 1993] Charles W Warren. *Fast path planning using modified A* method*. In [1993] Proceedings IEEE International Conference on Robotics and Automation, pages 662–667. IEEE, 1993. (Cited in page 12.)
- [Wilkes 1998] D Mitchell Wilkes, A Alford, Robert T Pack, T Rogers, RA Peters and Kazuhiko Kawamura. *Toward socially intelligent service robots*. Applied Artificial Intelligence, vol. 12, no. 7-8, pages 729–766, 1998. (Cited in page 14.)
- [Wolfe 2020] Benjamin Wolfe, Bobbie Seppelt, Bruce Mehler, Bryan Reimer and Ruth Rosenholtz. *Rapid holistic perception and evasion of road hazards*. Journal of experimental psychology: general, vol. 149, no. 3, page 490, 2020. (Cited in page 110.)
- [Xie 2021] Zhanteng Xie, Pujie Xin and Philip Dames. *Towards Safe Navigation Through Crowded Dynamic Environments*. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 4934–4940. IEEE, 2021. (Cited in page 18.)
- [Yanco 1998a] Holly A Yanco. *Integrating robotic research: A survey of robotic wheelchair development*. In AAAI Spring Symposium on Integrating Robotic Research, pages 136–141. Citeseer, 1998. (Cited in page 14.)
- [Yanco 1998b] Holly A Yanco. *Wheesley: A robotic wheelchair system: Indoor navigation and user interface*. Assistive technology and artificial intelligence, pages 256–268, 1998. (Cited in page 14.)
- [Zender 2007] Hendrik Zender, Patric Jensfelt and Geert-Jan M Kruijff. *Human- and situation-aware people following*. In RO-MAN 2007-The 16th IEEE

International Symposium on Robot and Human Interactive Communication, pages 1131–1136. IEEE, 2007. (Cited in pages 16 and 28.)

- [Zeng 2006] Qiang Zeng, Chee Leong Teo, Brice Rebsamen and Etienne Burdet. *Design of a collaborative wheelchair with path guidance assistance*. In Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006., pages 877–882. IEEE, 2006. (Cited in page 15.)