Human-aware Autonomous Vehicle Navigation in Pedestrian-rich Unstructured Environments

by

Mahsa Golchoubian

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2024

© Mahsa Golchoubian 2024

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Anne Spalanzani Professor, Grenoble Alpes University
Supervisor(s):	Nasser Lashgarian Azad Professor, Dept. of Systems Design Engineering, University of Waterloo
	Kerstin Dautenhahn Professor, Dept. of Electrical and Computer Engineering, University of Waterloo
Internal Member:	Catherine Burns Professor, Dept. of Systems Design Engineering, University of Waterloo
Internal-External Member:	Yash Vardhan Pant Assistant Professor, Dept. of Electrical and Computer Engineering, University of Waterloo
Internal-External Member:	Yue Hu Assistant Professor, Dept. of Mechanical and Mechatronics Engineering, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis incorporates content from peer-reviewed published or in-press papers. For all these publications, Mahsa Golchoubian was the principal investigator and first author, creating the first complete manuscript draft that her co-authors then commented on.

Journal Articles:

[J1] Mahsa Golchoubian, Moojan Ghafurian, Kerstin Dautenhahn and Nasser Lashgarain Azad, "Uncertainty-Aware DRL for Autonomous Vehicle Crowd Navigation in Shared Space", IEEE Transaction on Intelligent vehicles, 2024.

[J2] Mahsa Golchoubian, Moojan Ghafurian, Kerstin Dautenhahn and Nasser Lashgarain Azad, "Pedestrian Trajectory Prediction in Pedestrian-Vehicle Mixed Environments: A Systematic Review," in IEEE Transactions on Intelligent Transportation Systems, vol. 24, no. 11, pp. 11544-11567, Nov. 2023.

Conference Papers:

[C1] Mahsa Golchoubian, Moojan Ghafurian, Kerstin Dautenhahn and Nasser Lashgarain Azad, "Polar Collision Grids: Effective Interaction Modelling for Pedestrian Trajectory Prediction in Shared Space Using Collision Checks," 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC), Bilbao, Spain, 2023, pp. 791-798.

[C2] Mahsa Golchoubian, Moojan Ghafurian, Nasser Lashgarain Azad and Kerstin Dautenhahn, "Characterizing Structured Versus Unstructured Environments Based on Pedestrians' and Vehicles' Motion Trajectories," 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC), Macau, China, 2022, pp. 2888-2895.

[C3] Mahsa Golchoubian, Moojan Ghafurian, Nasser Lashgarain Azad, and Kerstin Dautenhahn, "What are Social Norms for Low-speed Autonomous Vehicle Navigation in Crowded Environments? An Online Survey," In Proceedings of the 9th International Conference on Human-Agent Interaction (HAI '21), 2021, pp. 148–156.

Abstract

Autonomous Vehicles (AVs) have the potential to enhance transportation safety, improve efficiency, and elevate quality of life. Despite significant advancements in AV technology, operating these vehicles in dynamic, crowded environments that requires frequent interaction with other decision-making agents remains challenging. A key example is the interaction between AVs and pedestrians. While most research has focused on these interactions in structured road settings, the complexity and diversity of AV navigation among pedestrians in unstructured environments (e.g., shared spaces, airport terminals) have been less explored. In such pedestrian-rich environments, AVs must be human-aware, meeting people's expectations while ensuring both their safety and comfort. At the same time, navigating these spaces requires reasoning about mutual interactions and accounting for the uncertainty in pedestrian behaviour. This thesis introduces a novel approach to address these challenges, presenting an integrated prediction and planning framework for AV navigation among pedestrians in unstructured shared environments.

The thesis is structured into two main phases: a design requirement study and an algorithmic development phase. Given the novelty of this application, the first phase focused on understanding the perceptions and preferences of pedestrians regarding AV behaviour in common interactive scenarios within unstructured settings. Additionally, we examined the unique aspects of pedestrian behaviour in these environments, identifying common behaviours AVs must manage and gathering existing datasets that better represent pedestrian behaviour in such settings. This study highlights the importance of considering uncertainty in pedestrian behaviour, shaping the direction of the development phase.

In the algorithmic development phase, we propose a novel proactive, uncertainty-aware Deep Reinforcement Learning (DRL) decision-making algorithm. This algorithm efficiently accounts for complex interaction effects with multiple pedestrians while maintaining reasonable computational time. The navigation algorithm is made proactive and farsighted by integrating the DRL motion planner with a data-driven pedestrian trajectory predictor. Our novel prediction model is designed to forecast pedestrian trajectories in highly interactive shared environments. It uses a collision risk metric to identify key interacting agents and encodes their effects through a newly engineered interaction feature which guide the learning process. During training, we prevented overconfident predictions and improved estimates of prediction uncertainty using an augmented loss function that incorporates uncertainty awareness.

Unlike other DRL algorithms in this area, our model's DRL motion planner accounts for prediction uncertainty, integrating it into the reward function to encourage the AV to

minimize collision probability with pedestrians over a prediction horizon. Additionally, the reward function design encourages socially aware behaviours, such as reducing speed during close encounters, respecting pedestrians' personal space, and adhering to social norms identified in our earlier design requirement study.

We trained our model in a simulation environment that contains realistic pedestrian trajectory behaviour in the presence of vehicles in shared spaces. The simulation results demonstrate that our uncertainty-aware DRL navigation framework outperforms state-of-the-art DRL crowd navigation and uncertainty-aware model predictive control (MPC) models, both in terms of efficiency and social behaviour aspects.

Overall, this thesis contributes to the advancement of socially-aware crowd navigation algorithms beyond human-sized mobile robots to autonomous vehicles operating as mobility aids among pedestrians in unstructured environments. It demonstrates how agent interactions can be effectively modelled within prediction and planning modules, and how uncertainty in these predictions can be integrated into a DRL-based motion planner.

Acknowledgements

This work would not have been possible without the support of many people. First and foremost, I would like to express my deepest gratitude to my supervisors, Prof. Nasser Lashgarian Azad and Prof. Kerstin Dautenhahn, for their guidance and support throughout my PhD. Their mentorship provided me with opportunities to strengthen my skills across various dimensions and grow as a researcher. Their trust and insightful discussions encouraged me to explore my field of interest and approach this work from diverse perspectives.

My heartfelt thanks go to Dr. Moojan Ghafurian, my mentor and co-author, whom I regarded as a third supervisor. Her readiness to help, positive attitude, and invaluable collaboration made this work possible. I am also deeply thankful to the members of my examining committee for their time, valuable insights, and constructive feedback. I would also like to thank Dr. Krzysztof Czarnecki and his research group for welcoming me to their weekly meetings, where I gained valuable knowledge through the discussions and benefited from their thoughtful feedback on my presentation.

I am grateful to my friends and fellow researchers at both the AIS and SIRRL labs for their support and feedback, and for fostering a fun and enjoyable working environment, making these years memorable. I feel fortunate to have shared time with them both in the lab and during our social gatherings.

Words cannot fully express my gratitude to my parents for instilling in me a passion for science and learning, and for their constant motivation and support throughout my educational journey. Their encouragement and guidance has been a constant source of strength. I am also thankful to my sister Parisa and Ali for helping me settle in Canada during the challenging early days of the COVID-19 pandemic and for their emotional support during that time. A special thank you goes to my younger sister, Dorsa, whose presence in Canada during the later years of my PhD was heartwarming and brought me immense joy.

Finally, my deepest gratitude goes to my beloved husband, Sajjad, for his unwavering love and support throughout these years. He has been a constant source of motivation, always encouraging me to step beyond my comfort zone and aim high. His presence and emotional support provided the strength I needed to overcome the challenges of this PhD journey, making this achievement possible.

Dedication

To Sajjad, Mom, and Dad

Table of Contents

Examining Committee ii

Author's Declaration iii

Statement of Contributions iv

Abstract v

Acknowledgements vii

Dedication viii

List of Figures xiv

List of Tables xxi

List of Abbreviations xxiii

1 Introduction 1

1.1 Motivation 4

1.2 Challenges for navigation among pedestrians 5

1.3 Research problem and objectives 7

1.4 Contributions 8

1.5 Thesis outline 10

2 Preliminary mathematical concepts **13**
2.1 Linear mixed effect models . . . 13
2.2 Unsupervised learning . . . 14
2.3 Reinforcement learning . . . 14

3 Related work **19**
3.1 Socially-aware navigation . . . 19
3.1.1 Evaluating social awareness of a navigation strategy . . . 22
3.2 Characteristics of motion behaviour in unstructured environments . . . 23
3.3 Decision-making in pedestrian-rich environment . . . 25
3.3.1 Reactive collision avoidance methods . . . 26
3.3.2 Proactive navigation methods . . . 27
3.3.3 Learning-based crowd navigation methods . . . 29
3.4 Prediction and planning integration in learning-based crowd navigation . . . 31
3.5 Pedestrian trajectory prediction . . . 34
3.5.1 Modelling pedestrian motion . . . 35
3.5.2 Interaction modelling . . . 38
3.5.3 Uncertainties in pedestrian trajectory prediction . . . 44
3.6 Planning under prediction uncertainty . . . 47
3.7 Summary of identified gaps . . . 49

4 Social norms in unstructured environments: a survey **51**
4.1 Research questions . . . 52
4.2 Method . . . 52
4.2.1 Procedure and measures . . . 54
4.2.2 Participants . . . 55
4.2.3 Data analysis approach . . . 55
4.3 Results . . . 56

4.4 Discussion 61
4.4.1 Limitations 63
4.5 Summary and conclusion 64

5 Characterizing unstructured environments 65
5.1 Research questions 66
5.2 Method 66
5.2.1 Existing trajectory datasets 67
5.2.2 Extracted trajectory features 68
5.2.3 Data analysis approach 72
5.3 Results and discussion 72
5.3.1 Limitations 78
5.4 Summary and conclusion 79

6 Pedestrian trajectory prediction 81
6.1 Problem formulation 83
6.2 Method 83
6.2.1 Model architecture 84
6.2.2 Interacting agent selection 86
6.2.3 Polar Collision Grid interaction feature 87
6.3 Experiments 88
6.3.1 Dataset 89
6.3.2 Evaluation metrics and baselines 91
6.3.3 Implementation details 92
6.4 Results and discussion 92
6.4.1 Quantitative evaluation 92
6.4.2 Qualitative evaluation 95
6.5 Uncertainty-aware Polar Collision Grid predictor 98

6.5.1 Evaluation metrics . . . 99
6.5.2 Results and discussion . . . 100
6.5.3 Limitations . . . 102
6.6 Summary and conclusion . . . 103

7 Vehicle motion planning **105**
7.1 Problem formulation . . . 107
7.2 Deep reinforcement learning method . . . 108
7.2.1 Network architecture . . . 109
7.2.2 Reward formulation . . . 112
7.3 Model predictive control method . . . 115
7.3.1 MPC with a hard constraint on minimum distance . . . 116
7.3.2 MPC with a soft constraint on personal space . . . 117
7.3.3 MPC with collision probability hard constraint . . . 117
7.4 Experiments . . . 118
7.4.1 Simulation environment . . . 118
7.4.2 Baselines and ablation models . . . 119
7.4.3 Evaluation metrics . . . 120
7.4.4 Implementation details . . . 120
7.5 Results and discussion . . . 121
7.5.1 Quantitative results . . . 121
7.5.2 Qualitative results . . . 126
7.5.3 Revisiting survey scenarios . . . 130
7.5.4 Limitations . . . 139
7.6 Summary and conclusion . . . 140

8 Conclusions **142**
8.1 Summary of research . . . 142
8.2 Contributions to original knowledge . . . 145
8.3 Limitations and future work . . . 146

References **150**

APPENDICES **174**

A Online survey **175**

A.1 Ethics clearance certificate 176

A.2 Online survey - Demographic questions 177

A.3 Online survey - Instructions for pedestrian role 178

A.4 Online survey - Instructions for passenger role 179

A.5 Online survey - Pre-survey questions 180

A.6 Online survey - Questions for pedestrian role 181

A.7 Online survey - Questions for passenger role 196

A.8 Online survey - Post-survey questions 211

B Letter of copyright permission **214**

B.1 Copyright permission for reuse of Fig. 5.1(a) and Fig. 5.1(b) 215

B.2 Copyright permission for reuse of Fig. 5.1(c) and 5.1(d) 216

B.3 Copyright permission for reuse of Fig. 5.1(e) 217

B.4 Copyright permission for reuse of Fig. 5.1(f) and Fig. 6.4 218

List of Figures

1.1 Examples of low-speed vehicles operating in unstructured/less structured environments shared with pedestrians. 2

1.2 Human-aware navigation lies as the intersection of two main field of motion planning and human-robot interaction. 3

1.3 Autonomous vehicles navigation pipeline. This thesis will focus on the local motion planner and the prediction module within the planner. 4

1.4 The design stages of the thesis, illustrating the flow and interconnections between the chapters. 10

1.5 The research questions covered in each chapter of the thesis. The main research questions and their corresponding sub-questions are cross-referenced to their definitions in the thesis for the reader's convenience. 12

2.1 Basic block diagram of Reinforcement Learning. 16

3.1 Main categories of proposed approaches for crowd navigation in the literature. 26

4.1 The results of the multiple choice question scenarios. (a) Pedestrian group's selected reactions about the cart approaching from behind in Scenario 3 (see Table 4.1). (b) Expectation of the passenger group in the same situation. "Ped" stands for pedestrian. (c) Participants' preferred side for the cart's operation in a corridor with an unidirectional flow of pedestrians as in scenario 4. (d) Participants' expectations of the cart when it gets behind a bidirectional flow of pedestrians as in scenario 7. 57

4.2 The results of the pre-survey and post-survey questions. 61

5.1 Images from the existing datasets used to characterize structured versus unstructured environments. These images are sourced either from the original papers introducing the datasets or from screenshots of dataset videos provided by the authors. See Appendix B for copyright permission. 69

5.2 Approach angle definition in a pedestrian-vehicle interaction. 70

5.3 The value of features extracted from pedestrians' trajectories in each dataset with 95 % confidence interval. 73

5.4 The value of features extracted from vehicles' trajectories in each dataset with 95 % confidence interval. 74

5.5 The value of features extracted from pedestrian-vehicle interaction in each data. 74

5.6 Datasets clustered in two groups based on the defined features. For each dataset, the portion of its data points in each cluster is reported. 75

5.7 The mean of the vehicle features within each cluster along with the 95% confidence intervals. 76

5.8 The mean of the interaction features within each cluster along with the 95% confidence intervals. 76

5.9 The mean of the pedestrian features within each cluster along with the 95% confidence intervals. 77

5.10 The process through which the results of this study can be used for developing trajectory prediction modules for autonomous vehicles (AVs). 78

6.1 The overall framework of our proposed pedestrian trajectory prediction. The embedded spatial feature, the pedestrian-pedestrian interaction and the pedestrian-vehicle interaction features are concatenated before being inputted into the LSTM module. 84

6.2 Time to Collision (TTC): the moment when the distance between two agents falls below a minimum acceptable threshold, assuming they continue following their current velocity profiles without any changes. 87

6.3 The process of building the interaction features for pedestrian i according to the neighbouring agent's time to collision (TTC) and approach direction. Eight discrete sectors around the target pedestrian are considered in this figure for the approach angles. Only neighbouring agents with a TTC lower than a threshold will be considered in the construction of the interaction feature. The interaction features for neighbouring vehicles and pedestrians are constructed separately. 89

6.4 The Hamburg Bergedorf Station site, where the HBS dataset was collected. The picture is sourced from [140] with proof of permission provided in Appendix B. 90

6.5 The original Social LSTM (a) selects interacting agents based on relative distance falling below a threshold and stores the hidden states of those interacting agents in a spatial grid, constructing the interaction tensor. In the modified version (b), we alter the original method's interacting agent selection by filtering out agents based on our Time-to-Collision (TTC) checking method while leaving the rest of the interaction tensor construction unchanged. 94

6.6 The predicted trajectories for our method compared to the baseline methods for a pedestrian-vehicle interaction scenario in the HBS test dataset. A marker with a larger size represents a later time step along the trajectory. The position of the vehicles at the last time step of the observation period is indicated with a square marker. Our method detects the target pedestrian A (indicated with start) to have interaction with vehicle B at this time step according to the TTC criteria. This identified interacting vehicle is indicated with a filled square marker at this time step. Considering this interacting vehicle, our method has predicted a trajectory that avoids the predicted collision while being close to the ground truth trajectory. 95

6.7 The predicted trajectories for our method compared to the deep-learning-based baseline methods in the HBS dataset. Our predictions overall are closer to the ground truth compared to the baselines. Case (c) is an example with a non-accurate prediction from our *Polar Collision Grid* method where a probable group interaction between Ped A and Ped B is mistakenly detected as a possible collision ending up with a prediction for Ped B that deviates to the right of the ground truth data. Interacting agents with a sample target pedestrian (shown with a star) are shown with a filled square marker. 96

6.8 Prediction of our Collision Grid method compared to Social LSTM for a pedestrian-pedestrian interaction scenario. Our method produces a prediction closer to the ground truth for pedestrians such as Ped C. This is done by detecting Ped C's interacting agents (Ped F and Ped B) more accurately while relying on time to collision for interaction modelling as opposed to the Social LSTM model that considers all agents closer to the target pedestrian than a threshold distance regardless of their moving direction which happens to be all the agents present in the scene here. Interacting agents for Ped C (target) are indicated with a filled square marker at time step t_{obs} while markers for non-interacting agents are unfilled. 97

6.9 Predicted trajectories of (a) the polar collision grid (PCG) model trained with only the negative log likelihood loss, compared to (b) the same model trained with an additional uncertainty loss, referred to as the Uncertainty-Aware PCG (UAW-PCG). The blue circles represent the 1σ standard deviation of the predicted distribution for the future position around the estimated mean (depicted in red). Agents' positions at the current time steps are marked with a star, and larger markers along the vehicle's trajectory highlight positions at later times. Here "Ped" and "Veh" stand for pedestrian and vehicle respectively. 102

7.1 We propose the use of predicted position and its associated covariance in the training of a deep reinforcement learning model with a novel reward function. Through this uncertainty-aware coupled prediction and planning approach, our model learns safe, foresighted, and smooth navigation behaviours among pedestrians in a shared space. 106

7.2 The crowd navigation problem represented as a graph with spatial and temporal edges. There are two types of spatial edges: one for pedestrian-autonomous vehicle interactions and another for human-human interactions. 108

7.3 Our integrated prediction and planning framework. The DRL network architecture is based on the model introduced in [122], with novel additions and modifications highlighted in orange. The inclusion of prediction uncertainties (covariance) is a key feature, integrated both as part of the observation state and in the formulation of our novel reward function. This covariance is obtained from the predictor module, which addresses pedestrian-vehicle interaction by incorporating the AV's state and an estimate of its future path as input. In our model, prediction covariance, alongside the states themselves, contributes to the human-human attention module. 109

7.4 Overview of the proposed framework. At each time step, current observations are input to the predictor, generating pedestrians' predicted positions and associated uncertainties over a horizon. These predictions inform collision probability calculations between the AV and each pedestrian's predicted path, contributing to the reward function of the DRL module. Utilizing this predictive data, the DRL motion planner is trained to generate a probability distribution of optimal actions, mitigating collision risk with both current and intended pedestrian paths. 110

7.5 Danger (a) is specified as a condition when the vehicle intrudes into the pedestrian's personal space (PS), shown as the yellow circle. Collision (b) occurs when the physically occupied spaces of the two agents, modelled as pink circles, overlap. 113

7.6 Mean and 95% confidence interval of difference model's path length over the test set scenarios. 124

7.7 Training convergence plots for varying goal achievement reward values while maintaining a constant collision penalty coefficient. 125

7.8 Comparison of two models trained with and without the speed-dependent (SD) danger penalty on an identical test scenario. The figure illustrates an autonomous vehicle (in black) navigating through a crowd of pedestrians (in red) towards a goal indicated by a star. It depicts the trajectories of both the AV and the pedestrians for the past 6 frames. A dashed yellow line encircles the current position of each pedestrian, representing their personal space. 127

7.9 Comparison of the performance of our DRL-based uncertainty-aware PCG model (UAW-PCG) with the uncertainty-aware MPC model on the same scenario as in Fig. 7.8. The mean trajectory and 1σ standard deviation of the pedestrians' future paths over the next 6 frames are depicted by yellow dots and blue ellipses, respectively. The AV's maximum sensor detection range is shown as a dashed black circle around its current position, with predictions being made only for pedestrians within this range. 128

7.10 Comparison of the autonomous vehicle's trajectory behaviour in an identical test scenario. This figure presents the trajectory of the AV using our proposed Deep Reinforcement Learning (DRL) model, which accounts for prediction uncertainties (UAW-PCG pred), against the baseline model that does not consider such uncertainties (PCG pred). The trajectories are illustrated over three consecutive frames. 129

7.11 The back interaction scenario observed in the CITR dataset [206] resembles scenario #5 from the survey, while the lateral interaction scenario aligns with scenario #6. The screenshots from the CITR dataset were extracted from its existing video recordings [210]. 131

7.12 Performance of our DRL model compared to the MPC model on a scenario from the CITR dataset, similar to the back approach scenario #5 of the online survey (Table 4.1). While the MPC model attempts to reach the goal by overtaking the crowd of pedestrians through heading changes, our model follows the pedestrian flow, aligning with the preferred behaviour indicated by the survey results. Despite its overtaking attempts, the MPC model could not reach the goal, resulting in a timeout. In contrast, our model successfully reached the goal in a perceived safe manner and within a reasonable time. 133

7.13 Performance of our DRL model compared to the MPC model and the human-driven trajectory on a lateral scenario from the CITR dataset, similar to the 90-degree approach scenario #6 of the online survey (Table 4.1). The collision avoidance strategy of the MPC involves more heading changes, while our DRL method adjusts speed, which is perceived safer according to the online survey results. This approach is also more in line with the human-driven trajectory but more cautious. 134

7.14 Performance of our DRL model in a custom scenario replicating Scenario #3 of the survey, using the ORCA model for pedestrian behaviour under two conditions. (a) When the robot is visible to pedestrians, resembling the case where pedestrians react to the vehicle by opening up a space for it, which was the most frequently selected pedestrian behaviour in the survey. In this case, the AV understands and overtakes the pedestrian. (b) When the robot is invisible to pedestrians, resembling the case where pedestrians do not react to the AV and follow their own paths. In this scenario, the AV will follow pedestrian A, complying with the expected behaviour of the AV from the survey. This demonstrates the AV's ability to adapt to different pedestrian behaviours. 136

7.15 Performance of our DRL model in a custom scenario replicating Scenarios 5 and 7 of the survey, using the ORCA model for pedestrian behaviour. In both cases, the AV follows the flow of pedestrians to reach its goal, as identified from the survey results. 137

7.16 Performance of our DRL model in custom-built scenarios replicating Scenarios 6 and 9 from the survey, using the ORCA model for pedestrian behaviour. In Scenario 9, random goals are assigned to the pedestrian at each time step to mimic the erratic behaviour of a child who eventually moves to one side. In Scenario 6, the AV adjusts its speed to avoid collisions, which is perceived as safer behaviour according to the survey results compared to changing its heading. When confronted with a child with high trajectory uncertainty, the AV decelerates and almost comes to a stop, resuming movement once the child reaches a safe distance, as expected based on the survey results. . 138

List of Tables

4.1 The main scenarios of the survey. Upper-case letters are behaviours for which a separate questionnaire was completed by the participants. The lower-case letters show the multiple-choice questions. 53

4.2 Designed questionnaire for scenario 5 with a continuous Likert scale from strongly disagree to strongly agree. 54

4.3 Thematic analysis results of pedestrians' reasoning in Scenario 3 where the cart is positioned behind a pedestrian with their left side occupied by other pedestrians and their right side by pylons. The table reports the frequency of mentions for each identified theme in the reasons provided for the chosen action. "Moving to left/right" are sub-categories under the broader choice of clearing the cart's path. 58

4.4 Thematic analysis results on passengers' reasons for choosing to stay behind person A in scenario 3. 59

4.5 Thematic analysis results on participants' reasons for their preferred operational side of the cart in a corridor in scenario 4. 60

5.1 The generalized linear model leading to best model fit according to the AIC criteria (AIC=72.83). "Ped" stands for Pedestrian. 75

6.1 Prediction errors for baseline methods compared to the proposed method and its variants on the HBS dataset. 93

6.2 The effect of our proposed interacting agent selection based on collision checks (used as a filter) on the performance of the Social LSTM model reported on HBS dataset. 94

6.3 Pedestrian trajectory prediction results. Bolded is the best. 100

7.1 DRL navigation and baseline results. Bold is the best. Underline is the second best. (SR: Success Rate, CR: Collision Rate, TR: Timeout Rate) . 121

7.2 One step average computational time of different models. 122

7.3 Sensitivity analysis of the reward value for goal achievement (r_g). The default value is denoted by (*). 125

7.4 Performance of our DRL navigation algorithm and baseline methods on CITR dataset. Bold is the best. 132

List of Abbreviations

ADE Average displacement error 91, 93, 100, 101, 144

AIC Akaike information criterion 75

AV Autonomous Vehicle v, 1–10, 20–22, 25, 33, 34, 48, 49, 51, 52, 64, 65, 81, 105, 107–116, 118–120, 122, 123, 125, 126, 130, 135, 139–146, 148, 149

CA Cellular Automota 35, 46

CITR Control and Intelligent Transportation Research Lab 131, 132, 139

CNN Convolutional Neural Network 36, 41

CR Collision Rate 120, 121, 125, 132

CVAE Conditional Variational Autoencoder 37, 45, 46

D-CBF Dynamic Control Barrier Function 47

DQN Deep Q-Networks 15

DRL Deep Reinforcement Learning v, vi, 9, 11, 29, 30, 32–34, 48–50, 98, 101, 105–108, 111, 112, 115, 118–124, 126, 130, 132, 135, 139–141, 143–149

DUT Dalian University of Technology 67, 68, 72, 78

FDE Final displacement error 91, 93, 100, 101, 144

fps frame per second 120

GAE Generalized Advantage Estimation 16

GAN Generative Adversarial Network 37, 45, 46

GCN Graph Convolutional Networks 30, 36, 43

GLM generalized linear model 74, 75, 79

GLMM Generalized linear mixed-effect model 14, 75

GMM Gaussian mixture model 45

GNN Graph Neural Network 30, 36, 42, 43

GRU Gated Recurrent Unit 36, 43, 108, 111

GT ground truth 99, 107, 119

HBS Hamburg Bergedorf Station 67, 68, 90, 92, 102, 107, 118, 120, 139, 148

HC Hannover Campus 67, 68

HE Heading error 91, 93

HIT Human Intelligence Task 55

HRI Human-Robot Interaction 2, 19

inD Intersection Drone Dataset 67, 72, 90

IRL Inverse Reinforcement Learning 29

LMM Linear mixed-effect model 13, 14, 55, 73

LR Linear Regression 91

LSTM Long Short-Term Memory 9, 29, 36, 37, 41, 43, 84–86, 88, 91, 92

MD Mahalanobis distance 99, 114, 117

MDP Markov Decision Processes 14, 15, 107

MHD Modified Hausdorff Distance 91, 93

MLP Multi-layer perceptron 29, 30

MPC model predictive control vi, 47, 48, 115–117, 119–122, 124, 126, 132, 141, 149

NLL negative log-likelihood 86, 98–101, 119

ORCA Optimal Reciprocal Collision Avoidance 26, 31, 32, 106, 112, 118, 132, 135, 139, 147

PCG Polar Collision Grid 85, 88, 93, 99–101, 103, 119, 120, 122, 123, 132

POMDP Partially Observable Markov Decision Process 29, 48

PPCG Pedestrian Polar Collision Grid 88

PPO Proximal Policy Optimization 16, 17, 33, 108

PS Personal Space 112–114, 116, 117, 119, 120, 123

PSR predictive state representation 32

ReLU rectified linear unit 85, 110, 111

RL Reinforcement Learning 13, 15, 16, 43, 149

RMS-prop Root Mean Squared Propagation 92

RMSE root mean square error 91

RNN Recurrent Neural Network 30, 31, 36, 43

RVO Reciprocal Velocity Obstacles 26

SD speed-dependent 113, 119, 122

SDD Stanford Drone Dataset 67, 72, 73, 75, 78–80, 89, 90

SE Speed error 91, 93

SFM Social Force Model 26, 27, 35, 37–40, 44

SR Success Rate 120, 121, 125, 132

ST-graph spatio-temporal graph 30, 31, 144

TR Timeout Rate 120, 121, 125, 132

TTC time-to-collision 40, 82, 86–88, 93, 96, 103, 143, 144, 146

UAW-PCG Uncertainty-Aware Polar Collision Grid 99–101, 104, 119, 120, 122

V-learning Value learning 30, 32, 48

VAE variational auto-encoders 45

VO Velocity Obstacle 26

VPCG Vehicle Polar Collision Grid 88

Chapter 1

Introduction

With the continuous evolution of technology, an increasing number of robots are being integrated into various aspects of our daily lives. Delivery robots now traverse sidewalks, service robots assist in hospitals, supermarkets, airport terminals, and restaurants, while prototypes of AVs undergo testing on roads. These deployments hold the promise of enhanced task efficiency, while AVs also aim to improve safety by mitigating human errors, which contribute to 94% of accidents [62], and providing enhanced accessibility to mobility for all. Across these applications, these robots and AVs must navigate dynamic environments shared with humans (e.g., pedestrians, human-driven vehicles), necessitating a priority on safe and appropriate interaction throughout their journey.

While significant advancements in autonomous vehicle technology have enabled the testing of autonomous trucks on highways and robotaxis being deployed in few cities in US and China without human safety operators, the capacity of current AVs to navigate complex urban environments necessitating high interaction with pedestrians and human-driven vehicles remains notably limited [43].

On regular roads, efforts are made to minimize vehicles interactions with vulnerable road users like pedestrians through spatial and temporal segregation of their operational zones via traffic lights and road markings. However, such interactions become unavoidable during the operation of low-speed vehicles serving as mobility aids in pedestrian-dominated, unstructured environments such as airport terminals and shopping malls (see Fig. 1.1(a) for an example), or off-road campus environments (see Fig. 1.1(b)). These interactions for an autonomous vehicles operation in such environments carry increased safety concerns compared to smaller service robots, owing to the larger size and higher speed of passenger-carrying vehicles.

Similarly, pedestrian-vehicle interaction in shared spaces—a concept in city planning where road curbs and surface markings are removed and traffic rules are relaxed to encourage negotiation [45, 147]— can become diverse and complex. This shared space approach aims to enhance traffic management and strike a balance between mobility needs and urban life quality (see Fig. 1.1(c) for an example).

Navigating autonomous vehicles in such environments presents challenges that highlight the current limitations of AVs in interacting with other traffic participants. Efficiently negotiating space among crowds of pedestrians while prioritizing safety for both pedestrians and passengers is critical for AVs operating in such settings. Therefore, this thesis aims to propose a navigation algorithm designed for safe and efficient operation of low-speed autonomous vehicles among pedestrians in shared or unstructured environments.

The challenge of autonomous agents navigating among crowds of pedestrians is commonly referred to in the literature as *human-aware navigation*, *social navigation* or *crowd navigation* [109, 171]. Human-aware navigation lies at the intersection of robot motion planning and Human-Robot Interaction (HRI), focusing on the interactions that occur during navigation [109]. Therefore, both navigation efficiency (e.g., path length, navigation time) and social considerations (e.g., perceived safety, acceptability) should be taken into account as shown in Fig. 1.2.

However, most of the existing work on this topic focuses on the social navigation of small mobile robots within crowds of pedestrians, rather than addressing the application of low-speed AVs in the same context. Unlike small mobile robots, the operation of AVs in such environments is more safety critical due to their higher speed and larger size as passenger-carrying vehicles. Consequently, this influences pedestrian behaviour around these vehicles

(a) Airport buggies (b) Vehicles on off-road campus environment (c) A shared space [199]

Figure 1.1: Examples of low-speed vehicles operating in unstructured/less structured environments shared with pedestrians.

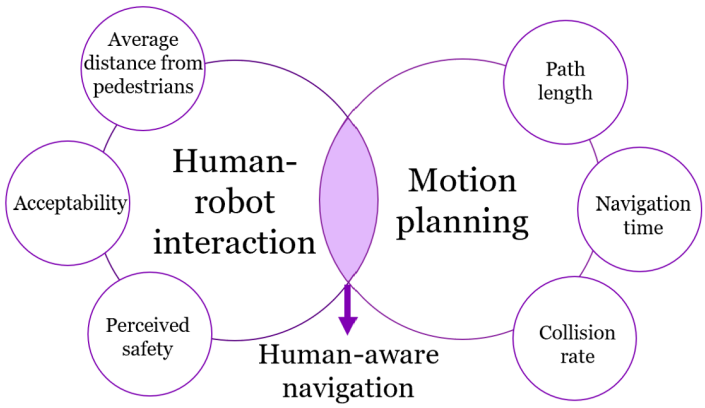

Figure 1.2: Human-aware navigation lies as the intersection of two main field of motion planning and human-robot interaction.

and their perceptions of the autonomous agent. These distinctions highlight the need to study the design requirements and develop crowd navigation algorithms specifically tailored for AVs among pedestrians in shared and unstructured environments. This unique aspect justifies dedicated attention, which is the primary goal of this thesis.

In contrast, in the field of autonomous vehicles, there are studies on how AVs interact with pedestrians. However, these studies have typically been confined to scenarios involving road crossings, where the interaction occurs at a 90-degree approach angle. This is a limited perspective because interactions between pedestrians and AVs in unstructured environments or shared spaces can be much more diverse, occurring from multiple directions and with different expectations. For instance, in relaxed traffic rule settings like shared spaces without defined priorities, or where pedestrians have dominance, such as in indoor environments (e.g., airport terminals), behaviours can vary widely. This includes individuals strolling leisurely, others rushing to a destination, or the presence of freely roaming children—all impacting risk tolerance levels and behaviour variety. Therefore, the operation of low-speed autonomous vehicles in such unstructured or less-structured environments has been identified as worthy of separate study within this thesis.

The overall navigation pipeline of an autonomous vehicle consists of three main modules: perception, planning, and control (see Fig. 1.3). The perception module utilizes the vehicle's sensor input to handle localization, object detection, and tracking. This information is then utilized in the planner module, which serves as the decision-making component within the pipeline.

The planner module can be further divided into several sub-modules, including the

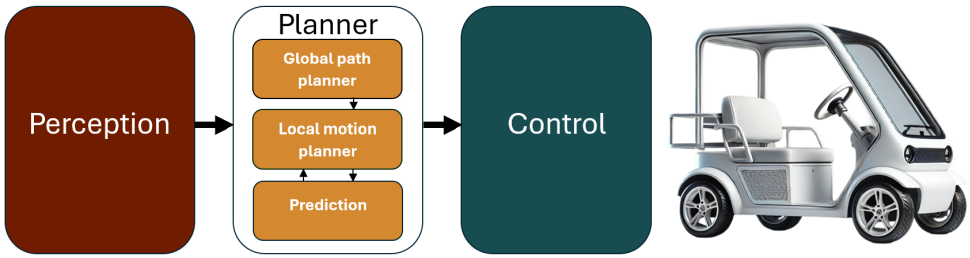

Figure 1.3: Autonomous vehicles navigation pipeline. This thesis will focus on the local motion planner and the prediction module within the planner.

global path planner, local motion planner, and prediction module. The global path planner, also known as the mission or route planner, determines the longer-term path based on the map and static obstacles in the environment. The output of the global planner provides key waypoints for reaching a destination over the long term.

The local motion planner is responsible for collision avoidance with dynamic obstacles and generates a trajectory considering the vehicle's kinematics and dynamics to follow the waypoints generated by the global path planner. Collision avoidance with dynamic obstacles in the local planner is facilitated through the prediction module, which collaborates closely with the local motion planner using perception data.

The output of the local motion planner, such as speed and heading angle at each time step, is then transmitted to the motion control module to produce lower-level actuator commands, such as throttle or steering angle, for the vehicle system.

In this thesis, our focus is solely on the planner module, specifically proposing a local motion planner and prediction module that collaborate closely to produce safe and efficient motion behaviour for the AV when navigating among pedestrians. Throughout this study, we assume that the positions of the vehicles and pedestrians have already been provided by the perception module in the form of x and y positions within a global frame.

1.1 Motivation

Consider an autonomous airport buggy operating in a busy terminal, tasked with transporting a few late arriving passengers from terminal A to E to catch their connecting flight. Given the urgency, the vehicle must reach the destination promptly while ensuring it avoids colliding with pedestrians in the crowded airport.

Overemphasizing safety could lead to overly cautious behaviour, where the vehicle stops frequently for pedestrians, even when they might have been willing to cooperate by making space. This overly conservative approach compromises efficiency and could cause passengers to miss their connecting flight. Conversely, driving too aggressively—such as sudden braking, unpredictable acceleration, or passing too closely—can startle people, make them uncomfortable, and increase the risk of collisions.

Therefore, there exists a balance between safety and efficiency, where safety encompasses not only physical avoidance of collisions but also ensuring psychological or perceived safety for pedestrians and passengers. It's important to avoid causing discomfort and behave in a natural, predictable manner that aligns with social norms to gain acceptance of automated systems among pedestrians.

It is important to emphasize that when we speak of striking a balance between safety and efficiency, we are not suggesting sacrificing safety for efficiency. Rather, we aim to avoid overly conservative behaviours that overlook the cooperative nature of pedestrians and the impact of vehicle presence on pedestrian behaviour, which can encourage pedestrians to create space willingly.

All these considerations of safety and efficiency stem from recognizing that pedestrians are not simply dynamic obstacles; they are decision-making agents who influence the decisions of nearby autonomous agents and are, in turn, influenced by the actions of these agents, which can be seen in changes to their trajectory behaviour.

Taking into account these mutual interaction effects for space negotiation purposes is crucial for achieving a balance between safety and efficiency in the design of human-aware and socially compliant navigation algorithms. Developing a deep understanding of these interaction effects remains a challenging topic in AV navigation among other decision-making agents.

In this thesis, we propose methods to more effectively capture these interaction effects to develop a safer and more efficient navigation algorithm for AV operation among pedestrians in shared, unstructured spaces.

1.2 Challenges for navigation among pedestrians

Autonomously navigating among other decision-making agents, in our case pedestrians, presents three main challenges:

1) **Ensuring Perceived Safety and Acceptability**: The first challenge involves accounting for the perceived safety and acceptability of the AV's motion behaviour. Since autonomous vehicles are designed for humans, considering their needs and expectations is critical. Many motion planning algorithms focus solely on path efficiency and physical collision avoidance. However, even if these algorithms are proven safe, deploying them near pedestrians may not be perceived safe if the vehicle gets too close or exhibits unnatural or unpredictable behaviour that does not align with social norms and people's expectations. Therefore, understanding people's expectations and incorporating considerations for perceived safety into the design is essential for developing an appropriate navigation algorithm.

2) **Managing Mutual Interaction Effects**: The second challenge involves considering mutual interaction effects, which are essential for balancing safety and efficiency. While autonomous agents account for the presence of pedestrians when planning their trajectories, they must also consider how their own actions can influence people's behaviour. This interaction often leads to cooperative collision avoidance and space negotiation. Addressing these mutual interaction effects requires the ability to predict other agents' behaviour based on their observed state and social cues. While human drivers naturally possess these skills, replicating them in AVs presents significant challenges. Currently, because of limited understanding of human behaviour, social cues, and mutual interaction effects, AVs often adopt overly conservative behaviours to prioritize safety, which can lead to unnecessary stops and inefficient navigation [167]. These conservative behaviours can confuse humans and be perceived as unpredictable by human drivers, leading to traffic congestion and accidents. In fact, 57% of AV crashes on roads are rear-end collisions caused by human drivers reacting to the AV's unexpected behaviour [173].

3) **Dealing with Uncertainties in Pedestrian Behaviours**: The third challenge involves managing inherent uncertainties in dynamic environments, particularly concerning pedestrians' behaviours, which are often influenced by unobservable states like intent. Effectively capturing interaction effects requires reasoning about the near future and predicting how the interactive scene will evolve. However, due to the uncertainties in other agents' behaviour, predictions become probabilistic. Therefore, balancing safety and efficiency hinges on appropriately considering and adapting to these uncertainties.

In this thesis, we focus on tackling each of these challenges with our proposed methods aimed at addressing them effectively. It is worth noting that the scope of this work is

limited to a single vehicle operating among crowds of pedestrians. Consequently, this thesis does not consider scenarios involving multiple vehicles, which would introduce the added complexity of vehicle-vehicle interactions.

1.3 Research problem and objectives

According to the challenges outlined in the previous section, we define our research questions as follows:

RQ1 What considerations are necessary for autonomous vehicle operation in pedestrian-rich, unstructured environments, given that people's expectations of the autonomous vehicle and the perceived safety of its behaviour may differ between structured and unstructured environments?

Given the less explored operation of AVs as mobility aids in pedestrian-crowded unstructured environments, we aim to study the design requirements for such systems. We will explore people's expectations of AV operation in these environments, focusing on motion behaviours that affect perceived safety. Furthermore, we will investigate how these behaviours can be incorporated into the design of the AV's navigation algorithm.

RQ2 What potential differences might exist in the motion behaviour of pedestrians and vehicles between structured and unstructured environments, and how should these differences inform the design of navigation algorithms for unstructured settings?

To further study the design requirements for an AV's operation in an unstructured environment, we aim to better understand such settings by examining potential differences in pedestrian and vehicle trajectories between unstructured and structured settings and how these differences should be considered in designing the navigation algorithm.

RQ3 How can interaction effects between different agents during navigation be effectively and computationally efficiently modelled in the decision-making algorithm of an autonomous vehicle?

Given the importance of modelling interaction effects for safe, natural, efficient, and intelligent motion behaviour, we seek to investigate effective ways of capturing

these interactions. In a shared environment, there exist both pedestrian-pedestrian and pedestrian-vehicle interactions, which differ in nature and influence. We aim to explore effective methods for encoding these interaction effects on both pedestrians' future trajectory behaviour and the AV's decision-making and motion planning.

RQ4 How can uncertainty in pedestrians' behaviour be captured and integrated into the decision-making algorithm of an autonomous vehicle navigating among a crowd of pedestrians?

Given the inherent uncertainty in predicting pedestrians' behaviour from an AV's perspective and the importance of prediction in motion planning for producing farsighted behaviours, we investigate how these prediction uncertainties can be accounted for in the selected motion planning algorithm. Our goal is to encourage safe yet efficient behaviour despite these uncertainties.

1.4 Contributions

Throughout this thesis, by addressing the specified research questions, we make the following main contributions to the body of knowledge:

1. **Developing crowd navigation algorithm for low-speed AVs operating in unstructured shared spaces.** While most of the literature focuses on either crowd navigation of small mobile robots in unstructured environments or the interaction of AVs with pedestrians in structured road environments at crossing areas, this thesis develops a crowd navigation framework designed for the operation of low-speed AVs among pedestrians in unstructured environments. Key considerations include the diversity of interactions, higher safety criticality due to the larger size and higher speed of the vehicle compared to small mobile robots, people's different expectations of the AV's behaviour in such settings, different pedestrian movement patterns and trajectory behaviours, and the datasets representing such environment.

2. **Identifying social norms for AV operation in unstructured environments.** Due to the novelty of the application, it is unclear if social norms from pedestrian-pedestrian interactions or vehicle operation on roads will transfer to such settings. Therefore, we investigated people's expectations of AV behaviour in these unstructured environments through a survey. This process identifies perceived safe behaviours and some social norms the vehicle should comply with in specific scenarios,

guiding the development of our algorithm in subsequent steps. This contribution specifically addresses the first research question (**RQ1**).

3. **Identifying differences in motion behaviours in structured vs. unstructured environments.** By investigating the trajectory features of different datasets, we identified characteristics that vary significantly among datasets and can be attributed to the nature of the environments in which they were captured, whether structured or unstructured. This quantitative investigation towards defining these environments is a novel approach, as opposed to the qualitative definitions found in the literature. It provides specific considerations for developing navigation algorithms for a target environment in line with addressing our second research question (**RQ2**).

4. **Designing an interaction-aware data-driven pedestrian trajectory prediction model.** Our proposed Long Short-Term Memory (LSTM) network is designed for predicting pedestrian trajectories in a shared space with vehicles, considering their interaction with both vehicles and other pedestrians separately. We propose the use of more engineered features and heuristics for effectively selecting interacting agents and encoding their effect on pedestrians' future trajectories, addressing the third research question (**RQ3**) for the trajectory prediction component of our navigation framework. Moving beyond reliance on just raw position and velocity data as seen in the literature, we demonstrate how these engineered features can guide the neural network in producing more accurate predictions. We further propose techniques for preventing overconfident predictions and generating more accurate estimations of prediction uncertainties as part of our response to the fourth research question (**RQ4**).

5. **Designing an uncertainty-aware DRL based crowd navigation algorithm.** We developed a DRL motion planner that integrates with the pedestrian trajectory prediction to effectively use prediction information, considering the uncertainties in these predictions in a novel way to learn safer, more efficient, and farsighted behaviours. The novel integration of the prediction and planner in our overall navigation architecture captures the mutual interaction effects, considering both the effect of pedestrians' presence on the vehicle's decisions and, vice versa, the effect of the AV's presence on the predicted future positions of pedestrians. This addresses the third research question (**RQ3**). By incorporating prediction uncertainty into the decision-making process of the DRL motion planner, we minimize the collision probability with pedestrians over a prediction horizon, addressing the fourth research question (**RQ4**).

6. **Constructing a realistic pedestrian trajectory simulation environment in the presence of vehicles.** We developed a simulation environment where the scenarios and pedestrians' motion behaviours are based on realistic trajectory data of pedestrians in the presence of vehicles in a shared space. We train and evaluate our navigation algorithm in this simulation environment, allowing us to learn behaviours that are more likely to be successful when deployed in the real world, as the simulation environment closely resembles possible real-world pedestrian behaviours. Our learned policy is compared with other state-of-the-art methods and evaluated in both realistic scenarios of this simulation environment and other custom-generated scenarios, using both social and efficiency metrics.

1.5 Thesis outline

This thesis comprises eight chapters. After reviewing related work, we first establish a clear understanding of the design requirements and the operational environment, with its specific considerations. The design requirements study unfolds in two stages: the first focuses on the AV's expected behaviour in unstructured environments in form of social norms (Chapter 4), while the second explores pedestrian behaviour in such settings, which must be incorporated into the design (Chapter 5). Following this two-stage design requirements study, we turn to the algorithmic development aspects of the navigation framework, which involves two stages: pedestrian trajectory prediction (Chapter 6) and vehicle motion

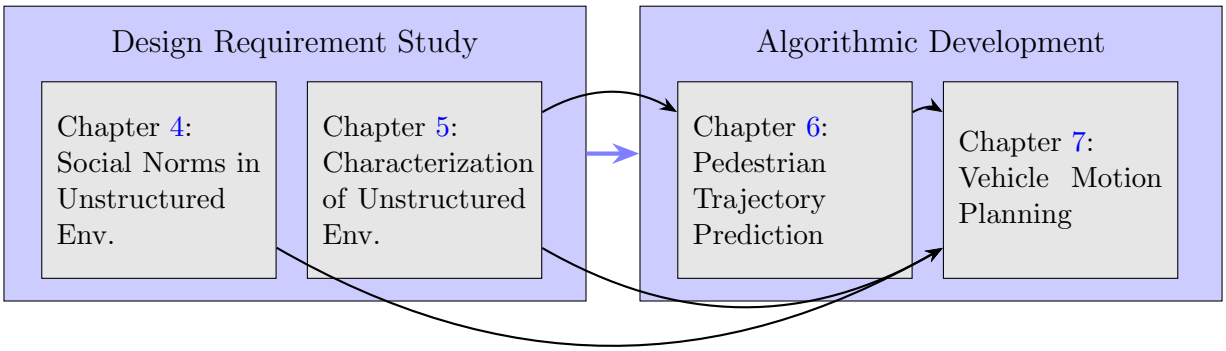

Figure 1.4: The design stages of the thesis, illustrating the flow and interconnections between the chapters.

planning (Chapter 7). The motion planning chapter integrates the findings from the design requirements study and the model built for prediction to develop our overall proposed algorithm. This connection across chapters, which guides the narrative of the thesis, is illustrated in Fig. 1.4. Additionally, figure 1.5 outlines the links between our research questions and the chapter in which they are addressed, serving as a guide throughout the thesis. Below is a summary of the thesis chapters:

Chapter 1 is the current chapter which introduces the motivation for the research, the challenges of the problem at hand, our research problem and objectives in the form of our main research questions, and our main contributions.

Chapter 2 introduces the fundamental mathematical concepts and methods used in the thesis.

Chapter 3 provides background and literature review of related topics, including socially-aware navigation and collision avoidance methods, trajectory prediction models, approaches for integrating prediction and planning in a unified framework, and consideration of uncertainty in motion planning. Throughout this review, we identify gaps that our proposed methods aim to fill. This chapter draws partially from our published review paper in [J2], and the related work sections of our other papers [J1], [C1], [C2], and [C3].

Chapter 4 presents our first step in understanding the navigation behaviour's required specifications by studying people's expectations and identifying perceived safe behaviours for the vehicle from a human-centred perspective. This chapter is based on our research published in [C3].

Chapter 5 describes our approach to better understand the unstructured nature of the environment in which the vehicle operates, focusing on specific considerations and differences in pedestrian trajectory behaviour between structured and unstructured environments. These insights inspire the design of the navigation algorithm in the subsequent chapters. This chapter is based on our published work in [C2].

Chapter 6 details our developed data-driven pedestrian trajectory prediction model as part of our designed navigation framework, emphasizing effective interaction modelling and accurate uncertainty estimation of the future pedestrian positions. This chapter is based on our research published in [C1] and [J1].

Chapter 7 presents the design of our DRL motion planner and its integration with the previously developed trajectory predictor into an overall uncertainty-aware navigation framework. It includes evaluations of our method in a simulation environment based on different metrics and the satisfaction of the previously identified expected behaviours. This chapter is based on our published journal paper [J1].

Chapter 8 concludes the thesis by summarizing the research, listing the contributions, and acknowledging the limitations of the work, which motivate future research directions.

Appendix A contains the ethics clearance certificate for the survey conducted in Chapter 4, as well as the survey questions.

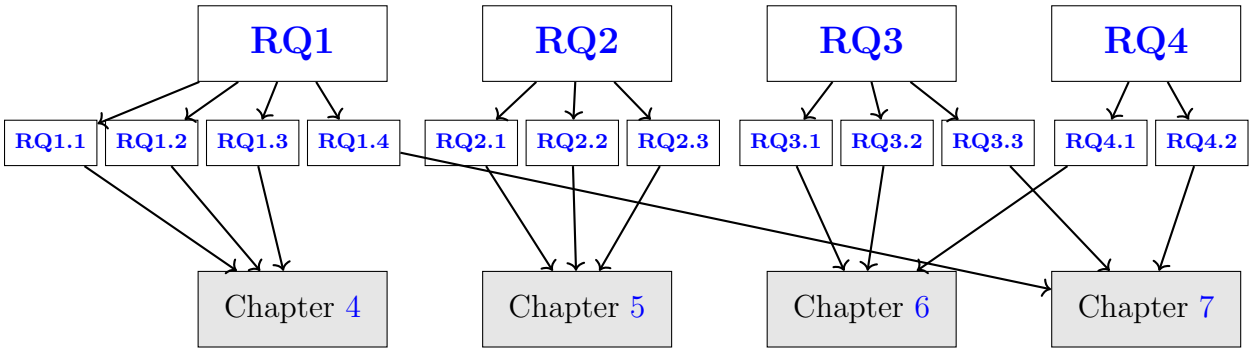

Figure 1.5: The research questions covered in each chapter of the thesis. The main research questions and their corresponding sub-questions are cross-referenced to their definitions in the thesis for the reader's convenience.

Chapter 2

Preliminary mathematical concepts

In this chapter, we briefly introduce some of the fundamental concepts the rest of the thesis is built upon. We first shortly provide a gentle introduction of the Linear mixed effect models used as a statistical analysis tool in Chapters 4 and 5. Next, we describe the k-means clustering method as part of the unsupervised machine learning techniques used in Chapter 5. Finally, we cover some of the main concepts in Reinforcement Learning (RL), which is the method used for our motion planner, in Chapter 7.

2.1 Linear mixed effect models

Linear mixed-effect models (LMMs) are extensions of linear regression models and are statistical models that consider both fixed and random effects. Fixed effects consist of levels that are the main interest of the experiment and remain the same when repeating the experiment. In contrast, random effects are not the primary interest and arise from the random selection of subjects, possibly from clusters of related units (potentially from higher levels not explicitly modelled). LMMs are particularly useful when there are repeated measures on the same unit in an experimental setting or when data is correlated due to the hierarchical grouping of subjects or units. The random effects in the model account for general variability among subjects and the correlations between them, enabling a variety of correlations to be modelled.

In general, a linear mixed effect model can be represented as [200]:

$$\mathbf{y}_i = X_i\beta + Z_i\mathbf{b}_i + \mathbf{e}_i, i = 1, 2, ..., n \tag{2.1}$$

In this equation $\mathbf{y}_i = (y_{i1}, y_{i2}, ..., y_{in_i},)^T$ is the n_i repeated measurements collected from subject i for the response variable y, $\beta = (\beta_1, ..., \beta_p)^T$ is the unknown vector of fixed effects, $\mathbf{b}_i = (b_{i1}, ..., b_{iq})^T$ is the unknown vector of random effects. Matrix X_i of size $n_i \times p$ and Z_i of size $n_i \times q$ are known design matrices, and $\mathbf{e}_i = (e_{i1}, ..., e_{in_i})^T$ is the unknown vector of random errors of the repeated measurements of individual i.

The fixed effects β are considered population-level parameters that are the same across all subjects, while the random effects $\mathbf{b}_i$ are individual-level parameters accounting for between-individual variation in the population level parameters.

Generalized linear mixed-effect models (GLMMs) are extensions of LMMs that also accommodate non-normal distributions of the response variable, such as binomial, Poisson, or gamma distributions.

2.2 Unsupervised learning

Unsupervised machine learning algorithms perform inference using input data without any provided labels or outputs. K-means clustering is a popular and simple algorithm within the unsupervised learning category that searches for a fixed number of groups or cluster k in a dataset, such that the datapoints in each group are more similar to each other than the ones from separate clusters. The goal of the k-means clustering is, therefore, to predict the centroid of these clusters and the label for each data point specifying its assigned cluster. The process is done by repeatedly alternating between (a) assigning points to a cluster given the current centroid and (b) updating the position of the centroid based on the current data point assignments until convergence.

2.3 Reinforcement learning

Sequential decision-making problems can be formulated as Markov Decision Processess (MDPs). An MDP is defined by the tuple (S, A, P, R, γ). In this formulation, at each discrete time t, $s \in S$ is the current state of the system. The agent takes action $a \in A$ by following a policy $\pi(a|s)$ at state s. Through this action a at state s, the agent transitions to the next state s' according to the state transition probability function $P(s'|s, a)$ and receives a scalar reward r as feedback for the action taken. The optimality of the sequence of actions, resulting in a trajectory, is measured through the return $R = \sum_{k=0}^{\infty} \gamma^k r_k$, which is the accumulated future reward discounted by $\gamma \in [0, 1)$, balancing the importance of

immediate vs. long-term rewards. A value function numerically specifies how good it is for the agent to be in a given state (State Value Function $V^{\pi}(s)$) or to execute a given action at a given state (State-Action Value Function or Q-function $Q^{\pi}(s,a)$), in terms of the expected return the agent will receive by following a policy π starting from that state or state-action pair [179].

$$\begin{aligned} V^{\pi}(s) &= \mathbb{E}_{\pi}[R|s_0 = s] \\ Q^{\pi}(s,a) &= \mathbb{E}_{\pi}[R|s_0 = s, a_0 = a] \end{aligned} \tag{2.2}$$

In RL, which follows the MDP formulation, an agent learns by interacting with a simulated environment and collecting experiences in the form of tuples (s, a, r) as shown in Fig. 2.1. The goal in RL is to learn an optimal policy, π^*, for the agent that maximizes the expected return, or in other words, the value function.

The optimal value function ($V^*(s) = \max_{\pi} V^{\pi}(s)$) associated with the optimal policy satisfies the Bellman equation (Eq. 2.3), based on which the optimal policy can be derived according to Eq. 2.4. As specified in this equation, deriving the optimal policy from a learned value function requires knowledge of the state transition function of the environment and the resulting state of the environment s' after executing the action a at state s.

$$V^*(s) = \max_{a} \sum_{s',r} P(s', r|s, a)[r + \gamma V^*(s')] \tag{2.3}$$

$$\pi^* = \arg\max_{a} \sum_{s',r} P(s', r|s, a)[r + \gamma V^*(s')] \tag{2.4}$$

Model-based RL methods attempt to use the agent's collected experience to estimate the state transition function and then use planning to derive the optimal policy based on the learned model. On the other hand, model-free methods directly learn an estimation of the value function, Q-function, the policy function, or a combination of them.

Within model-free methods, there are two main categories: Q-learning (action-value methods) and policy optimization methods. In Q-learning methods, the action value function is learned, and the best action is selected based on the estimated Q-function ($a = \arg\max_{a} Q_(s, a)$). Deep Q-Networks (DQN) is a popular method in this category [132]. On the other hand, policy optimization methods learn a parameterized policy directly, allowing action selection without necessarily requiring a value function. A value

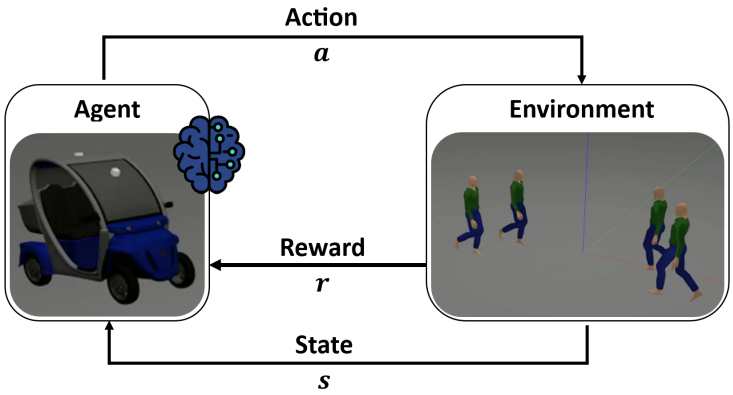

Figure 2.1: Basic block diagram of Reinforcement Learning.

function can still be learned alongside the policy function to guide the learning process, but it is not essential. In these methods, the parameters of the policy $\pi(a|s,\theta)$ are often learned using policy gradient methods for maximizing a scalar performance measure $J(\theta)$. Therefore, the policy parameters θ, are updated based on the gradient of $J(\theta)$ respect to those parameters according to Eq. 2.5. Examples of policy optimization methods include Asynchronous Advantage Actor-Critic (A3C) [131] and Proximal Policy Optimization (PPO) [166].

$$\theta_{t+1} = \theta_t + \alpha \nabla J(\theta_t) \tag{2.5}$$

Methods that approximate both the value and policy functions are often called actor-critic networks, where the actor network is the learned policy function, and the critic network is the learned value function.

Proximal Policy Optimization

PPO is a policy optimization method within the model-free RL category that maximizes the performance by optimizing a surrogate objective function [166]. PPO is an on-policy algorithm capable of handling both discrete and continuous action spaces. The policy parameters are updated based on a clipped surrogate objective function given in Eq. 2.6. In this equation, ϵ is a small hyperparameter limiting the changes of the policy in each update and $\hat{A}^{\pi_{\theta_t}}(s,a)$ is an estimation for the advantage function using Generalized Advantage Estimation (GAE) according to Eq. 2.7. In this equation λ is the GAE parameter.

$$L^{clip}(\theta) = \mathbb{E}_{s,a\sim\pi_{\theta_t}}\left[\min\left(\frac{\log\pi_\theta}{\log\pi_{\theta_t}}A^{\log\pi_{\theta_t}},\ \text{clip}\left(\frac{\log\pi_\theta}{\log\pi_{\theta_t}}, 1-\epsilon, 1+\epsilon\right)\hat{A}^{\pi_{\theta_t}}(s,a)\right)\right] \tag{2.6}$$

$$\hat{A}^{\pi_\theta}(s_t, a_t) = \sum_{l=0}^{N}(\gamma\lambda)^l\left(r_{t+l} + \gamma V^\pi(s_{t+l+1}) - V^\pi(s_{t+l})\right). \tag{2.7}$$

In addition to the objective for the policy, PPO also incorporates an objective for the value function, which is learned simultaneously. This objective is expressed as a squared error loss:

$$L^{value}(\theta) = \mathbb{E}_{s_t}\left[(V_\theta(s_t) - \hat{R}_t)^2\right]. \tag{2.8}$$

Moreover, the loss function of PPO includes an entropy regularization term to encourage exploration:

$$L^{entropy}(\theta) = \mathbb{E}_{a\sim\pi_\theta}\left[H\left(\pi_\theta(s)\right)\right], \tag{2.9}$$

where $H(\pi_\theta(s))$ represents the entropy of the policy π_θ.

The overall loss function is defined as a linear weighted combination of these three objectives, as specified in Eq. 2.10.

$$L^{PPO} = -L^{clip} + c_1 L^{value} - c_2 L^{entroy} \tag{2.10}$$

Algorithm 1 presents the pseudo-code for the PPO algorithm.

Algorithm 1 Proximal Policy Optimization (PPO)

Initialize the parameters θ for the policy and value network
Duplicate policy $\pi_{\theta_{\text{old}}} \leftarrow \pi_\theta$.
for iteration $= 1, 2, \ldots$ **do**
    Collect trajectories using current policy:
    for $t = 1$ to N **do**
        Take action a_t according to $\pi_\theta(a_t|s_t)$
        Observe reward r_t and next state s_{t+1}
    end for
    Compute advantages $A(s, a)$ using Generalized Advantage Estimation
    Compute the return estimates $\hat{R}_t$
    Optimize policy network:
    Compute ratio $r(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$
    Compute clipped surrogate objective:

$$L^{clip}(\theta) = \mathbb{E}_t \left[\min\left(r(\theta) \cdot A(s,a),\ \text{clip}(r(\theta), 1-\epsilon, 1+\epsilon) \cdot A(s,a)\right)\right]$$

    Compute value loss:

$$L^{value} = \mathbb{E}_t \left[(V_\theta(s_t) - \hat{R}_t)^2\right]$$

    Combine losses (include entropy regularization):

$$L^{PPO} = -L^{clip} + c_1 L^{value} - c_2 L^{entroy}$$

    Update θ using gradient descent on L^{PPO}
    Update $\pi_{\theta_{\text{old}}}$ to match current policy π_θ ($\pi_{\theta_{\text{old}}} \leftarrow \pi_\theta$)
end for

Chapter 3

Related work

In this chapter, we will review the literature and state-of-the-art methods related to the three challenges that we pointed out in Chapter 1 for navigation among pedestrians.

3.1 Socially-aware navigation

With the increasing number of service robots and autonomous vehicles being introduced, the field of *human-aware navigation* or *socially-aware navigation* has gained significant interest in recent years. Robots and vehicles operating close to people need to be aware of humans in their proximity and treat them as social entities, considering their safety and preferences [171]. According to the definition provided by Francis et al., "A socially navigating robot interacts with humans, achieving its movement objectives while adjusting its behaviour to ensure that the experience of nearby agents is either maintained or improved" [63].

When research in this area focuses on the independent socially-aware navigation task itself, rather than incorporating additional specific social considerations that arise from an assistive or collaborative role of the robot, this type of navigation is sometimes referred to as "*crowd navigation*" [171].

Human-aware or socially-aware navigation lies at the intersection of robot motion planning and HRI, focusing on the interactions that occur during navigation [109]. Social navigation initially emerged as part of the human-robot interaction research domain [171]. Therefore, most papers in this area focus on mobile robots, though social aspects of autonomous vehicles (e.g., cars and shuttles) sharing space with pedestrians have recently

gained research interest [171]. This includes studies on vehicle-pedestrian interaction, with a comprehensive survey on pedestrian behaviour studies involving autonomous vehicles and AV-pedestrian communication methods provided in [147, 154]. Therefore, recently, the terms "*social*," "*human-aware*," and "*socially aware*" have begun to be used in the autonomous vehicle area as well. within this area, some have studied pedestrian-aware navigation [99, 152]. Another group of studies has focused on autonomous vehicles' cooperation with human-driven vehicles during navigation [181, 189, 57, 167], a field that can also be categorized as part of socially-aware navigation research [171].

However, most human-aware vehicle-related studies are conducted for interactions occurring on the road, overlooking the diversity and complexity of interactions that can happen between pedestrians and vehicles in unstructured environments, along with the differing behavioural expectations that may arise in such settings. This less-explored area of socially-aware navigation for autonomous vehicles in unstructured environments motivated the formation of our first research question (**RQ1**). Among the few works in this context is Kabtoul et al.'s method, which generates proactive manoeuvres for AVs around pedestrians in shared spaces, considering both frontal and lateral interactions using dynamic channels [99]. However, their approach overlooks the different expectations pedestrians may have regarding acceptable and perceived safe behaviour of the AV. Additionally, by relying on simulated pedestrian behaviour based on the social force model, it does not consider the potential variations in pedestrian behaviour specific to these contexts. These aspects, which may influence the algorithm, are addressed in this thesis through answering the provided research question (**RQ1**).

Several survey papers have extensively reviewed articles on socially-aware navigation, categorizing them based on different aspects they consider [109, 128, 63, 171]. Within the context of socially-aware navigation, a robot's situational awareness has been studied from different perspectives, including its awareness of surrounding agents and the social norms people typically follow [171]. The robot's awareness of surrounding agents, which could include trajectory prediction or interaction consideration, will be discussed in detail in the next sections. Here, we focus on literature examining robots' awareness of common social norms in navigation. These could include considering other agents comfort and preferences, behaving naturally and legibly, and adhering to cultural norms during navigation. These three aspects, referred to as comfort, naturalness, and sociability [109], have been studied in the human-aware navigation literature and will be reviewed here.

Comfort is defined as not causing stress or annoyance in humans during the robot's navigation. Therefore, human comfort and preferences are addressed under this category in socially-aware navigation. Proxemics, defined as the study of personal space and the distance individuals keep from each other in social situations [158], is one of the main

elements considered for robots' socially-aware navigation. For modelling proxemics, some researchers have considered a circular area around humans that the robot should avoid [150, 22]. Others have used a Gaussian function to model this area [172, 194]. Factors such as gender or the human's status (sitting, standing, or moving) have also been considered in formulating this space, as seen in [142] and [126], respectively. In a non-Gaussian formulation for personal space, Bera et al. [15] and Jiang et al. [90], in separate works, consider human emotion as a factor influencing this personal space. The social distances depend also on the robot or the vehicle's size as well as its speed [44]. Studies have shown the strong effect of the autonomous vehicles' speed [165] and distance [23] on pedestrians' crossing decisions, and it has been suggested that these criteria may have a more dominant effect than the vehicle's size [51]. In this thesis, we will model pedestrians' personal space as a circle and estimate its size using real datasets on pedestrian-vehicle interactions, fine-tuning it in simulations. Additionally, we will consider the effect of vehicle speed during close passes on pedestrian discomfort and study other aspects of comfortable and perceived safe motion behaviours through a survey.

Naturalness refers to the robot or AV having similar decision-making behaviours to humans. Designs considering this aspect attempt to make the robots' paths smooth, i.e., without sharp deviation angles and velocity changes. This includes avoiding zigzag motions to improve the predictability of the vehicle's movement [134]. Specific algorithms have also been used in mimicking human navigation behaviours to achieve the most natural-looking paths and collision avoidance techniques for a robot [59, 106, 187]. In autonomous vehicles, often explicit visual and auditory communication means are added for making the vehicle's movements and intentions clearer and more predictable to the other road users [50]. However, the concept of naturalness is more than just the smoothness of the trajectory or behaving similarly to humans in communicating intent. This concept can be expanded to the autonomous agent making decisions and actions in the same way humans do in an interaction [109]. Reasoning about others' behaviour in interaction and planning a path considering the mutual effects of trajectories on one another will lead to more natural navigation behaviours. Mathematically modelling these interactions in the motion planner of the AV is the approach we are taking in this thesis for increasing the naturalness of the AV's behaviour when navigating among pedestrians.

Sociability refers to the robot's compliance with social and cultural norms and conventions during navigation. One of the well-known norms in navigation is passing on the right (left) and overtaking on the left (right) for countries that have right-hand (left-hand) traffic, and human-aware navigation algorithms have been developed to comply with these norms [126]. Johnson and Kuipers designed a navigation algorithm for a robot wheelchair that learned this right-hand norm from observing pedestrians' behaviours [92]. To induce

the right-hand norm in their socially aware motion planning, Chen et al. defined an extra term in the reward function of their reinforcement learning algorithm for breaking the symmetry of the left and right passage in favour of the direction compatible with the norm [35]. Kirby et al. [104] coded social conventions such as keeping on the right side of a hallway as a constraint in their navigation algorithm. The social force model was extended in [155] by enforcing these social conventions to be followed. These social conventions were also introduced in the strategy selection when using a geometric reactive navigation approach [155].

Penalizing the robot's motion when heading straight towards a human who is also moving towards the robot is another norm formulated in [102]. Other examples of social norms are when a person excuses oneself and asks for a way through a blocked path. This social behaviour was tested on a robot in [161] while comparing it to two other non-social behaviours of stopping and waiting for the path to get cleared and performing a simple obstacle avoidance manoeuvre. Shahrezaie et al. extracted some social rules for robot motion in public environments such as museums through subjective analysis of unstructured interview data and suggested rules for robots' behaviour to be considered social [169]. Due to the less-studied application of low-speed autonomous vehicles operating in unstructured environments and the unclear transferability of some social norms in these interactions, this thesis takes steps toward identifying potential social norms for AV operation in such contexts, addressing our first research question (**RQ1**) informed by this research gap.

3.1.1 Evaluating social awareness of a navigation strategy

Methods for evaluating the social awareness of a robot's or AV's navigation behaviour can be divided into two broad categories: qualitative and quantitative measures [171]. Qualitative measures rely on non-numeric data to describe the social awareness of the autonomous agent, while quantitative measures aim to objectively evaluate these aspects through numerical data.

In qualitative approaches, the trajectory or speed profile of the robot is typically analyzed to draw subjective conclusions [4, 136]. This is often achieved by providing and analyzing screenshots of the simulation environment at different time steps [75, 122], showcasing the progress of the trajectory and its interactions, and comparing it to other methods [31, 150]. The legibility of the motion and the acceptance of the behaviour are usually subjectively investigated through user studies [129, 143]. In these studies, people's perceptions and expectations are analyzed using data collected from questionnaires, open-ended questions, or interviews [168].

Quantitative approaches numerically measure the performance of the navigation policy. Many of these quantitative measures in the literature focus on navigation metrics such as success rate, collision rate, and path length, as navigation efficiency is a part of social navigation [35, 31, 30]. Additionally, other quantitative metrics for measuring human discomfort or the naturalness of the trajectory have been used, focusing more on the socially-aware aspect of navigation.

Within the discomfort measure category, metrics include human-robot distance, the number of personal space intrusions, time spent in the personal space of humans, and defined safety indexes (e.g., Social Individual Index, Social Group Index), among others [122, 186, 107].

Within the naturalness measure category, several similarity metrics are used to compare the robot's trajectory to a human's trajectory, especially when the robot is designed to mimic human behaviour [60]. These measures, such as average displacement error and final displacement error, are common in pedestrian trajectory prediction [3] but have also been used for socially-aware navigation when a reference trajectory of a human in the same situation is available [48]. Other metrics in this category are irregularity measures, which count the number of unnecessary heading changes [193]. In our work, we will use both quantitative and qualitative measures to evaluate both the navigation efficiency and the social awareness aspects of our developed navigation framework.

3.2 Characteristics of motion behaviour in unstructured environments

When designing the behaviour of an autonomous vehicle for operation in a specific environment, it is important to study the unique characteristics of the target environment and the relevant factors that should be considered. In our case, this involves examining unstructured environments, its definition, and the unique motion behaviours of pedestrians in such settings.

Many researchers in the area of motion prediction and planning are using the terms structured and unstructured environments for addressing two distinct application areas for their automated systems [32, 175, 24, 74, 148]. Different definitions have been proposed in the literature for classifying these environments into structured and unstructured. In one type of definition that is focused on vehicular traffic environments, the distinction is made based on the existence of lane marks and strict traffic rules (e.g., see [97, 101, 105, 86]). Within this definition, a structured environment is believed to be a place with clear lane

marks and road dividers and where traffic participants adhere to strict traffic rules [97]. On the other hand, in an unstructured environment, such marked lanes or road dividers do not exist and no strict traffic rules are followed by drivers and pedestrians. In unstructured environments, all traffic participants are required to pay more attention due to the reduced traffic rules [97]. Pedestrianized spaces and parking lots are considered as unstructured environments under this definition [97, 101].

In another definition, the coherency of the movement pattern has been introduced as the aspect that distinguishes structured from unstructured environments [139, 91, 160]. In this definition, an environment is called structured when a main motion track can be identified which is usually caused by a condition in the environment, such as the pedestrian crosswalks on roads. In contrast, in unstructured environments, pedestrians can move freely in different directions [139]. In this definition, unlike the previous one, a pedestrianized environment can be either structured or unstructured based on the existing motion patterns.

Although these two definitions overlap in some aspects, there are also some differences that make it difficult to specify whether an environment is structured or unstructured solely based on these qualitative definitions. Despite the differences in the two definitions provided above, it is agreed in the literature that the behaviours of traffic participants in these two types of environments are different. This becomes important for an autonomous vehicle that needs to predict the future motions of other traffic participants [113, 1], as this prediction can be influenced by the type of environment in which the vehicle is operating.

Many real-world pedestrian trajectory datasets have been collected from different environments, which are used as a benchmark for trajectory prediction algorithms. While the nature of the environment and vehicles' presence can have a considerable influence on the trajectory profile of pedestrians, no clear classification has been made between the different datasets based on quantitative measures such as the pedestrians' and vehicles' trajectory features. This gap has informed our second research question (**RQ2**), aimed at gaining a better understanding of pedestrian motion behaviour characteristics in unstructured environments.

Different datasets have been compared in terms of the complexity of the human trajectories for the prediction task by introducing some indicators in [6]. However, the focus of this comparison was only on pedestrians' behaviour and, therefore, mostly included datasets in which pedestrians are the only agents present without being influenced by vehicles.

Robicquet et al. have detected different classes of navigation styles for pedestrians in a dataset based on the different ways pedestrians handle their interaction in a collision

avoidance situation [159]. However, it is still unclear how these navigation styles can also be affected by the type of the environment.

Furthermore, traffic participants' spatial and movement behaviour in different shared space layouts is studied in [13], considering both vehicles and pedestrians. However, the number of behavioural features studied was limited to mean speed, number of interactions, and yield ratio.

Also, in another study focusing on pedestrians, pedestrians' behaviour in a shopping mall was classified into three categories of "walking straight", "finding a way", and "walking around" in [138], showing that roaming around is a common trajectory behaviour that can be seen in a shopping mall environment. However, the analysis was restricted to one dataset with only pedestrians moving in a mall. Nevertheless, we get inspiration from some of the behavioural measures used in these papers (e.g., path efficiency [6], stop ratio [138]) for building the quantitative motion features that can describe unstructured vs. structured environments in this thesis.

3.3 Decision-making in pedestrian-rich environment

The vision of deploying mobile robots among pedestrians in public spaces to accomplish various tasks has attracted significant research interest over the past few decades. A crucial capability for these robots is collision avoidance with pedestrians. Various approaches to collision avoidance have been proposed in the literature, which can be categorized as *reactive*, *proactive*, and *learning-based* methods for autonomous navigation in crowded environments.

Reactive methods rely solely on the current state of the environment, responding to immediate changes as they occur. In contrast, *proactive* methods are defined as approaches that predict how the environment will evolve and take early actions based on these anticipated changes. *Learning-based* approaches aim to learn these dynamics from data or through interaction in a simulation.

The literature on each of these categories will be discussed separately in the following sections. As we delve into this section, our focus will be on the independent socially-aware navigation task of the robot or AV among a crowd of pedestrians. Therefore, we will predominantly use the term "*crowd navigation*" instead of "*socially-aware navigation.*"

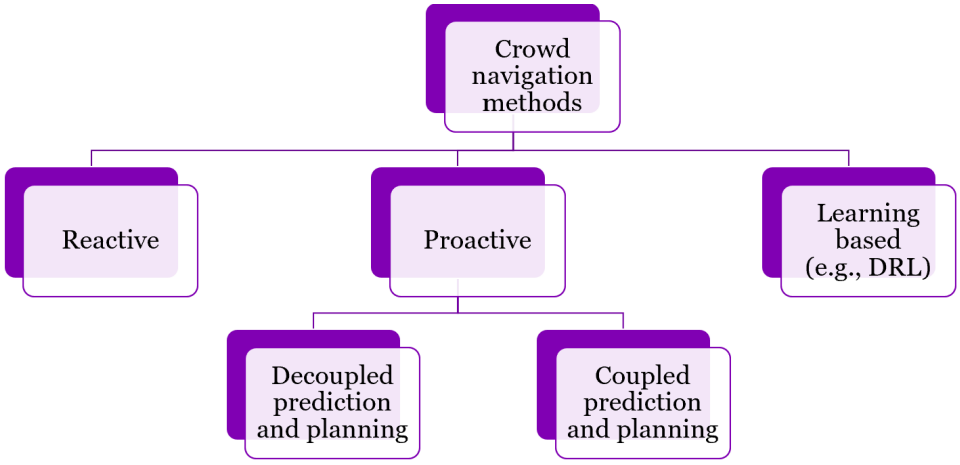

Figure 3.1: Main categories of proposed approaches for crowd navigation in the literature.

3.3.1 Reactive collision avoidance methods

Earlier methods proposed for crowd navigation were based on reactive collision avoidance methods, which try to decide the next collision-free move of the robot solely based on the current state of the agents and the one-step interaction rule.

One notable approach in this category is the Reciprocal Velocity Obstacles (RVO) [191] and its derivatives, including the Optimal Reciprocal Collision Avoidance (ORCA) [190] algorithms. These methods have become popular in multi-agent collision avoidance in robotics. They are based on the concept of a Velocity Obstacle (VO), which is the set of all velocities that will cause a collision with another agent at some point in the future if that nearby agent maintains its current velocity. In this way, a velocity obstacle for each nearby agent will be generated, based on which a safe velocity vector within the feasible collision-free space will be selected for the robot. The VO concept and ORCA model continue to be utilized in recent research, whether as a means to simulate other agents' behaviour in the pedestrian simulation environment [31, 30, 122, 36], or as external knowledge to enhance safety of learning-based methods [149, 201].

Another commonly employed reactive navigation and collision avoidance method in the literature is the Social Force Model (SFM) [79]. The SFM models agent interactions using a repulsive force for collision avoidance and introduces an attractive force towards the goal position to facilitate navigation to the destination. The SFM remains actively utilized in the literature on crowd navigation, often in combination with other methods [65].

Reactive-based algorithms have become popular due to their simplicity, making them

computationally efficient for real-time applications. However, their reliance solely on the current state of agents makes them somewhat short-sighted, potentially resulting in oscillatory and unnatural trajectories [61, 108]. Furthermore, these algorithms are built on assumptions such as reciprocity (e.g., agents in collision sharing equal responsibility for avoiding the collision), which may not always hold true, especially in scenarios involving non-cooperative or distracted pedestrians.

The SFM also has been shown to produce very conservative behaviours, as reflected in the circuitous path it takes (often resulting in large detours) [18], and could lead to robot freezing problems in crowded spaces [123].

Due to the shortcomings of reactive-based methods, proactive collision avoidance methods were proposed for the crowd navigation problem. These approaches can leverage predictions of crowd movement, which will be discussed next.

3.3.2 Proactive navigation methods

In contrast to reactive collision avoidance algorithms, which have a short-sighted perspective, proactive navigation algorithms incorporate reasoning about the future evolution of the dynamic surroundings. They employ prediction algorithms to generate more farsighted behaviours, integrating predictions regarding other agents' trajectories into the decision-making process. Navigation algorithms falling within this category are often referred to as trajectory-based methods [36], which can be categorized into two classes of decoupled and coupled prediction and planning approaches.

Decoupled prediction and planning algorithms

Decoupled prediction and planning algorithms first predict all agents' trajectories and then attempt to plan the robot's path in the remaining unoccupied space [188, 54, 14, 10, 144]. For example, in [10], probabilistically safe trajectories for the robot are generated through a chance-constrained optimization problem given the prediction distribution of other agents.

However, decoupled approaches may lead to the robot freezing problem in crowded settings, where the prediction renders most of the area unsafe to traverse [184]. Assuming that unconstrained uncertainty growth in prediction leads to most of the space becoming unsafe, some research has focused on more accurate human motion models [14] to limit individual agents' predicted uncertainty, which has shown to improve navigation performance to some extent [54]. However, it has been shown that even with perfect prediction,

the robot freezing problem is still unavoidable with certain configurations or levels of crowd density [183].

The issue with decoupled prediction and planning approaches, contributing to the robot freezing problem, lies in neglecting the fact that pedestrians are not merely dynamic obstacles but decision-making agents. They adapt their paths based on interactions with other agents and participate in joint collision avoidance, also influenced by the robot's actions [183, 128]. In certain scenarios, the robot might anticipate cooperative agents creating space for it [183], a behaviour commonly observed in human-human interactions. Thus, addressing the robot freezing problem relies on appropriately modelling cooperation and interaction effects, recognizing the mutual influence between both agents [183]. Consequently, a coupled prediction and planning approach is essential for generating a collision-free and freezing-free path in crowded spaces which will be discussed next.

Coupled prediction and planning algorithms

In the category of proactive navigation algorithms, coupled prediction and planning approaches jointly plan feasible trajectories for all agents [108, 183, 110], giving as output a cooperative collision avoidance solution. In this line, Kretzschmar et al. modelled the navigation problem as a joint mixed probability distribution of the trajectories of all agents present [108].

Trautman et al. introduced the use of Interacting Gaussian Processes to solve cooperative collision avoidance by modelling trajectories as Gaussian Processes and formulating the dependencies between trajectories through an interaction potential function, which addresses the collision avoidance requirement [183].

In [178], the concept of cooperative collision avoidance is extended beyond simply coupling trajectories to also include coupling agents' intention distributions, thereby modelling the agents' preferred trajectory adjustments during interactions.

Game theory is another method used for modelling cooperation among multiple decision-making agents and jointly planning all agents' trajectories in a crowd navigation problem. Fridovich-Keil et al. used differential games to navigate a mobile robot between two humans and execute interactive collision avoidance manoeuvres [64].

However, implementing this joint planning approach in crowded settings, coupled with the need for frequent updates to compensate for deviations in other agents' paths from the jointly planned trajectory, imposes a computational time burden, making them challenging for real-time implementation [36]. This has motivated the introduction of learning-based methods in the crowd navigation literature, which will be reviewed next.

3.3.3 Learning-based crowd navigation methods

To address the challenge of modelling mutual interaction effects in crowd navigation while maintaining low computational time for real-time implementation, numerous DRL approaches have emerged in recent years. These approaches aim to learn interaction effects either from real-world data i.e., Inverse Reinforcement Learning (IRL) [180, 59, 106, 187], or by training the agent directly in a simulated environment using DRL methods to learn the interaction effects [36, 58, 35, 31, 122]. This approach allows the computational burden to be shifted to a training phase conducted offline. Subsequently, the trained model can be efficiently executed for real-time inference.

In DRL methods proposed for crowd navigation, interaction effects necessary for collision avoidance are learned implicitly through training. Various network architectures have been proposed in the literature to capture these effects, which will be reviewed in the following subsection.

Modelling interaction effects in DRL crowd navigation methods

One of the first DRL algorithms for multi-agent collision avoidance was proposed by Chen et al., [36]. They formulated a decentralized two-agent collision avoidance problem as a Partially Observable Markov Decision Process (POMDP). Their DRL solution involved training a value network to estimate the time to goal and deriving the optimal action based on this value function and an estimation of the nearby agents' transition probabilities. The DRL state was defined as a joint configuration of the robot's state and a single nearby agent. To extend this to more than two agents, they suggested a min-max solution where the selected action from a discrete set corresponds to the highest lower joint state value with respect to each neighbour's projected state. However, this approach could face scalability issues with a larger number of agents [36].

To better scale and manage a variable number of neighbouring agents, Everett et al. [58] proposed anew architecture using a LSTM network. This architecture sequentially inputs the neighbours' information in reverse order of their distance to the ego robot to ensure that more attention is paid to closer neighbours. The output of the LSTM cell was then concatenated with the robot's own state to form a joint state vector. This joint state vector was processed by a Multi-layer perceptron (MLP) to eventually output both a value and a policy, represented as a discrete probability distribution over actions [58].

However, closer agents are not necessarily the most important ones to pay attention to, as designed through the reverse distance order LSTM feeding of their network architecture

[58]. Other factors, such as the approaching direction and speed of the surrounding agents, also play a critical role in how much attention the robot should pay to each particular neighbour to adjust its path accordingly. Therefore, the concept of self-attention mechanisms for encoding pairwise interactions was used by Chen et al. [31]. In their method, a weighted sum of the embedded states of all the neighbours is calculated as part of a pooling module to produce the interaction feature. This interaction feature, along with the robot's state, goes through an MLP network to estimate the state value in a Value learning (V-learning) crowd navigation algorithm. The attention scores in the weighted sum of the embedded states of the neighbours are calculated through a neural network structure that takes as input each individual's embedded state vector along with the mean vector of all other agents, eventually passing through a softmax function to produce each individual's attention score [31]. In their algorithm, they also accounted for human-human interactions in addition to the common first-order human-robot interaction by augmenting each human's state with a local map that captures the neighbouring agents of that human [31]. This approach expanded the robot's ability to predict and reason, as human-human interactions can indirectly affect the robot's decisions.

Instead of using a local map to approximate the indirect effect of human-human interaction on the robot's navigation policy, Chen et al. proposed the use of a relational graph to fully model human-human interactions alongside human-robot interaction [30], employing the concept of Graph Neural Networks (GNNs). Viewing each agent in the environment as a node in the graph, they inferred relations between the agents, represented as edges in the graph, and utilized Graph Convolutional Networks (GCN) to calculate pairwise interaction features. Pairwise relationships were inferred through a similarity function between different nodes for each edge in the graph, stored as a matrix. Through message passing mechanisms between nodes, each node aggregates its neighbours' information, which is weighted by the inferred relations from the previously constructed matrix. This mechanism computes each node's state representation, embedding information about its interactions with others. The state representation then serves as input to the value estimator, modelled as an MLP, which was used within a model-based DRL framework to calculate the robot's best action [30].

Liu et al. [123] extended the idea of modelling interaction effects in a crowd navigation problem with graphs by incorporating a temporal dimension, resulting in spatio-temporal graphs (ST-graphs). This modification enabled capturing the changes in the graph over time. They employed Recurrent Neural Networks (RNNs) to encode the temporal patterns inherent in all agents' movements and how the robot's interaction with other agents evolves over time. Separate RNNs were introduced for encoding the features of spatial and temporal edges, the outputs of which were fed into an attention module (scaled dot-product

attention [192]) to assign attention weights to each spatial edge and output a weighted sum of the spatial edges. This interaction feature was eventually concatenated with the robot's node feature and passed through another RNN network to output both value and policy within a model-free actor-critic structure [123].

Liu et al.'s model architecture in [123] originally focused solely on human-robot interaction. However, in their subsequent work [122], they expanded their approach by incorporating human-human spatial edges into their ST-graph. Additionally, they introduced separate attention layers for encoding human-human and human-robot interactions. This extension allowed them to consider indirect interaction effects on the robot's navigation.

While learning-based methods in crowd navigation problems excel in capturing complex interaction effects through experience gained in simulation environments, it remains unclear how far-sighted their learned policies are and whether they can be classified as proactive navigation algorithms. In other words, there are limited works on incorporating prediction into learning-based navigation algorithms, which is beneficial for capturing interaction effects over a longer horizon. This led us to the formation of the third research question of this thesis on effective ways of capturing interaction effects in learning-based navigation methods (**RQ3**). We will investigate this in both the trajectory prediction of nearby pedestrians and the mutual interaction effects important for motion planning, as well as the integration of prediction and decision-making in learning-based navigation methods towards development of proactive learning-based navigation algorithm. To this end, we will first review existing work on integrating prediction and planning in learning-based methods and identify gaps in this integration for robustly capturing long-horizon interaction effects. Subsequently, we will delve deeper into pedestrian trajectory prediction models and examine how interaction effects are addressed in related research.

3.4 Prediction and planning integration in learning-based crowd navigation

The introduction of learning-based methods in crowd navigation was motivated by their capability to implicitly learn complex interaction effects with feasible execution times during inference. This initiative addressed computational challenges encountered by conventional non-learning coupled prediction and planning methods in crowded environments while also mitigating the issue of shortsighted and unnatural behaviours observed in simplified reactive approaches like ORCA. In this context, we revisit the learning-based crowd navigation literature, focusing on the integration of prediction and planning. This exploration aims

to understand how prediction can contribute to making these methods proactive and more farsighted and to identify the unexplored areas in this integration that inspired the research for this thesis.

Numerous strategies have been proposed within the category of model-free DRL for crowd navigation, that involves learning a value function (V-learning) through interactions within a simulated environment [36, 35, 31]. However, a significant challenge encountered by these methods in determining actions based on the value function is the unknown state transitions of the dynamic environment (see Eq. 2.4). To tackle this challenge, some approaches rely on assumptions about pedestrian behaviour, such as assuming a known and deterministic dynamic for humans. These assumptions include utilizing a constant velocity model to predict the pedestrian's next state [36, 35] or assuming that all agents adhere to the ORCA model (with access to the ground truth of the simulation environment) [31]. It's crucial to acknowledge that these assumptions may not consistently align with real-world scenarios, particularly in crowded and highly interactive settings. Furthermore, pedestrians' behaviour is inherently stochastic and cannot be reliably modelled deterministically.

These assumptions regarding the state transition function of other nearby agents were relaxed in [58, 123] through the integration of a policy network alongside the value network within an actor-critic structure for the DRL network. This approach relies on the actor network to directly generate actions rather than deriving the optimal action from the value network, thereby mitigating the need for assumptions about the behavioural models of other agents.

As another approach to address the challenge posed by the unknown state transition probability in dynamic environments within V-learning approaches, Chen et al. adopted a model-based approach in DRL utilizing a prediction model [30]. In their framework, a prediction module for nearby agents is jointly learned to serve as the state transition probability function in selecting the best action sequence using the concurrently learned value function. Employing a d-step planning method [137], the authors conduct a rollout of the agents' state d-steps into the future using their prediction model and value estimator to determine the optimal action through a Monte Carlo tree search approach [30]. This integration of prediction into the DRL policy generator (motion planner) is achieved and inspired by the structure of a model-based DRL.

Using the same d-step planning approach, Matsumoto et al., combined a novel predictive state representation (PSR) model within the structure of the DRL framework [127]. The PSR in their architecture serves as the model that predicts the sate of the pedestrians in the next state based on the current state as well as the robot's action as an input.

However, to manage computational costs associated with the tree search, these meth-

ods are constrained to consider only a discrete action space [30, 127]. Moreover, their prediction model is designed to predict solely the next state and does not consider future states over a horizon. Consequently, they rely only on the current state for predictions, overlooking the impact of the agent's trajectory history on their future positions, which becomes particularly important for longer-horizon predictions.

The DenseCAvoid method, introduced in [164], represents another crowd navigation algorithm leveraging predictions of pedestrians' next states. However, it uniquely incorporates this prediction into the framework of a policy gradient DRL method, specifically PPO, which directly learns a policy network to generate the robot's actions. In their approach, they integrate the prediction within the DRL network by treating the next-step prediction as a novel observational state [164]. It's important to note that the prediction network utilized in their model has been previously trained on real-world data [29].

Liu et al., also combine a previously trained data-driven pedestrian prediction network with an actor-citric DRL structure for the planner and train it using PPO [122]. However, unlike previous models, instead of relying solely on the predicted sate for the next step [164, 30, 127], Liu et al. utilize the forecasted trajectory of pedestrians over a prediction horizon when integrating prediction and DRL planning [122]. They integrate all the prediction states into the DRL's observation space and also incorporate the predicted pedestrian trajectory into the reward function. The latter (which was a novel aspect not previously seen in the formulation of the papers with a relatively similar approach [164]) ensures that the robot is penalized not only for getting too close to the pedestrians' current position but also for intersecting with the predicted future trajectory, ensuring pedestrian comfort. Employing an actor-critic structure with PPO policy-based learning, they effectively handle continuous actions, distinguishing themselves from previous works that solely rely on discrete actions. Integrating pedestrians' predictions into DRL learning has demonstrated the ability to facilitate longsighted robot behaviours [122].

However, their model does not incorporate the impact of the robot's presence and actions on the predicted trajectory of pedestrians, assuming instead that the robot is invisible to humans during prediction and focusing solely on human-human interactions [122]. This limitation in the existing literature and state-of-the-art methods has informed other aspects of our third research question (**RQ3**).

Moreover, another notable gap in the existing DRL crowd navigation literature is the lack of consideration for how uncertainties in pedestrian trajectory predictions affect the learned policy. This gap informed the fourth research question of this thesis (**RQ4**). Addressing prediction uncertainties is especially valuable in scenarios where pedestrians might suddenly enter the robot's or AV's field of view without many past observable states, or

for stationary pedestrians who might unexpectedly initiate movement. Such situations common in unstructured environment, are where data-driven predictions are prone to significant errors that the planner should mitigate for . Therefore, in the next section, we will review pedestrian trajectory prediction models and examine how interactions effects and uncertainties in future predictions are addressed in the literature at the prediction layer. Following this, we will explore the limited existing research on incorporating prediction uncertainty in DRL-based planning algorithm when using these two modules in tandem.

3.5 Pedestrian trajectory prediction

Modelling pedestrians' motion behaviours to predict their future trajectories is critical for autonomous vehicles navigating in shared environments. Accurate prediction involves integrating multiple sources of information, including agents' past trajectories and inter-agent interactions (e.g., how one pedestrian's movements influence another's). Safe autonomous vehicle operation hinges on understanding how the vehicle's presence and actions affect pedestrians' trajectories.

Many pedestrian trajectory prediction models have been proposed that include only pedestrian-pedestrian interaction without considering any vehicle in their model. These models can hardly be used for AV algorithms as they do not capture the effect of the vehicle's motion on the pedestrian's decisions and motions. Among the papers that include the pedestrian-vehicle interaction in the pedestrian's motion model, a great focus has been put on pedestrian road crossing behaviour [154, 217] that takes into account the strict structure of the conventional roads.

Given our focus on pedestrian trajectory behaviour in the presence of vehicles within an off-road, unstructured environment, here we only focus on related literature for pedestrian trajectory prediction and motion modelling that includes pedestrian-vehicle interaction without restricting it to a structured road environment. Specifically, our review in this section emphasizes three key perspectives: 1) The pedestrian trajectory prediction and motion models proposed, 2) The methodologies employed for modelling interactions between pedestrians and vehicles within these frameworks, 3) The strategies employed for modelling uncertainty regarding pedestrians' future positions.

3.5.1 Modelling pedestrian motion

Various methods in the literature model and predict pedestrian trajectories. These approaches typically fall into two categories: *expert-based methods*, which rely on hand-crafted rules, and *data-driven methods*, which learn from real-world trajectory data. Some methods combine elements from both categories. Below, we provide a brief review of the literature in each category.

Expert-based models

Expert-based models involve a human-crafted explicit rule that governs pedestrians' state transition. Within this category are models that use equations of motion from laws of physics for propagating forward the state of the agent. A well-studied model in this category is the SFM model [79], that uses the second law of motion for simulating pedestrians' motion toward their goal position.

In the dynamic equation of SFM, the motion of the pedestrian is formulated to be under the influence of multiple forces with the basic ones being an attractive force to the destination and repulsive forces from both obstacles and other agents (e.g., pedestrians or vehicles) in the environment. The summation of these forces specifies the acceleration vector of a pedestrian according to Newton's second law of motion [207, 157]. While the underlying dynamic equation of SFM with only the attraction force towards the destination can be used as a basic motion model for pedestrians, the social forces from other agents are an additional factor in the SFM for interaction modelling. Therefore, some papers have used the SFM as an interaction model in a multi-layer architecture on top of a simple motion planner that calculates the agent's shortest path to its destination considering only the static obstacles in the environment [9, 94, 96].

Other expert-based models based on kinematic equations are the constant velocity and constant acceleration models implemented in [216] for pedestrians while using a Kalman Filter state estimation. Within this category Zhang et al, model the vehicle's motion using a constant turn rate and constant velocity formulation. In their model, the two agents' motions are made coupled by transforming the pedestrians' coordinate into the vehicle's moving frame [216].

The Cellular Automota (CA) method is another approach commonly used in simulation environments for generating agent's trajectories as consecutive occupied cells over time in a discretized world [42, 119]. The transition between different cells in CA is governed by some hand-crafted rules and pre-specified probabilities.

Data-driven models

Instead of relying on explicitly defined dynamic equations or hand-crafted rules for describing the pedestrians' motions, many researchers have worked on learning these behavioural motion patterns from data (e.g., [153, 177, 7]). This approach is followed in the literature with the goal of modelling pedestrians' more complex movements that cannot be captured by a hand-tuned function. For this purpose, different function approximates have been fitted to the human trajectory data ranging from linear regression models (e.g. [98, 66]) to more complex deep learning models (e.g., [17, 215, 120]).

Due to the sequential nature of pedestrian trajectory, many of the models proposed in this area are based on RNN and variants such as LSTM networks and Gated Recurrent Unit (GRU), that can capture the dependencies between time series data [17, 41, 66]. The output of the popularly used LSTM cells, called the *hidden state*, encodes the trajectory history of each agent.

For the purpose of trajectory prediction, usually, an encoder-decoder architecture is used [153, 103, 52, 156] where the states of the pedestrian, mainly consisting of their position coordinates will first get encoded into a hidden state as part of the encoder module. Then the extracted feature of each pedestrian's trajectory along with all other related features (e.g., the hidden states of other agents or scene features) will go through a decoder module to output the predicted states of the pedestrian's future trajectory. The way the encoded features of different agents are shared with each other in the architecture for modelling interaction is discussed in the next section. As another method, a transformer network is used in [212] for pedestrian trajectory prediction.

While the above methods apply sequential models, it should be pointed out that non-sequential models such as the Convolutional Neural Networks (CNNs) have also been used for trajectory prediction [215, 25]. CNN is used in [25] for extracting trajectory patterns and for predicting the next piece of trajectory among the alternative paths within a range of possible smooth turns.

Others have used CNN in combination with LSTM to capture the dependencies between the spatial features in the form of agent-agent and agent-scene interaction along with the temporal features encoded by LSTM [34, 198, 29, 39, 33, 28].

An alternative approach followed in the literature for capturing spatial and temporal dependencies simultaneously in a trajectory prediction problem is the use of GNN [120, 27, 133, 125, 218, 121].

A variant of GNNs called GCN has also been used for trajectory prediction of pedestrians in a mixed environment of multi-class agents [130, 151, 176]. In these models, the

feature of each node and the structure of the graph's connections in the form of an adjacency matrix are inputted to a multi-layer convolutional network to extract common patterns in the trajectories.

Creating multiple plausible future trajectories through repeated sampling of generative models is another common approach followed in the literature for considering the multi-modality nature of human trajectory. This is done through the use of Generative Adversarial Networks (GANs) in [221, 78, 56, 111, 198, 83] with different architecture for the generator and the discriminator but all consisting of LSTM layers. The other generative model used in the literature is the Conditional Variational Autoencoders (CVAEs) which formulates the future trajectory of each agent conditioned on its past trajectory and a latent variable that can be sampled multiple times [197, 80, 66, 16, 39, 218]. This latent variable is learned during the training phase by providing both the trajectory history and the future trajectory of an agent.

Combination of data-driven and expert-based models

In a few papers, deep learning models have been combined with expert-based models for predicting pedestrians' future trajectories [93, 116, 115]. Johora et al. [93], have shown that the combined approach outperforms each of the pure methods. In their architecture, the expert-based methods consist of the three interacting modules of free-flow trajectory planning, the social force model, and the game-theory decision layer. This combination is executed in parallel with the deep learning module. The output of the deep learning-based module is then used when no conflict between the predicted trajectories of agents is detected. In case of a detected conflict in the generative trajectories of the deep learning model, the output of the expert-based model will be implemented [93].

Moreover, in the data-driven deep learning models proposed in [116, 115], some physics-based dynamic constraints on the motion of agents are applied at the near-end layers of the model to guarantee the feasibility of the generated trajectories. In [115], the agent's actions are predicted through a network and a kinematic cell, which takes the predicted actions and the current state as inputs to generate the agent's next state.

In a couple of expert-based models (e.g., SFM), the parameters of the model are estimated by using real-world trajectory data [9, 96, 146, 209, 42]. In these methods, parameters such as safety distances, range of repulsive forces, payoff matrices of the game [96] or the social factor and perception zones in the SFM [146] are calibrated through a process to get the minimum difference between the ground-truth trajectories in the dataset and the simulated ones using the model.

3.5.2 Interaction modelling

The presence of a moving vehicle close to a pedestrian can highly affect the pedestrian's motion behaviour and is an important factor to consider when modelling or predicting the pedestrian's future trajectory [215, 56]. This effect, referred to as the pedestrian-vehicle interaction effect [215], has been modelled in many different ways in the literature depending on the type of model used for generating the pedestrians' trajectory, such as being expert-based or data-driven.

From a broad perspective, these interaction effects can be categorized to be modelled either explicitly or implicitly. In the explicit interaction modelling approaches (e.g., see [205, 9, 94, 81]), the effect of the vehicle on a pedestrian's motion is forced through some clear terms in the formulation of the pedestrian's motion such as through explicit forces in the SFM (e.g., [146, 219]). On the other hand, data-driven modules often account for these interactions implicitly by using the vehicle's trajectory as another input to the model along with the target pedestrian's own trajectory (e.g., see [17, 27]). As these models are trained on datasets collected from real-world interaction scenarios, it is expected that the model will learn the interaction already embedded in the data.

More details on the interaction models proposed in the literature in each sub-category will be discussed next.

Explicit interaction modelling

The original SFM model, proposed by Helbing and Molnar [79], only considers the social interactions between pedestrians. Many papers have extended that model by introducing additional forces for modelling pedestrian-vehicle interaction [207, 20]. The way these additional forces are formulated is the main difference between these extended social force models.

In these models, each vehicle exerts a repulsive force on the pedestrian based on its relative distance and approach direction. The effect of the relative interaction distance is captured in the so-called *decaying function* [209] which is usually selected to be an exponential function of the distance [205, 8]. Another function which is used in some of the SFM model formulations is called the *anisotropy function* [209, 208]. This function formulates the effect of the different interacting directions on the magnitude of the repulsive force. Such a function will consider, for example, the higher influence that a pedestrian walking toward the vehicle will experience compared to another pedestrian that is walking away from the vehicle when both of them are positioned at the same distance from the vehicle, which means having the same value for the decaying function [209, 208].

Some papers have modelled the vehicle with a circular shape same as how pedestrians are modelled in SFM but with a larger radius [146, 145]. In these models, the asymmetrical effect of the vehicle's influence is accounted for in the anisotropy function or considering the approach angle and its change rate [146, 145]. On the other hand, some others have modelled a vehicle with an ellipse and considered the asymmetry of the vehicle's influence by calculating the pedestrian's relative distance based on its angular position with respect to the ellipse [8, 219, 9].

In [205], one foci of the ellipse is placed at the vehicle's middle rear and the other at an extended front point based on speed, accounting for the increased danger at the front and extending the danger zone with speed. A nearby pedestrian experiences a repulsive force perpendicular to the ellipse's tangent line in their model [205].

Others have used a fixed ellipse which encloses the vehicles and repulsive force is modelled as an exponential function of the distance between the vehicle and the pedestrian, where the distance depends on the relative orientation of the pedestrian with respect to the vehicle [219, 8, 9]. Yang et al., define three areas around the vehicle — namely the front area, body area, and back area — and propose a different formulation for the lateral and longitudinal forces of the vehicle on the pedestrian in each area [208]. Of course, the front area always has the highest magnitude of forces due to being more dangerous. A rectangular shape contour with an extended triangular shape in the front is used for the vehicle in [209]. The vehicle-pedestrian interaction force is calculated as a function of the minimum distance between the pedestrian and the contour. The magnitude of the force is also adjusted according to the walking direction of the pedestrian.

The repulsive force in the SFM increases when the agents get closer to each other and almost vanishes when they get far away [79]. Because of this, short-range repulsive forces of the SFM can cause frequent urgent detours when it comes to pedestrian-vehicle collision avoidance. Therefore, many papers have combined the SFM with other collision avoidance strategies [8, 9, 157, 146, 145]. In this regard, a long-range collision avoidance method is proposed in [8, 9] for handling potential collisions that might last by just following the SFM formulation. Within this method, potential conflicts are predicted by geometrically projecting the shadow of the pedestrian in the direction of the car. The minimum speed and direction change for avoiding collisions are then calculated by solving an optimization problem.

For long-range conflicts, defined as conflicts happening in more than 2 seconds in [157], the authors consider a force that keeps the pedestrian in a safe zone. This is done by defining a force perpendicular to the car which models the pedestrian's tendency to walk parallel to the car in a shared space environment.

SFM is used alongside a proposed decision model in [146, 145], which gets activated whenever a conflict is predicted with the vehicle by referring to the time-to-collision (TTC) parameter. The decision model decides on an action based on (a) the type of the interaction being from the front, back, or lateral, and (b) the expected crossing order between the pedestrian and the vehicle. These actions are implemented through some newly introduced forces in the SFM and can create a sharp turn, a stop, or a step-back action. For no-conflicting interactions, the simple SFM in which the vehicle just has a repulsive force on the pedestrian is used.

SFM's ability to directly map perception to action is used in [94, 95, 96] for solving simple reactive interactions. However, for handling more complicated interactions which require deciding among different alternative actions, a game-theoretic layer is added on top of the SFM. The decided action at the game-theoretic layer is executed by the force-based layer. Later, this model was proposed to be used beside a deep learning model and perform as a safe backup for the data-driven model [93]. Within this structure, whenever the predicted trajectories of the data-driven model end up in a collision, the output from the combined SFM and game theory is used instead, ensuring a collision-free trajectory prediction [93].

Game theoretic methods used in the literature for considering the effect of interacting agents' actions on one another also fall into this category as they solve for a coupled decision for the interacting agents by searching the Nash equilibrium [94, 95, 96, 82, 93]. These calculations are based on a payoff matrix that explicitly encodes the effect of one agent's decision on the other. For example, the game-theoretic layer in the proposed architectures of [94, 95, 96], models the interaction between a pedestrian and a vehicle as a Stackelberg sequential leader-follower game [196] where a payoff matrix is designed for possible discrete actions of *continue*, *decelerate*, or *deviate* for the pedestrian, and *continue* and *decelerate* for the vehicle. The final optimal actions are executed via the social force layer to generate the coupled trajectories.

Other heuristics such as TTC have also been used in combination with the Social force model [219, 145] to keep track of interaction between pedestrians and vehicles. In [219] the TTC governs the principal behind the vehicles behaviour in interaction with pedestrians, where the vehicle breaks whenever the TTC is lower than a threshold. In [145] another decision layer on top of the SFM gets activated when the TTC between agents gets below a threshold.

Implicit interaction modelling

In data-driven deep learning methods for pedestrian trajectory prediction, interaction effects are encoded through a separate module. This module uses the trajectory information of nearby agents along with the target agent's data to implicitly learn the impact of interactions on the target's trajectory through data. Various models, like pooling mechanisms and graph neural networks, are used to combine these trajectories and will be reviewed separately.

Pooling trajectories of interacting agents

One of the common ways for pooling hidden states of all surrounding agents is by constructing an occupancy grid map centred at the target agent's position [17, 78, 39]. In these occupancy maps, the hidden state of all agents residing in the same grid cell will be summed together, building a tensor that embeds information of all interacting agents that can impact the future trajectory of the target pedestrian. This tensor input will be used along with the target agent's own spatial hidden state as the two main inputs to the LSTM network used for the trajectory prediction.

While the above-mentioned grid map used for constructing the tensor input has a rectangular shape in [17, 78], Cheng et al., have used a circular polar occupancy map where the orientation and the distance of the agents with respect to the target pedestrian specify the occupied cells [39].

In the above-mentioned papers, grid maps used for social pooling are constructed locally around each target agent. A global version of such spatial feature maps at each time step is proposed and used in [198] by dividing the bird's-eye view of the scene into discrete cells. In this constructed map the feature encoding of each agent is placed in a tensor based on the agent's position coordinate. These 2D tensors at each time step go through a CNN while the temporal dependencies between these spatial maps over time are also analyzed by using a separate LSTM network.

In [28], two maps are created for each agent: the *Horizon map* and the *Neighbour map*. The Horizon map pools embeddings from prioritized interactions in a semi-elliptical region in front of the agent, while the Neighbour map does so in a larger elliptical area around the agent, considering the closest nearby agents. These maps go through a Convolutional network to capture local dependencies, and the outputs are concatenated with the target agent's embedding to predict its future trajectory [28, 29].

While the construction of a tensor input based on an occupancy grid map for encoding interactions are mainly based on relative distances [17, 78, 39], Cheng et al, have proposed a tensor input that relies on collision probability with other agents [41]. However, their

collision probability is again formulated as a function of the distance between the pedestrian and the vehicle considering their approach direction and the elliptical shape of the car.

Instead of using occupancy grid maps to store the hidden states of nearby agents in a tensor format that maintains relative positions, some methods directly embed the relative distances between the target agent and its neighbors. These distances are then aggregated using softmax [56] or a max pooling layer [215, 76, 221] to construct the interaction feature. In other words, these methods use the relative distance values directly rather than relying on pedestrian positions on an occupancy map.

Two different types of pooling modules are proposed in [221], one for homogeneous agents and the other for heterogeneous agents. The pedestrian-vehicle interaction effect is considered in the heterogeneous pooling module in which the average output of the vehicle's homogeneous pooling module is combined with each pedestrian's output from its own homogeneous pooling module.

The relative importance of surrounding agents when aggregating their state can also be specified with an attention mechanism [197], and transformers with attention mechanisms have been applied to trajectory forecasting tasks [118, 213]. While the attention coefficients of the surrounding agents are mainly constructed based on their relative Euclidean distance to the target pedestrian, other social features such as the bearing angle between the two agents and their speculation distance were also used [111].

Most models in the literature treat the dependency of a future trajectory on an agent's past states and social interactions separately. However, Yuan et al. introduced *Agent-Former*, a transformer model that jointly captures temporal and social dependencies by flattening trajectory features across both agent and time [214]. To maintain the temporal and agent-specific identities within the sequence, they modified the original transformer's attention mechanism, replacing positional encoding with time encoding and introducing agent-aware attention, which distinguishes features of the same agent from those of others. This modification enables the state of one agent at a specific time to directly influence the future state of another agent at a different time.

Message passing through spatial edges of graph neural network

In GNNs for pedestrian trajectory prediction, spatial edges connect nodes to model interactions between agents and their impact on each agent's future position. The process of aggregating information from connected nodes, called message passing, updates the node's representation. An attention mechanism is typically used to encode the relative importance of these connections for the target agent.

Two nodes or agents are connected by an edge in the graph and are considered to interact based on criteria that vary across different papers. A popular criterion is the

spatial distance where two agents are connected with a spatial edge in the graph if they are within a certain distance from each other [27, 115, 66]. Then, the features of connected neighbours to an agent will be used with the agent's own feature for the calculation of its future location. While there are papers that rely on a criterion for specifying connecting edges, some just start with a fully connected graph [83], in other words accounting for all agents in the scene.

In a novel approach proposed by Li et al., the existence of an interaction edge in a GNN model is decided through a RL framework on top of the prediction model [116]. In their RL formulation, the action is specified as making the edge between the two nodes on or off, and the reward is defined based on how close the prediction of the selected graph connection is to the ground truth trajectory.

Su et al., have proposed different connectivity graphs for interaction consideration in their GNN structure namely *View graph*, *Direction graph* and *Rate graph* [176]. In the *View graph* only agents present in a target pedestrian's field of view are connected to that pedestrian's node in the graph and are considered to have an interaction effect on the pedestrians. For *Direction graph* and *Rate graph*, these edges are decided based on whether the two agents' moving direction has a crossover [176].

Some utilize glsGCN to directly operate on the graph [27, 130, 151, 176]. These models construct an adjacency matrix to store connected edges in the graph, with weights proportional to inverse relative speed [151] or relative distances [27, 176] of interacting agents. In GCN structure, attributes of connected nodes are aggregated using this matrix, combined with the agent's own attribute to predict its future path, considering the agent's interactions with others.

As opposed to GCN, other GNN methods have used RNN such as LSTMs to encode the sequential dynamic features of the edges in the graph. In the later, the attributes of the edge are relative measurements such as relative distance between the two connected traffic agents in the graph [120, 125, 218, 133, 115, 66, 121, 83], relative velocity [133, 115, 66] or relative heading/yaw angel [133, 115]. These edge features are sent to an RNN (e.g., LSTMs [120, 218, 25], GRU [133, 115]) to encode the information of the spatial interactions between the agents and to be used along other embedded information (e.g., the temporal edge) to output the predicted trajectory in subsequent layers.

Another feature utilized as edge attributes in the literature is the type of interacting agents. This is represented by a unique encoder in [120, 125, 121], incorporated as a feature of the spatial edge. Mo et al. [133] also employed edge type, specified as a discrete indicator based on the types of agents interacting with each other.

In implicit methods of encoding interaction effects in learning-based approaches, much

of the focus is on using only relative distance information. When speed and heading direction are included, they are typically treated as raw inputs, leaving it to the network to infer their indirect effect on interaction. In contrast, some expert-based modelling approaches pay considerable attention to incorporating speed, orientation and approach direction in the explicit formulation of interaction effects, as discussed in methods like the new variants of SFM model, reviewed in Section 3.5.2. Recognizing the potential for better-engineered features to capture interaction effects in learning-based methods and to advance this area of research, we formulated the third research question of this thesis **RQ3**, which will be elaborated on in chapter 6 and chapter 7.

3.5.3 Uncertainties in pedestrian trajectory prediction

In reality, predicting pedestrians' future trajectories involves uncertainties due to limited observable information, making assigning probabilities to future positions more realistic. While some methods, like the social force model, simulate trajectories deterministically, others in the literature consider the probabilistic nature of pedestrian trajectories.

In existing literature, uncertainty in pedestrian trajectory prediction is typically addressed through three main approaches: 1) predicting parameters of a distribution function over future states, 2) employing generative models to produce multiple trajectories for each pedestrian, and 3) various other methods, each discussed separately.

Predicting the parameters of a distribution function over pedestrians' future positions

Many data-driven methods have been proposed for predicting the probability distribution of the pedestrian's next state [120, 41, 130, 17, 151, 80, 56, 176]. In these models, instead of outputting the exact values of the pedestrian's next position, the network is trained to output the parameters of a bi-variate Gaussian distribution over the x and y position coordinate of the pedestrian's future trajectory [120, 41, 130, 17, 151, 125, 121, 28]. Therefore, the output of the final layer of their network has five components consisting of two means, two standard deviations, and a correlation value $(\sigma_x, \sigma_y, s_x, s_y, \rho)$.

Some other models have used higher dimensions for the Gaussian distribution [80]. Li et al. predict both pedestrians' and vehicles' trajectories, taking into account the kinematic constraint of a vehicle's motion [115]. Therefore, as part of their network, they predict the vehicle's next input action as a multi-variant Gaussian distribution and use the vehicle's kinematic formulation to output the vehicle's next position given its current position and the predicted action distribution [115].

Pedestrians' future positions are predicted as a Gaussian mixture model (GMM) in [56]. In this mixed model, an extra parameter representing the weighting of each Gaussian component is also added to the output of the network besides the parameters of the Gaussian distributions themselves. Instead of Gaussian distribution, Su et al., proposed the use of a Cauchy distribution for the predicted position of pedestrians. As a result, they predict the location and scale parameters of a Cauchy distribution $(m_x, m_y, \gamma_x, \gamma_y)$ [176].

Using generative models to produce multiple trajectories

Another approach in the literature involves directly predicting position coordinates while considering the uncertainty of pedestrian trajectories by generating multiple outputs. This method acknowledges the multi-modality nature of pedestrians' trajectories, implying the existence of several acceptable future paths based on past trajectories [198, 221, 218, 16]. This is achieved using generative models such as variational auto-encoders (VAE) [197, 218, 16] or GAN [221, 198, 111].

In the use of CVAEs, each agent's future position is predicted given a history of its trajectory and a sample from a latent variable [197, 218, 16]. This latent variable itself is learned during the training phase by taking as input both the pedestrian's previous and future trajectories. When using the trained network for prediction, multiple samples from the latent variable will let producing multiple outputs for the trajectory of the pedestrian.

Trajectron is a trajectory prediction model introduced in the literature that combines a CVAE with recurrent models and spatio-temporal graphs [88]. The decoder takes samples of the latent variable from the CVAE framework and outputs the parameters of a GMM, effectively addressing the multi-modality of predicted trajectories. This model was further enhanced in its *Trajectron++* version by incorporating environmental information through semantic maps and considering dynamical constraints, which are crucial for trajectory prediction of agents with non-holonomic motions (e.g., vehicles) [163]. In *Trajectron++*, the decoder outputs the parameters of a bivariate Gaussian distribution over the agents' control actions, which are then mapped to positions using the agents' dynamical equations, thereby generating dynamically feasible trajectory predictions.

The ability of GAN to produce multiple outputs is leveraged in [221, 198, 111] as another way of satisfying the multi-modality nature of pedestrians' future trajectories. This is done in [221, 198, 111] by inputting a random noise, sampled from a Gaussian distribution, along with the other agent's encoded features to the generator module of the GAN. These different generated paths will then be evaluated in the discriminator of the GAN structure. In the discriminator module of the GAN in [39], a bi-variate Gaussian distribution is fit to these different samples produced by the generator and a joint probability distribution is created for storing these trajectories. The final predicted trajectory is then proposed to

be the most likely prediction of the joint probability density function.

Each of these generative models has its own strengths and weaknesses. GAN's are usually prone to the mode collapse problem where the network only generates one mode of trajectory with high variance, which does not thoroughly satisfy the multimodality criteria. CVAE models do not suffer from this issue by explicitly modelling multimodality, yet they present a challenge in defining additional loss criteria for proper training [220].

Therefore, for ensuring multimodality of the predicted trajectories, some models have combined the strength of both methods by proposing a hybrid CVAE-GAN model which adds the adversarial training benefit of GANs to the CVAE models [220, 114]. By replacing the GAN's generator module with a CVAE structure, the hybrid model is encouraged to learn a latent space that cannot be discriminated from the latent space of the real trajectories, and the overall model can produce more diverse and realistic trajectories [220, 114].

Other methods for modelling uncertainty

There are few other proposed methods in the literature that account for uncertainties of pedestrians' future positions [119, 7, 177, 116]. In the CA framework used in the [119], a probability is assigned to each possible reaction of the pedestrian towards a vehicle that suddenly ran into the pedestrian critical gap acceptance. Four reactions are considered for this model, namely going forward, rushing forward, waiting, or stepping backward [119].

In [7] the next position of the pedestrian is predicted as a probability distribution over the current position and speed of the person and some other parameters such as the risk of collision with a nearby vehicle.

Through the Bayesian Network framework used in [177], the probability of the selected action is calculated given the current position of the agents and its discrete latent variable. By sampling an action from this distribution, the current state of the agent can be propagated to the next time step.

Other methods such as adding noise to the predicted state of the agent are among the simple ways of accounting for the uncertainty of the pedestrian's predicted state [116].

In this thesis, we will address uncertainty in pedestrian trajectory behaviour by predicting the parameters of a distribution and incorporating these uncertainties into our learning-based navigation algorithm. We will also explore how to handle these uncertainties when integrating prediction and planning modules. To this end, the next section will briefly review related work on planning under prediction uncertainty, highlighting research gaps where we can contribute new knowledge and where our more detailed research questions are inspired.

3.6 Planning under prediction uncertainty

A crucial consideration in utilizing predicted pedestrian trajectories for robot motion planning is the inherent stochastic nature of pedestrian behaviours. It's important to acknowledge that the predicted trajectory may differ from the actual behaviour that will unfold in the future. To address this challenge, one potential strategy involves frequent updates to predictions as new information about the pedestrians' state becomes available over time. However, even prediction errors for a single step ahead, can adversely impact the actions generated by the downstream motion planner.

Different optimization methods for incorporating prediction uncertainty in navigation algorithms have been proposed in the literature. In this regard, some methods aim to ensure collision avoidance by considering the worst-case behaviour of other agents [203]. In [89], a MPC combined with a Dynamic Control Barrier Function (D-CBF) is used to enable a mobile robot to avoid colliding with a pedestrian. Their method employs a Kalman filter to predict the motion of other agents, with the increasing uncertainty of these agents' future positions modelled as a parameterized ellipse [89]. However, while these robust optimization methods provide provable safety guarantees, they tend to be overly conservative, often overreacting to unlikely future threats.

To address the issue of overly conservative behaviour, probabilistic optimization methods using chance constraints have been proposed for robot navigation in crowded, uncertain environments [5, 10, 2]. These methods aim to keep the probability of safety constraint violations within a specified confidence level. In [10], a combination of a Gaussian process with sampling-based reachability computation was employed to account for the multiple possible behaviours of surrounding dynamic agents. Given the higher complexity of evaluating chance constraints for multi-modal outputs of machine learning-based trajectory prediction algorithms, Ryu et al. [162] proposed using a distributionally robust chance constraint in a MPC formulation for robot crowd navigation. In their model, an ambiguity set is constructed from Monte Carlo samples of a machine learning-based multi-modal human trajectory prediction model [163], and an optimal control input is determined by ensuring constraint satisfaction for all distributions within the ambiguity set [162].

However, in the case of multi-modal prediction, minimizing a cost function over the entire planning horizon can still lead to conservative behaviours. To address this, contingency planning frameworks have been proposed, primarily for autonomous vehicles interacting with human-driven vehicles on roads. These frameworks plan trajectories conditioned on different possible outcomes of the prediction model [77, 46, 11, 37, 135].

In this context, some works have divided the planning problem into short-term and long-

term planning by defining a branching time over the planning horizon [77, 46, 11, 135]. This approach assumes that the ambiguity in the intent of other agents will be resolved by the branching time, allowing the AV to plan according to the determined intent of the other agent [135]. A shared trajectory is calculated to ensure safety for all possible intents of the surrounding agents until the branching time. After this point, the trajectory branches into different contingent plans for each possible intent of the other agents. The problem is formulated as an MPC optimization problem combined with a Bayesian filter to continuously update the AV's belief about the intent of other agents in [135].

Another line of work for planning under prediction uncertainty models the problem as a partially POMDP [85, 202, 124] and plans in the belief space based on a discrete set of other agents' intentions. However, the solutions to these problems become computationally intractable as the problem scales.

While various optimization methods for planning under the uncertainty of other agents' behaviour have been proposed in the literature, especially in combination with MPC methods, the incorporation of diverse predictions for pedestrians' trajectories and their associated uncertainties is limited in existing DRL-based crowd navigation methods. Monte Carlo tree search or DRL-based decision-making algorithms proposed for crowd navigation typically rely solely on the deterministic prediction of future pedestrian states, neglecting the nuanced uncertainties that exist in these predictions [30, 127, 122].

Among the few DRL crowd navigation algorithms aiming to accommodate the uncertainty of predicted pedestrian behaviour, Li et al., introduced a novel approach that integrates multiple plausible predictions generated by a multimodal prediction algorithm within the V-learning method [117]. In their methodology, the value function is assessed across all potential future predictions, with the optimal action chosen from a discrete set. However, despite considering multiple predictions, each anticipated trajectory entails a deterministic future position for pedestrians without assigning probabilities to these predictions. Moreover, the utilization of different predictions is confined to the value function calculation and does not extend to the formulation of the state or reward functions [117].

Katyal et al., proposed a DRL crowd navigation model that adapts its policy based on prediction errors from a pre-trained model, interpreting these errors as indicators of novel pedestrian behaviours [100]. When these errors exceed a threshold, the system switches from a fast navigation policy to a more cautious one, with the robot operating at a lower speed. This adaptation aims to avoid potential collisions, but reducing speed may not always be the best strategy for handling uncertainties in pedestrian behaviour.

To address this gap in considering prediction uncertainty within DRL crowd navigation methods, which inspired the **RQ4** of this thesis, we propose integrating uncertainties

associated with pedestrians' predicted trajectories directly into the observation state and reward formulation of a DRL algorithm that is coupled with the prediction model. Through a novel reward design, we can penalize the agent for actions that lead to a collision probability surpassing a predefined threshold, thereby encouraging the trained policy to minimize collision probabilities across the prediction horizon. This approach ensures that during training, the robot learns safe behaviour by accounting for the inherent uncertainties in pedestrians' predicted trajectories.

3.7 Summary of identified gaps

The discussion of related work in this chapter led us to identify gaps and areas where we could contribute to existing knowledge. This informed the research questions addressed in this thesis.

While there is a substantial body of literature on human-aware navigation, most studies focus on small mobile robots operating among humans. These robots are not highly safety-critical due to their small size and low operating speed. Research on human-aware navigation for AVs, whose operations near pedestrians are more safety-critical, has begun to gain interest, primarily concentrating on pedestrian-vehicle interactions in crossing scenarios. Few methods have been proposed for low-speed AV operations outside road structures in unstructured environments using expert-based models. Therefore, to the best of our knowledge, our work is the first DRL-based crowd navigation method developed for low-speed vehicle operations in shared unstructured spaces, with potential applications as mobility aids in pedestrian-rich environments such as airport terminals and shopping malls.

The identified gap in human-aware navigation for AVs in such settings motivated us to explore people's expectations of AV operations in these environments and what behaviours during these interactions would be perceived as safe and supportive of pedestrian comfort. This formed our first research question (**RQ1**).

We also identified a gap in understanding how pedestrian motion behaviour differs between unstructured and structured environments. This motivated us to investigate these differences to better understand our targeted operational environment and its unique characteristics, especially since no quantitative characterization of unstructured environments was available in the literature. This gap informed the second research question of this thesis (**RQ2**).

In developing the navigation algorithm, handling the complex interactions between the vehicle and multiple pedestrians in crowded settings is a significant challenge. Various

methods for interaction modelling have been proposed in the literature, but finding effective ways to capture these interactions still remains an open research question. Given the importance of algorithm scalability in crowded settings and the need to keep computational time bounded, we pursued learning-based approaches. However, we identified a gap between proactive navigation algorithms relying on trajectory prediction and DRL methods. Most DRL crowd navigation methods do not rely on prediction, limiting the foresight of their learned policies concerning interaction effects. This gap informed our third research question (**RQ3**). Among the few DRL methods that do consider prediction, we identified another gap in accounting for prediction uncertainty in the decision-making process. This led to the research question of how to incorporate uncertainty in pedestrian behaviour into the DRL navigation framework, considering this uncertainty in both pedestrian trajectory prediction and DRL-based motion planning (**RQ4**).

Additionally, in reviewing trajectory prediction models to inform the motion planner, we noticed a lack of data-driven models for unstructured environments shared with pedestrians. Most existing models are trained on datasets containing only pedestrians or data collected from road environments where pedestrian-vehicle interactions are minimal. Furthermore, among the proposed data-driven models in the literature, we identified room for improvement in the effective selection of interacting agents and the encoding of interaction effects. This informed other detailed aspects of our third research question (**RQ3**). Most methods in this area rely solely on distance measures or raw velocity data, overlooking the importance of approach direction and engineered features from velocity and distance vectors and heuristics that could better guide data-driven models in learning the interaction effects. This aspect will also be addressed in the design of our pedestrian trajectory prediction module.

Overall, we identified multiple gaps in the literature that informed the research questions addressed in this thesis. These gaps and our proposed solutions will be discussed in the subsequent sections.

Chapter 4

Social norms in unstructured environments: a survey

In this chapter, we investigate the design requirements for the perceived safe behaviour of a low-speed AV operating among pedestrians in unstructured environments through an exploratory online survey. The results from this study provided insights into factors to consider in the design of our decision-making algorithm.

Well-known social norms govern human-human interactions, while traffic rules regulate human-vehicle and vehicle-vehicle interactions on the roads. However, it is unclear if and how these conventions will translate to the operation of AVs in unstructured environments, which differ from typical road situations. Many studies in the literature have examined pedestrians' interactions with AVs in crossing zones and how AVs can communicate with pedestrians in the absence of a driver [154]. However, to the best of our knowledge, these studies have focused solely on structured road environments. It remains unclear whether the same results would apply to smaller, lower-speed AVs operating in unstructured, crowded environments, such as golf cart-shaped vehicles in airport terminals or shopping malls. In these unstructured, crowded environments, where no specific lanes or signs separate operational areas and no official regulations exist, behaviours might differ. Additionally, expectations might vary between a small AV (e.g., golf cart) and a small robot interacting with people in the same environment due to differences in size and speed.

Motivated by these exploratory questions, we conducted an online study at the beginning of the project to understand people's expectations and preferences regarding the behaviour of a low-speed AV, introduced in the survey as a golf cart, operating among pedestrians [67]. The COVID-19 pandemic necessitated online studies. As a result, we

used sketched interaction scenarios and asked participants about their opinions and feelings towards different possible behaviours of an autonomous golf cart in each scenario. This study helped us identify conventions and social norms that people expect the AV to follow, thereby making its operation among humans more acceptable.

4.1 Research questions

The detailed research questions for the conducted study in this chapter, as part of the thesis's first main research questions (**RQ1**), are as follows:

RQ1.1 Are there general patterns in the way people (assuming the roles of pedestrians and passengers) expect the cart to behave when it is operating in an unstructured environment like a shopping mall or an airport terminal?

RQ1.2 For feeling safe, do people (assuming the role of pedestrians) expect the cart to follow rules that are similar to road traffic rules for passing and overtaking despite the environment being unstructured?

RQ1.3 How safe and comfortable do people feel when the cart tries to negotiate space with them for passing or crossing? Are they willing to share their space with these carts and let them pass? What do they think about the priorities in these shared-space environments?

As the nature of this study is exploratory, we did not formulate any hypotheses.

4.2 Method

We designed an online study in which we described and sketched nine common interaction scenarios that could occur between an autonomous golf cart and one or more pedestrians, including scenarios such as passing, overtaking, crossing, and turning. The details of each scenario are shown in Table 4.1.

In each interaction scenario, we considered possible behaviours for either the cart or the pedestrian in the sketched scene. The possible designed behaviours for each scenario differed and were limited to the cart's path selection relative to the crowd, the direction

Table 4.1: The main scenarios of the survey. Upper-case letters are behaviours for which a separate questionnaire was completed by the participants. The lower-case letters show the multiple-choice questions.

No.	Scenario	Behaviours	No.	Scenario	Behaviours
1	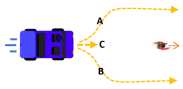	A. Overtaking on the left. B. Overtaking on the right. C. Following the pedestrian.	2	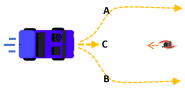	A. Passing on the left. B. Passing on the right. C. Wait for the pedestrian to clear the path.
3	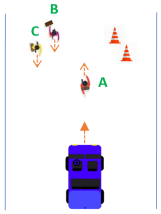	For pedestrians in place of A: a. Clearing cart's path by moving to right. b. Clearing cart's path by moving to left. c. Keep walking straight until a place where the cart can overtake me. d. Paying no attention e. Others For passengers: a. Expecting A to clear the cart's path. b. Expecting B and C to clear the cart's path. c. Following A until getting to a place where it could overtake. d. Others	4	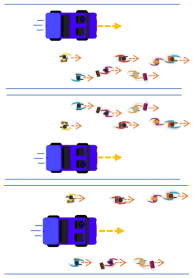	a. Moving on the left. b. Moving on the right. c. Moving in the middle. d. No preference. e. Others
5	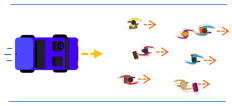	A. Following the flow of pedestrians. B. Requesting a way through	6	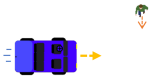	A. Yielding to the pedestrian by stopping. B. Yielding to the pedestrian by slowing down. C. Adjusting path to the left or right while keeping the same speed. D. Taking the priority to cross by just continuing the path without decelerating. E. Using visual and auditory signal for showing the intention to cross first and taking the priority to cross.
7		a. Following group A. b. Opening a passage in the middle. c. Opening a passage on the left side of the corridor. d. Opening a passage on the right side of the corridor. e. Others	8	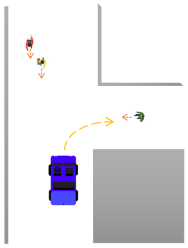	A. Showing turn intention implicitly only through movement. B. Showing turn intention using flashing lights. C. Showing turn intention by projecting the future path on the floor.
9		A. Stopping and waiting for the child to get into a safe zone. B. Decelerating and changing path to the right or left for avoiding the child.	Pre		a. Passing on the left. b. Passing on the right. c. No preferred direction.

Table 4.2: Designed questionnaire for scenario 5 with a continuous Likert scale from strongly disagree to strongly agree.

Statement
I feel comfortable with this behaviour of the cart
I think this behaviour is dangerous
I think this behaviour is appropriate
I would not have expected this behaviour
I would trust a cart that shows this behaviour
The cart did not seem to be aware of me (the pedestrian) when it showed this behaviour

the cart passes or overtakes a pedestrian or other choices such as deceleration, stopping and path changing for collision avoidance, as well as option of the cart requesting a pass.

We also studied the effect of people's roles in the scene on the perceived safety of the cart's behaviour. We considered two different roles: a pedestrian walking close to the cart, or being a passenger on the cart.

Including people's roles and the cart's behaviour choices, we had a mixed factorial design in each of our scenarios. The participant's role was a between-subject factor with two levels (pedestrian and passenger), while the cart's behaviour was a within subject factor with the number of available choices ranging from 2 to 5 depending on the scenario (Table 4.1).

4.2.1 Procedure and measures

After giving consent and answering demographic questions (i.e., age, gender, country of residence), the participants were randomly assigned to either the role of a pedestrian or a passenger on the cart. The participants were requested to imagine themselves in the scene in the role they were given and answer the questions from that perspective. An instruction page gave each participant a general view about the concept of the passenger-carrying golf cart operating autonomously in a shopping mall or an airport terminal among the crowd, along with some real-world images of these golf carts in an indoor environment. The golf cart was described as an electric vehicle with a max speed of 9.5 mph.

The participants then answered two pre-survey questions about a human-human interaction scenario, asking about their preferred direction for passing or overtaking another person. The participants then went through the 9 designed interaction scenarios one by one. These scenarios are shown in Table 4.1. For scenarios 1, 2, 5, 6, 8, and 9, the participants completed the perceived safety sub-scale of the Godspeed questionnaire for each

of the cart's behaviours. For scenario 9, the subscale of perceived intelligence was also provided. In scenario 5, along with the perceived safety subscale, we also added six other Likert scale questions shown in Table 4.2.

Participants evaluated scenarios 3, 4, and 7 in Table 4.1 in the form of a multiple-choice questions where the participants responded by choosing all the behavioural options they preferred. In the multiple-choice questions, we also asked participants to explain why they selected a specific option. We used multiple choice questions for these scenarios instead of the Godspeed questionnaire, as we were interested in scenario-specific answers, focusing on the norms related to the operation of the vehicle.

The survey ended with a set of questions, where the participants were asked about their general concerns towards the operation of these carts in the described environments, as well as their specific thoughts on the interaction of the cart with children. We included this question in order to probe participants' possible concerns regarding vulnerable populations, in this case restricted to children. We also asked which approach direction may be most concerning when the car is getting close to them. At the end of the study, the participants saw a *thank you* page and were given an end code to submit their Human Intelligence Task (HIT) on Amazon Mechanical Turk.

4.2.2 Participants

We recruited 131 participants on Amazon Mechanical Turk. The recruitment was limited to US and Canadian workers with an approval rate of 95% or higher based on at least 50 HITs. Fifteen participants failed the attention check question. Therefore, the final data was based on 116 participants' responses (49 Female, 66 Male, 1 Unknown; age range: 18-74). 57 participants were randomly assigned to the role of pedestrian, while the other 59 participants were assigned to the role of passenger. The survey took on average 24 minutes to complete and participants were paid 2 USD upon completing the study and a pro-rated amount if they wished to stop earlier. The study received Ethics clearance from the University of Waterloo Research Ethics Committee (see Fig. A.1).

4.2.3 Data analysis approach

Quantitative data: LMMs were used for studying significant differences among the perceived safety of the distinct behaviours for the cart in each scenario and between the two different roles. We took potential confounding factors into account and also searched for potential interaction effects in the data. For coming up with one single number for the

perceived safety of each behaviour, we averaged the scores of the three questions in the perceived safety subscale of the Godspeed questionnaire while reversing the score of the last question for making it consistent with the other two questions. For multiple choice questions, we performed a Chi-square test to check if there was a significant difference between the selected choices and, in case of a significant difference, performed pairwise comparisons.

Qualitative data: A thematic analysis was conducted to examine participants' reasoning for their choices in scenarios 3 and 4, as captured in the open-ended response sections of the questionnaire. Following the step-by-step guidelines for thematic analysis outlined in [21], I manually generated initial codes based on participants' stated reasons for their selected behaviours. These codes were then organized and grouped into relevant themes, with associated codes collated under each theme. After reviewing and refining the themes, I assigned descriptive labels to them. Finally, the frequency of each theme in participants' responses was counted. The results were categorized based on the selected choices and are presented in Tables 4.3-4.5.

4.3 Results

In our single pedestrian-cart passing and overtaking scenarios (scenarios 1 and 2), the passenger group perceived passing on the right much safer than the cart waiting for the pedestrian to clear the path ($se = 0.21, t = 3.44, p < 0.01$). We did not observe any significant difference in the pedestrian group.

Figure 4.1 shows the result of the three multiple choice scenarios (scenarios 3, 4, and 7). In scenario 3, where the cart gets blocked behind a single pedestrian, significant differences were found between the preferred choices in both groups ($Pedestrians : {X^2}_{(3,N=59)} = 45.34, p < 0.001; Passengers : {X^2}_{(3,N=57)} = 100.68, p < 0.001$). For the pedestrian group, the "clearing the cart's path by moving to the right" behaviour was selected significantly more than "moving to the left" ($p < 0.001$) and also "Keep walking straight" ($p < 0.001$). In the same scenario, the passenger group selected the behaviour of "staying behind the pedestrian" significantly more than expecting pedestrian A ($p < 0.001$) or pedestrians B and C ($p < 0.001$) to clear the cart's path.

Based on thematic analysis of participants' reasons for their chosen behavioural options, we identified distinct themes for each choice. For example, among participants in the role of pedestrians, four main themes emerged for the decision to clear the cart's path in scenario 3: 1) Feeling uncomfortable being followed, 2) Avoiding the risk of getting hit, 3)

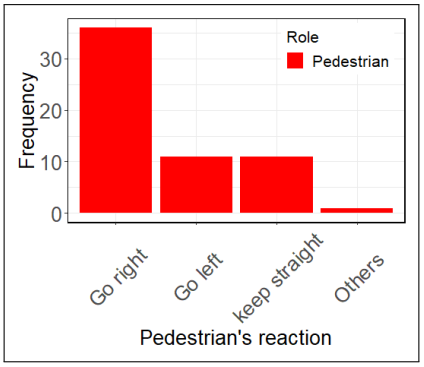

(a) scenario 3 - pedestrians

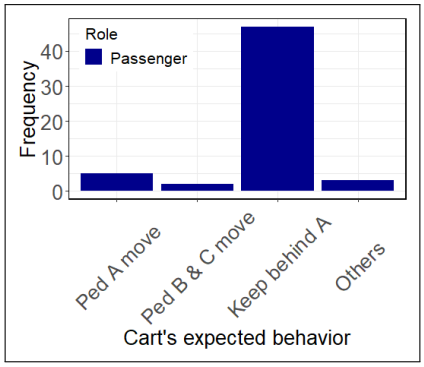

(b) scenario 3 - passengers

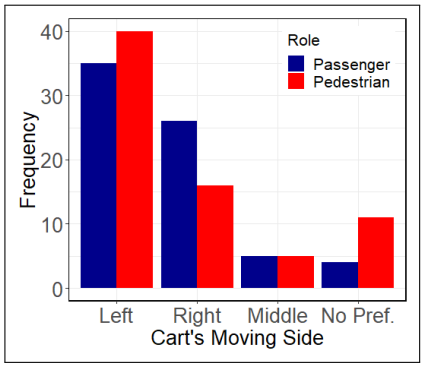

(c) scenario 4

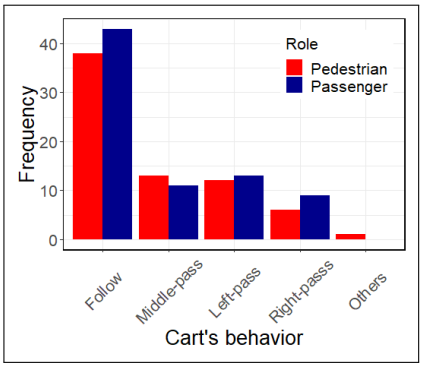

(d) scenario 7

Figure 4.1: The results of the multiple choice question scenarios. (a) Pedestrian group's selected reactions about the cart approaching from behind in Scenario 3 (see Table 4.1). (b) Expectation of the passenger group in the same situation. "Ped" stands for pedestrian. (c) Participants' preferred side for the cart's operation in a corridor with an unidirectional flow of pedestrians as in scenario 4. (d) Participants' expectations of the cart when it gets behind a bidirectional flow of pedestrians as in scenario 7.

Acknowledging the cart's right of way, and 4) Allowing a faster vehicle to overtake. These themes and the corresponding number of responses for each are summarized in Table 4.3. The identified themes for participants' reasons for moving either left or right when clearing the cart's path, as well as for the alternative behaviour of continuing to walk straight, are also detailed in the table.

Participants in the pedestrian group mainly chose to clear the cart's path as they were uncomfortable with and anxious about the cart's presence behind them, and they found it

Table 4.3: Thematic analysis results of pedestrians' reasoning in Scenario 3 where the cart is positioned behind a pedestrian with their left side occupied by other pedestrians and their right side by pylons. The table reports the frequency of mentions for each identified theme in the reasons provided for the chosen action. "Moving to left/right" are sub-categories under the broader choice of clearing the cart's path.

Choice	**Themes**	**#**	**Choice**	**Themes**	**#**
Clearing cart's path	Uncomfortable being followed	7	Moving left	Viewing pylons as a tighter restriction	5
	Risk of getting hit	6		Feel safer joining other pedestrians	2
	Cart's right of way	2		Traffic rules (passing on right)	1
	Letting faster vehicle overtake	7	Moving right	Not disturbing other pedestrians	14
				Traffic rules (faster vehicle on left)	5
				Clear of people	2
				Intuitive	2
Keep walking straight	Priority rule	2			
	No room on either side	7			

safer for themselves to clear the cart's path to reduce the risk of getting hit. For those who chose to move to the right, the most frequently mentioned theme was avoiding disturbing other pedestrians walking in the opposite direction on the left. Furthermore, the theme of allowing faster vehicles to overtake on the left was the second most frequently cited reason for this choice (Table 4.3). On the other hand, for participants in the role of a passenger in the same scenario, we identified several themes for their decision to stay behind the pedestrian. The most frequently mentioned themes included the belief that staying behind was the safest option and that the cart should not rely on people getting out of its way. Another recurring theme was the lack of free space to take other actions, along with concerns that pedestrians might not notice the cart. Participants emphasized that the cart should remain behind the pedestrians until the left path cleared (see Table 4.4).

In a unidirectional flow of pedestrians in a corridor (scenario 4), the cart moving on the left side of the corridor was preferred significantly more than moving on the right ($p < 0.01$) and moving in the middle ($p < 0.001$) by the participants in the pedestrian group as shown in Fig. 4.1(c). In the passenger group, moving in the middle was significantly less preferred as compared with moving on the left ($p < 0.001$) and moving on the right ($p < 0.001$). However, we did not observe any significant difference between the left and right passage of the cart in the passenger group. Results are summarized in Fig. 4.1(c). Obeying the traffic rule that faster vehicles drive on the left emerged as the most frequently mentioned theme for preferring the cart's left passage in the provided reasons (see Table 4.5). Thematic analysis of the reasons also suggested that the participants' concern about the cart passing

Table 4.4: Thematic analysis results on passengers' reasons for choosing to stay behind person A in scenario 3.

Themes	**Mentions**
Safest option	22
The cart shouldn't rely on people getting out of way	11
No free space for other actions	9
Bothers no one	6
Pedestrian might not realize the cart	6
Wait for the left path to get clear	4
Acting like a pedestrian in the same situation	3
Pedestrians' right of way (priority)	3

from the middle was due to the high risk of this option.

In a bidirectional flow of pedestrians (scenario 7), there was a strong preference in both groups (passengers and pedestrians) for the option where the cart follows the flow of pedestrians walking in the same direction as the cart (group A), compared to all other options of opening a passage through the flow of pedestrians whether from the right ($p < 0.001$), left ($p < 0.001$), or through the middle ($p < 0.001$). Results are shown in Fig. 4.1(d). No significant difference was found among the other choices (i.e., left, right, and middle passages).

When the cart gets behind an unidirectional flow of pedestrians (scenario 5), the results suggest that requesting a way through is perceived significantly less safe than when the cart just follows the flow of pedestrian ($se = 0.16, t = -3.81, p < 0.001$). Participants in the pedestrians group perceived both behaviours less safe than the passengers ($se = 0.18, t = -3.07, p < 0.01$).

Further, the passengers found the cart's request for a path through significantly more dangerous ($p < 0.001$) and inappropriate ($p < 0.01$), as compared with following the flow of the pedestrians, and rated to be more uncomfortable within that situation ($p < 0.05$). Passengers also did not expect the cart to request a way through ($p < 0.001$) and their trust level decreased significantly when the cart requested for a path, as compared to when it followed the flow ($p < 0.001$). We did not see any significant difference between the preferences of the pedestrians, except that they were not expecting the cart to request a way through compared to staying behind them ($p < 0.01$). For all the additional questions, except for the one on awareness, the pedestrian group's ratings were significantly lower than the passenger group's ratings when the cart's behaviour was to follow the crowd.

Comparing the perceived safety of the different behaviours in a ninety-degree approach scenario (scenario 6), the cart's lateral path change was perceived as significantly less safe

than yielding to the pedestrian whether by a complete stop ($se = 0.11, t = 16.36, p < 0.001$) or by slowing down ($se = 0.11, t = 7.53, p < 0.001$). Signaling the pedestrian while taking the priority had a significant positive effect on the perceived safety of the cart compared to the same behaviour with no signaling ($se = 0.11, t = -6.82, p < 0.001$).

The effect of using explicit cues for communicating the intent, compared to implicit cues, was also observed in the turn scenario (scenario 8). Using both flashing lights and a projected path significantly and positively affected the safety ratings ($se = 0.14, t = 10.84, p < 0.001$ and $se = 0.14, t = 6.4, p < 0.001$, respectively). While the passenger group rated flashing lights to be safer than the projected path ($se = 0.14, t = 4.43, p < 0.001$), we did not find a difference in pedestrians group's perception.

When encountering a child (scenario 9), the cart's complete stop was perceived not only significantly safer ($se = 0.13, t = 11.05, p < 0.001$), but also more intelligent ($se = 0.12, t = 12.27, p < 0.001$) compared to the behaviour where the cart just adjusts its path to avoid the child. Moreover, passengers had a safer feeling than pedestrians towards the cart's complete stop in this scenario ($se = 0.19, t = 2.36, p < 0.05$).

The results of our pre-survey and post-survey questions are shown in Fig. 4.2. In human-human interaction, there was a significant difference in the preferred passing direction ($X^2{}_{(2,N=116)} = 94.88, p < 0.001$). Participants had a clear preference for passing another person on the right compared to passing on the left ($p < 0.001$). The highest general concerns about the operation of the cart were (a) the presence of children and (b) the cart's approaching sound not being noticeable, respectively. The results also showed that participants were concerned about the cart approaching from behind more than the front

Table 4.5: Thematic analysis results on participants' reasons for their preferred operational side of the cart in a corridor in scenario 4.

Choice	**Themes**	**Mentions**
The cart moving on left	Traffic rule: Faster vehicle on the left/overtaking on the left	37
	Negative feeling about the cart going through the middle	8
	More intuitive	4
	Similar to the structure of sidewalk-road	3
The cart moving on right	Traffic rule: the cart should keep on right	7
	Negative feeling about the cart going through the middle	7
	Safer	3
Selecting both left and right	Lower risk of hitting someone	3
	Safer to have pedestrian on one side	3
	Better to keep pedestrians and the cart separate	3
	Uncomfortable with the cart in the middle	2

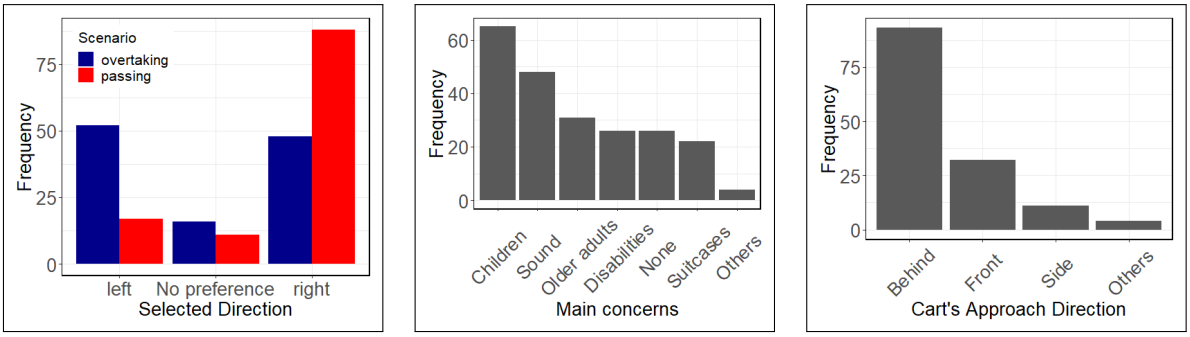

(a) Human-Human interaction (b) Main Concerns (c) Approach Direction

Figure 4.2: The results of the pre-survey and post-survey questions.

($p < 0.001$) and side ($p < 0.001$) approach directions as shown in Fig. 4.2(c).

4.4 Discussion

We conducted an online study to understand people's preferences regarding the behaviour of an autonomous, low speed cart moving among pedestrians. By showing participants sketches of possible interaction scenarios between a passenger cart and pedestrians in an unstructured environment, and providing different options for the cart's behaviour in each scenario, we assessed people's expectations of the cart's behaviour. Participants were randomly assigned to two groups, i,e. passengers of the cart or pedestrians. We also studied if the traffic rules are expected to be obeyed by the cart despite the unstructured nature of the environment.

While our human-human interaction scenario suggested that participants had clear preferences about the passing direction, we did not observe significant differences between participants' preferences for passing direction in cart-pedestrian interaction (scenario 2). However, when more than one pedestrian was involved, participants showed a significant tendency to obey traffic rules of passing on the right and overtaking on the left. Speed difference seems to be a major factor governing people's preferences for the overtaking direction, as the participants mentioned the common rule on the road of higher speed vehicles moving on the left for their reasons (Table 4.5). This can also explain the non-significant difference for overtaking direction in the human-human interaction scenario as humans' speeds are not much different from each other.

Participants' seemed to prefer following traffic rules even in unstructured environments,

as shown in our scenarios. For example, the cart taking the left side passage was preferred, which can be similar to the traffic rule that requires faster vehicles to move on the left. Even those who chose the right passage for the cart were referring to a traffic rule that cars generally drive on the right side, viewing the scenario more as a passing situation rather than overtaking. Also, the participants in the passenger group appeared to have a general belief that the cart should not rely on people getting out of its way. Instead, it was perceived the safest for the cart to stay behind the pedestrian/s until the left path gets clear. However, in the pedestrian group, most participants chose to clear the cart's path despite the presence of obstacles on both sides. Reasons provided by participants showed that participants in the pedestrian group felt uncomfortable about being followed by the cart and they found it safer for themselves to clear the cart's path to reduce the risk of getting hit.

So, by introducing a passenger-carrying cart even in an unstructured environment, participants seem to still view the scene as a structured environment (e.g., a road) and prefer the cart to follow the traffic rules, which can increase the perceived safety of the cart (Addressing **RQ1.1** and **RQ1.2**). It is likely that participants have shown this preference as the operation of a motorized vehicle in these unstructured environment is relatively new and the most familiar interactions similar to that are the encounters that happen on the roads. So participants might be carrying over their expectation from the roads to these unstructured environments.

Further, regarding participants' preferences about the behaviour of the cart when negotiating space with pedestrians (**RQ1.3**), we found that there is a high tendency to separate the vehicle's path from the pedestrians' path as much as possible for reducing the risk of injuring someone (scenario 4). This was shown in participant's stated reasons for not selecting the middle passage for the cart in scenario 4 (Table 4.5). Also, people were mostly expecting the cart to just follow the flow of pedestrians without disturbing them. This was the case for both a unidirectional flow of pedestrian (scenario 5) and a bidirectional flow of pedestrian (scenario 7). The responses to the scenarios which simulated situations where the cart gets blocked behind pedestrian(s) (scenarios 3, 5 and 7) have shed some light on how people think about the priorities in these unstructured environments (Addressing **RQ1.3**). The responses showed that people may not necessarily give priority to the vehicle by clearing the path for the cart despite its higher operational speed.

Also, as the presence of children was the mostly mentioned concern people had, and also because children's movements, if moving freely, are highly unpredictable, it might be better for the cart to do a complete stop when encountering children (also suggested by scenario 9). The use of some type of alert was also frequently recommended by the participants.

Overall, based on our results we can extract the following tentative norms, which may serve as design suggestions for a navigation algorithm of an autonomous, passenger-carrying cart operating in an unstructured, crowded environment, with the ultimate goal to help increase the perceived safety and acceptability of the cart.

- Obeying common traffic rules of overtaking on the left and passing on the right when interacting with a crowd.
- Keeping the cart's path separate from other pedestrians as much as possible.
- Using explicit auditory and visual signal for communicating intent.
- Using velocity change instead of lateral deviation of the path for collision avoidance.
- Coming to a complete stop instead of bypassing when encountering a child that is running about.

As the above mentioned behaviours were perceived safer by the participants in both passenger and pedestrian groups, paying attention to these behaviours could help improve the cart's acceptance in the crowd.

4.4.1 Limitations

This study clearly had multiple limitations. Due to the COVID-19 situation at the time we were not able to conduct any in-person studies with a real golf cart. Our sketched-based scenarios are likely to have affected the results, as imagining situations based on the sketches would be different from realistic interactions. Also, in those sketches used, it was difficult to provide rich contextual information on the nature of the unstructured environments, the speed and size of the cart, its specific design and communication features (e.g. lights), and most importantly three-dimensional and temporal information which may be crucial in real encounters of passengers or pedestrians with an autonomous golf car. This could raise doubts on whether the participants had indeed imagined themselves in a truly unstructured environment and also in the correct role by just seeing the top-view of the abstract sketches of the scene, despite clearly pointing to a specific pedestrian in that scene. Future research needs to investigate if the results of this study would translate to real world interactions while also including a wider range of interaction scenarios. It would be also interesting to further study how much the autonomous nature of the cart influenced the responses by comparing it to a human-driven cart in the same scenarios. The type of preferences and social norms studied in this work are also highly cultural dependent and this may limit our results to US and Canada, or countries that follow a right-hand traffic rule. Despite these limitation, this preliminary study served as a starting point for

understanding people's expectation of these carts in an unstructured environment, in order to design autonomous, low-speed vehicles that can successfully operate in crowded areas.

4.5 Summary and conclusion

We studied people's expectations regarding how an autonomous, low-speed, golf-cart type vehicle should navigate in an unstructured, crowded environment through an online survey. Our results suggested that the cart is not expected to disturb the flow of the pedestrians to get through the crowd, but, whenever it wants to pass a group of pedestrians, it is expected to follow common traffic rules. While following this norm was considered to be important in the interaction of the cart with a group of pedestrians, we did not find any preference for passing or overtaking direction in a single pedestrian-cart interaction, especially when there was space on both sides of the pedestrian. Results also suggested that pedestrians may mostly react to a cart behind them by clearing its path as they do not feel comfortable about the cart's presence behind their backs. This was also consistent with participants' most concerned approach direction of the cart, i.e. being from behind. Furthermore, in a 90-degree approaching interaction, participants perceived speed adjustment by the cart safer than any heading change for collision avoidance.

While these results are not entirely surprising and need to be confirmed in future more realistic scenarios and settings, they form a useful starting point for our quest to understand the norms that people expect a golf cart to follow when operating in unstructured and crowded environments.

The findings from this study inspired the design of our decision-making algorithm to encourage the AV to exhibit behaviours similar to those expected by humans, as identified in this online survey. Specifically, these findings guided the formulation of the reward function for our DRL-based decision-making algorithm.

Additionally, we will rely on common pedestrian reactions observed in the survey scenarios to create similar behaviours and evaluate the performance of our algorithm later in Chapter 7. This evaluation will determine if the behaviour generated by our algorithm aligns with the expected behaviour identified in the online survey.

After exploring the expectations for the vehicle's behaviour in this chapter as part of our design requirements study, we will shift focus to pedestrians in the next chapter, specifically examining their common behavioural patterns in unstructured environments. This will provide a better understanding of the characteristics of the environment in which our AV will operate.

Chapter 5

Characterizing unstructured environments

Before delving into the design of the navigation algorithm, we conduct a preliminary study to better understand the unstructured environment in which our AV will operate and its key distinctions from structured environments. This study marks the second stage of our design requirements phase, focusing on analyzing pedestrians' trajectory behaviours across various environment types to identify potential differences between structured and unstructured settings.

Many data-driven trajectory prediction models have been proposed for the use of autonomous transportation systems. These models are usually trained on datasets that contain real-world trajectories of pedestrians and vehicles in an environment. Depending on the type of the environment these datasets have been captured from, the trained predictive model will also learn the motion behaviour of traffic participants within the same type of environment. Paying attention to this point is crucial when it comes to generalizing the trained data-driven motion predictor to other similar environments.

However, while it is valuable to classify the datasets based on the type of environment (i.e., unstructured versus structured), it is hard to do this classification based on the different qualitative definitions provided in the literature, as reviewed in section 3.2. Therefore, here we investigate how to quantitatively define the type of environments through evaluating quantitative features for classifying environments to unstructured and structured, by focusing on the traffic participants' behaviours embedded in the trajectories.

Based on the existing literature, we assume that traffic participants have similar behaviours in each type of environment, and based on this similarity we will suggest a re-

lationship between the extracted quantitative features and the structured/unstructured nature of the environment.

5.1 Research questions

This chapter explores three key sub-research questions that are related to the second main research question (**RQ2**) of this thesis:

RQ2.1 What are the notable differences between pedestrians' and vehicles' trajectory behaviours in structured versus unstructured environments?

RQ2.2 How can these potential differences quantitatively characterize structured versus unstructured environments, complementing qualitative distinctions?

RQ2.3 How can existing datasets on agents' (i.e., pedestrians and vehicles) trajectories be categorized based on these quantitative measures?

Answering these questions will aid us in understanding the unique characteristics of trajectory behaviours in unstructured environments that should be considered in the design of the navigation algorithm. Additionally, it will help us categorize the existing dataset of pedestrian trajectory data based on the nature of the environment. This classification is crucial as we progress in predicting pedestrian trajectories, guiding the selection of appropriate datasets for training our data-driven predictor. Ultimately, our goal is for the trained model to effectively capture pedestrian behaviour patterns in unstructured environments in the presence of vehicles, which is the focus of this thesis.

5.2 Method

To classify environments that include both pedestrians and vehicles into structured and unstructured categories, we first gather several datasets commonly used in pedestrian prediction that contain both pedestrians and vehicles. We then compare these datasets based on a few extracted trajectory features, such as mean speed and trajectory variability, using two methods: (a) k-means clustering to determine if the datasets can be divided into two clusters representing structured vs. unstructured environments and (b) a generalized linear model to investigate if our assumptions about the effect of features on identifying the type of environment hold true. Through this investigation, we propose quantitative measures for distinguishing between the two types of environments.

5.2.1 Existing trajectory datasets

There are a few datasets that include both pedestrians and vehicles which are analyzed in this paper, discussed below.

Stanford Drone Dataset (SDD) is a dataset that has been captured from the campus of Stanford University and consists of eight different scenes [159]. The dataset includes trajectories of around 19k agents in different classes, from pedestrians, bicyclists, and skateboarders to cars, carts, and buses. Here, we have analyzed the trajectories of only pedestrians, cars and carts within four scenes of this dataset (hyang, coupa, nexus, death circle) that include a relatively higher percentage of car/cart compared to the other scenes. In this dataset, we filtered out the annotations with "lost" value of 1 which indicates annotations outside the view screen.

Intersection Drone Dataset (inD) contains trajectories of more than 11K traffic participants, including pedestrians, vehicles, and bicycles [19]. The dataset was recorded from four different German intersections, two of which were used in the analysis of the current chapter.

Dalian University of Technology (DUT) dataset contains trajectories of pedestrians and vehicles interacting with each other in two locations on the campus of Dalian University of Technology in China with close to 2K pedestrian and 65 vehicle trajectory data [206]. One of the locations is an intersection without a traffic light. The other location is a shared space near a roundabout where agents can move freely.

Hamburg Bergedorf Station (HBS) dataset is captured from a shared space street in Germany near a busy railway station [140]. It includes trajectory data of 1115 pedestrians, 331 cars and 29 cyclists.

Hannover Campus (HC) dataset was captured from a shared street in a university campus in Germany [38]. It contains trajectory data of 282 pedestrians, 15 vehicles and 40 cyclists.

All the above datasets have been captured from stationary cameras or drones. In each of these datasets, the position track of different agents is provided frame by frame. These positions were all transformed to a world cartesian coordinate before being used for analysis.

Three of the above datasets (i.e. HBS, HC and the roundabout scene of the DUT dataset) are collected from shared spaces environments which are defined as urban environments where the curbs, road surface markings and traffic signs have been removed to encourage different traffic participants to share the same space and negotiate the priority based on social rules [147].

The secondary use of the above datasets for this study has received ethics clearance from the University of Waterloo Research Ethics Committee.

5.2.2 Extracted trajectory features

Pedestrian motion features

The features that we extracted from each individual pedestrian's trajectory were: (a) mean speed of pedestrian, (b) stop fraction of pedestrian, (c) variability of pedestrian's trajectory, (d) path efficiency of pedestrian, (e) pedestrians' path orientation entropy, (f) density of pedestrians, and (g) density of standing pedestrian. The annotated x-y positions of agents are already available in these datasets. Using this positional data over time, we developed a programming script that calculates the specified features from the trajectory data.

The first three features are related to the speed profile of the pedestrian. We calculated the velocity vector from the pedestrian's position vector in consecutive frames knowing the frame rate through which the dataset was captured. For the mean speed feature, we calculated the average speed of the pedestrian during its whole trajectory. For the stop fraction, we computed the percentage of time the pedestrian is in a stop along its whole trajectory. We have considered a pedestrian to be in a stopping position when their speed gets lower than 0.5 m/s [138]. Variability of trajectory is defined as the trace of the covariance matrix of the speed [138] and is calculated according to the equation below:

$$tr(cov(\overrightarrow{V})) = \frac{1}{n}\sum_{t}(v_t - \bar{v})^2 \tag{5.1}$$

In this equation, $\overrightarrow{V}$ is a vector of speeds in different time steps t, ($\overrightarrow{V} = [v_1, v_2, ..., v_n]$), $\bar{v}$ is the average speed and n is the total number of time steps that the pedestrian is present in the scene.

We calculated path efficiency as the ratio of the distance between the endpoints of the trajectory over the trajectory length and indicates how close the path is to a straight line [6]. Some of the datasets have a wider perspective of the capturing scene and therefore in those

© 2019 IEEE
(a) DUT - intersection [206]

© 2019 IEEE
(b) DUT - roundabout [206]

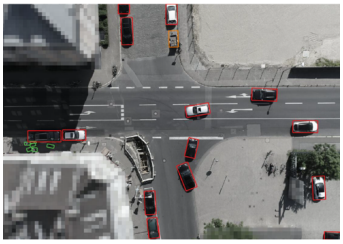

© 2020 IEEE
(c) inD - scene 1 [19]

© 2020 IEEE
(d) inD - scene 4 [19]

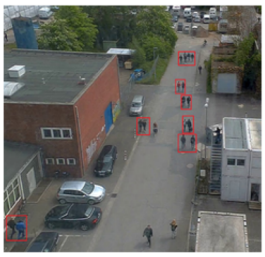

© 2019 IEEE
(e) HC [38]

(f) HBS [140]

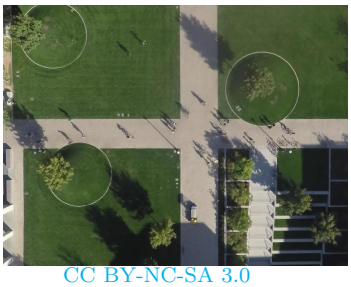

CC BY-NC-SA 3.0
(g) SDD - hyang [159]

CC BY-NC-SA 3.0
(h) SDD - death circle [159]

CC BY-NC-SA 3.0
(i) SDD - nexus [159]

CC BY-NC-SA 3.0
(j) SDD - coupa [159]

Figure 5.1: Images from the existing datasets used to characterize structured versus unstructured environments. These images are sourced either from the original papers introducing the datasets or from screenshots of dataset videos provided by the authors. See Appendix B for copyright permission.

scenes, the pedestrians' trajectories can deviate from a straight line due to some geometrical constraints in the scene, such as the buildings. To compensate for that in the calculation of the path efficiency, we divide the trajectories into smaller sections, each lasting 4.8 seconds, and the path efficiency is calculated using the cumulative value of the smaller sections of the trajectory within the formulation. These smaller pieces of the trajectory that last 4.8 seconds are called trajlet in [6] and are commonly used in trajectory prediction as the observation interval of each agents' trajectory for future position prediction.

For the path orientation feature, we calculate the angle of the straight line connecting the start and endpoint of each trajlet relative to a global horizon. However, the exact number of path orientations is not important and cannot be compared between the different datasets due to the inconsistency of the orientation of their global Cartesian frame. In fact, what is important is the diversity of the path orientation within each environment that could represent the pedestrians' level of freedom in selecting their movement direction. For measuring this diversity, we calculate the entropy of the path orientation for each dataset and keep it as a feature.

All the above features are calculated for pedestrians who are not stationary throughout the whole dataset. Stationary pedestrians are considered as those with a stop fraction of more than 90% and are excluded from the above feature calculations as those features are not meaningful for stationary pedestrians, and the calculations are prone to high error for pedestrians that have a speed of more than 0.5 m/s only at 10% of their whole trajectory. However, for calculating the density of the pedestrians, all the moving and standing pedestrians are being considered and their total number is divided by the area of the scene from which the dataset is being captured. As the stationary pedestrians were ignored in the first 5 features, we added a last feature that specifies the density of standing pedestrians in the dataset. These two density features are calculated for each pedestrian as the average

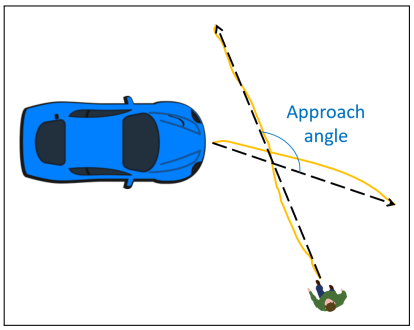

Figure 5.2: Approach angle definition in a pedestrian-vehicle interaction.

density while the pedestrian is present in the scene since the density of the crowd could affect each individual's motion behaviour.

Vehicle motion features

From the vehicle trajectory data, we extracted: (a) the mean speed of the vehicle, (b) stop fraction of the vehicle, and (c) variability of vehicle's trajectory.

We defined and calculated these features in the same way as in the pedestrian features. The only difference is in the stop fraction feature, where a vehicle is considered to be in a stop when its speed is less than 1 m/s. Same as in the pedestrians feature, here we also exclude vehicles that have a mean speed of less than 0.5 m/s throughout their whole presence which is used as an indication for a parked vehicle.

Pedestrian-vehicle interaction features

The way pedestrians and vehicles interact in a structured vs. unstructured environment can also be very different. In an unstructured environment, as there are no pre-specified paths or lane marks present [101], pedestrians and vehicles can approach one another from different directions and since no strict rules apply in unstructured environments [97], different behaviours can be observed in terms of which class of agent (pedestrians or vehicles) gets the priority when their paths cross one another. Therefore, (a) entropy of approach angle, (b) priority (percentage of pedestrian priority), and (c) ratio of vehicles to pedestrians were the three features that we defined for capturing these differences.

For extracting these behaviours as interaction features, we first focus on those pairs of pedestrian-vehicle that get closer to each other than a threshold that we have selected to be 4 meters here. Based on the speed and size of the vehicles, this four meters was a rational distance choice below which calculating the approach angle was justifiable.

For each pair of pedestrian-vehicle, we then calculate their approach angle as the angle between the straight line connecting the endpoints of each of the two agent's trajectories while they are in interaction (Fig. 5.2). We will then calculate the entropy of the approach angle distribution within each dataset as a measure of the approach angle diversity. As another feature, we calculate the ratio of pedestrians gaining priority over vehicles to the total number of times a path crossing situation happened between a vehicle and pedestrians during their interaction. We defined priority in terms of who got to first cross the intersection point of the two crossing trajectories. We also include the ratio of vehicles to

pedestrians in each frame as a feature, with the average of this ratio being calculated over all the frames of the dataset.

5.2.3 Data analysis approach

Having the above-extracted features for each dataset, we apply k-means clustering to find two clusters that could represent structured vs. unstructured environments for the data in the combined datasets. By using the k-means clustering technique we have the flexibility of choosing the number of clusters to be two, for representing the two types of environments that we wanted to distinguish between. To have all features together, we used the average of the vehicle features for each dataset and added them to the interaction and pedestrians features.

We have eliminated the outliers by using the $1.5 \times IQR$ (Interquartile range) method. Therefore, the data points that are $1.5 \times IQR$ below the 25% percentile or above the 75% percentile are excluded.

As the next step, we evaluate each feature for its significant effect on the clustering result. For that we fit a generalized linear model to the data, using the pedestrian features as the fixed effect and the label of the environment from the clustering result as the output. This was done using the *glm* function in RStudio. The results of these analyses are presented in the next section.

5.3 Results and discussion

The mean and 95% confidence intervals for each feature are shown in Figs. 5.3, 5.4, and 5.5. Note that the different scenes of each dataset are plotted separately. According to Fig. 5.3, pedestrians in the SDD dataset have higher stop fraction and trajectory variability with a lower path efficiency compared to the other datasets. In Fig. 5.3(f), it is observed that the DUT dataset has the highest pedestrian density as it has been captured from a small congested area, but the density of standing pedestrians is higher in SDD according to Fig 5.3(g). The highest vehicle mean speed is from the inD dataset as it is being captured from a typical road environment (Fig 5.4). The inD dataset also has the highest ratio of vehicles to pedestrians among all the datasets (Fig. 5.3(c)) and also the highest percentage of vehicle priority (Fig. 5.3(b)).

The clustering results for each dataset based on these thirteen defined trajectory features is shown in Fig. 5.6. This figure shows the percentage of the dataset that goes into

each of the two clusters, A and B. For majority of the dataset, all the data points stand unanimously within one cluster.

According to this clustering result, cluster A which contains the four different scenes of the SDD dataset represents trajectory behaviours that happen in a campus environment and mainly off the road. All the other datasets fall into the second cluster, cluster B, which contains mostly road environments having either a conventional structure with designated operational areas for pedestrians and vehicles or having the design of a shared street.

By assigning all the points in a dataset to a label that represents the label of the majority of the points in that dataset from the k-means clustering, we have then compared the values of each feature between the two clusters in Figs. 5.7-5.9. In these figures, the mean value of each feature is shown with the 95% confidence interval. The statistical analysis using LMMs [12] for each feature separately while considering the dataset as a random effect shows that pedestrians' stop fraction ($se = 0.005, t = -91.06, p < 0.001$),

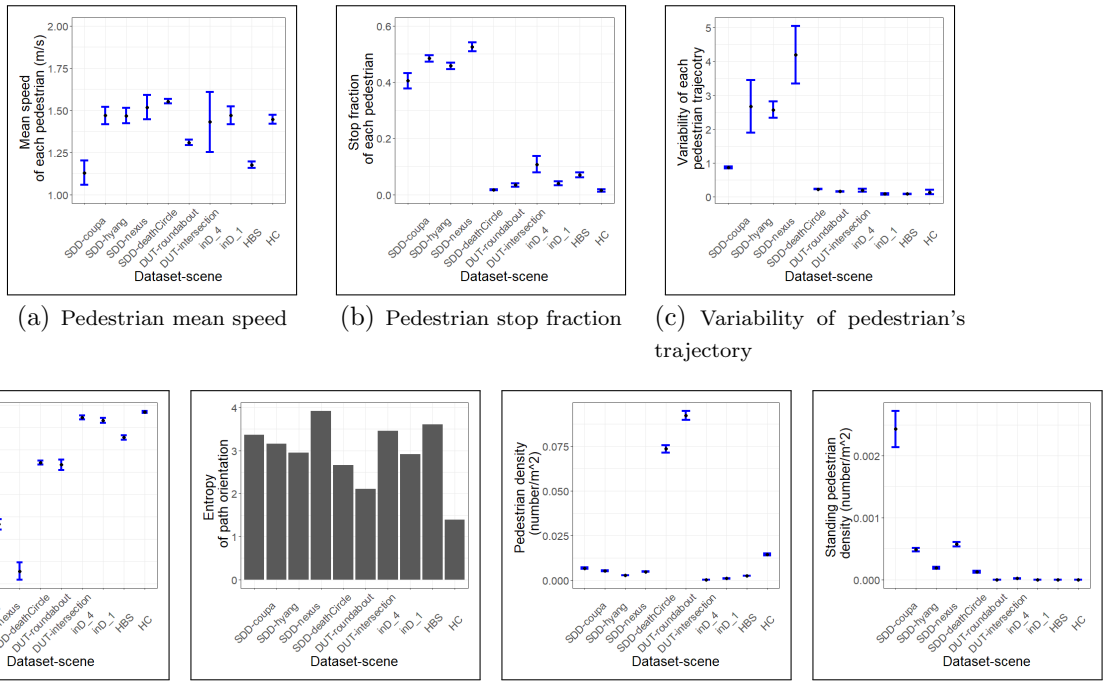

(a) Pedestrian mean speed (b) Pedestrian stop fraction (c) Variability of pedestrian's trajectory

(d) Pedestrian path efficiency (e) Entropy of path orientation (f) Pedestrian Density (g) Standing Pedestrian Density

Figure 5.3: The value of features extracted from pedestrians' trajectories in each dataset with 95 % confidence interval.

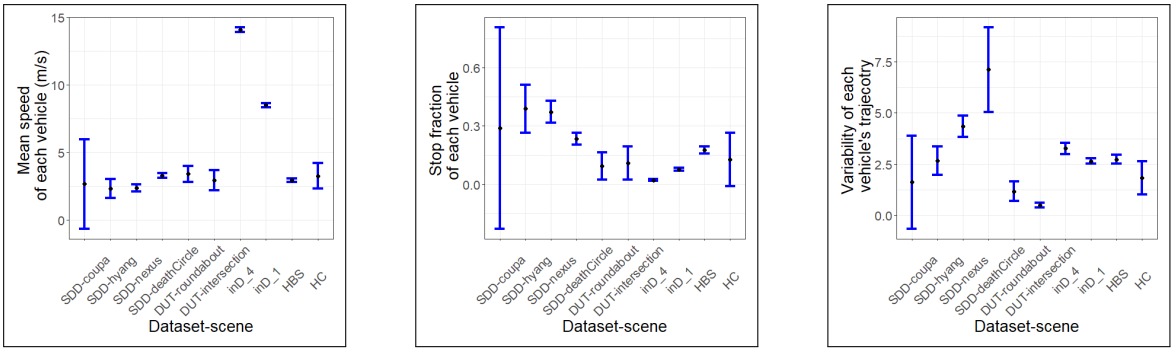

(a) Vehicle mean speed (b) Vehicle stop fraction (c) Variability of Vehicle 's trajectory

Figure 5.4: The value of features extracted from vehicles' trajectories in each dataset with 95 % confidence interval.

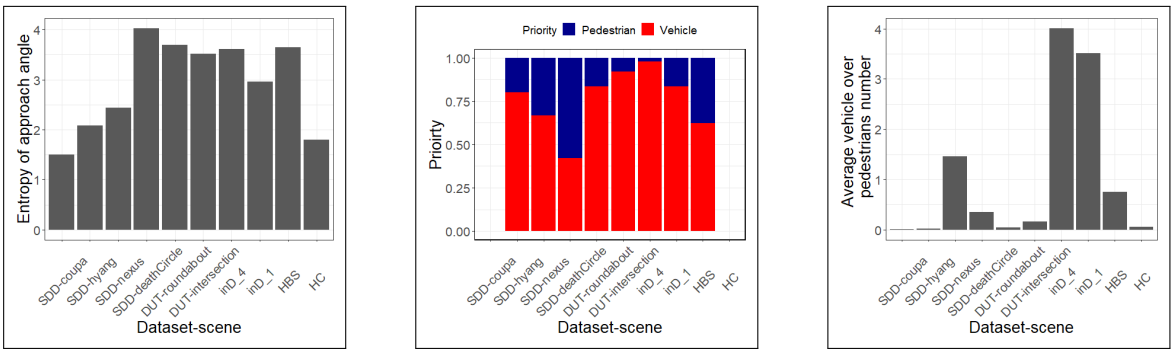

(a) Entropy of approach angle (b) Priority in path crossing (c) Ratio of vehicles to pedestrians

Figure 5.5: The value of features extracted from pedestrian-vehicle interaction in each data.

the variability of pedestrians' trajectory ($se = 0.069, t = -30.11, p < 0.001$), and the pedestrians' path efficiency ($se = 0.023, t = 4.731, p < 0.05$) are significantly different among the two clusters. The path efficiency is significantly lower in cluster A compared to cluster B. In addition, the density of standing pedestrians has significantly different values in the two clusters ($se = 1.82e - 5, t = -30.86, p < 0.001$). Among the vehicle-related features, the stop fraction of the vehicle ($se = 0.051, t = -4.157, p < 0.05$) is significantly different between the two clusters.

For considering the effect of pedestrians' features all together, we fitted a generalized linear model (GLM), using maximum likelihood estimation for predicting the label of

Table 5.1: The generalized linear model leading to best model fit according to the AIC criteria (AIC=72.83). "Ped" stands for Pedestrian.

Fixed effects	**Estimate**	**Std. Error**	**z value**	**Pr(> |z|)**
Ped mean speed	2.897	1.129	2.567	0.01 (*)
Ped variability	-11.288	1.778	-6.349	<0.001 (***)
Ped average density	-0.736	0.323	-2.279	0.023 (*)

clusters. All the features were centred to have a mean of zero for this analysis To avoid collinearity, pedestrian features that only had no or low correlation (< 0.3) with each other were selected as fixed effects, and to improve model fit, the combination resulting in the least Akaike information criterion (AIC) was considered, which is reported in Table 5.1. The result of the GLM shows that all three fixed effects of mean speed, variability, and average pedestrian density had a significant effect on the clustering result. Using other combinations of non-correlating features, other pedestrian features not considered in the model of Table 5.1 also had significant effects on predicting the clustering result. We also added the dataset as a random effect to the model (to take the effect of specific locations/datasets into account) and fit a GLMM. This led to pedestrian path efficiency being the only significant fixed effect.

This difference may be because we only have one dataset representing cluster A, so by including the dataset as a random effect, we are also accounting for the differences between the two environments. As we were not able to find more datasets representing an unstructured environment similar to SDD, future work is needed to further investigate the differences between the two environments once more datasets become available.

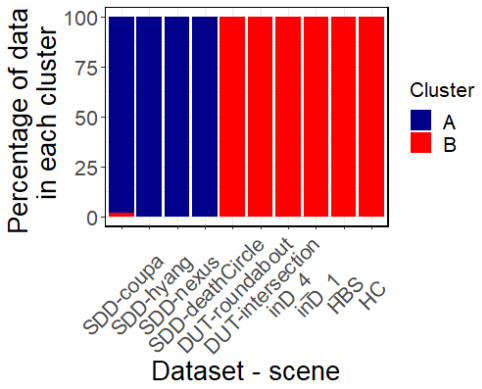

Figure 5.6: Datasets clustered in two groups based on the defined features. For each dataset, the portion of its data points in each cluster is reported.

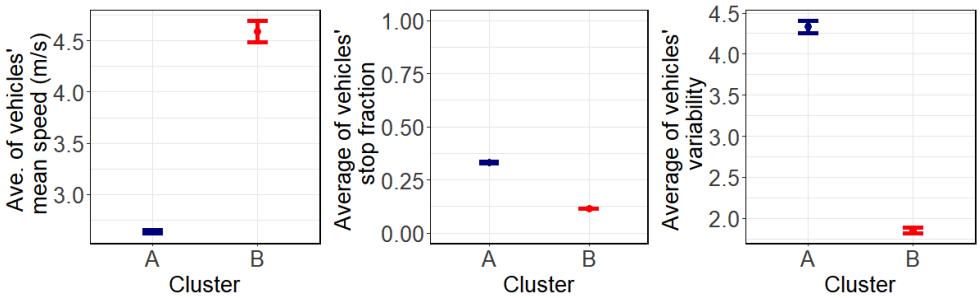

Figure 5.7: The mean of the vehicle features within each cluster along with the 95% confidence intervals.

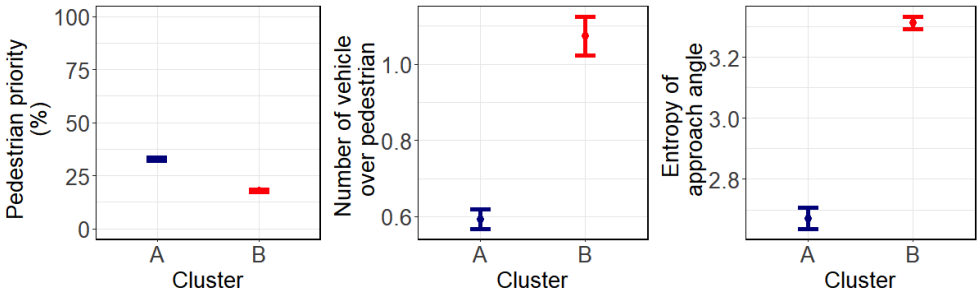

Figure 5.8: The mean of the interaction features within each cluster along with the 95% confidence intervals.

General Comparison of Environmental Features: The environments in cluster A, overall, are more dominated by pedestrians than vehicles, as can be seen from their lower vehicle over pedestrian ratios compared to cluster B (Fig. 5.8). The higher stop fraction in pedestrians' trajectory profile in cluster A, as seen in Fig. 5.9, can be a usual behaviour observed on-campus where people might linger in the environment or regularly come to a stop along the path when meeting a friend. These behaviours will also cause high variability in the speed profile. The lower path efficiency of pedestrians in cluster A is equivalent to more deviations from a straight path which can be the trajectory behaviour of a pedestrian that is roaming around without having an exact destination point in mind. This behaviour is unlikely to happen for pedestrians who are crossing the road or are walking on the sidewalk of an urban area towards their destination.

Therefore, based on our results, we can point out some similarities between the trajectory features observed in cluster A and the characteristics of an unstructured environment. Slightly higher entropy of path orientations in cluster A (Fig. 5.9) is an indication of higher diversity in the path directions, which happens when pedestrians are not restricted to follow designated paths under strict rules and are free to move anywhere. This was one

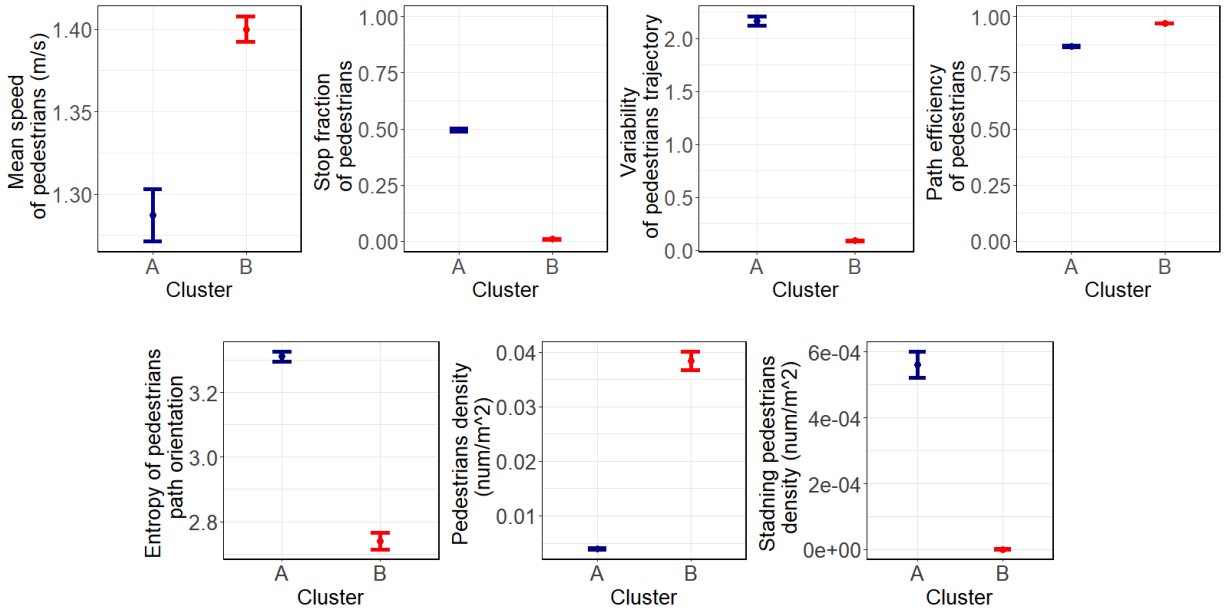

Figure 5.9: The mean of the pedestrian features within each cluster along with the 95% confidence intervals.

of the main elements of the qualitative description of an unstructured environment.

We can also describe an unstructured environment by the lower speed and higher variability of vehicles trajectory in those environments (Fig. 5.7), as there are no strict rules specifying the priorities or rights of way, and vehicles should pay more attention as pedestrians can be present everywhere. This could cause regular stops for the vehicles, resulting in slightly higher numbers for the stop fraction feature.

The effect of having no strict traffic rule in cluster A, as an example of an unstructured environment, can also be observed in the higher percentage of times that the vehicle yields and gives priority to pedestrians in the interactions happening in cluster A compared to cluster B (Fig. 5.8). The pedestrian's dominance in the scenes of cluster A and their environment not having a road structure could have affected these priority behaviours in interaction which could be a common behaviour seen in unstructured environments.

The number of standing pedestrians per unit area in cluster A is higher than in cluster B, showing that we should expect more standing pedestrians in environments similar to the ones in cluster A or in other words in unstructured environments. This is somehow also connected to having higher stop fraction within these environments.

Other features that could be added to the definition of an unstructured environment,

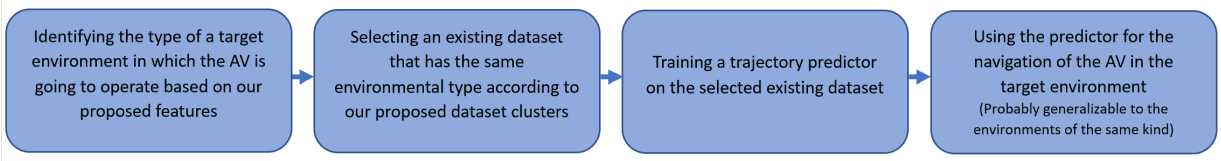

Figure 5.10: The process through which the results of this study can be used for developing trajectory prediction modules for autonomous vehicles (AVs).

suggested by our results, are moving pedestrians having slightly lower speed with a higher probability of coming to a stop along the path. This means that we should expect more variability in the pedestrians' trajectory in an unstructured environment. These behavioural features are important to consider when predicting a pedestrian's trajectory in different environments, as well as in the context of the vehicle's decision-making.

5.3.1 Limitations

Our work had limitations. The number of existing datasets that include both pedestrians and vehicles is limited especially when it takes place outside the road structure, which is closer to the definition of an unstructured environment. An unstructured environment similar to the SDD dataset can also be found in other areas, such as in airport terminals and shopping malls, for which we could not find any data which also includes vehicles. Therefore, cluster A in our results only contains one major dataset (SDD) with its different scenes. Although having more datasets that could have fallen under cluster A would have strengthened our argument about the clusters representing unstructured vs. structured environments, there is still a clear difference between the SDD dataset and all the other datasets, which are believed to be a result of the different nature of the environments.

The other limitation comes from the differences between the field of view and the geometry of the environments where the data were captured, with some datasets (e.g. DUT) being more focused on a small area as opposed to the others that have covered a broader region (e.g. in, SDD). This would affect the length of the trajectory for each agent present in the dataset, with shorter trajectories probably showing fewer variations in the trajectory profile due to limited information available from them. A larger dataset that contains more examples of unstructured environments can be investigated in the future, which, to the best of our knowledge, is not presently available.

5.4 Summary and conclusion

We studied the differences in the trajectory features of five existing datasets containing both pedestrians and vehicles. After dividing the datasets into two clusters based on 13 different features related to agents' motion trajectories, we analyzed the effect of these features on distinguishing between the unstructured and structured nature of the captured environments. Among the two resulting clusters, one primarily included road-like environments, which we named the structured environment. The other cluster contained scenes from a campus environment, happening mainly outside of a road structure, and was interpreted as an unstructured environment. Therefore, we found some differences in the trajectory features of pedestrians in structured versus unstructured environments.

Based on our findings, we discussed the differences in the motion behaviours of vehicles and pedestrians in structured versus unstructured environments. Features such as stop fraction, variability of pedestrian trajectories, and the density of standing pedestrians significantly differed between the two clusters, addressing **RQ2.1**. These features, can serve as quantitative measures for distinguishing between these environments, as demonstrated through our clustering process and the GLM model we fitted, addressing **RQ2.2**.

Our results suggest that an unstructured environment can be characterized as a place where pedestrians have more variability in their trajectories, are more likely to deviate from a straight path, and are more likely to stop regularly along their path compared to when they are moving in a structured environment. Therefore, one can conclude that pedestrian behaviours in unstructured environments are associated with greater uncertainties. Additionally, a higher number of standing pedestrians can be observed in unstructured environments, as these areas typically offer open spaces for pedestrians to linger and interact. The increased presence of standing pedestrians with unclear next intentions also contributes to the greater uncertainties in such environments.

Considering these trajectory features, one can use the related dataset for a similar type of environment to train a pedestrian trajectory predictor, learning the behavioural elements likely to be seen in that type of environment. Fig 5.10 shows how our work can be used in the process of developing a trajectory predictor for an autonomous vehicle.

The datasets identified in each cluster helped address **RQ2.3**. According to the clustering results of our five datasets, which we mapped to the environment nature based on the qualitative descriptions of un/structured environments, scenes from the SDD dataset belonged to the unstructured cluster, while other datasets, including those from shared spaces, were clustered in the structured environment group.

While shared environments can qualitatively fall under the definition of unstructured

environments, our different classification output for the datasets from this space might be due to the small portion of the shared environment being captured in these datasets, resulting in more straight-line paths of pedestrians crossing vehicles' paths. Furthermore, we made the simplifying assumption that environments could be divided into two categories: structured and unstructured. However, this binary classification may oversimplify the reality, as environments likely exist on a continuum of structure rather than as discrete categories. For example, introducing a third category for "less structured" environments might better capture the nature of shared spaces. Nonetheless, with the limited datasets available and based on our investigation, the two-group classification yielded the best clustering results, as reported in this chapter. Overall, due to the limited datasets we had, especially from unstructured environments containing both pedestrians and vehicles, further studies are needed for a more reliable categorization of the datasets into structured and unstructured environments.

Although, based on our clusters, SDD is the only dataset representing an unstructured environment, we do not limit ourselves to this dataset for training our pedestrian trajectory predictor. This decision is mainly because of the limited number of close interactions happening between vehicles and pedestrians in the SDD dataset. Therefore, we will still consider datasets captured from shared space environments as options for possible behaviour in unstructured environments and consider the existence of various close interactions between pedestrians and vehicles as another important factor in deciding the dataset to use for training our pedestrian trajectory predictor in the next chapter.

The main takeaway from this chapter's study is the importance of considering uncertainties when predicting pedestrian trajectories in unstructured environments, as revealed through the higher trajectory variability, increased stop fractions, and greater number of standing pedestrians. Consequently, accounting for these uncertainties will be an important aspect of our design in the subsequent phases of our algorithmic development.

Having explored specific pedestrian behavioural patterns in unstructured environments, we conclude our two-stage design requirements studies discussed in the previous chapters as part of our responses to **RQ1** and **RQ2**. We will now proceed to the algorithmic development phase, leveraging the information gathered thus far, starting with pedestrian trajectory prediction, which will be our next focus.

Chapter 6

Pedestrian trajectory prediction

Having established the design requirements and gained a deeper understanding of the target environment and pedestrian behaviour in previous chapters, we now transition to the technical development phase of our navigation algorithm. We begin by developing a pedestrian trajectory prediction algorithm for interactive scenarios in shared spaces with vehicles. This chapter focuses on the design of this predictor, while the next chapter will cover its integration with the motion planner within the overall navigation framework.

To investigate one of the main research questions of this thesis—effective ways of capturing interaction effects in the AV decision-making algorithm—we focus on how these interaction effects, comprising both pedestrian-pedestrian and pedestrian-vehicle interactions in our target environment, can be captured in the prediction module.

Accurate pedestrian trajectory prediction requires the integration of different sources of information. This includes the past trajectory history of each agent as well as the between-agent interaction effects (e.g., how one pedestrian's movement may affect another pedestrian's). Due to the complexity of pedestrians' motion in interactive settings, many deep-learning models have been proposed in the literature for predicting pedestrian trajectories while considering the interaction effects (e.g.,[3, 76, 204, 163, 112, 39]) as discussed in Chapter 3.5.

In deep learning approaches for modelling interaction effects, which is the method used in our prediction module in this chapter, three main research questions raises as subquestions under two of the main research questions in this thesis (**RQ3** and **RQ4**):

RQ3.1 How should interacting agents be specified to account for their impact on trajectory prediction methods?

RQ3.2 How to effectively formulate the effect of these interacting agents on the pedestrians' future trajectories, in learning-based prediction methods?

RQ4.1 How can the uncertainty in pedestrians' behaviour be captured in the prediction model?

Most of the existing literature in this area has focused primarily on proposing methods to address the second question, while offering simplified solutions to the first question. For example, some studies consider agents as interacting if they are within a certain threshold distance of each other [3, 17, 39, 112, 27, 115, 66], or include all agents in the scene by aggregating their trajectory information to model interaction [56, 215, 83]. Although some methods have proposed learning the importance of each neighbouring agent through weightings in attention layers [56, 111], having a measure to select interacting agents initially could better guide training and accelerate convergence. For example, in classical heuristic-based methods, pedestrians' motions when encountering vehicles are designed to be governed by time-dependent collision risk features such as TTC [146, 145] or minimum gap acceptance [42, 119]—defined as the time remaining between the pedestrian and the nearest approaching vehicle.

Getting inspired by these methods, here, we are proposing the use of TTC for specifying interacting agents (both pedestrians and vehicles) that can influence a target pedestrian's future trajectory.

Moreover, in response to the second question about formulating the possible effect of the interacting agents on a target pedestrian's trajectory, most of the existing models rely on relative distance features such as pooling the information of all interacting agents in a fixed positional grid map around the target pedestrian [3, 17, 204, 39, 112], or aggregating the information of the neighbouring agents using their relative distance to the target pedestrian [215, 76, 221, 111, 56]. In all these distance-dependent methods of capturing interaction, the effects of approach directions and speed are missing, while these two factors can indeed have a high impact on pedestrians' motion behaviour [84, 195]. In other words, two agents could be close, in terms of relative distance, but are not necessarily influencing each other's trajectory, if walking in parallel or getting away from each other.

Few other models have considered raw velocity information as an additional agent state when generating predictions [111, 115, 133]. However, velocity data can be used to derive more interpretable engineered interaction features that can better guide the learning process.

Therefore, in response to the second question, our approach leverages the information embedded in velocity vectors to calculate collision probabilities and define a polar collision

grid map. This map, serving as an interaction feature, captures the influence of agent interactions on a pedestrian's future trajectory. We believe that a neural network model can more effectively learn interaction-influenced motion patterns when provided with interpretable, engineered features as input.

6.1 Problem formulation

Pedestrian trajectory prediction is defined as predicting the future positions of pedestrians given a short observed history of their own and the close-by agents' (pedestrians and vehicles) trajectories [3]. The position of pedestrians and vehicles at time step t along their trajectories are denoted as $X_t^{ped} = [\mathbf{s}_t^1, \mathbf{s}_t^2, ..., \mathbf{s}_t^{n_p}]$ and $X_t^{veh} = [\mathbf{s}_t^1, \mathbf{s}_t^2, ..., \mathbf{s}_t^{n_v}]$, where "ped" stands for pedestrian and "veh" stands for vehicle, and n_p and n_v are the numbers of pedestrians and vehicles in the scene respectively. The x-y position coordinate of pedestrian or vehicle i at time t in this formulation is defined as $\mathbf{s}_t^i = (x_t^i, y_t^i)$. The velocity vector of each agent denoted as $\mathbf{v}_t^i = (v_{x_t}^i, v_{y_t}^i)$ can also be calculated using the position difference between successive time steps, knowing the timestamp between the captured frames.

Given the observed trajectories of all agents for the time steps between $0 < t \leq T_{obs}$, we predict the trajectory of the pedestrians across the future time period of $T_{obs}+1 < t \leq T_{pred}$ as $\hat{Y}_t = [\hat{\mathbf{s}}_t^1, \hat{\mathbf{s}}_t^2, ..., \hat{\mathbf{s}}_t^{n_p}]$ as stated in Eq. 6.1, where F is our neural network function. The ground truth trajectory of pedestrians for any time step in this prediction period is denoted as $Y_t = [\mathbf{s}_t^1, \mathbf{s}_t^2, ..., \mathbf{s}_t^{n_p}]$. The predicted position coordinates $\hat{\mathbf{s}}_t^i$ are derived from the predicted Gaussian distribution output of the model $\tilde{\mathbf{y}}_t^i$, which will be described later in Eq. 6.7.

$$\hat{Y}_{t+1:t+T_{pred}} = F(X_{t-T_{obs}:t}^{ped}, X_{t-T_{obs}:t}^{veh}) \tag{6.1}$$

6.2 Method

Here, we rely on a neural network as the function F in Eq. 6.7 that outputs the pedestrian trajectory for a prediction horizon, taking the observed trajectories of pedestrians and vehicles as input.

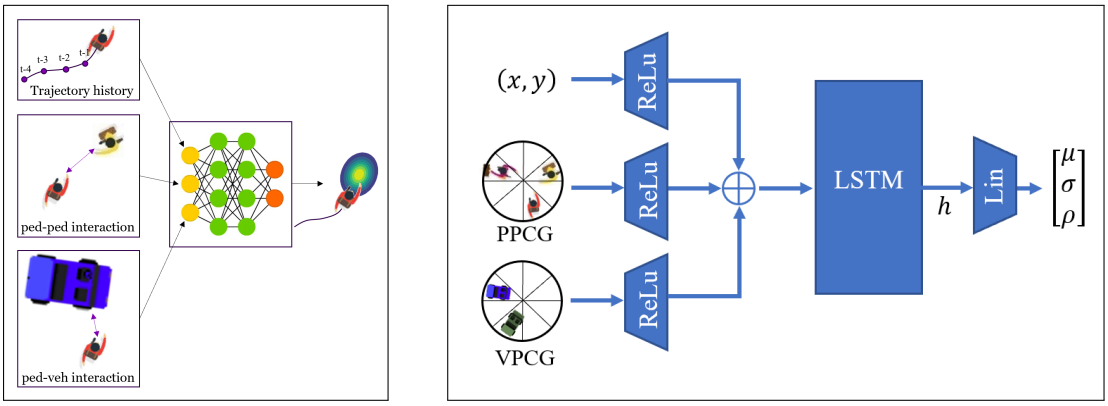

Figure 6.1: The overall framework of our proposed pedestrian trajectory prediction. The embedded spatial feature, the pedestrian-pedestrian interaction and the pedestrian-vehicle interaction features are concatenated before being inputted into the LSTM module.

6.2.1 Model architecture

Our overall framework for predicting pedestrians' trajectories in an environment shared with both pedestrians and vehicles is shown in Fig. 6.1 in more detail. The movement pattern in the pedestrian's trajectories is learned through an LSTM network. The network receives three types of input features for each pedestrian: (1) the spatial embedded feature of the pedestrian, (2) the embedded pedestrian-pedestrian interaction feature, and (3) the embedded pedestrian-vehicle interaction feature. These features are aggregated together and used as the input to the LSTM at each timestep.

Similar to other works [3, 56], instead of predicting the actual future positions, our model will predict the parameters of a bi-variant Gaussian distribution of possible future positions to account for prediction uncertainty, which is important in unstructured environments. Therefore, the output of our LSTM goes through a linear layer to predict the means (μ_x, μ_y), standard deviations (σ_x, σ_y) and the correlation coefficient (ρ) of the bi-variant Gaussian distribution. The details of each sub-component and the loss function used for training are described below.

Spatial embedded feature

The effect of each pedestrian's own trajectory history on the predicted next position is captured through this feature, where the spatial coordinate of the pedestrian is embedded

through a single layer with rectified linear unit (ReLU) non-linearity as the activation function. In the following equation, this layer is denoted as $\phi_e(.)$ with W_e being its embedding weights.

$$e_t^i = \phi_e(\Delta \mathbf{s}_t^i; W_e) \tag{6.2}$$

The spatial input will be provided as displacements $\Delta \mathbf{s}_t^i = (x_t^i - x_{t-1}^i, y_t^i - y_{t-1}^i)$ which is the relative distance between pedestrian i's current time step t and the previous time step $t-1$.

Interaction embedded features

How interactions with other agents (i.e., pedestrians and vehicles) can influence a pedestrian's future trajectory is captured through the interaction feature input. These interactions are captured through a Polar Collision Grid (PCG) map in the form of a tensor that will be discussed in more detail in Section 6.2.3. We have separate tensors for encoding pedestrian-pedestrian (PPCG) and pedestrian-vehicle interaction (VPCG), as a vehicle can influence a pedestrian's trajectory differently from how another neighbouring pedestrian does. Each of these two tensors will be processed by going through an embedding layer with ReLU nonlinearity (ϕ_p for ped-ped and ϕ_m for ped-veh[1] interaction) before being used as an input to the LSTM network.

$$p_t^i = \phi_p(PPCG_t^i; W_p) \tag{6.3}$$

$$m_t^i = \phi_m(VPCG_t^i; W_m) \tag{6.4}$$

Prediction module

The three inputs of spatial feature (e_t^i), ped-ped interaction feature (p_t^i) and ped-veh interaction feature (m_t^i) will be concatenated and used as input to the LSTM cell with W_l weights.

$$h_t^i = LSTM(h_{t-1}^i, concat(e_t^i, p_t^i, m_t^i); W_l) \tag{6.5}$$

[1]Ped stands for pedestrians. Veh stands for vehicle

The output of the LSTM will then go through a linear layer with W_o embedding weight to output the parameters of the bi-variant Gaussian distribution. The predicted position can then be generated by sampling from this distribution.

$$\tilde{\mathbf{y}}_t^i = [\mu_x^i, \mu_y^i, \sigma_x^i, \sigma_y^i, \rho^i]_t = W_o h_t^i \tag{6.6}$$

$$\hat{\mathbf{s}}_{t+1}^i \sim \mathcal{N}(\tilde{\mathbf{y}}_t^i) \tag{6.7}$$

Loss function

Given that the model is predicting the parameters of a bi-variant Gaussian distribution ($\hat{Y}_t^i$) over the next positions, a negative log-likelihood (NLL) loss function is used for learning the parameters of the model as follows:

$$Loss = -\sum_{i,t} log(P(\mathbf{s}_{t+1}^i | \tilde{\mathbf{y}}_{t+1}^i)) \tag{6.8}$$

6.2.2 Interacting agent selection

We rely on collision risks to indicate each target pedestrian's nearby interacting agents. Only these specified interacting agents will then be considered for constructing the interaction feature of each target pedestrian. This method is opposed to other approaches followed in the literature that consider either all the agents present in the scene or rely only on distance for specifying interacting agents and building the interaction feature. The intuition is that at each time step, a pedestrian's trajectory will be affected only by those agents that are in a collision course with it if the pedestrian keeps its current trajectory trend.

We use the feature of TTC as a proxy for measuring the collision risk with others. Therefore, at each time step, having the position and velocity vectors we calculate the time left for the agents to get closer to each other than a minimum comfortable distance (d_{min}), which we define as collision (see Fig. 6.2). This time between agent i and j with position vectors $\mathbf{s}$ and velocity vectors $\mathbf{v}$ is calculated according to Eq. 6.9. The notations used are for the information at time step t, but we have removed the subscript t for simplicity.

$$TTC^{ij} = \frac{-(\mathbf{d}_{rel} \cdot \mathbf{v}_{rel}) \pm \sqrt{(\mathbf{d}_{rel} \cdot \mathbf{v}_{rel})^2 - |\mathbf{v}_{rel}|^2(|\mathbf{d}_{rel}|^2 - d_{min}^2)}}{|\mathbf{v}_{rel}|^2} \tag{6.9}$$

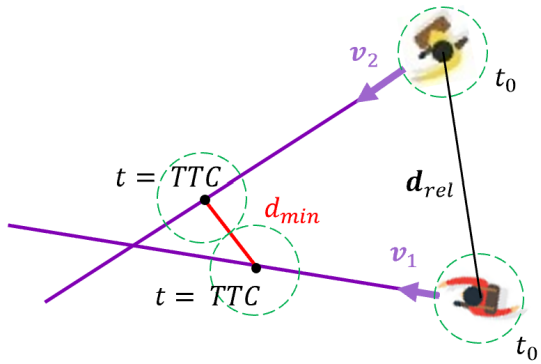

Figure 6.2: Time to Collision (TTC): the moment when the distance between two agents falls below a minimum acceptable threshold, assuming they continue following their current velocity profiles without any changes.

In the above equation, the vectors for relative distance and the relative velocity between agents i and j are denoted as $\mathbf{d}_{rel} = \mathbf{s}^i - \mathbf{s}^j$ and $\mathbf{v}_{rel} = \mathbf{v}^i - \mathbf{v}^j$ respectively. Moreover, (.) denotes the dot product of two vectors and $|\mathbf{d}|$ is the size of the vector $\mathbf{d}$.

Using Eq. 6.9, all the agents with a valid TTC of greater than zero and less than a specified threshold of TTC_{thre} will be considered for constructing the interaction features, as discussed in the next section. For any agent that might already be within d_{min} distance of the target pedestrian, a TTC of zero will be considered.

6.2.3 Polar Collision Grid interaction feature

At each time step, having the specified interacting agent and their TTC, we construct the interaction feature in the form of a polar grid centred at each target pedestrian's current position and in the direction of its current velocity vector. We follow a similar approach to how we relied on collision risk for specifying interacting agents for building the interaction effect feature.

With the justification that a pedestrian follows relatively similar evasive manoeuvres for agents in the same collision course in terms of approaching angle, we consider the angle between the velocity vectors of two potentially colliding agents as another feature for encoding interaction and learning the trend of the evasive behaviour. Therefore, we discretize the angular space around each pedestrian's current velocity vector to a specified number of sectors (n_{sector}), and for each conflicting agent, we calculate the discretized

sector where its approach angle lies on. The angles are calculated with respect to the target pedestrian's current velocity vector in a counter-clockwise direction. This process results in a polar grid map around each pedestrian's current velocity vector while specifying the presence of each interacting agent in each cell of this polar grid.

To convert this polar grid map into an interaction feature that also includes collision risk information, we store the $TTC_{thr} - TTC$ value of each interacting agent in its corresponding cell. If multiple agents occupy the same cell, we use only the information of the agent with the closest time to collision (i.e., the riskiest agent) in that cell, as this is the first interacting agent the target pedestrian is likely to react to. In other words, we aggregate the information of the multiple interacting agents by focusing on the riskiest agent in each approach angle cell which is the agent with the lowest TTC or the highest value of $TTC_{thr} - TTC$. The resulting grid will be called the Polar Collision Grid (PCG), which is a tensor of size $1 \times n_{sector}$ calculated according to Eq. 6.10 for each pedestrian i at time step t.

In this equation, $\mathbb{1}_n[\theta]$ is a binary indicator function to specify if θ (the angle between vector $\mathbf{v}_t^i$ and $\mathbf{v}_t^j$) is in the nth cell of the grid, and N^i is the set of all interacting agent of pedestrian i.

$$PCG_t^i(n) = \max_{j \in N^i}(\mathbb{1}_n[\angle(\mathbf{v}_t^i, \mathbf{v}_t^j)](TTC_{thr} - TTC^{ij})) \tag{6.10}$$

This tensor is constructed separately for pedestrians and vehicles that are interacting with a pedestrian where the former is referred to as Pedestrian Polar Collision Grid (PPCG) and the latter as Vehicle Polar Collision Grid (VPCG). Fig. 6.3 visually shows the process of creating these interaction features for pedestrian i in the middle for a sample scenario. In this figure pedestrians F and G, despite being close to pedestrian i, are not included in the interaction feature since they are not causing any collision risks.

These interaction features will then be embedded according to Eq. 6.3 and Eq. 6.4 and used as an input to the LSTM.

6.3 Experiments

We have evaluated the effectiveness of our proposed method by first comparing the overall methods with a couple of baseline methods in the literature. Then we focus on the effectiveness of each of our introduced modules by conducting ablation studies.

6.3.1 Dataset

In this thesis, we are focused on predicting pedestrian trajectories in an unstructured environment shared with vehicles. Our method examines the effects of both pedestrian-pedestrian and pedestrian-vehicle interactions on the trajectory prediction of pedestrians, necessitating a dataset that includes both pedestrians and vehicles for training and evaluation. Additionally, we need the trajectories of the dataset to closely resemble behaviours in unstructured environments. In other words, we aim to train our model in settings where interactions are less regulated by traffic rules and more governed by negotiation and social etiquette, as these interaction patterns are more complex to model compared to behaviours forced by traffic rules or emerged from strict environmental structures (e.g, crosswalks).

The SDD dataset was the only one identified in our unstructured environment cluster in Chapter 5. However, the SDD dataset contains very limited instances of close interactions between pedestrians and vehicles. Therefore, we expanded our dataset selection in this

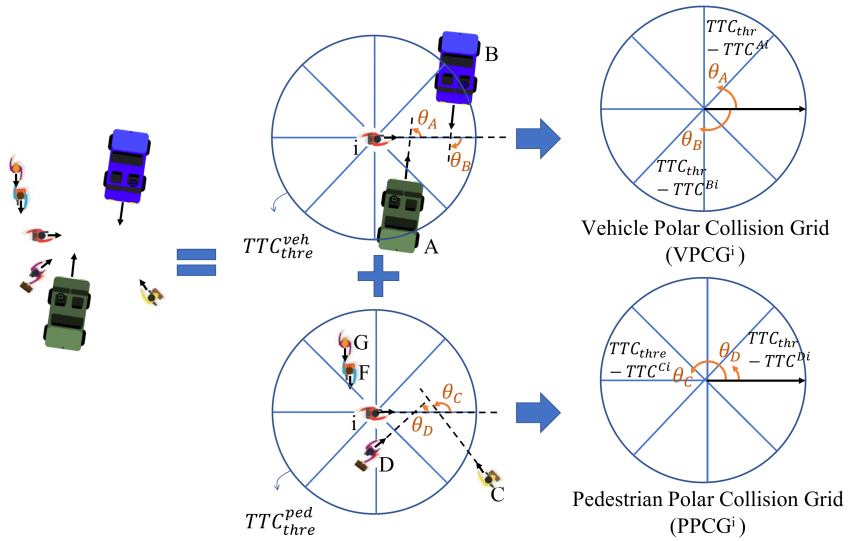

Figure 6.3: The process of building the interaction features for pedestrian i according to the neighbouring agent's time to collision (TTC) and approach direction. Eight discrete sectors around the target pedestrian are considered in this figure for the approach angles. Only neighbouring agents with a TTC lower than a threshold will be considered in the construction of the interaction feature. The interaction features for neighbouring vehicles and pedestrians are constructed separately.

chapter to also include shared spaces, where traffic rules and road markings are alleviated, and different types of agents are encouraged to share the same space [147]. This aligns closely with the qualitative definition of an unstructured environment. Therefore, here we use the HBS dataset [140], which, among the shared space datasets analyzed in Chapter 5, has the highest number of close interactions between pedestrians and vehicles. Its motion features are also closest to those of the SDD and furthest from the inD (an example of a dataset captured in a strictly structured road environment). This is evidenced by features such as pedestrian and vehicle mean speed, vehicle stop fraction, and priority of path crossing, as shown in Figures 5.3-5.5.

The real-world trajectory data in the HBS dataset has been captured from a shared space close to a busy train station in Germany from a bird's eye view (see Fig. 6.4), and includes highly interactive scenarios between pedestrians and vehicles. This dataset has already been used in other trajectory prediction methods in the literature [41, 39, 96]. There are 338 vehicles, 1115 pedestrians, and 22 cyclists in the 30-minute video collected from this area, and the trajectory data are available at 2 Hz. The dominant number of pedestrians in this dataset also suits our method, which is focused on predicting only pedestrians' trajectories. Same as [41, 140, 40], we use the first 10 minutes (31% of the dataset) as the test set and the remainder for training. Here we observe 6 time-steps covering 3 seconds of the trajectory and predict the next 6 time steps (3 seconds) of the trajectory, similar to [41, 40].

Figure 6.4: The Hamburg Bergedorf Station site, where the HBS dataset was collected. The picture is sourced from [140] with proof of permission provided in Appendix B.

6.3.2 Evaluation metrics and baselines

In line with the literature, we use the following metrics to evaluate the method and report the prediction errors.

- Average displacement error (ADE): The Euclidean distance between predicted and ground truth positions at each time step averaged over all predicted time steps.
- Final displacement error (FDE): The Euclidean distance between predicted and ground truth positions at the final predicted time step.
- Modified Hausdorff Distance (MHD): The largest distance between any two points on predicted and ground truth trajectories (without considering time step alignment).
- Speed error (SE): The root mean square error (RMSE) between the speed of the predicted trajectory and the ground truth speed at each time step.
- Heading error (HE): The RMSE between the heading of the predicted trajectory and the ground truth heading at each time step.

We compare the performance of our method against the following baseline methods:

- Linear Regression (LR): A linear regression model for each position dimension individually.
- Vanilla LSTM: A naive LSTM considering only trajectory history without the interaction effects from pedestrians and vehicles.
- Social LSTM: The LSTM model with the social pooling layer proposed by Alahi et al. [3] considering only pedestrian-pedestrian interaction.

The first two baseline models, unlike the Social LSTM model and our proposed Polar Collision Grid model, do not account for interaction effects from other agents in the target pedestrian's predicted trajectory. Therefore, comparing their performance with ours should highlight the impact of including interaction effects. Of these two baseline models, the Linear Regression model uses a simpler approach, while the Vanilla LSTM relies on a neural network to capture complex patterns in pedestrian trajectory behaviour, demonstrating the advantages of using a data-driven model.

On the other hand, both our Polar Collision Grid model and the Social LSTM [3] consider interaction effects, using grid-based methods to capture them. However, while the Social LSTM uses a square grid based on relative position data as its interaction encoding input, our model incorporates additional information—such as speed, approach angle, and collision risk—within a polar grid structure. Comparing the performance of our model with that of the Social LSTM should therefore reveal the effectiveness of our proposed interaction features.

6.3.3 Implementation details

The embedding layer for the spatial inputs and the two interaction features all have a size of 64. We used a hidden state dimension of 128 for the LSTM module. The number of sectors (n_{sector}) for the polar collision grid was set to 8. A TTC_{thre} of 9 seconds with a d_{min} of 0.7 meters was used for the pedestrian-pedestrian interaction features. These values were 8 seconds and 1 meter for the pedestrian-vehicle interaction feature, respectively. These values were selected after further tuning within a reasonable range according to the physical dimensions of the agents and the common personal space of humans [72] as well as the agents' minimum distance and speed ranges in the HBS dataset. Other hyperparameters of the model were also tuned by testing different values and selecting the one that resulted in the lowest error. A training batch size of 10 was used for 200 epochs. We used the Root Mean Squared Propagation (RMS-prop) optimizer with a learning rate of 0.001. The Nvidia GeForce RTX 2080 Ti GPU was used for the training and testing processes.

6.4 Results and discussion

6.4.1 Quantitative evaluation

The performances of the trained models on the HBS test set are summarized in Table 6.1. Since the output of the deep-learning models is a distribution over the predicted position, for evaluation, same as [215], we generated 20 samples for each model and reported the best result. Lower numbers are better. The best values among the different models for each metric are shown in boldface. The results show that our proposed method outperforms the baseline methods on all the evaluation metrics except for heading angle error. These results demonstrate the advantage of formulating interaction according to collision probability and approaching directions instead of only relying on relative distances. According to Table

Table 6.1: Prediction errors for baseline methods compared to the proposed method and its variants on the HBS dataset.

Models \Metrics	**ADE (m)**	**FDE (m)**	**MHD (m)**	**SE (m/s)**	**HE (°)**
Linear regression	0.696	1.238	2.995	0.390	44.0
Vanilla LSTM	0.305	0.676	2.855	0.240	32.7
Social LSTM	0.309	0.677	2.852	0.244	**31.5**
P-PolarCollisionGrid	0.304	0.664	2.811	0.235	32.6
V-PolarCollisionGrid	0.305	0.669	2.827	**0.232**	32.9
PolarCollisionGrid (PCG) (Ours)	**0.295**	**0.648**	**2.791**	0.235	31.7

6.1, Social LSTM has slightly higher prediction errors than the Vanilla LSTM, which does not consider interactions. This is consistent with the result reported by others for these two models using other datasets [76, 215].

Ablation study: To more clearly show the benefit of our proposed interaction modules for both pedestrian-pedestrian and pedestrian-vehicle interaction, we conducted an ablation study where we took out the interaction modules in a step-wise process and compared prediction errors with and without those modules in place.

Following this, we have trained our model by taking out the pedestrian-vehicle interaction feature and only keeping the collision grid associated with the pedestrian-pedestrian interaction. We called this the *P-PolarCollisionGrid*. We repeated the same process this time by taking out the pedestrian-pedestrian interaction feature and only keeping the pedestrian-vehicle polar collision grid and called the model *V-PolarCollisionGrid*. The prediction errors of these two models along with our complete model *PolarCollisionGrid PCG* are also reported in Tabel 6.1.

The results show that adding each of the interaction modules positively improves the prediction accuracy of the model in terms of ADE, FDE, MHD and heading error (HE). This indicates the effectiveness of our interaction module in capturing the influence of potentially conflicting agents on the predicted trajectory of a pedestrian. However, the lowest speed error is for the model that only considers pedestrian-vehicle interaction. This could be due to the high influence that vehicles have on pedestrians' speed adjustment along their trajectory especially when a pedestrian stops for giving way to a vehicle.

Interacting agent selection: For more specifically analyzing the effect of our proposed method for the selection of interacting agents based on TTC, we compared two versions of the Social LSTM baseline model. The first version is the original one, where the interaction feature is constructed using agents closer than a specified distance. In the second version, the neighbouring agents are first filtered based on our proposed TTC crite-

Table 6.2: The effect of our proposed interacting agent selection based on collision checks (used as a filter) on the performance of the Social LSTM model reported on HBS dataset.

Models \Metrics	**ADE (m)**	**FDE (m)**	**MHD (m)**	**SE (m/s)**	**HE (°)**
Social LSTM	0.309	0.677	2.852	0.244	**31.5**
Social LSTM +Filtered interacting agents	**0.298**	**0.658**	**2.827**	**0.234**	32.3

ria, and then the original social pooling module [3] is constructed considering those agents with collision risks (see Fig. 6.5).

The performance evaluation of these two versions in Table 6.2 shows that filtering the agent to consider only the ones in a conflict course with the target agent will actually

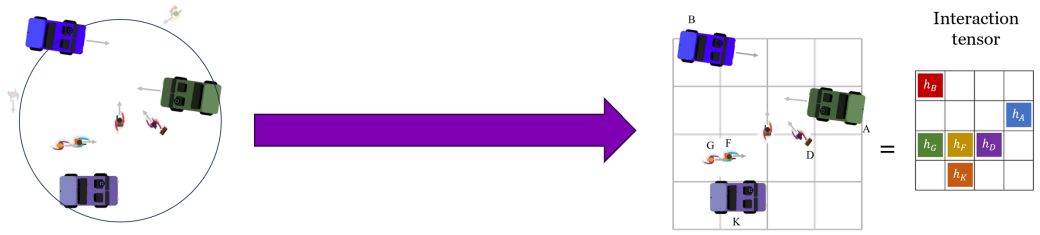

(a) Original SocialLSTM

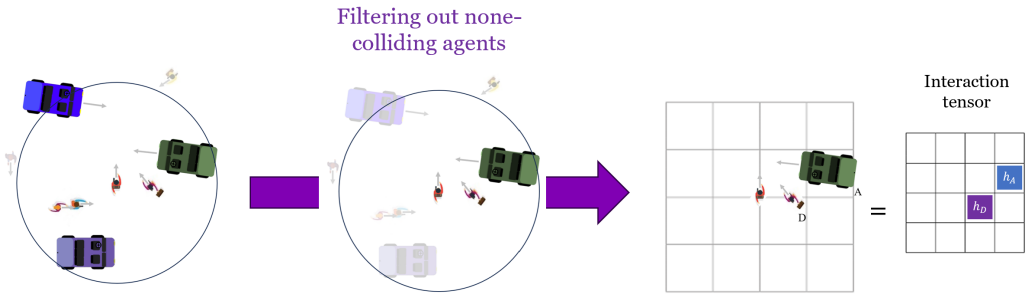

(b) Social LSTM with filtered interacting agent

Figure 6.5: The original Social LSTM (a) selects interacting agents based on relative distance falling below a threshold and stores the hidden states of those interacting agents in a spatial grid, constructing the interaction tensor. In the modified version (b), we alter the original method's interacting agent selection by filtering out agents based on our Time-to-Collision (TTC) checking method while leaving the rest of the interaction tensor construction unchanged.

improve the prediction accuracy. According to this result, removing the information of those agents that do not produce a collision risk on the target agent seems to be not only harmless but also beneficial, as the target agent's trajectory will not be affected by those agents. Therefore, our proposed method of selecting interacting agents based on collision risk seems to outperform the distance-based methods by concentrating on those agents that actually make the target agent adjust its path to avoid possible collisions.

6.4.2 Qualitative evaluation

To qualitatively analyze the predicted trajectory of our model compared to other baselines, we plotted the prediction outputs for a couple of interactive scenarios in the test set as shown in Fig. 6.6-6.8.

In the scenario shown in Fig. 6.6, our method detects vehicle B as an interacting agent with pedestrian A in the last observed timestep (t_{obs}), meaning that a collision will

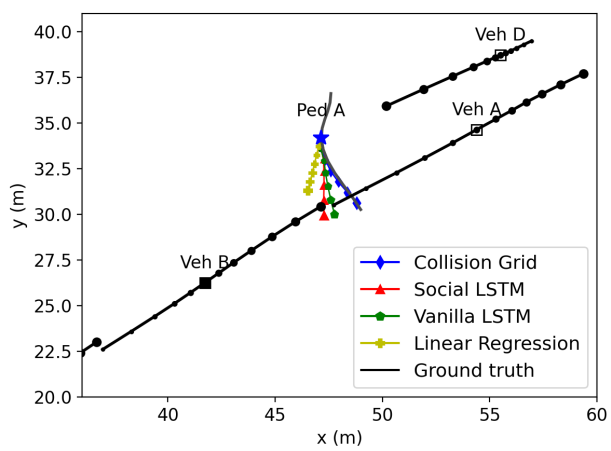

Figure 6.6: The predicted trajectories for our method compared to the baseline methods for a pedestrian-vehicle interaction scenario in the HBS test dataset. A marker with a larger size represents a later time step along the trajectory. The position of the vehicles at the last time step of the observation period is indicated with a square marker. Our method detects the target pedestrian A (indicated with start) to have interaction with vehicle B at this time step according to the TTC criteria. This identified interacting vehicle is indicated with a filled square marker at this time step. Considering this interacting vehicle, our method has predicted a trajectory that avoids the predicted collision while being close to the ground truth trajectory.

take place if the two agents continue their velocity as it is at time t_{obs}. Having identified this interaction, our method then considers vehicle B's information on TTC and approach angle when predicting pedestrian A's future trajectory. Therefore, as shown in Fig. 6.6, our method predicts an avoidance manoeuvre which is close to the ground truth. This is

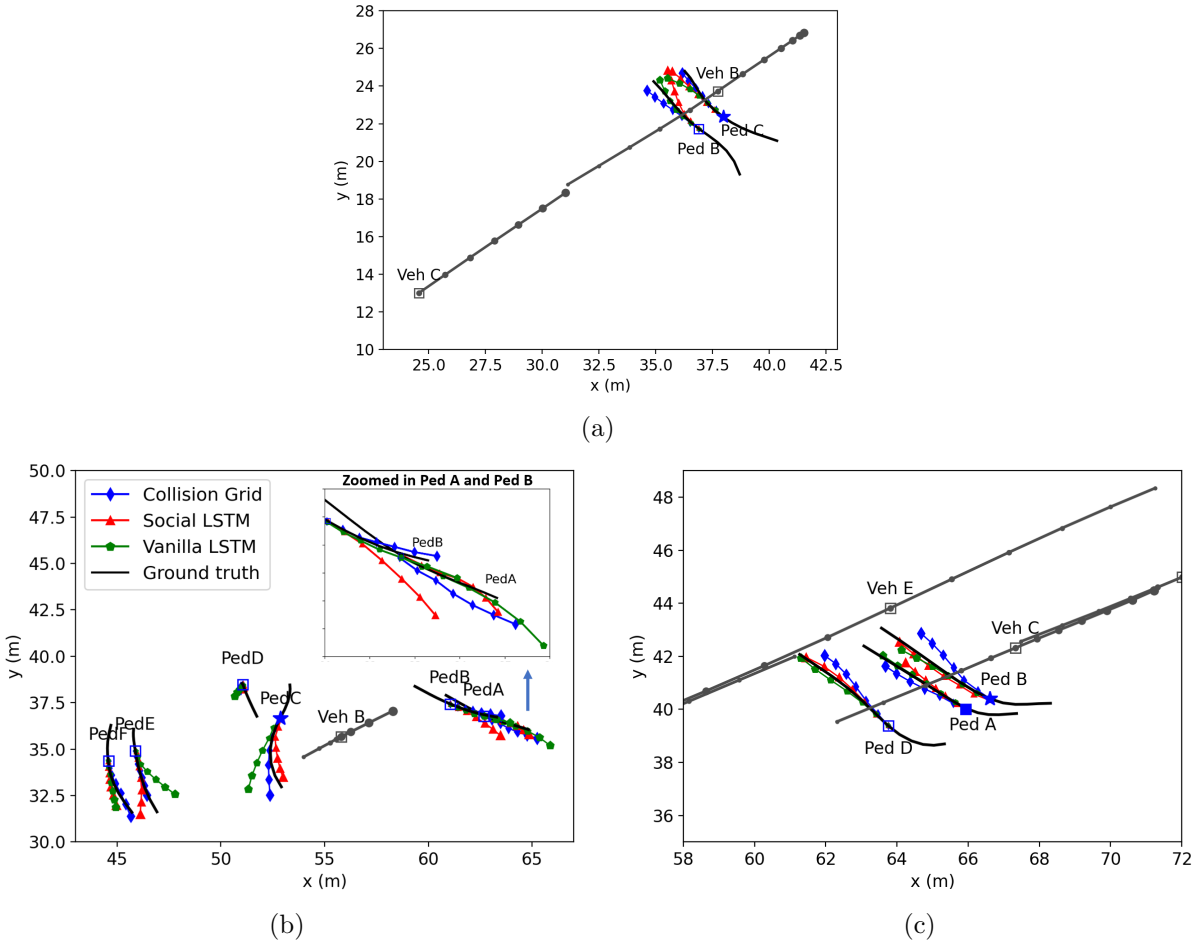

Figure 6.7: The predicted trajectories for our method compared to the deep-learning-based baseline methods in the HBS dataset. Our predictions overall are closer to the ground truth compared to the baselines. Case (c) is an example with a non-accurate prediction from our *Polar Collision Grid* method where a probable group interaction between Ped A and Ped B is mistakenly detected as a possible collision ending up with a prediction for Ped B that deviates to the right of the ground truth data. Interacting agents with a sample target pedestrian (shown with a star) are shown with a filled square marker.

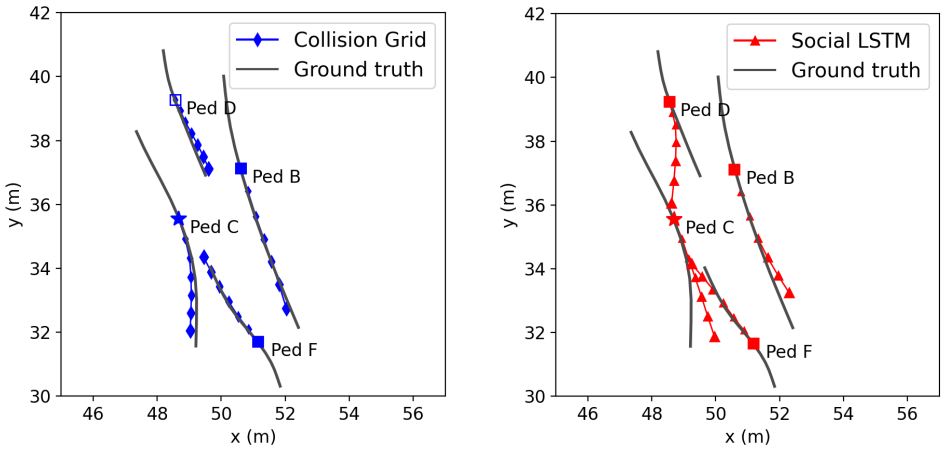

(a) (b)

Figure 6.8: Prediction of our Collision Grid method compared to Social LSTM for a pedestrian-pedestrian interaction scenario. Our method produces a prediction closer to the ground truth for pedestrians such as Ped C. This is done by detecting Ped C's interacting agents (Ped F and Ped B) more accurately while relying on time to collision for interaction modelling as opposed to the Social LSTM model that considers all agents closer to the target pedestrian than a threshold distance regardless of their moving direction which happens to be all the agents present in the scene here. Interacting agents for Ped C (target) are indicated with a filled square marker at time step t_{obs} while markers for non-interacting agents are unfilled.

while the Social LSTM method, which does not consider the pedestrian-vehicle interaction, predicts a trajectory that ends up in a collision with the vehicle at the last time step of the prediction window. This supports the effectiveness of our pedestrian-vehicle interaction module.

Fig. 6.8 shows another example focused on pedestrian-pedestrian interactions and compares the prediction of our Collision Grid method with the predictions of the Social LSTM method that also models the interaction between pedestrians. Focusing on Pedestrian C as a target agent, which is indicated with a star marker at time step t_{obs}, we show all its neighbours at the same time step with a square marker. In each of the methods, the neighbouring agents that will be considered for constructing the interaction feature are indicated with a filled square marker. According to Fig. 6.8 while the Social LSTM considers all the agents present in this scene in building the interaction feature, our method only detects pedestrians F and B as interacting agents for pedestrian C and therefore relies only

on these pedestrians' information for constructing the interaction feature. This results in a more accurate prediction for Pedestrian C's future trajectory with less deviation from the ground truth compared to the Social LSTM method. This shows the effectiveness of our proposed interacting agent selection process that focuses specifically on neighbouring agents that can actually affect the target agent's trajectory due to being in a potential collision course, as opposed to other methods such as Social LSTM that consider all agents that are closer than a distance to the target agent regardless of their moving direction. This latter approach can introduce extra not useful information for interaction consideration that can lead to an undesired prediction as seen in Fig. 6.8 for Social LSTM.

More qualitative examples of our method's prediction compared to other baseline methods can be found in Fig. 6.7 for other scenarios. There are cases where our method has more deviation from the ground truth compared to other methods. Fig. 6.7(c) provides an example in which two pedestrians (Ped A and Ped B) are walking together as a group. For this case our method has a false positive collision detection and predicts a trajectory for Ped B that is slightly to the right of the ground truth, getting away from Ped A, instead of predicting close parallel paths for these two pedestrians. However, despite the poor prediction in these particular cases, our algorithm produces less error on average compared to other baselines' predictions, as seen in Table 6.1. In the future, the prediction of our algorithm can be further improved by adding a group detection to filter out possible false collision detection in the model.

6.5 Uncertainty-aware Polar Collision Grid predictor

As discussed in Chapter 5, pedestrian trajectory behaviours are associated with higher uncertainties in unstructured environments, which is important to consider in both prediction and planning modules of our navigation framework. In the predictor module presented so far in this chapter, we have taken the initial steps to account for these uncertainties in our forecasts by predicting the parameters of the Gaussian distribution. However, since we will consider the prediction uncertainty by using the covariance of the predicted trajectory in the DRL motion planner in the next chapter, accurately predicting the covariance, along with the actual x and y positions of the data sampled from it, is essential to prevent prediction errors from negatively impacting our decision-making later on.

Models trained solely with the NLL loss are prone to providing overconfidence predictions, where the output distribution aligns precisely with the ground truth predicted position, with a sharp distribution similar to a Dirac delta distribution [87]. This im-

plies a low standard deviation for the predicted distribution or, in other words, very low uncertainty.

To overcome this problem, we train our previous Polar Collision Grid (PCG) predictor model with a revised loss function. In this modified version, in addition to the common NLL loss, we also add a novel uncertainty loss aimed at improving covariance prediction without compromising accuracy in mean prediction (Eq. 6.11).

Employing the combined loss in Eq. 6.11 directs the training process not only to maximize the log probability of the ground truth position but also to consider the distance between the ground truth point and the predicted distribution. This encourages predictions to have a higher standard deviation for the same predicted mean, aiming to encompass the ground truth.

$$\begin{aligned} \text{Loss} &= \text{Loss}_{NLL} + \text{Loss}_{Uncertainty} \\ &= -\sum_{i,t} \log(P(s^i_{t+1}|\tilde{\mathbf{y}}^i_{t+1})) + W \sum_{i,t} d^{MD}(s^i_{t+1}, \tilde{\mathbf{y}}^i_{t+1}) \end{aligned} \tag{6.11}$$

In this equation, d^{MD} represents the Mahalanobis distance (MD) (see Eq. 6.12) calculated between the ground truth (GT) position of the pedestrians and the predicted Normal distribution. The uncertainty loss function is introduced as an addition to the original NLL loss function, with a weight denoted as W.

To demonstrate the advantages of incorporating this uncertainty-aware loss function for more precise predictions of the distribution of pedestrians' future positions, we compare the original PCG model, trained exclusively with the NLL loss, with the same model trained using the updated loss function detailed in Eq. 6.11. We refer to our improved predictor model as Uncertainty-Aware Polar Collision Grid (UAW-PCG).

$$d^{MD}(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x} - \mathbf{y})^\top \mathbf{\Sigma}^{-1} (\mathbf{x} - \mathbf{y})} \tag{6.12}$$

6.5.1 Evaluation metrics

Given this modified pedestrian trajectory predictor's emphasis on generating less overconfident predictions by improving the estimation of the covariance matrix for future positions, for this part, we assess our predictor module using the measures recommended in [87] as follows:

Table 6.3: Pedestrian trajectory prediction results. Bolded is the best.

Model	**ADE (m)**	**FDE (m)**	**NLL**	ΔESV_1	ΔESV_2	ΔESV_3
PCG	$0.343_{\pm 0.15}$	$0.753_{\pm 0.33}$	$6.323_{\pm 14.88}$	**-0.231**	-0.464	-0.415
UAW-PCG	$\mathbf{0.333}_{\pm 0.14}$	$\mathbf{0.732}_{\pm 0.32}$	$\mathbf{0.537}_{\pm 3.06}$	0.440	**0.089**	**-0.012**

- Average Displacement Error (ADE): Average Euclidean distance between the predicted and the ground truth position over the prediction horizon.
- Final Displacement Error (FDE): Average Euclidean distance between the predicted and the ground truth position at the final step of the prediction horizon.
- Mean Negative Log Likelihood (NLL): Mean NLL of the ground truth future positions under the predicted distribution outputted by the predictor model.
- Delta Empirical Sigma Values (ΔESV_i): The difference between the empirical proportion of the ground truth positions residing in the i-th standard deviation (e.g., 1σ, 2σ, 3σ) of the anticipated distribution and the same proportion derived from an ideal Gaussian ($\Delta ESV_i : \sigma_{pred,i} - \sigma_{ideal,i}$, where $\sigma_{ideal,1} \approx 0.39$, $\sigma_{ideal,2} \approx 0.86$, $\sigma_{ideal,3} \approx 0.99$). Negative and positive values of ΔESV_i indicate overconfident and under-confidence in the predicted distribution accordingly [87].

6.5.2 Results and discussion

Quantitative results

Table 6.3 presents the results of the Polar Collision Grid model, trained with our proposed combination loss, in comparison to the same model trained solely with the NLL loss.

Upon comparing the two trained models, it is evident that the model trained exclusively with NLL exhibits overconfident predictions, as indicated by its negative ΔESV_i values (e.g., ΔESV_3 = -0.415). In contrast, the model trained with the combination loss yields a ΔESV_3 value closer to zero (-0.012), which is considered ideal. Additionally, larger positive values are observed for ΔESV_1, indicating underconfidence for the uncertainty-aware predictor. However, underconfidence in predictions is preferable compared to overconfident predictions when using the predictor results for planning. This approach provides a wider safety margin and promotes a more conservative behaviour.

Notably, the model exclusively trained with NLL loss demonstrates a higher standard deviation in the NLL loss measure on the test set when compared to the model trained

with the combination loss. This difference also highlights the overconfident predictions of the original model. Despite occasionally achieving very low NLL values towards the lower end of the NLL distribution over the test cases (6.323 ± 14.88) the model struggles to generalize effectively to unseen data in the test set.

Therefore, the incorporation of the uncertainty loss has proven beneficial, enabling the model to better estimate the uncertainty (covariance matrix) of its predicted distribution for the future positions of pedestrians. This enhancement is accompanied by a slight improvement in the ADE and FDE measures, with ADE decreasing by 1 cm and FDE decreasing by 2.1 cm. In Table 6.3, ADE and FDE values are reported for the predicted mean to facilitate a fair comparison. Therefore, these values may differ slightly from those reported for PCG model in Table 6.1, which reported metrics on the best 20 samples from the predicted distribution.

Qualitative results

To conduct a qualitative analysis of the predicted trajectories and visually assess the overconfidence prediction of the original model, we have generated visualizations of the predictions for a sample case from the test set, as depicted in Fig. 6.9. This figure illustrates that, despite potential deviations between the predicted mean and the ground truth in certain cases, the Uncertainty-Aware PCG (UAW-PCG) assigns a probability to the occurrence of the ground truth. In contrast, the original model tends to overlook this probability, especially in scenarios involving stationary pedestrians who may unexpectedly initiate movement in subsequent time steps, like the case of Ped D in Fig. 6.9. For such cases the UAW-PCG excels in anticipating higher uncertainties in the next time step, capturing the likelihood of sudden movements. Consequently, for enhancing pedestrian safety in DRL training for collision avoidance, we rely on predictions from the UAW-PCG.

Another notable observation is that the UAW-PCG model grasps the pattern of uncertainty propagation across the prediction horizon. It consistently predicts distributions with higher standard deviations for time steps farther into the future. This aligns with the correct understanding that prediction uncertainty naturally increases for longer horizons. This crucial pattern is not distinctly evident in the overconfident predictions of the model trained solely with NLL loss function.

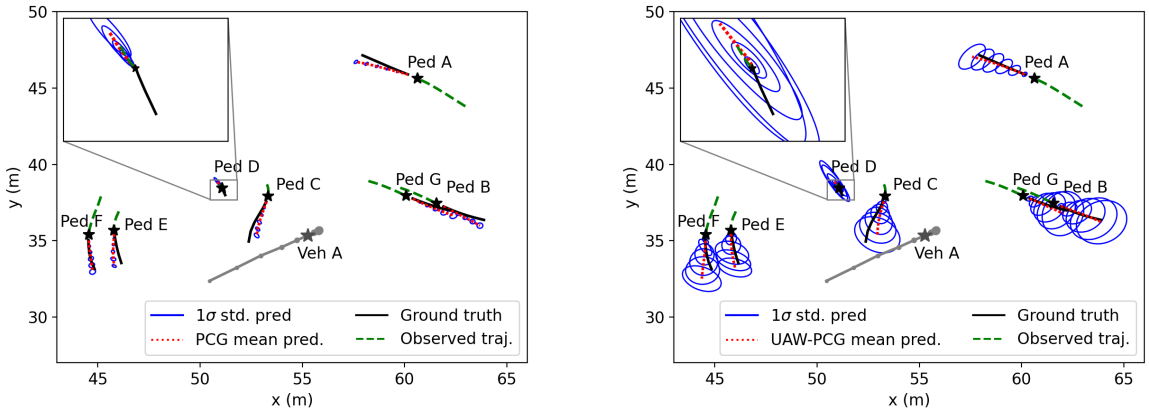

(a) PCG

(b) UAW-PCG

Figure 6.9: Predicted trajectories of (a) the polar collision grid (PCG) model trained with only the negative log likelihood loss, compared to (b) the same model trained with an additional uncertainty loss, referred to as the Uncertainty-Aware PCG (UAW-PCG). The blue circles represent the 1σ standard deviation of the predicted distribution for the future position around the estimated mean (depicted in red). Agents' positions at the current time steps are marked with a star, and larger markers along the vehicle's trajectory highlight positions at later times. Here "Ped" and "Veh" stand for pedestrian and vehicle respectively.

6.5.3 Limitations

Our pedestrian trajectory prediction model has some limitations. The goal was to capture the interaction between pedestrians and vehicles to predict future pedestrian trajectories, particularly in unstructured and shared environments where the lack of segregation between pedestrian and vehicle spaces increases interaction complexity. As a result, we focused on datasets representing unstructured environments with highly interactive scenarios between pedestrians and vehicles. However, only a few such datasets are available, and some are limited in the number of data points and interactive trajectories. Consequently, we trained and tested our algorithm on one dataset (HBS) from a shared space, which provided highly interactive scenarios between pedestrians and vehicles. Our model demonstrated superior performance compared to other baselines in these scenarios. Nevertheless, as more datasets from unstructured environments containing both pedestrians and vehicles become available, our model could be further trained and tested to improve its generalizability.

Additionally, in order to concentrate on patterns derived solely from trajectory data, we did not account for the influence of environmental structures or static obstacles. We believe, however, that these factors could be incorporated into our current model by introducing the environment map as an additional input.

Lastly, to capture the uncertainty in pedestrian behaviour, we predicted the parameters of a bivariate Gaussian distribution as the output. However, our approach focused on unimodal predictions, despite the potential benefits of adopting multimodal predictions, as demonstrated in some existing literature [211, 49, 56, 163], to capture multiple plausible pedestrian intents. The decision to use unimodal predictions was influenced by the complexity of modelling intent, which can span a large continuous space in unstructured environments. Moreover, pedestrian intent is highly context-dependent and influenced by key locations in the area, factors that were not explicitly encoded in our model. To mitigate this limitation, we aimed to improve the prediction of uncertainty in pedestrian behaviour. Specifically, we worked on providing less overconfident predictions of the covariance matrix to better reflect the uncertainty in their intent and subsequent trajectories.

6.6 Summary and conclusion

This chapter focused on the design of the pedestrian trajectory predictor model and how to effectively capture the interaction effects between both pedestrian-pedestrian and pedestrian-vehicle interactions in the prediction module. This work addresses the third primary research question of the thesis concerning the capture of interaction effects in the prediction component (**RQ3**).

We proposed using TTC as an effective heuristic for selecting interacting agents in response to the first sub-research question of this chapter (**RQ3.1**). The information of the selected interacting agents was then used to construct the interaction feature, which served as an input for the prediction model. For this construction, we proposed the PCG map. This map uses collision-risk information, including TTC and the approach direction of the interacting agents, to encode the interaction effects.

Our results showed that filtering neighbouring agents to focus only on those that could affect the target pedestrians by posing a collision risk allows better capture of the interaction effects on the predicted trajectories. By directly using information on the approaching risk direction and the criticality of the situation using TTC, we were able to predict corresponding changes in the trajectory more accurately than baseline methods, addressing the second sub-research question of this chapter (**RQ3.2**). According to the results, we believe

that providing the prediction network with more relevant engineered features helps deep-learning-based prediction methods better learn the patterns underlying the trajectories, leading to predictions closer to the ground truth.

Furthermore, by complementing the loss function of the deep learning prediction model with a novel uncertainty loss term, we were able to provide more accurate covariance predictions as a measure of the pedestrian's behaviour uncertainty, improving the measure of overconfidence predictions. The trained model with this modified loss function, which we called UAW-PCG, was an attempt to answer the last sub-research question of this chapter (**RQ4.1**) related to **RQ4**.

In the next chapter, we will demonstrate how our proposed prediction model integrates with the motion planner and how it can help improve the decision-making algorithm. This integration aims to make the decision-making process safer and more farsighted by also accounting for the uncertainties associated with these predictions.

Chapter 7

Vehicle motion planning

In this chapter, we focus on developing our DRL-based motion planner for the navigation of low-speed AVs in unstructured environment and its integration with the pedestrian trajectory prediction module. Our goal is to develop an overall navigation framework that is safe and efficient for the AV's operation among pedestrians, bringing together the information and models developed in the previous chapters. Throughout the design of this DRL motion planner, we aim to answer the following research questions:

RQ3.3 How can the interaction effects between the AV and pedestrians be effectively and proactively captured in the AV's decision-making and motion planning processes based on DRL?

RQ4.2 How can the uncertainties in the predicted trajectories of other agents be incorporated into the formulation of the DRL motion planner, and what are the positive effects of this incorporation?

RQ1.4 How can the perceived safe behaviours for the AV, discovered in the online survey of Chapter 4, be encouraged in the design of the DRL-based motion planner and to what extent does the behaviour of the trained model align with these perceived safe behaviours in similar scenarios?

These research questions will be explored as part of the design of the DRL motion planner and through the evaluation of the trained model on test scenarios, as well as on scenarios similar to those in the online survey.

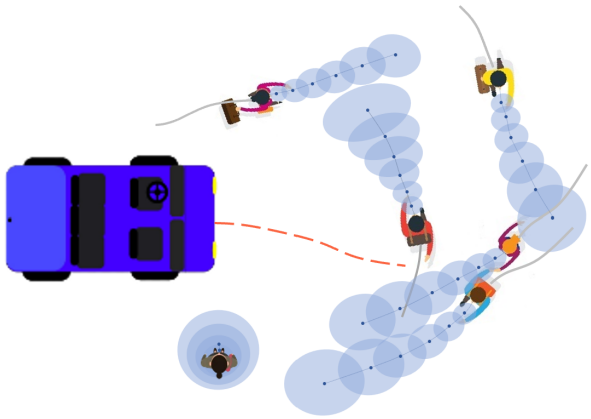

Figure 7.1: We propose the use of predicted position and its associated covariance in the training of a deep reinforcement learning model with a novel reward function. Through this uncertainty-aware coupled prediction and planning approach, our model learns safe, foresighted, and smooth navigation behaviours among pedestrians in a shared space.

We will be focusing on the DRL method for our motion planner design due to the strength of these models in learning complex interaction effects happening when navigating among other decision-making agents through interaction in a simulation environment while ensuring a fast computational time during inference. Due to these advantages in recent years, several crowd navigation algorithms based on DRL have been proposed [36, 31, 122, 201].

However, there are several limitations to the existing DRL models designed for crowd navigation: (i) They often rely solely on the current state of all agents, neglecting the benefits of incorporating predicted human trajectories or, when considering predicted trajectories, they ignore the uncertainties of pedestrians' behaviour, resulting in a potentially misleading reliance on prediction; (ii) Current reward functions have limitations in incorporating pedestrian comfort, particularly regarding the distance the robot should keep from pedestrians and the impact of robot's speed at close passes on pedestrian discomfort. (iii) Simulation environments used for training focus on robots close in size to pedestrians and make strong assumptions about pedestrian behaviours based on the ORCA model. These assumptions include cooperative behaviour, specifically sharing equal responsibility for collision avoidance, deterministic behaviours, resolving collisions immediately, and exhibiting abrupt trajectories or manoeuvre that may not align with human dynamics, ultimately failing to reflect real-world behaviours.

To overcome these limitations, we introduce an uncertainty-aware DRL-based coupled

prediction and planning model designed for navigating a low-speed AV within a shared space with pedestrians (Fig. 7.1). To address the first limitation of not accounting for prediction or their uncertainties, we enhance the model by incorporating not only the predicted future trajectories of pedestrians but also accounting for the uncertainties associated with these predictions in the formulation of a model-free DRL framework. To tackle the second limitation regarding pedestrian comfort, we propose a novel reward function component for the DRL model. This addition aims to encourage the AV to learn improved social behaviour. Finally, the last limitation concerning simulation environments is addressed by training our model in a simulation environment derived from realistic pedestrian-vehicle interactions in a shared space. This simulation environment is constructed using data from the HBS dataset [140, 141] to ensure a more realistic representation of pedestrian behaviours in shared spaces.

7.1 Problem formulation

The problem of navigating from a start point to a goal position among a crowd of pedestrians can be formulated as a sequential decision making process using the MDP notation defined through the tuple (S, A, P, R, γ). Let w_t be the AV's current state and s_t^i be the i-th pedestrian's current state. The state of the AV consists of its current position (p_t^{AV}), velocity, heading angle, size (specified as radius r^{AV}), preferred velocity and goal position ($w : [p_x, p_y, v_x, v_y, \theta, r^{AV}, v_{pref}, g_x, g_y]$).

We define pedestrian i's current state as its current position s_t^i: (p_x^i, p_y^i). Considering the pedestrians' position history over the last H time steps ($s_{t-H:t}^i$), we also predict the pedestrians' future position and its covariance over the next K time steps as $\hat{s}_{t+1:t+K}^i$ and $\hat{\Sigma}_{t+1:t+K}^i$ with the GT predictions being denoted as $s_{t+1:t+K}^i$. Therefore, we define the state of the MDP, $s_t \in S$ as the joint state of the AV and the n pedestrians states including their current and predicted state and covariance: $s_t^{jn} : (w_t, s_t^1, \hat{s}_{t+1:t+K}^1, \hat{\Sigma}_{t+1:t+K}^1, ..., s_t^n, \hat{s}_{t+1:t+K}^n, \hat{\Sigma}_{t+1:t+K}^n)$. Pedestrians modelled as circles have a radius denoted by r^{ped}.

Considering a unicycle model for the AV, we define the action $a_t \in A$ as the speed and heading change $a_t = [v_t, \Delta\theta]$. Taking the action at each time step the AV receives a reward $r_t \in R$ and the AV and the whole environment moves to the next step according to an unknown state transition probability $P(s_{t+1}^{jn}|s_t^{jn}, a_t)$. The objective is to find the optimal policy $\pi(a_t|s_t^{jn})$ that maps the joint state to an action while maximizing the cumulative discounted reward ($\sum_{k=0}^{\infty} \gamma^k r_{t+k}$) with a discount factor of γ. The navigation process will continue until (a) the AV reaches the goal, (b) collides with any pedestrians or (c) the

simulation exceeds the allowed time given to the AV to reach the goal. Figure 7.4 depicts an overview of the proposed framework, with detailed descriptions of each module provided in subsequent sections.

7.2 Deep reinforcement learning method

A model-free DRL network is used as the AV's decision-making algorithm. We built upon the DRL network architecture proposed in [122], which represents interactions among agents using edges in a spatio-temporal graph structure (see Fig. 7.2). This architecture captures both the direct impact of pedestrian-AV interactions and the indirect influence of pedestrian-pedestrian interactions on AV's decision-making (Fig. 7.3). These effects are modelled and aggregated through attention mechanisms. The temporal correlation of the graphs at consecutive time steps, referred to as temporal edges, is captured using a GRU. The GRU outputs features to both a value network and a policy network, both of which are trained together using PPO. To model the limited sensor range of the AV, all pedestrians within a specified range are considered to construct the spatial edges. The detailed process within the network architecture is explained next.

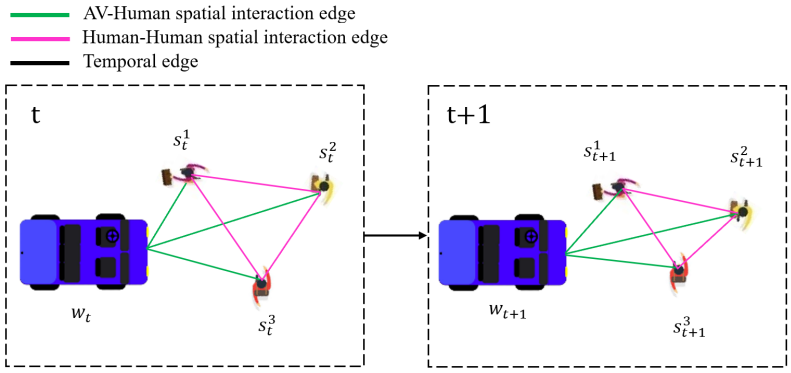

Figure 7.2: The crowd navigation problem represented as a graph with spatial and temporal edges. There are two types of spatial edges: one for pedestrian-autonomous vehicle interactions and another for human-human interactions.

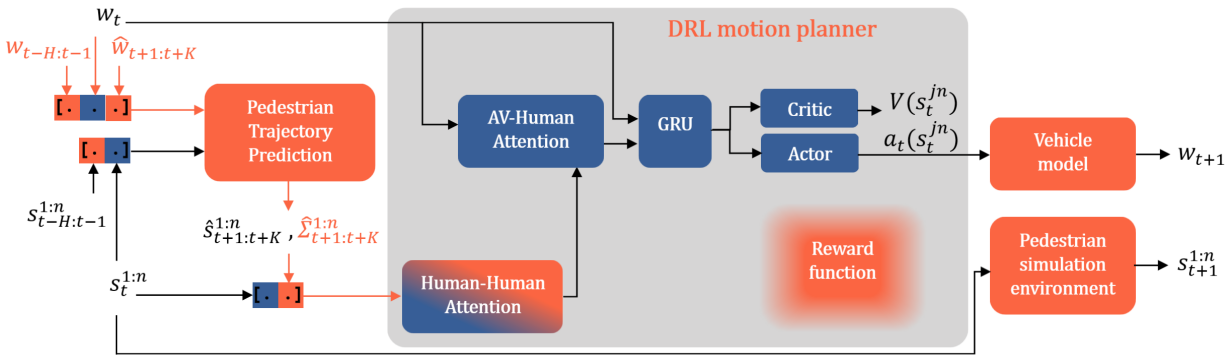

Figure 7.3: Our integrated prediction and planning framework. The DRL network architecture is based on the model introduced in [122], with novel additions and modifications highlighted in orange. The inclusion of prediction uncertainties (covariance) is a key feature, integrated both as part of the observation state and in the formulation of our novel reward function. This covariance is obtained from the predictor module, which addresses pedestrian-vehicle interaction by incorporating the AV's state and an estimate of its future path as input. In our model, prediction covariance, alongside the states themselves, contributes to the human-human attention module.

7.2.1 Network architecture

The indirect effect of interactions between pedestrians on the AV's motion is captured by encoding the human-human spatial edges and feeding them into an attention layer. The spatial edge consists of the information about pedestrians' current and future positions, as well as the covariance matrix of their predicted positions, all concatenated together. The attention mechanism functions similarly to the scaled dot-product attention introduced in [192] and stated as Eq. 7.1. Within this equation, the attention scores applied to the value V are computed using key K and query Q, and normalized by d_k, which is the dimension of the query Q and value V.

$$Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d_k}})V \tag{7.1}$$

For this problem, K, Q, and V are built by passing the embedded spatial edge through a linear layer as stated in Eq. 7.2:

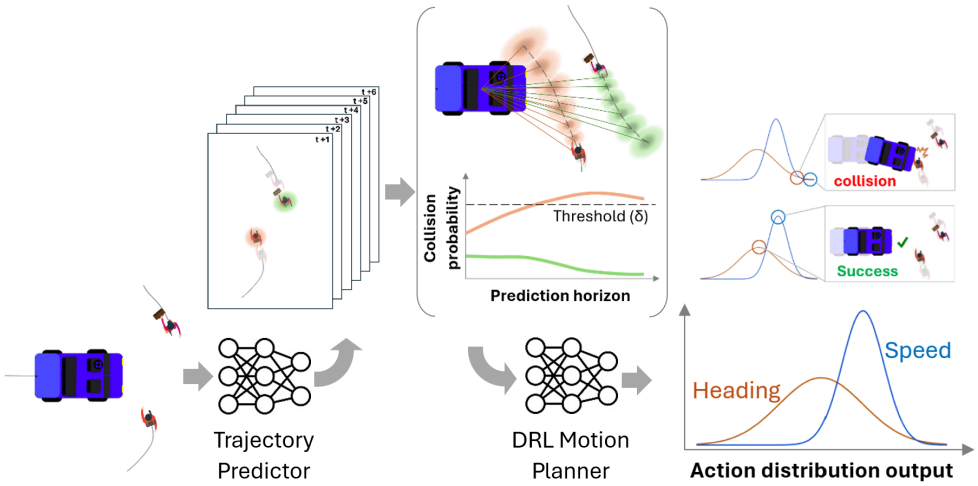

Figure 7.4: Overview of the proposed framework. At each time step, current observations are input to the predictor, generating pedestrians' predicted positions and associated uncertainties over a horizon. These predictions inform collision probability calculations between the AV and each pedestrian's predicted path, contributing to the reward function of the DRL module. Utilizing this predictive data, the DRL motion planner is trained to generate a probability distribution of optimal actions, mitigating collision risk with both current and intended pedestrian paths.

$$
\begin{aligned}
g_t^{HH} &= \phi_g([\hat{s}^1_{t:t+K}, \hat{\Sigma}^1_{t+1:t+K}, ..., s^n_{t:t+K}, \hat{\Sigma}^n_{t+1:t+K}]) \\
Q_t^{HH} &= [g_t^{HH}]^T W_Q^{HH} \\
K_t^{HH} &= [g_t^{HH}]^T W_K^{HH} \\
V_t^{HH} &= [g_t^{HH}]^T W_V^{HH}
\end{aligned}
\tag{7.2}
$$

In this equation, ϕ_g is a two layer neural network with ReLU nonlinearity. The human embedding, u_t^{HH} is then obtained using the multi-head scaled dot-product attention with 8 heads.

The AV-Human attention then captures the interaction effect between the AV and pedestrians by again using the scaled dot-product attention formulation of Eq. 7.1 where K this time is the embedded state information of the AV, and V and Q are the linear embedding of the weighted human feature output of the human-human attention (u_t^{HH}) as stated in Eq. 7.3. The number of attention heads is set to 1 for the AV-Human attention. In

Eq. 7.3, ϕ_v and ϕ_q are single layer neural networks with ReLU nonlineararity. Throughout this attention process, a binary mask is used to maintain attention only towards visible pedestrians within the vehicle's sensor range.

$$\begin{aligned} u_t^V &= \phi_v(w_t) \\ Q_t^{VH} &= [\phi_q(u_t^{HH})]^T W_Q^{VH} \\ K_t^{VH} &= [u_t^V]^T W_K^{VH} \\ V_t^{VH} &= [\phi_q(u_t^{HH})]^T W_V^{VH} \end{aligned} \tag{7.3}$$

The twice-weighted human feature output of the AV-Human attention, u_t^{VH} is then concatenated with the AV's embedded state u_t^V and passed through the GRU to encode the temporal edge and the time dependency between the spatial interaction over time as per Eq. 7.4. The output of the GRU then goes through two separate fully connected networks with *Tanh* nonlinearity functioning as the actor and critic networks to create the value $V(s_t^{jn})$ and the features of the policy network $\pi(a_t|s_t^{jn})$. Modelling the policy as a normal distribution, the actor features are then used to generate the parameters of the distribution from which the actions can be sampled.

$$h_t = GRU(h_{t-1}, (u_t^V, u_t^{VH})) \tag{7.4}$$

While our DRL network's overall architecture is based on the one proposed in [122], our model differs in four key aspects, as depicted in Fig. 7.3.

Firstly, our predictor, developed in Chapter 6 not only generates the covariance for the predicted position of pedestrians but also integrates it as part of the joint state input for the DRL policy generator, affecting the attention module that captures agent interactions. The predicted covariance also contributes to formulating the reward function during training.

Secondly, unlike the integrated prediction and planning model proposed in [122], our pedestrian trajectory predictor incorporates the AV's state as input, capturing the influence of the AV on the predicted pedestrian trajectory. We use the AV's most recent action to project its trajectory with a constant velocity model, providing an estimated trajectory for the AV over the next K steps ($\hat{w}_{t+1:t+K}$). This projection serves as a simple estimation of how pedestrians are likely to perceive the AV's trajectory in the near future. Subsequently, this estimated future trajectory, along with the current AV's observed trajectory,

is inputted into the pedestrian trajectory predictor module to model pedestrian-AV interactions within the predictor. It is important to note that the constant velocity projection of the AV's trajectory is solely utilized within the pedestrian trajectory prediction module for calculating AV-pedestrian interaction features. Based on these predictions, the DRL module outputs only the best action for the AV at the current time step. Once this single action is executed and the environment transitions to the next time step, the entire process repeats with the newly available observations.

Thirdly, we adopt a unicycle kinematic model, which more accurately represents the non-holonomic kinematics of vehicles, as opposed to the holonomic kinematic model utilized in [122] for training the DRL network.

Finally, our model is trained within our own developed novel simulation environment with pedestrian behaviour based on realistic human trajectories in a space shared with vehicles rather than assuming pedestrians comply with the ORCA model.

7.2.2 Reward formulation

In the design of the reward function for the DRL network, the AV is given a high reward when it reaches the goal point within a predefined acceptable radius outlined in the goal set S_{goal}. Conversely, the system penalizes the AV not only for collisions with pedestrians but also for intruding into their Personal Space (PS), referred to as "danger cases" (see Fig. 7.5).

In all other instances, the AV is rewarded for its progress towards the goal ($r_{progress}$), penalized for large actions (r_{action}), and for coinciding with pedestrians' future predicted positions (r_{pred}). The reward function is outlined in Eq. 7.5, where d_{min} represents the minimum distance between the AV and all pedestrians. This is computed by subtracting the sum of the agents' radii from the Euclidean L^2-norm of the difference between their positions, where $\|.\|$ represents the Euclidean norm ($d_{min} = \min_i(\|p^{AV} - s^i\| - r^{AV} - r^{ped})$.

$$r(s_t^{jn}, a_t) = \begin{cases} r_g & \text{if } s_t^{AV} \in S_{goal} \\ r_c & \text{if } d_{min} < 0 \\ r_{danger} & \text{if } d_{min} < PS \\ r_{progress} + r_{pred} + r_{action} & \text{otherwise} \end{cases} \tag{7.5}$$

To model pedestrians' discomfort, we explore two versions of danger penalties. In the first version, we adopt a widely used penalty formulation from the literature. Here, the AV

incurs a linearly increasing penalty as its proximity to a pedestrian falls below a specified threshold—the closer the AV gets, the higher the penalty [122, 123, 31]. However, this approach falls short in addressing the influence of the AV's speed on pedestrians' comfort zones [26, 44], neglecting the individuals' tendency to maintain a larger distance from a vehicle travelling at higher speeds.

To address this limitation, we incorporate the AV's speed as a crucial factor in the second version of the danger penalty, referred to as speed-dependent (SD). We introduce the AV's speed as a determinant in calculating the penalty, as outlined in Eq. 7.6. This modification is important, particularly when modelling interactions in a mixed human-vehicle environment, recognizing the elevated risk linked to vehicles compared to smaller robots.

$$r_{danger} = \begin{cases} r_c(1 + \frac{v_{AV}}{v_{max}})(\frac{PS - d_{min}}{PS}) & \text{if speed depend. (SD)} \\ r_c(\frac{PS - d_{min}}{PS}) & \text{otherwise} \end{cases} \tag{7.6}$$

The progress reward ($r_{progress} = (-d^t_{goal} + d^{t-1}_{goal})$), is aimed at guiding the AV towards the goal position, with d_{goal} representing the distance of the AV from the goal position. The action reward (r_{action}), as stated in Eq. 7.7 is to penalize the AV for large heading changes,

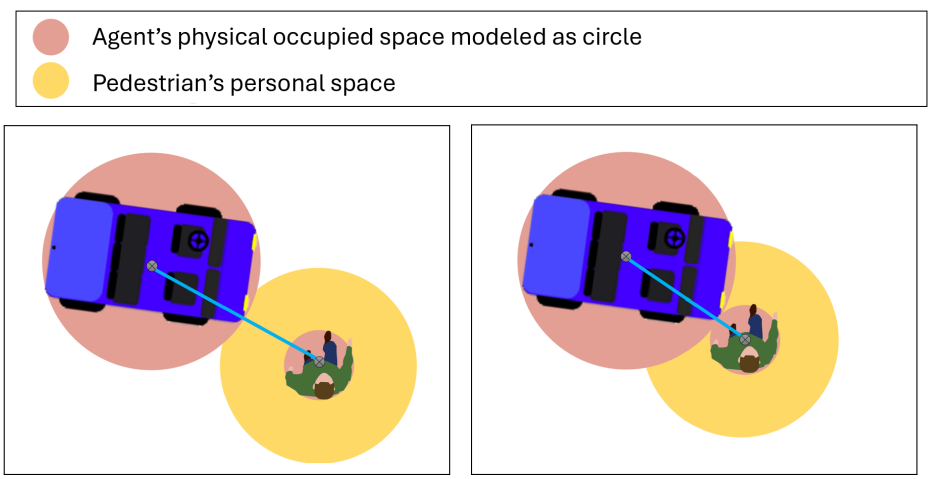

(a) Danger (b) Collision

Figure 7.5: Danger (a) is specified as a condition when the vehicle intrudes into the pedestrian's personal space (PS), shown as the yellow circle. Collision (b) occurs when the physically occupied spaces of the two agents, modelled as pink circles, overlap.

in line with the findings from our online survey, as well as backward motion, to encourage smoother trajectories. In this equation, $|u|$ is the absolute value of u, and $\mathbb{1}(u < 0)$ is an indicator function that equals 1 if $u < 0$, and 0 otherwise.

$$\begin{aligned} r_{action} &= -\alpha \cdot r_{heading} - \beta \cdot r_{speed} \\ r_{heading} &= a_{heading}^2 \\ r_{speed} &= |a_{speed}| \cdot \mathbb{1}(a_{speed} < 0) \end{aligned} \tag{7.7}$$

The prediction reward (r_{pred}) involves penalizing the AV if its position coincides with the pedestrians' intended future position, as per their predicted trajectory. This aspect of the reward function was initially introduced in [122]. However, in contrast to their deterministic prediction, our approach incorporates the uncertainties associated with predicted trajectories embedded in the covariance matrix of the predicted distribution when formulating the prediction reward. Consequently, we impose a penalty on the AV when the collision probability between its current state and the predicted position distribution of pedestrian i exceeds a threshold δ, based on Eq. 7.8 [53]. In this equation, V_s is the volume of the sphere of radius $r^{ped} + r^{AV} + PS$. We indicate the occurrence of this condition using the indicator $\mathbb{1}_{t+k}^i$ (Eq. 7.9), for all the pedestrians' predicted state distribution over the next K time steps, with a decreased penalty for collisions in steps further into the future as per Eq. 7.10.

$$\begin{aligned} &P_{t+k}^i(collision) \approx \\ &V_s \frac{1}{\sqrt{det(2\pi\hat{\Sigma}_{t+k}^i)}} \exp(-\frac{1}{2} d^{MD}(p_t^{AV}, \hat{s}_{t+k}^i, \hat{\Sigma}_{t+k}^i)^2) \end{aligned} \tag{7.8}$$

$$\mathbb{1}_{t+k}^i = \begin{cases} 1 & \text{if } P_{t+k}^i(collision) > \delta \\ 0 & \text{otherwise} \end{cases} \tag{7.9}$$

$$\begin{aligned} r_{pred}^i(s_t^{jn}) &= \min_{k=1,\dots,K} (\mathbb{1}_{t+k}^i \frac{r_c}{2^k}) \\ r_{pred}(s_t^{jn}) &= \min_{i=1,\dots,n} r_{pred}^i(s_t^{jn}) \end{aligned} \tag{7.10}$$

7.3 Model predictive control method

Due to the higher safety criticality of a low-speed autonomous vehicle in a crowded environment, linked to their larger size and speed compared to small mobile robots, we conduct a comparison between our DRL model and methods from optimal control, where collision avoidance can be formulated as a constraint. In this context, we adopt the MPC method as a baseline, ensuring a fair comparison as it can be integrated with the pedestrian trajectory predictor in a manner similar to our DRL model. In the tests of the MPC navigation algorithm, the entire DRL motion planner block depicted in Fig. 7.3 will be substituted with the MPC-based motion planner. Since we are using the same trajectory predictor for pedestrians' future states as in our DRL algorithm, the prediction horizon of the MPC is set to match the prediction length of our predictor, which is six timesteps, with each timestep being 0.5 seconds.

We adopt the MPC formulation used for crowd navigation in [2], where the base optimization problem is defined by minimizing the following cost:

$$\min_{p_{1:K}, a_{0:K-1}} J_K(p_K^{AV}) + \sum_{k=0}^{K-1} J_k(a_k, p_k^{AV}) \tag{7.11}$$

The cost in this formulation consists of a stage cost and a terminal cost, as defined in Eq. 7.12 and Eq. 7.16 respectively.

$$J_k(a_k, p_k^{AV}) = J_k^a(a_k) + J_k^p(p_k^{AV}) + J_k^{ED}(p_k^{AV}, s_k^{1:n}) \tag{7.12}$$

The stage cost consists of three components: a control input cost defined in Eq. 7.13, a normalized target distance cost outlined in Eq. 7.14, and a pedestrian closeness cost specified in Eq. 7.15. The control input cost aims to minimize the usage of control signals. The normalized target distance cost serves a similar purpose as the progress-to-goal component in the reward function of the DRL formulation, while the pedestrian closeness cost penalizes the AV for getting close to pedestrians. In these equations, Q_a, Q_p, and Q_{ED} are the weight coefficients.

$$J_k^a(a_k) = a_k^\top Q_a a_k \tag{7.13}$$

$$J_k^p(p_k^{AV}) = Q_p \left(\frac{\|p_k^{AV} - p_{goal}\|}{\|p_0^{AV} - p_{goal}\|} \right)^2 \tag{7.14}$$

The pedestrian closeness cost is defined as the sum of the inverse of the Euclidean distances between the AV and all pedestrians within the AV's field of view.

$$J_k^{ED}(p_k^{AV}, s_k^{1:n}) = Q_{ED} \sum_{i=1}^{n} \frac{1}{d_{i,k}^{ED}(p_k^{AV}, s_k^i)^2} \tag{7.15}$$

Finally, within the terminal cost, the AV is penalized if it fails to reach the goal position by the end of the prediction horizon. The prediction horizon for the MPC optimization problem is determined by the prediction length of the predictor, which is currently set to 6 steps.

$$J_K(p_K^{AV}) = J_k^p(p_K^{AV})|_{k=K} \tag{7.16}$$

This thesis addresses the optimization problem outlined in Eq. 7.11, as a baseline, by integrating two distinct constraints for collision avoidance adapted from [2]: the uncertainty-unaware and uncertainty-aware constraints. For the uncertainty-unaware MPC, two versions of hard and soft constraints on the minimum distance will be considered. All the different versions of the MPC baseline will be discussed in the following subsections.

7.3.1 MPC with a hard constraint on minimum distance

Under the uncertainty-unaware constraint, Eq. 7.17 establishes a minimum threshold for the Euclidean distance between the AV and each pedestrian, using the mean of the predicted distribution of pedestrian positions within the prediction horizon. Using this as a hard constraint to complement Eq. 7.11, we define our first MPC baseline method, as outlined in Eq. 7.18.

$$d_{k,i}^{ED}(p_k^{AV}, \hat{s}_k^i) \geq (r^{ped} + r^{AV} + PS) \tag{7.17}$$

$$\begin{aligned} \min_{p_{1:K}, a_{0:K-1}} \quad & J_K(p_K^{\text{AV}}) + \sum_{k=0}^{K-1} J_k(a_k, p_k^{\text{AV}}) \\ \text{subject to} \quad & p_0 = p(0) \\ & a_{\min} \leq a_k \leq a_{\max} \\ & d_{k,i}^{ED}(p_k^{AV}, \hat{s}_k^i) \geq (r^{ped} + r^{AV} + PS) \end{aligned} \tag{7.18}$$

7.3.2 MPC with a soft constraint on personal space

In the second version of the uncertainty-unaware MPC baseline, we relax the portion of the hard constraint on the minimum distance related to respecting personal space. Here, the constraint on avoiding PS is made soft by introducing a slack variable sv and adding it to the control input vector to be minimized ($a_t = [v_t, \Delta\theta, sv]$). To ensure this soft constraint applies only to personal space violations while still guaranteeing physical collision avoidance, we consider an upper bound for the slack variable as a new constraint added to the MPC formulation. Therefore, the optimization problem becomes as follows:

$$\begin{aligned} \min_{p_{1:K}, a_{0:K-1}} \quad & J_K(p_K^{\text{AV}}) + \sum_{k=0}^{K-1} J_k(a_k, p_k^{\text{AV}}) \\ \text{subject to} \quad & p_0 = p(0) \\ & a_{\min} \leq a_k \leq a_{\max} \\ & d_{k,i}^{ED}(p_k^{AV}, \hat{s}_k^i) \geq (r^{ped} + r^{AV} + PS - sv) \\ & sv < PS \end{aligned} \tag{7.19}$$

7.3.3 MPC with collision probability hard constraint

For the uncertainty-aware version of our MPC baseline, a minimum threshold δ is set for the collision probability, considering both the covariance of predicted future pedestrian positions and their means (see Eq. 7.8). This is translated into a minimum Mahalanobis distance (MD), as stated in Eq. 7.20 [2]:

$$d_{k,i}^{MD}(p_k^{AV}, \hat{s}_k^i, \hat{\Sigma}_k^i)^2 \geq -2\ln(\sqrt{\det(\pi\hat{\Sigma}_k^i)}\frac{\delta}{V_s}) \tag{7.20}$$

Therefore, the optimization problem of the uncertainty-aware MPC is defined as follows:

$$\begin{aligned} \min_{p_{1:K}, a_{0:K-1}} \quad & J_K(p_K^{\text{AV}}) + \sum_{k=0}^{K-1} J_k(a_k, p_k^{\text{AV}}) \\ \text{subject to} \quad & p_0 = p(0) \\ & a_{\min} \leq a_k \leq a_{\max} \\ & d_{k,i}^{MD}(p_k^{AV}, \hat{s}_k^i, \hat{\Sigma}_k^i)^2 \geq -2\ln(\sqrt{\det(\pi\hat{\Sigma}_k^i)}\frac{\delta}{V_s}) \end{aligned} \tag{7.21}$$

7.4 Experiments

7.4.1 Simulation environment

The majority of crowd navigation algorithms undergo training in a simulation environment where humans are controlled by the ORCA model [31, 30, 122, 36]. While this simulation environment can be effective for applications involving small mobile robots in pedestrian crowds, our focus in this study is on the navigation of a low-speed AV in shared spaces with pedestrians. Modelling the heterogeneous interactions in such environments, particularly distinguishing between pedestrian-vehicle and pedestrian-pedestrian interactions, poses a challenge within the ORCA simulation framework. Moreover, the assumption that pedestrians consistently adhere to the ORCA model and actively participate in collision avoidance, as implied by taking half the responsibility, is not always valid, for instance, in cases involving distracted and non-cooperative pedestrians.

Hence, we employ our custom-built simulation environment, developed based on a real-world HBS dataset that has captured pedestrians' and vehicles' trajectories in a shared space [140, 141]. This approach allows us to train our AV in an environment that simulates real pedestrian behaviour. The same dataset was utilized to train the data-driven pedestrian trajectory prediction module in Chapter 6.

We extracted 310 scenarios from this dataset corresponding to the 331 vehicles initially present. Some vehicles that appeared toward the end of the dataset were excluded due to their short and insufficiently captured trajectories. Each extracted scenario is dedicated to one of the existing vehicles, tracing its trajectory from the initial frame upon entering the scene until it exits. The vehicle's positions at the beginning and end of each scenario are designated as the starting and goal positions within our simulation environment. The 310 scenarios are split into training, testing, and validation sets, with proportions of 64%, 20%, and 16%, respectively.

The behaviours of pedestrians in each scenario are simulated based on real-world data, while the actions of the AV are governed by the DRL policy network. The extracted scenarios involve highly interactive situations, requiring the AV to decelerate to yield to pedestrians or accelerate to pass a pedestrian to avoid collisions. In the most crowded scenarios, the overall scene can include up to 60 pedestrians over an approximately 1200 m^2 area (density of ≈ 0.05 pedestrian/m^2). The simulation environment, integrated into a gym environment, has been made publicly available as a benchmark for autonomous vehicle crowd navigation training[1].

[1] https://github.com/Golchoubian/UncertaintyAware_DRL_CrowdNav

7.4.2 Baselines and ablation models

Ablation models for DRL navigation method

For the entire integrated DRL navigation algorithm, we compare five models. Unless otherwise specified, all models are trained with the speed-independent penalty for danger cases. The first model (DRL. No pred) does not utilize the prediction output in DRL training. The second model (DRL. No pred, SD.) also omits prediction but is trained with a speed-dependent (SD) danger penalty to better capture pedestrians' discomfort at higher vehicle speeds. The third model (DRL. PCG pred) employs the prediction output of the Polar Collision Grid (PCG) predictor, trained exclusively with a NLL loss function. The fourth model (DRL. UAW-PCG pred) utilizes the prediction output of the Uncertainty-Aware Polar Collision Grid (UAW-PCG). Finally, the last model (GT pred) represents a hypothetical model with access to the ground truth (GT) future prediction, serving as an upper bound for the performance of models using deterministic prediction outputs. Additionally, we have included the output of all our measures on the trajectory of the actual human-driver for the ego vehicle in the scenario, as per the real dataset (Human-driver).

MPC baseline method for social navigation

We employ three versions of the MPC-based social navigation model proposed in [2] and detailed below.

MPC. dist. hard.: This MPC model, known as ED-MPC-EDC in [2], incorporates a hard constraint on the minimum Euclidean distance between the AV and each pedestrian. This constraint ensures that the distance exceeds the sum of their two radii and the pedestrian's PS.

MPC. dist. soft.: This second MPC model, referred to as ED-MPC-AEDC in [2], represents an adaptive version of the first model. In this case, the constraint for avoiding intrusion into the PS of pedestrians is made soft by introducing a slack variable as part of the control inputs to be minimized. Simultaneously, we ensured the model upholds a hard constraint for physical collision avoidance by incorporating an upper bound on the slack variable.

MPC. col. hard.: This third MPC model, denoted as ED-MPC-MDC in [2], is an uncertainty-aware MPC. It imposes a hard constraint on keeping the collision probability with the current and predicted position distribution of pedestrians below a threshold δ (set at 10% here) as per Eq. 7.20.

MPC. col. hard. PCG: This model is similar to the previous one (MPC. col. hard.) but employs the original PCG predictor instead of the uncertainty-aware predictor. This

model aims to showcase the effectiveness of the UAW-PCG prediction model over the original PCG, not only within the context of the DRL model but also for an MPC motion planner that takes prediction uncertainties into account.

7.4.3 Evaluation metrics

To assess the trained DRL navigation policy, we incorporate measures for both navigation efficiency and social aspects related to pedestrian comfort. Concerning navigation efficiency metrics, we provide information on the proportion of test set scenarios that successfully reach the goal (Success Rate (SR)), collide with any pedestrian (Collision Rate (CR)), or experience a timeout (Timeout Rate (TR)).

Regarding the social aspects, we employ measures such as intrusion ratio, minimum intrusion distance, and intrusion speed. Intrusion ratio is defined as the percentage of time throughout the entire navigation duration in which the AV intrudes into the personal space of any pedestrian. Furthermore, we define minimum intrusion distance as the smallest distance recorded during these intrusion instances and intrusion speed as the AV's speed at that distance. We report the mean and standard deviation of these measures across the scenarios in the test set.

7.4.4 Implementation details

We maintained the network architecture of the DRL model consistent with [122]. Pedestrians are represented by circles with a radius of 0.3 meters and a personal space (PS) of 1 meter beyond their physical dimension. The AV is modelled with a circle of a 1-meter radius. The maximum speed of the AV is set at 15 km/h, slightly below the allowable speed limit of 20 km/h in the shared space area where the HBS dataset was collected [140], from which we extracted our scenarios. Additionally, we limit the heading change of the AV output from the policy to 0.2 rad/sec. The simulation environment operates with a time step of 0.5 seconds, equivalent to the HBS dataset's frame rate of 2 frame per second (fps). Each scenario in the simulation environment is configured with a maximum allowable time of 15 seconds more than the total time for the driver to reach the goal point in the real dataset. This designated time serves as the timeout for the DRL training.

The reward for reaching the goal and the penalty for collisions, denoted as r_g and r_c in Eq. 7.5, are set to +10 and -20, respectively. The AV is assumed to have a range of view extending up to 15 meters. Lastly, the minimum collision probability threshold, δ, is set to 0.1 for both the DRL and MPC models. This value, used in the literature [182],

Table 7.1: DRL navigation and baseline results. Bold is the best. Underline is the second best. (SR: Success Rate, CR: Collision Rate, TR: Timeout Rate)

Model	**SR**	**CR**	**TR**	**Nav. time (s)**	**Path length (m)**	**Intrusion ratio (%)**	**Intrusion dist. (m)**	**Intrusion speed (m/s^2)**
MPC. dist. hard.	0.84	0.07	0.09	$15.29_{\pm 5.68}$	$47.94_{\pm 6.78}$	$1.55_{\pm 4.77}$	$0.72_{\pm 0.27}$	$2.47_{\pm 1.88}$
MPC. dist. soft	0.84	0.07	0.09	$15.24_{\pm 6.02}$	$47.81_{\pm 6.71}$	$2.02_{\pm 5.63}$	$\mathbf{0.79}_{\pm 0.25}$	$2.79_{\pm 1.76}$
MPC. col. hard.	0.88	0.05	0.07	$15.49_{\pm 5.82}$	$48.73_{\pm 8.16}$	$\underline{0.07}_{\pm 0.50}$	$0.02_{\pm 0.00}$	$4.16_{\pm 0.00}$
MPC. col. hard. PCG	0.84	0.09	0.07	$14.88_{\pm 5.13}$	$47.67_{\pm 6.65}$	$\mathbf{0.00}_{\pm 0.00}$	-	-
DRL. No pred	0.84	0.07	0.09	$\mathbf{13.41}_{\pm 3.49}$	$46.25_{\pm 6.07}$	$3.79_{\pm 5.75}$	$0.68_{\pm 0.27}$	$2.12_{\pm 1.49}$
DRL. No pred. SD.	$\underline{0.90}$	**0.02**	0.09	$15.24_{\pm 5.15}$	$46.65_{\pm 6.52}$	$2.35_{\pm 4.04}$	$0.68_{\pm 0.26}$	$\mathbf{1.84}_{\pm 1.44}$
DRL. PCG pred	$\underline{0.90}$	0.05	$\underline{0.05}$	$\underline{14.50}_{\pm 4.96}$	$\mathbf{45.41}_{\pm 5.69}$	$1.80_{\pm 3.67}$	$0.64_{\pm 0.24}$	$\underline{1.99}_{\pm 1.74}$
DRL. UAW-PCG pred (Ours)	**0.95**	$\underline{0.03}$	**0.02**	$14.62_{\pm 4.40}$	$\underline{46.06}_{\pm 6.00}$	$3.11_{\pm 5.38}$	$\underline{0.74}_{\pm 0.21}$	$2.00_{\pm 1.28}$
DRL. GT pred	0.95	0.02	0.03	$14.47_{\pm 4.16}$	$45.60_{\pm 5.74}$	$2.24_{\pm 3.56}$	$0.68_{\pm 0.25}$	$2.62_{\pm 1.39}$
Human-driver	1.00	0.00	0.00	$16.10_{\pm 5.57}$	$45.83_{\pm 6.59}$	$2.54_{\pm 3.93}$	$0.62_{\pm 0.29}$	$2.07_{\pm 1.66}$

was selected without tuning for optimized performance. However, reducing this collision probability threshold may lead to a higher occurrence of infeasible solutions in the MPC model. All DRL models undergo training for 4×10^7 time steps. The initial learning rate is set to 4×10^{-5} and undergoes linear decay after 5000 time steps.

7.5 Results and discussion

7.5.1 Quantitative results

The results of the trained DRL motion planner and its ablations, as well as the MPC baseline model, are presented in Table 7.1.

Effectiveness of deep reinforcement learning algorithm: The DRL models demonstrate superior performance compared to MPC models, exhibiting an overall higher success rate and a lower collision rate. Additionally, the DRL model tends to follow more direct paths to the goal, resulting in a reduced path length and lower navigation times in comparison to MPC models (refer to Table 7.1). It is worth noting that since the MPC model does not include a term equivalent to the speed-dependent danger penalty used in the "DRL. No pred. SD" model, a fair comparison of MPC with the DRL models in Table 7.1 should exclude this variant.

The average computational time required for each model to determine an optimal action

Table 7.2: One step average computational time of different models.

Model	**Computational time (sec)**
DRL. No pred	$0.01_{\pm 0.00}$
DRL. UAW-PCG pred (Ours)	$0.04_{\pm 0.02}$
MPC. col. hard.	$0.10_{\pm 0.08}$

at every time step is detailed in Table 7.2. The findings reveal that the DRL model operates more than twice as fast as the MPC model. This is particularly crucial for real-time implementation in crowded spaces, where satisfying a multitude of constraints (one for each visible pedestrian) can lead to extended convergence times for MPC, or even infeasibility in highly dense spaces. Consequently, the DRL model, which offloads computation time to the training phase and exhibits faster performance during inference, holds a distinct advantage in terms of computational efficiency.

Collisions in the MPC model arise either when it encounters infeasible solutions in crowded scenarios or when the predicted collision probability falls below the set threshold in the constraint for the MPC model with a minimum collision probability constraint (MPC. col. hard.). The latter cause of collisions can potentially be linked to the fact that Eq. 7.8 is an approximation and is subject to some error. Introducing a soft constraint to relax the hard constraint associated with avoiding pedestrians' personal space in the (MPC. dist. soft.) model can mitigate the occurrence of infeasible solutions. However, this comes at the cost of increased intrusion into pedestrians' personal space, as reported in Table. 7.1.

The MPC model with a hard constraint defined on the minimum collision probability exhibits the lowest collision rate of 0.05 among all MPC models. This highlights that the MPC model, which considers the uncertainty of pedestrians' next states, is better equipped to avoid intrusion into pedestrians' personal space (as indicated by the intrusion ratio in Table. 7.1), ultimately resulting in a reduced collision rate. However, the success of this MPC model depends on accurate uncertainty predictions. This is evident in the higher collision rate observed in a similar MPC model that, instead of utilizing the predictions of the UAW-PCG, relies on the overconfident predictions of the original PCG (MPC. col. hard. PCG).

Effectiveness of speed dependent danger penalty: In the two cases of the DRL model where no prediction is used, the model trained with a speed-dependent (SD) danger penalty exhibits a significantly reduced collision rate and intrusion ratio compared to the model trained with a speed-independent danger penalty in the reward function (Table 7.1). Additionally, the speed of the AV during the intrusion of pedestrians' personal space

for the DRL model trained with a speed-dependent penalty is lower than the speed of the AV in intrusion cases in the baseline model (DRL - No pred). Given that the only difference between these two models lies in the danger penalty, the results suggest that incorporating the impact of the AV's speed during intrusion on people's discomfort into the reward function contributes to learning a policy that better respects people's PS and reduces collision occurrences in the simulation.

Effectiveness of prediction: Consistent with the findings in [122], incorporating pedestrians' predicted trajectories during the training of the DRL model has a notable impact. This is evident in the increased success rate, reduced collision rate and timeout rate, as well as a lower intrusion ratio, as detailed in Table. 7.1. The comparison between the DRL model that does not consider prediction (DRL - No pred) and the one with prediction (DRL, PCG pred) and (DRL, UAW-PCG pred) highlights these improvements.

However, this enhancement comes at the expense of increased computational time for generating predictions, as shown in Table 7.2. Nonetheless, the computational time remains significantly below the 0.5-second time step at which the simulation operates.

Effectiveness of the uncertainty-aware DRL model: Our proposed model, which considers uncertainties in pedestrians' predicted trajectories (DRL - UAW-PCG), exhibits a 40% decrease in collision rate and 60% decrease in timeout rate, along with an overall 5.5% increase in the success rate, compared to the baseline model that neglects prediction uncertainties (DRL - PCG pred), as detailed in Table. 7.1. Furthermore, our uncertainty-aware prediction model is observed to maintain a greater distance from pedestrians when intruding into their personal space. The hypothetical DRL model assuming access to ground truth predictions without considering prediction uncertainty outperforms the same model relying on the predicted mean of the PCG module. This observation emphasizes that the performance degradation in the uncertainty-unaware model (DRL - PCG pred) is attributed to predictions deviating from ground truth trajectories, emphasizing the importance of considering uncertainties. Consequently, our model, which accounts for these uncertainties, demonstrates performance closer to the hypothetical model with access to ground truth (GT pred) and even surpasses it on certain safety metrics such as intrusion distance and intrusion speed. This superiority is derived from its more cautious behaviour, stemming from reliance on the predicted distribution of pedestrian future positions rather than single future position points.

Comparison to human driver trajectory: The average path length of our proposed model is slightly higher than that of the human driver in the dataset while achieving a lower navigation time (Table 7.1). This indicates that the average velocity of the AV in our model was higher while deviating slightly more from the straight path to the goal compared

to the human driver in the same scenario. It is important to note that these differences in the path length of our DRL model, compared to the human-driven trajectory, are not significant when contrasted with the path length of the MPC model. The MPC model's path length is statistically longer than the overall path length of the human driver ($se = 0.548, t = -4.121, p < 0.001$), as illustrated in Figure 7.6. This difference is attributed to the MPC model primarily deviating from a straight path for collision avoidance, whereas our model's strategy prioritizes speed adjustment, resulting in a trajectory closer to that of the human driver.

It is important to note that the intrusion ratio in the human driver case is not zero due to the errors introduced by modelling the vehicle with a circle. This results in the simulation environment detecting intrusion from the side when the vehicle passes a pedestrian despite maintaining an acceptable distance with its rectangular shape in the real-world data. The conservative assessment with a circular vehicle shape model contributes to this detection of intrusion for the human driver.

Reward sensitivity analysis: The values in the reward function were selected based on the literature [122, 123] and manually adjusted to enhance convergence and strike a balance between safety and efficiency. Here, we demonstrate this trade-off between safety and efficiency in the final learned policy through a sensitivity analysis of the goal achievement reward value (r_g) relative to the collision penalty value (r_c).

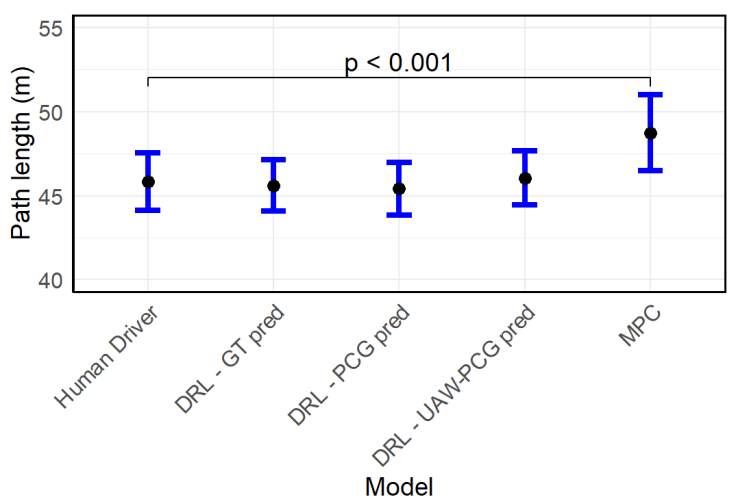

Figure 7.6: Mean and 95% confidence interval of difference model's path length over the test set scenarios.

Table 7.3: Sensitivity analysis of the reward value for goal achievement (r_g). The default value is denoted by (*).

Model	**SR**	**CR**	**TR**	**Nav. time(s)**	**Path length (m)**	**Intrusion ratio (%)**	**Intrusion dist. (m)**	**Intrusion speed (m/s^2)**
1	0.93	0.00	0.07	$15.11_{\pm 5.02}$	$45.66_{\pm 6.00}$	$2.17_{\pm 4.30}$	$0.64_{\pm 0.22}$	$1.75_{\pm 1.47}$
10*	0.95	0.03	0.02	$14.62_{\pm 4.40}$	$46.06_{\pm 6.00}$	$3.11_{\pm 5.38}$	$0.74_{\pm 0.21}$	$2.00_{\pm 1.28}$
100	0.90	0.07	0.03	$13.28_{\pm 4.13}$	$47.09_{\pm 6.72}$	$4.25_{\pm 7.17}$	$0.62_{\pm 0.28}$	$2.26_{\pm 1.11}$

With the default value of r_g set to 10 for the produced results of Table 7.1, we also trained and evaluated models with 1/10 and 10 times the default value. The training convergence plots are shown in Figure 7.7, and quantitative results for test scenarios are presented in Table 7.3. Reducing r_g to 1 results in zero collision rate at the expense of a slightly reduced success rate due to increased timeouts. This happens because, with a lower r_g compared to the default collision penalty of -20, the model prioritizes collision avoidance over reaching the goal, leading to longer navigation times on average and demonstrating a more conservative behaviour.

Conversely, significantly increasing the reward value for reaching the goal to 100 results in a policy that prioritizes goal achievement at the expense of collision avoidance. In this case, the AV rushes to reach the goal, resulting in shorter navigation times but riskier behaviour, as evidenced by higher intrusion rates, increased speed during intrusions, and shorter intrusion distances in Table 7.3. This behaviour ultimately leads to a higher colli-

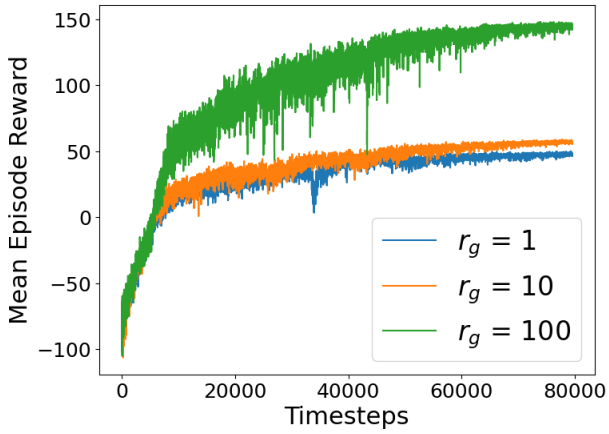

Figure 7.7: Training convergence plots for varying goal achievement reward values while maintaining a constant collision penalty coefficient.

sion rate and a decrease in the overall success rate. The trends revealed in this sensitivity analysis underscore the critical role of the reward function and how adjusting the relative reward values can assist in achieving an acceptable balance between safety and efficiency.

7.5.2 Qualitative results

The behaviours of the different DRL models, along with the best MPC model, are compared in the same scenario from the test set in Fig. 7.8-7.9. The baseline model, which does not use predictions (DRL - No pred), ends up getting too close to a pedestrian and ultimately colliding with it. The model trained with a speed-dependent danger penalty successfully avoids any collision in the same scenario by learning to slow down when approaching pedestrians, yet it still intrudes into a pedestrian's personal space at frame 16 (see Fig. 7.8). In contrast, our proposed uncertainty-aware DRL model, utilizing future predicted states of pedestrians and the associated uncertainties, successfully reaches the goal in this scenario without intruding into any pedestrian's personal space. This is achieved by predicting high uncertainties for a pedestrian in frame 11, prompting the AV to slow down from that frame (see Fig. 7.9).

The comparison between our uncertainty-aware DRL model and the best-performing MPC model in the same scenario is illustrated in Fig. 7.9. Although both models successfully reach the goal, the MPC model encounters infeasibility over a couple of frames, including frame 11. Despite not colliding with any pedestrian, this results in a highly unnatural planned path for that time step, posing a potential risk of collision. Another noticeable difference in the output behaviour of these two models is that the MPC model avoids the crowd in front by deviating to one side, while the DRL model achieves collision avoidance through speed adjustment, primarily maintaining a straight line towards the goal. This distinction is supported by the average lower path length of DRL methods compared to the MPC models, as shown in Table 7.1 and Fig 7.6. This behavioural difference is observed despite both models being constrained to the same upper and lower bounds for speed and heading change, and both models being encouraged to keep the heading change low (heading change control appears in the stage cost of MPC as well as the reward function of the DRL model). Besides the fact that speed adjustment behaviour better aligns with the path taken by the human driver in the same scenario from the real dataset, it is noteworthy that for collision avoidance, speed adjustment has been shown to be perceived safer by people compared to heading change, as shown by our survey in Chapter 4.

In Fig. 7.10, the performance of our uncertainty-aware model (DRL - UAW-PCG pred) is compared with its baseline, the uncertainty-unaware counterpart model (DRL - PCG

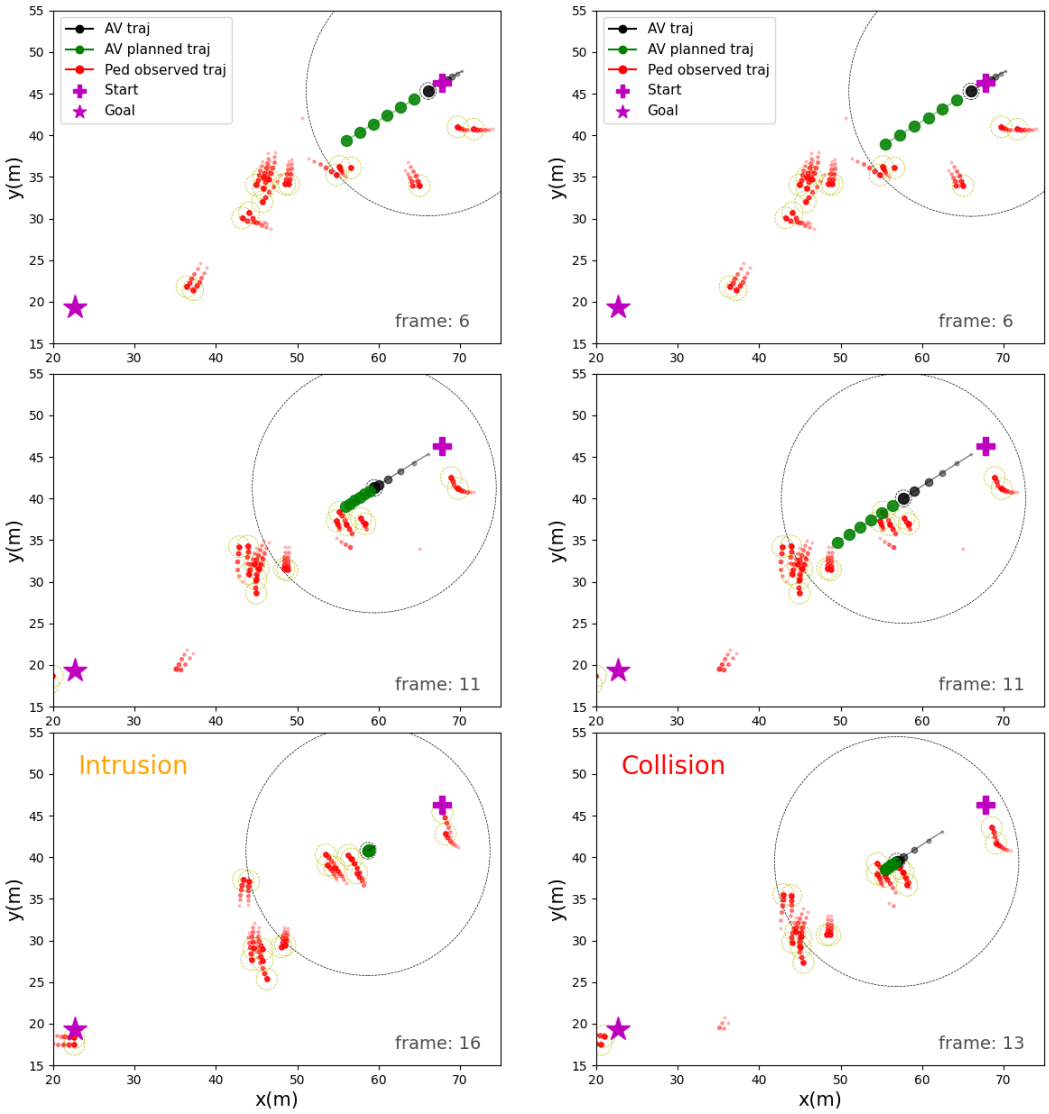

(a) DRL - No pred. SD. (Ours)

(b) DRL - No pred model

Figure 7.8: Comparison of two models trained with and without the speed-dependent (SD) danger penalty on an identical test scenario. The figure illustrates an autonomous vehicle (in black) navigating through a crowd of pedestrians (in red) towards a goal indicated by a star. It depicts the trajectories of both the AV and the pedestrians for the past 6 frames. A dashed yellow line encircles the current position of each pedestrian, representing their personal space.

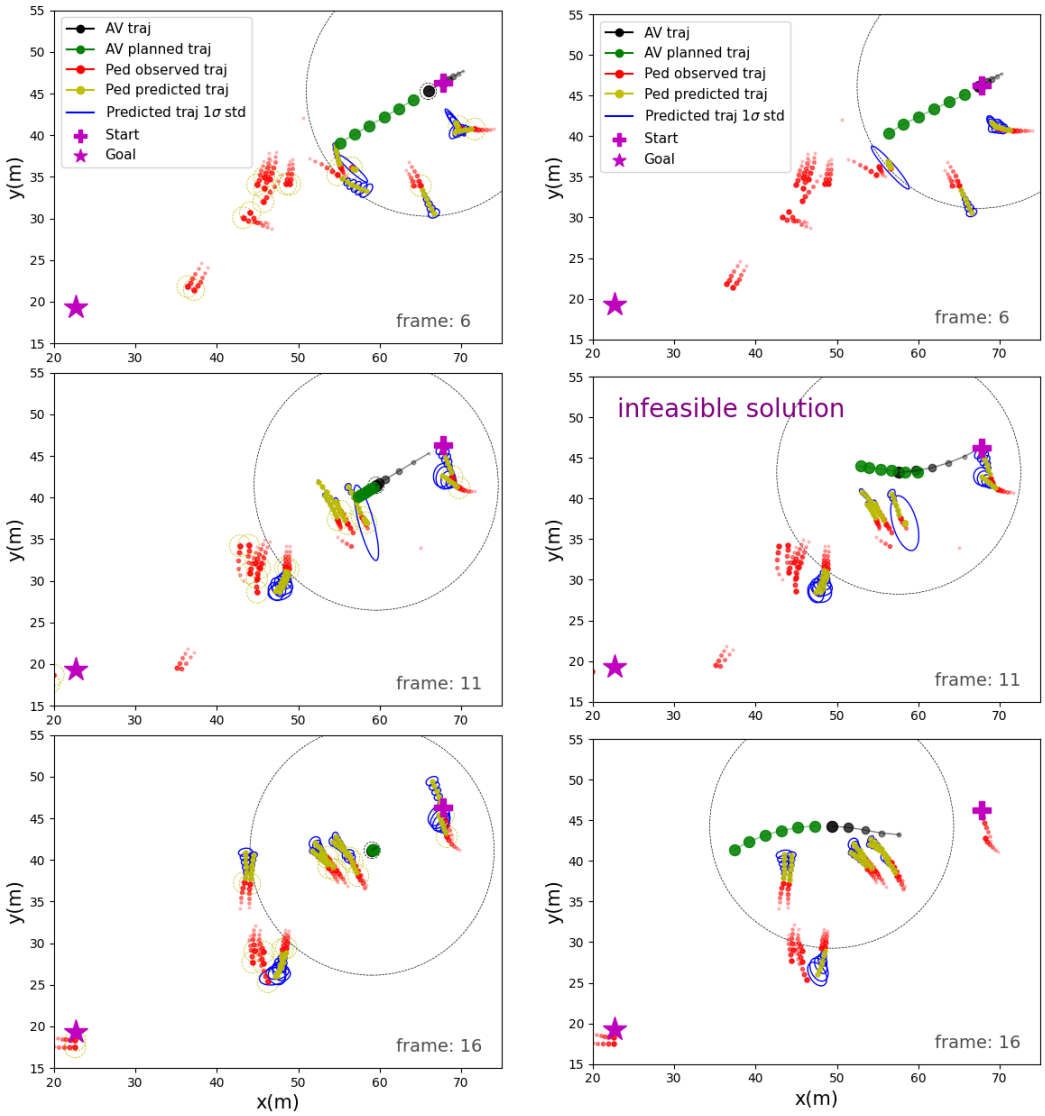

(a) DRL. UAW-PCG pred model (Ours)

(b) MPC. col. hard. model

Figure 7.9: Comparison of the performance of our DRL-based uncertainty-aware PCG model (UAW-PCG) with the uncertainty-aware MPC model on the same scenario as in Fig. 7.8. The mean trajectory and 1σ standard deviation of the pedestrians' future paths over the next 6 frames are depicted by yellow dots and blue ellipses, respectively. The AV's maximum sensor detection range is shown as a dashed black circle around its current position, with predictions being made only for pedestrians within this range.

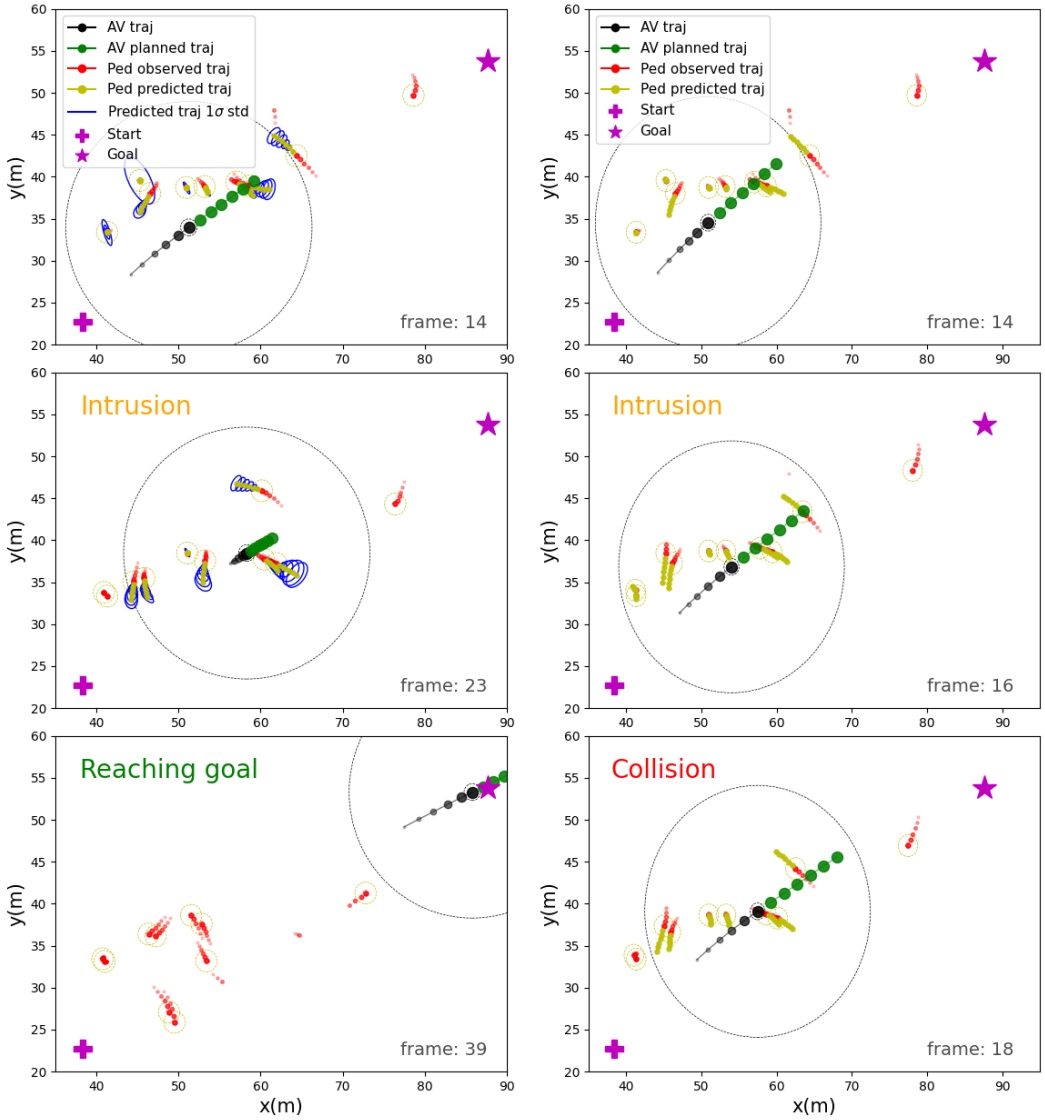

(a) DRL - UAW-PCG pred model (Ours) (b) DRL - PCG pred model

Figure 7.10: Comparison of the autonomous vehicle's trajectory behaviour in an identical test scenario. This figure presents the trajectory of the AV using our proposed Deep Reinforcement Learning (DRL) model, which accounts for prediction uncertainties (UAW-PCG pred), against the baseline model that does not consider such uncertainties (PCG pred). The trajectories are illustrated over three consecutive frames.

pred), in another scenario from the test set. By predicting uncertainties associated with pedestrians' future positions, our model maintains a safe distance from pedestrians on its left in this scenario through a gradual deviation to the right. In contrast, the uncertainty-unaware model, relying only on the actual predicted future position, causes the AV to drive too close to the pedestrian, intrude into their personal space and finally colliding with one. In this scenario, Our model yields to the pedestrian that the baseline model has collided with. However, due to early acceleration after the stop, the AV in our model slightly intrudes into the pedestrian's personal space from behind after they have passed. Nevertheless, unlike its counterpart, our model manages to reach the goal successfully. This scenario serves as an example of our model's ability to output more cautious behaviour by considering uncertainties associated with pedestrians' future trajectories.

7.5.3 Revisiting survey scenarios

With the trained DRL model in hand, we now aim to evaluate its performance in scenarios similar to those in the online survey of Chapter 4. This evaluation will help us understand how the generated trajectory behaviours of the AV align with the perceived safe behaviours identified based on the participants' responses in the survey.

It should be emphasized that throughout our navigation algorithm design, we focused only on some of the social norms identified in the online survey of Chapter 4, such as relying more on speed changes rather than heading changes for collision avoidance, and avoiding behaviours that require pedestrians to make way for the vehicle, in compliance with the perceived safe behaviours identified in the survey. We did not design the AV to follow the right-handed rule during passing or overtaking. This selection was mainly governed by the availability of similar scenarios in the existing real-world dataset we used for training our DRL model. For example, there were not enough real-world datasets of pedestrian-vehicle interactions containing situations where passing or overtaking on the left or right would be an option. In fact, other expected behaviours identified in the survey, such as following the group of pedestrians or relying on speed change rather than heading change, would alleviate the need for frequent overtaking where right-handed rules become important. Therefore, we focused only on replicating a portion of the scenarios in the survey, including scenarios 3, 5, 6, 7, and 9 of Table 4.1, and evaluating the performance of the algorithm on them. The inclusion of other social norms and their evaluation will be left to future work. The focus of scenario 8 was on explicit intent communication, which was out of the scope of the current design.

Evaluation of similar survey scenarios in the CITR dataset

Interaction situations similar to scenario 5 of the survey in Table 4.1 can be found in the Control and Intelligent Transportation Research Lab (CITR) dataset [206]. The CITR dataset contains trajectory data of controlled interaction scenarios between a golf cart and eight pedestrians in four main scenarios: 1) Front interaction (four videos), 2) Back interaction (four videos), 3) Unidirectional lateral interaction (eight videos), and 4) Bidirectional lateral interaction (ten videos), with each video being around 20 seconds long and captured from a bird's-eye view. We did not use this dataset for the original training of our model due to the limited data points. However, we decided to evaluate our trained model on scenarios from this dataset, particularly the back interaction and lateral interaction scenarios, as they are similar to scenarios 5 and 6 of the survey (see Fig. 7.11).

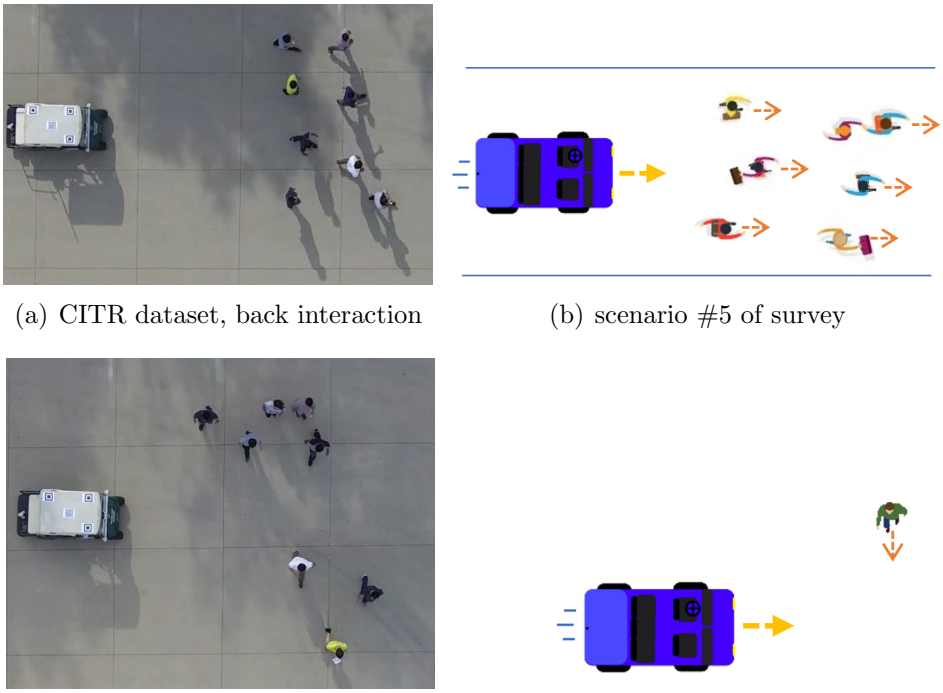

(a) CITR dataset, back interaction

(b) scenario #5 of survey

(c) CITR dataset, lateral interaction

(d) scenario #6 of survey

Figure 7.11: The back interaction scenario observed in the CITR dataset [206] resembles scenario #5 from the survey, while the lateral interaction scenario aligns with scenario #6. The screenshots from the CITR dataset were extracted from its existing video recordings [210].

Table 7.4: Performance of our DRL navigation algorithm and baseline methods on CITR dataset. Bold is the best.

Model	**SR**	**CR**	**TR**	**Nav. time (s)**	**Path length (m)**	**Intrusion ratio (%)**	**Intrusion dist. (m)**	**Intrusion speed (m/s^2)**
MPC. col. hard.	0.20	0.60	0.20	15.00$\pm$0.00	30.00$\pm$0.00	**2.86**$\pm$5.71	0.00$\pm$0.00	**0.00**$\pm$0.00
DRL. PCG pred	0.80	0.20	**0.00**	13.25$\pm$6.77	**17.82**$\pm$5.98	5.71$\pm$11.43	**0.92**$\pm$0.03	3.39$\pm$0.68
DRL. UAW-PCG pred	**1.00**	**0.00**	**0.00**	**11.50**$\pm$6.19	18.07$\pm$5.75	13.05$\pm$9.30	0.77$\pm$0.26	2.15$\pm$1.61
Human Driver	1.00	0.00	0.00	4.90$\pm$2.06	18.15$\pm$5.80	30.24$\pm$18.63	0.72$\pm$0.24	3.16$\pm$0.37

In this regard, we kept one video from each of the interaction scenarios (exceptionally two from the bidirectional lateral interaction as it contained more videos compared to other scenarios) for testing and used the remaining videos for fine-tuning our DRL model on the CITR dataset for 3×10^6 time steps. We also used the same training set to fine-tune the prediction model on the CITR dataset.

We performed the same fine-tuning process for the DRL and PCG baseline models. Table 7.4 shows the quantitative results of our model's performance compared to two of the main baselines we had before on the five scenarios of the test set.

Fig. 7.12 qualitatively demonstrates our algorithm's performance on the back interaction scenario, which replicates scenario 5 from the survey. Fig. 7.13 illustrates the performance on a lateral scenario similar to scenario 6 of the survey. Compared to the MPC model, our model exhibits trajectory behaviours more aligned with preferred behaviours identified in the survey, such as following the flow of pedestrians from behind in scenario 5 (Fig. 7.12) and showing a tendency towards speed adjustment for collision avoidance rather than heading changes in scenario 6 (Fig. 7.13). In Fig. 7.13, our model decelerates to avoid collisions with laterally approaching pedestrians, while the MPC model (also designed to minimize heading changes) moves to the right to avoid pedestrians while accelerating towards its goal.

Building custom scenarios using the ORCA model

While we were able to evaluate our algorithm's performance on two scenarios from the CITR dataset that resemble those in our online survey, not all scenarios from the survey had similar counterparts in the CITR or other datasets (e.g., scenarios 3 and 7, Table 4.1). To simulate these custom scenarios and assess our algorithm's performance in interactive environments, we opted to use the ORCA model for pedestrian behaviour simulation, as real-world pedestrian trajectory data was unavailable for these scenarios. Furthermore,

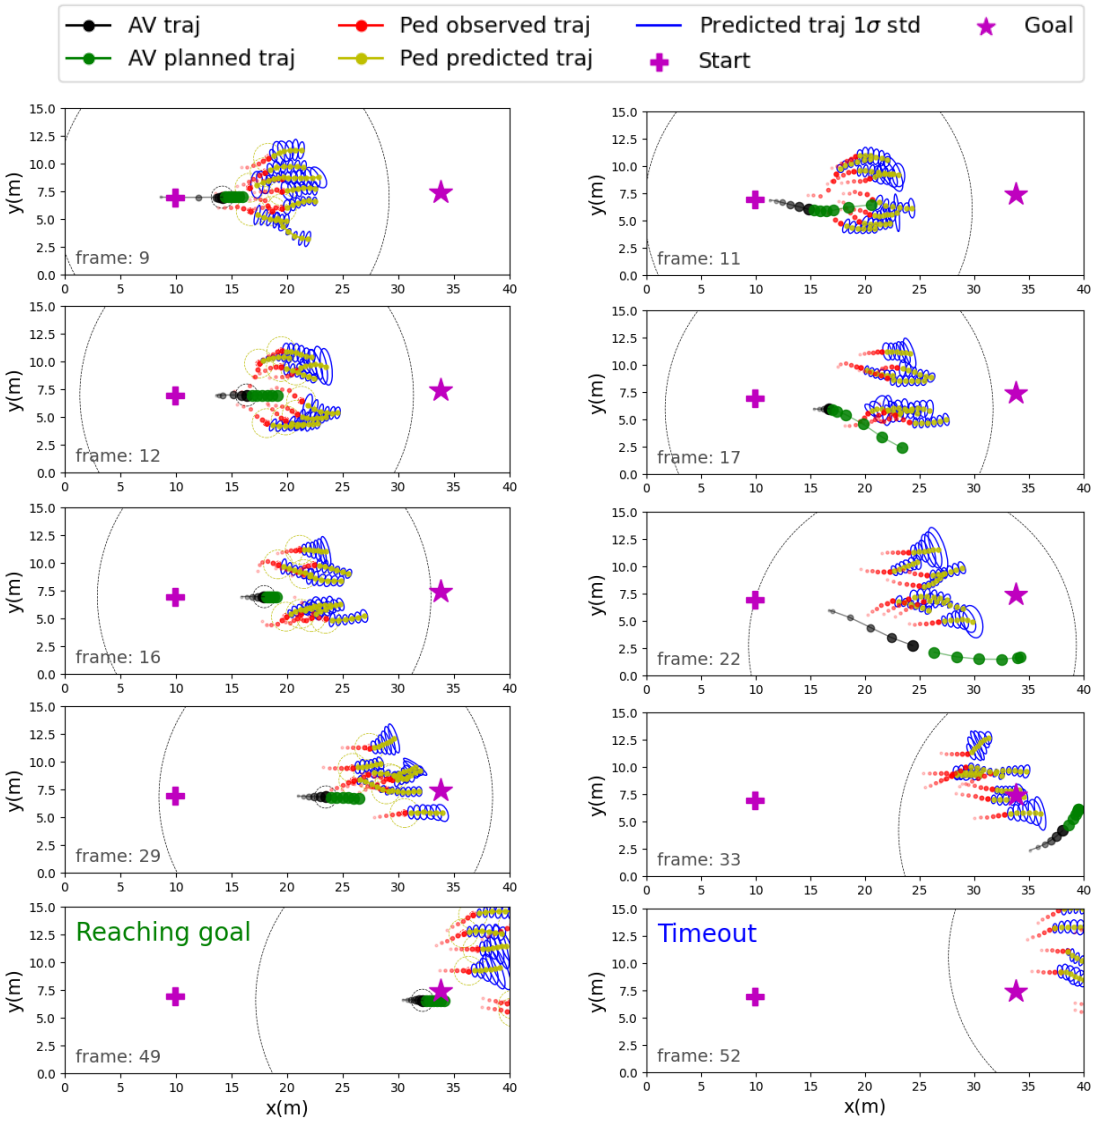

(a) DRL. UAW-PCG pred

(b) MPC - col. hard.

Figure 7.12: Performance of our DRL model compared to the MPC model on a scenario from the CITR dataset, similar to the back approach scenario #5 of the online survey (Table 4.1). While the MPC model attempts to reach the goal by overtaking the crowd of pedestrians through heading changes, our model follows the pedestrian flow, aligning with the preferred behaviour indicated by the survey results. Despite its overtaking attempts, the MPC model could not reach the goal, resulting in a timeout. In contrast, our model successfully reached the goal in a perceived safe manner and within a reasonable time.

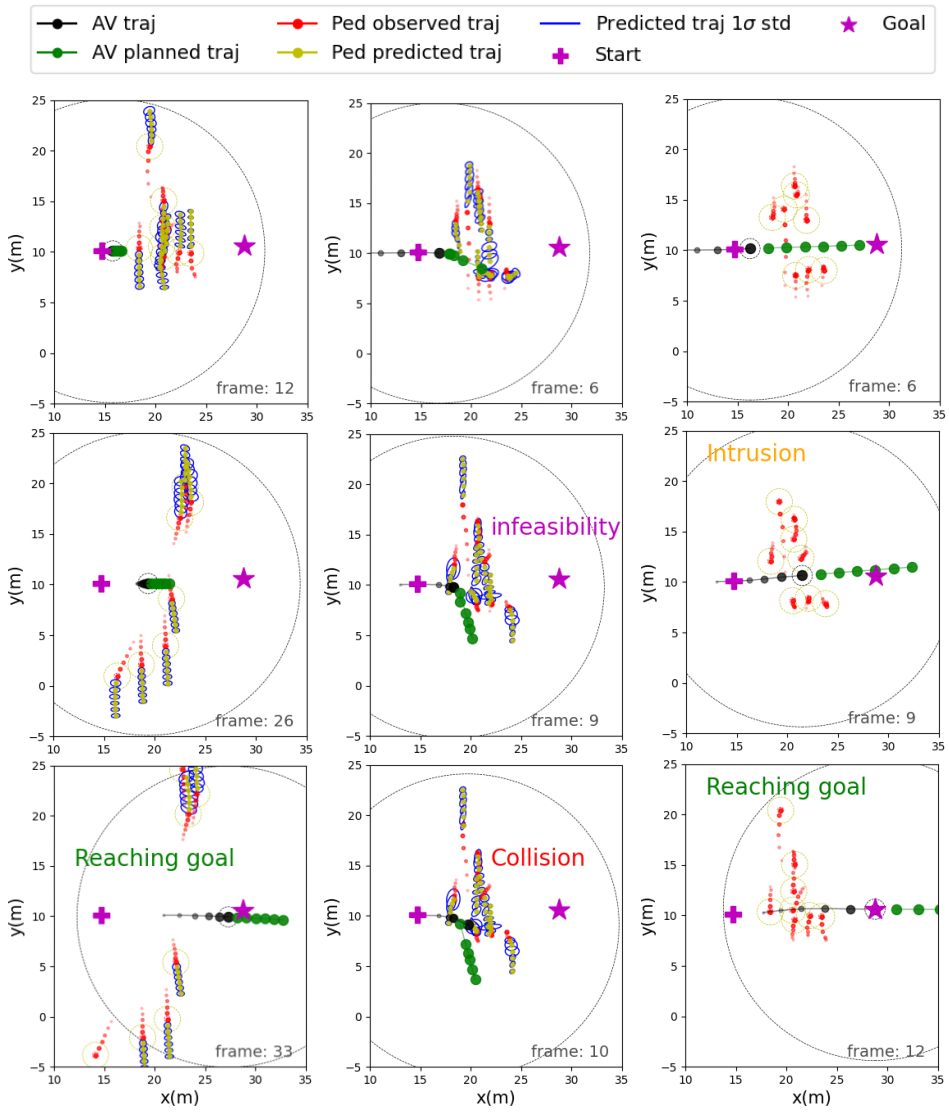

(a) DRL. UAW-PCG pred (b) MPC - col. hard. (c) Human Driver

Figure 7.13: Performance of our DRL model compared to the MPC model and the human-driven trajectory on a lateral scenario from the CITR dataset, similar to the 90-degree approach scenario #6 of the online survey (Table 4.1). The collision avoidance strategy of the MPC involves more heading changes, while our DRL method adjusts speed, which is perceived safer according to the online survey results. This approach is also more in line with the human-driven trajectory but more cautious.

our survey results provided insights into how pedestrians would react to AVs in specific scenarios (e.g., scenario 3). By implementing the AV as an agent within the ORCA framework, with the ability to switch its visibility to other pedestrians on and off, we can test our vehicle's decision-making algorithms against various pedestrian behaviours identified in the survey, such as creating space for the AV versus continuing along their intended path in scenario 3.

Figures 7.14 - 7.16, qualitatively demonstrate the performance of our DRL navigation framework in custom-created scenarios resembling scenarios 3, 5, 6, 7, and 9, utilizing the ORCA behaviour model for pedestrians.

Figure 7.14 replicates scenario 3 from the survey and is shown in two different cases. In the first case, the vehicle is made invisible to pedestrians, simulating option (a) from Table 4.1 where pedestrians "keep walking straight until a place where the cart can overtake," thus not reacting to the vehicle (Fig. 7.14(b)). In this scenario, our decision-making algorithm successfully follows the pedestrian without causing harm, aligning with the expected behaviour of the AV as perceived by survey participants in the role of AV passengers.

In the second case, the vehicle remains visible to pedestrians in the ORCA model, prompting pedestrians to react by moving away from the vehicle's path, similar to behaviour option (b) in the online survey where pedestrians "clear cart's path by moving to the right" (Fig. 7.14(a)). This option was the most selected by survey participants, acting as pedestrians. Here, the AV reacts to the cleared path and overtakes the pedestrian while progressing towards its goal. It's important to note that the ORCA simulation environment does not model static obstacles. Therefore, we could not simulate the red cones as visualized in scenario 3 of the survey. Nonetheless, we successfully simulated two potential behavioural responses in this scenario and qualitatively demonstrated how our AV adapts to different pedestrian behaviours—either not opening space or opening space for the AV—based on our trained DRL navigation policy.

For the custom simulation of scenario 7 from the survey, our results visualized in Fig. 7.15(b) demonstrate that the behaviour of the AV, guided by our decision-making algorithm, aligns with the most preferred option of "following group A" (Table 4.1) according to the survey results.

Similar observations apply to scenario 5, where the AV follows the pedestrian flow without requesting a passage (see Fig. 7.15(a)), and in scenario 6, where our navigation policy leans towards "yielding to the pedestrian by slowing down," the preferred behaviour identified by survey participants, rather than adjusting its path left or right (see. Fig. 7.16(a)).

Scenario 9, replicated in Fig. 7.16(b), showcases the capability of our uncertainty-aware trained model in managing situations involving erratic and unpredictable child behaviour,

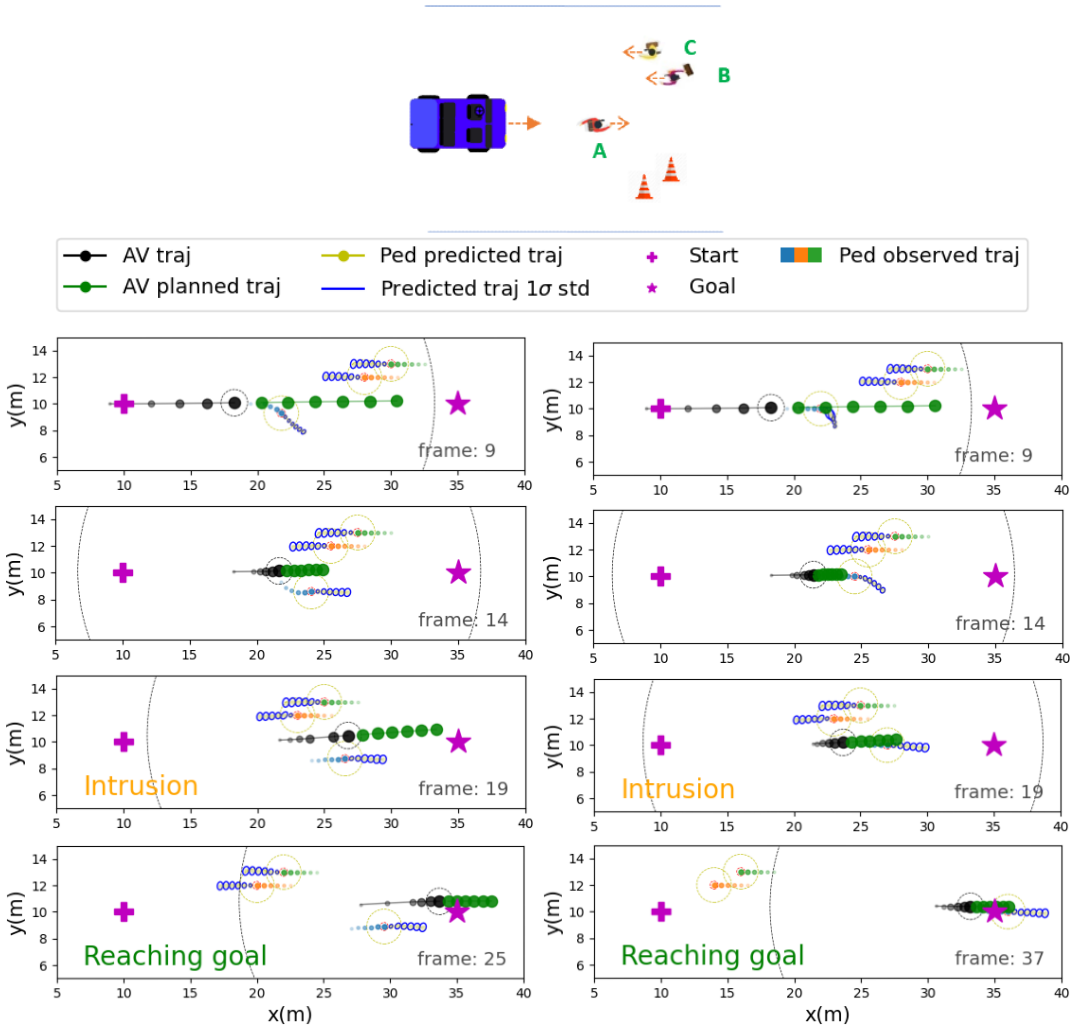

(a) Survey - scenario 3, visible AV

(b) Survey - scenario 3, invisible AV

Figure 7.14: Performance of our DRL model in a custom scenario replicating Scenario #3 of the survey, using the ORCA model for pedestrian behaviour under two conditions. (a) When the robot is visible to pedestrians, resembling the case where pedestrians react to the vehicle by opening up a space for it, which was the most frequently selected pedestrian behaviour in the survey. In this case, the AV understands and overtakes the pedestrian. (b) When the robot is invisible to pedestrians, resembling the case where pedestrians do not react to the AV and follow their own paths. In this scenario, the AV will follow pedestrian A, complying with the expected behaviour of the AV from the survey. This demonstrates the AV's ability to adapt to different pedestrian behaviours.

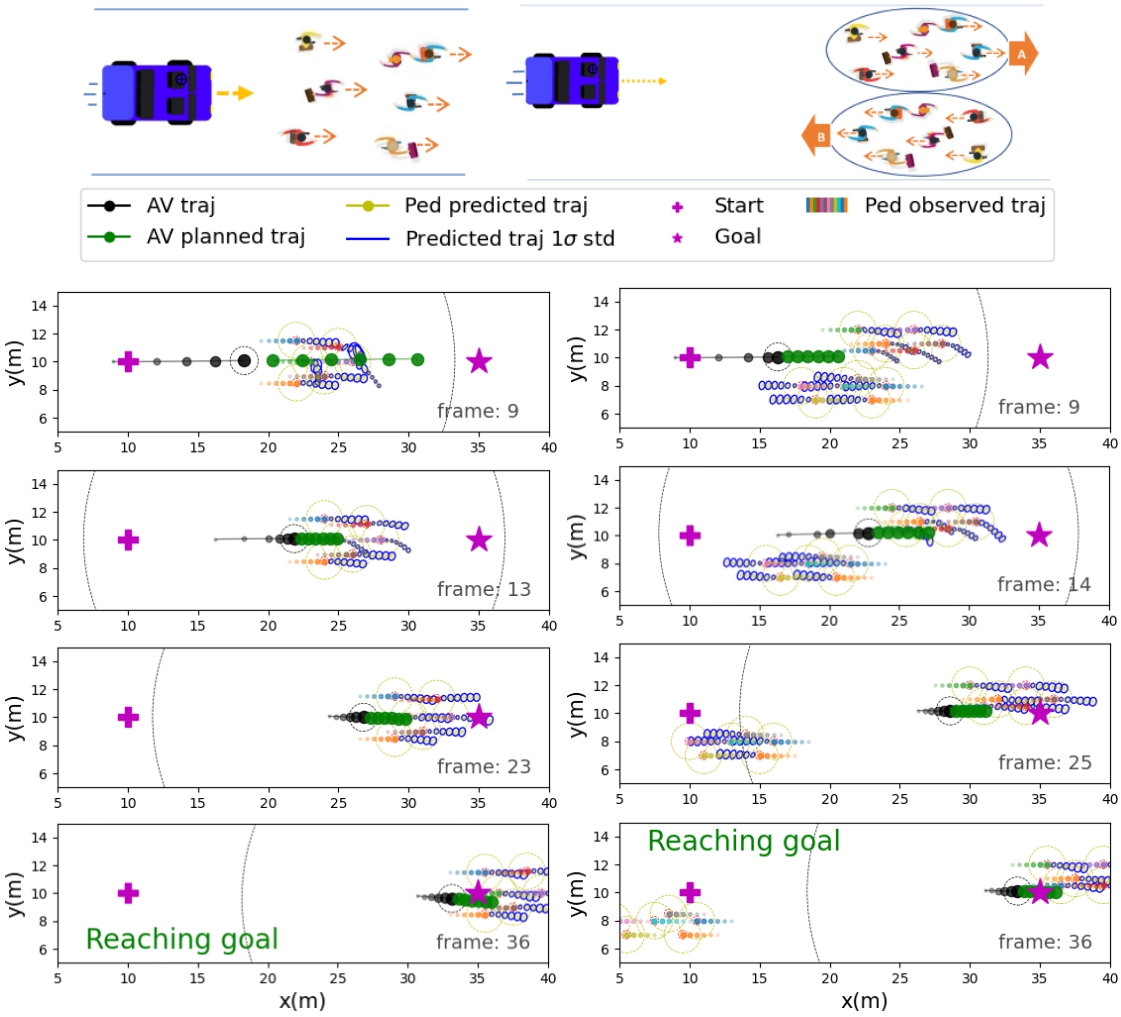

(a) Survey - scenario 5

(b) Survey - scenario 7

Figure 7.15: Performance of our DRL model in a custom scenario replicating Scenarios 5 and 7 of the survey, using the ORCA model for pedestrian behaviour. In both cases, the AV follows the flow of pedestrians to reach its goal, as identified from the survey results.

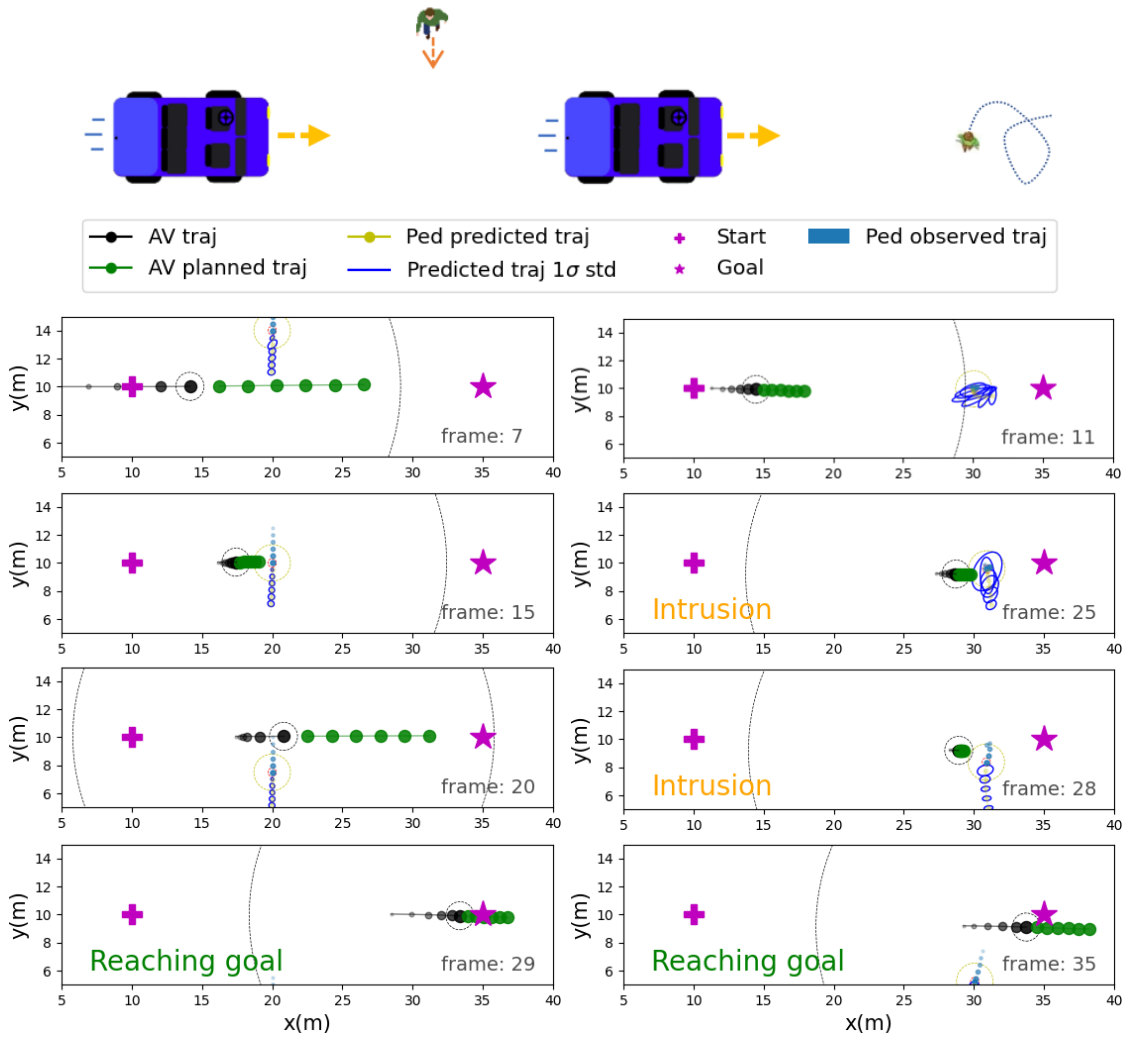

(a) Survey - scenario 6

(b) Survey - scenario 9

Figure 7.16: Performance of our DRL model in custom-built scenarios replicating Scenarios 6 and 9 from the survey, using the ORCA model for pedestrian behaviour. In Scenario 9, random goals are assigned to the pedestrian at each time step to mimic the erratic behaviour of a child who eventually moves to one side. In Scenario 6, the AV adjusts its speed to avoid collisions, which is perceived as safer behaviour according to the survey results compared to changing its heading. When confronted with a child with high trajectory uncertainty, the AV decelerates and almost comes to a stop, resuming movement once the child reaches a safe distance, as expected based on the survey results.

which can commonly occur in unstructured environments. We simulated this scenario by randomly assigning a nearby goal position for a few time steps before eventually assigning a goal further away, simulating the child gradually moving out of the AV's path. The AV's generated trajectory in this scenario shows it slowing down and almost stopping near the child until the child moves away, which aligns with the expected behaviour identified in the survey (Fig. 7.16(b)). Similar to scenario 6, the AV adjusts its speed rather than changing its heading, consistent with the perceived safe behaviour from the survey results.

While our DRL decision-making policy was not specifically trained for scenarios involving erratic pedestrian behaviour, the results in Fig. 7.15(b) qualitatively demonstrate our algorithm's potential. Despite occasional intrusions in Frames 25 and 28 (shown in Fig. 7.15(b)), the algorithm shows a collision-free behaviour when handling scenarios involving children and the high uncertainty associated with their behaviour.

7.5.4 Limitations

Our proposed method and simulation environment, while effective, have limitations. Due to the low operational speed of vehicles in the crowded unstructured environments that we targeted, we used a unicycle kinematic model for the vehicle in the simulation environment, which is an improvement over the holonomic kinematic model often used in the literature for crowd navigation algorithms [123, 122, 30, 31]. However, future work can explore the effect of more sophisticated kinematic models that better resemble real-world vehicles. Additionally, we modelled the AV as a circle in the simulation to facilitate distance calculations, potentially resulting in a larger than necessary distance during side passes due to the vehicle's rectangular shape. This could be addressed in the future by incorporating an elliptical model for the AV's shape.

Our model was exclusively tested on the HBS dataset, and the limited number of scenarios from the CITR dataset. This was primarily due to the limited availability of datasets depicting pedestrian and vehicle trajectories closely interacting within unstructured or shared spaces. Nevertheless, we partially demonstrated the generalizability of our algorithm by showcasing its performance on some custom-built scenarios based on ORCA that our trained model had not encountered during training.

Our work is limited in its direct application to environments with numerous static obstacles, as we mainly focused on interactions with other decision-making agents in open, dynamic spaces and did not explicitly incorporate static obstacles such as walls and physical structures into our algorithm. Our motion planner concentrated exclusively on local motion planning and managing pedestrian interactions within dynamic environments while leaving

the consideration of static obstacles to the global motion planner. However, future research could integrate environmental maps into both prediction and planning modules to better handle static obstacles directly.

Lastly, the simulation environment we proposed replays realistic human trajectories but lacks responsiveness to various AV behaviours as humans would react in the real world. The scenarios extracted from the HBS dataset, used to build our simulation, already include vehicles. Therefore, the replayed pedestrian trajectory data inherently reflects some reactive behaviours to the human-driven vehicles in the data, such as yielding. However, these reactions may not remain realistic, particularly if the behaviour of the AV driven by the DRL algorithm in the simulation deviates from the trajectory of the human driver in the original data. Nonetheless, the lack of pedestrian reactions to every possible AV behaviour has an advantage: it encourages the trained policy to minimize reliance on pedestrian cooperation for collision avoidance. This reduces the need for pedestrians to adjust their paths in real-world scenarios. In fact, training a DRL algorithm in a simulation environment with reactive agents can lead to the agent learning aggressive behaviours—forcing others to clear its path as it recognizes they will react to its actions [123].

Our approach avoids aggressive behaviours by not relying on pedestrian cooperation or reactions. Notably, our trajectory predictor, which takes the planned trajectory of the AV as input, can generate predictions that inherently account for pedestrian cooperation. However, it can also revise these predictions in subsequent time steps if cooperation is not observed, allowing the planner to adapt accordingly. As a result, our planner is trained to operate independently of pedestrian cooperation while remaining capable of recognizing and leveraging instances when pedestrians make space for it. This independence ensures robust performance and can enable even better results when pedestrian cooperation occurs.

7.6 Summary and conclusion

In this chapter, we developed a DRL-based decision-making algorithm to serve as the motion planner within our navigation framework, which incorporates our previously developed trajectory prediction algorithm. We integrated prediction uncertainties into the policy learning of the model-free DRL algorithm. The future positions of nearby pedestrians, along with their covariance matrix (representing the uncertainty in the predicted trajectories), were input into the DRL motion planner, which was trained with a novel reward function. This reward function was designed to promote perceived safe behaviours by encouraging the vehicle to maintain a safe distance from pedestrians' current positions, minimize the collision probability with their predicted future positions over a prediction

horizon, and consider pedestrians' comfort during close approaches. To simulate realistic pedestrian behaviour in the presence of vehicles, we trained our model in a simulation environment derived from real-world pedestrian trajectories in a shared space.

In addressing **RQ3.3**, which focuses on incorporating interaction effects into the motion planner design, we considered both sides of the mutual interaction effects through a novel integration of our trajectory predictor and motion planner modules. On one hand, we accounted for the influence of pedestrians on the AV motion planning by incorporating both their current and future trajectories into the AV's DRL motion planner, which is modelled as a spatio-temporal graph. On the other hand, we considered the effect of the AV's presence on pedestrians' predicted future trajectories—an aspect often overlooked in similar crowd navigation methods—by inputting a projection of the most recent planned vehicle trajectory into the pedestrian trajectory prediction module. This overall approach demonstrated its effectiveness by enabling the algorithm to proactively avoid predicted pedestrian positions in the near future, as supported by the results of the ablation studies.

Moreover, compared to MPC models, which serve as an example of optimal control methods for crowd navigation, we showed that our model exhibits faster computational times, superior performance, and trajectories closely resembling those of human drivers in similar scenarios.

RQ4.2 which examines the effect of other agent's behavioural uncertainty on the AV's decision-making, was addressed by incorporating the covariance matrix of pedestrians' predicted trajectories into the DRL state space. A probabilistic penalty was also added to the reward function, based on the collision probability with future predicted positions of pedestrians, considering their uncertainty. The effectiveness of this approach was validated through the ablation study, which showed a lower collision rate, reduced timeout rate, and increased minimum distance to pedestrians for our uncertainty-aware DRL motion planner.

The design of the DRL's reward function also incorporated insights from our online survey, particularly by penalizing heading changes to encourage the trained policy to rely more on speed adjustments, along with other social considerations such as respecting personal space and reducing speed during close passes, addressing **RQ1.4**.

Furthermore, in addressing **RQ1.4**,we revisited the survey scenarios by replicating them in simulation, demonstrating that the behaviour generated by our trained model aligns with the perceived safe behaviours identified in the survey, such as relying on speed adjustments rather than heading changes for collision avoidance. This step linked our designed process back to the initial survey conducted in this thesis, showcasing the designed algorithm's performance relative to the survey findings.

Chapter 8

Conclusions

8.1 Summary of research

This thesis presented a human-aware integrated prediction and planning framework for the autonomous navigation of a low-speed vehicle in an unstructured, pedestrian-rich environment. The applications range from mobility aids in airport terminals to AV operations in highly dynamic and interactive shared spaces. The first objective of this thesis was to better understand the design requirements for this less-explored application by studying perceived safe behaviours for AVs in shared and unstructured spaces and exploring the connection between pedestrian trajectory behaviour and the structured or unstructured nature of the environment. Proactive navigation, which entails more natural and farsighted behaviour, requires understanding how pedestrian trajectories will evolve. Therefore, the second goal was to develop a model for predicting pedestrian trajectories, considering the effects of interactions with both vehicles and other pedestrians in unstructured and shared environments. Finally, the third objective was to develop a computationally efficient motion planner that integrates this predictive information into the decision-making process, accounting for interaction effects, the identified design requirements, and uncertainties in the predicted pedestrian trajectories.

Chapter 4 presented our investigation into the preferred and perceived safe behaviours, in the form of social norms, for an AV to comply with when operating in an unstructured, pedestrian-dominated environment. This study addressed the first research question of this thesis (**RQ1**) and was conducted through an online survey with 116 participants. Through this survey, we explored people's expectations for the vehicle's navigation behaviour in close proximity to pedestrians in a few sketched interactive scenarios, with participants taking

either the perspective of a pedestrian near the AV or a passenger in the AV. This helped us take a human-centred approach to the design of the navigation algorithm. This first step of studying the design requirements was particularly important given that this application is less-explored in the literature. Most previous work has focused either on the AV's interaction with pedestrians in the strict structure of the road or on small mobile robots operating among pedestrians without any safety criticality due to their small size. Key findings from the survey, in response to **RQ1**, revealed that the AV's speed adjustment for collision avoidance was perceived safer than heading changes. Additionally, that the vehicle was expected to follow the flow of pedestrians rather than forcing them to make space when there was not enough room to overtake. These design requirements were later encoded in the reward function construction of the DRL motion planner.

In Chapter 5, we investigated the possible connections between pedestrian and vehicle trajectory features and the nature of the environment being structured or unstructured. This study aimed to identify key differences in trajectory features and behaviour patterns that are important to consider and model in the navigation algorithm as outlined in **RQ2**. Additionally, we aimed to determine which dataset best represents behaviours in an unstructured environment for training our data-driven prediction model. This initiative, not only provided a better understanding of the target environment for our design, but also prompted us to take steps toward quantifying the definition of an unstructured environment, which was missing in the literature. In this context, we extracted thirteen trajectory features, such as mean speed, stop fraction, and variability of pedestrians' trajectories, from available datasets including both pedestrians and vehicles. We then clustered the dataset scenes into two groups based on these features, with each cluster representing structured vs. unstructured spaces. The cluster that included more road environment scenes was considered to be the structured environment. Our results showed that features such as variability of pedestrian trajectories, stop fraction, path efficiency, and the density of standing pedestrians were significantly different between the two clusters. These features can be used to quantify unstructured environments or classify existing datasets. Key findings included a higher variability in pedestrian trajectories in unstructured environments, and common stop-and-go behaviours indicated by a higher stop fraction. These findings, besides the higher density of standing pedestrians that could suddenly start walking, highlighted the high uncertainty in the predicted trajectory of pedestrians in these environments, encouraging us to consider prediction uncertainty in the decision-making later on.

Chapter 6 focused on developing our deep learning-based pedestrian trajectory prediction model, addressing the overlooked problem of effectively identifying interacting agents in existing data-driven pedestrian trajectory models and contributing to the answer for **RQ3** of this thesis. For this purpose, we proposed using the TTC as a heuristic for distin-

guishing interacting agents, as opposed to the literature's focus on distance. Solely focusing on distance could easily lead to overlooking agents that are farther away but headed toward the target agent, which can significantly impact the target agent's future trajectory. Through our experimental results, we demonstrated how this simple yet effective heuristic-based method can filter out neighbouring agents with non-critical TTC, thereby improving prediction accuracy even for the existing prediction models without changing the structure of the rest of the model. This improvement was evidenced by a 3.5% decrease in ADE for the Social LSTM model when filtering interacting agents based on our proposed criterion. The second main problem addressed in this chapter, in line with **RQ3**, was the construction of an effective interaction feature for pedestrian trajectory prediction models in unstructured environments. We proposed our polar collision grid model, which accounts for both TTC and approach angle information for constructing the interaction feature separately for pedestrians and vehicles interacting with the target pedestrian. We showed that this novel interaction feature reduces prediction errors by 4.5% in ADE and 4.3% in FDE compared to baseline methods.

In Chapter 7, we developed our DRL-based motion planner, which integrates with our pedestrian trajectory predictor through a novel flow of information between them. In this setup, the planner considers the prediction outputs, taking into account the effect of pedestrians on the AV's decisions. Simultaneously, the prediction model uses the most recent planned trajectory of the AV to account for the AV's presence on pedestrians' future trajectories, thereby capturing the mutual interaction effect. This approach addresses **RQ3** at both the prediction and planning layers. This approach is novel because, in the state-of-the-art models proposed for crowd navigation, pedestrian trajectory predictions are usually generated without considering the robot's presence. The DRL motion planner learns interaction effects through its ST-graph architecture, trained in a realistic simulation environment we developed using real-world pedestrian trajectory data. One of the main problems we addressed in the motion planner design was the consideration of prediction uncertainty in the DRL's decision-making process, addressing **RQ4**. This aspect has been overlooked in DRL-based crowd navigation methods in the literature, but is important due to the higher presence of standing pedestrians, high stop fraction, and trajectory variability in unstructured environments, as discovered in Chapter 5. To address this gap and fulfill **RQ4** of this thesis, we included prediction covariance matrices as part of the state definition of the DRL, which gets involved in interaction modelling. We also designed a novel reward function that encourages the AV to minimize its collision probability with the predicted future positions of pedestrians over a prediction horizon, promoting farsighted behaviours. Within the reward function design, we also emphasized respecting pedestrian personal space and reducing speed during close passes as social-aware aspects of navigation. The

aspect of speed in close approaches has often been neglected in this area of research, but is important for AV operation given its higher speed and larger size, which we have encoded in our proposed algorithm. Furthermore, we encouraged minimal heading changes in line with the perceived safe behaviours identified in the survey results of Chapter 4.

Our overall designed uncertainty-aware DRL navigation framework showed superior performance, with a 40% decrease in collision rate and a 15% increase in minimum distance to pedestrians, reflecting the social considerations of the navigation algorithm, compared to the state-of-the-art model that does not account for prediction uncertainty. Additionally, our approach outperforms model predictive control methods that incorporate the same prediction uncertainties in terms of both performance and computational time while producing trajectories closer to those of human drivers in similar scenarios. The evaluation of the algorithms on scenarios similar to the survey scenarios in Chapter 4 revealed compliance with the expected and perceived safe behaviours identified from the survey results.

All in all, our developed human-aware navigation method, along with our contributions, addresses the key challenges of navigating among pedestrians in a dynamic, uncertain, and highly interactive unstructured environment. A list of our contributions is provided in the next section, followed by a discussion of the limitations of our work and the remaining aspects that need to be addressed in future research.

8.2 Contributions to original knowledge

In this thesis, and throughout the development of our autonomous navigation algorithm, we made the following contributions to the field, addressing the main research questions listed in Chapter 1.

1. The research focused on the novel application of autonomous vehicles (AV) navigating in pedestrian-rich, unstructured and shared spaces, addressing the complexities and safety concerns of such environments.

2. The thesis identified social norms for AV compliance in unstructured environments through an online survey with recruited participants (Addressing **RQ1**).

3. This work proposed quantitative trajectory features of pedestrians and vehicles that differ between structured and unstructured environments, emphasizing the importance of considering uncertainty in pedestrian behaviour for navigation algorithm design specially for operation in unstructured environments (Addressing **RQ2**). These

features can classify datasets with both pedestrian and vehicle trajectories based on the nature of the environment.

4. The research designed a novel data-driven pedestrian trajectory predictor model for shared spaces, featuring:
   - An effective interacting agent selection process based on Time-to-Collision (TTC) heuristics (Addressing **RQ3**).
   - A novel feature for encoding interaction effects based on approaching angle and TTC information (Addressing **RQ3**).
   - A novel loss function to improve the accuracy of forecasting the covariance of future positions, mitigating overconfident predictions (Addressing **RQ4**).
5. The thesis proposed a novel uncertainty-aware Deep Reinforcement Learning (DRL) motion planner with the following features:
   - Integration of the DRL motion planner with the pedestrian trajectory prediction model to capture mutual interaction effects (Addressing **RQ3**).
   - Utilization of prediction uncertainty in the DRL-based motion planner to minimize collision probability over a prediction horizon (Addressing **RQ4**).
   - Introduction of a novel reward function prioritizing pedestrian comfort by encouraging the AV to reduce speed during close approaches and comply with discovered social norms (Addressing **RQ1**).
6. This work constructed a novel simulation environment (gym-environment) with realistic pedestrian trajectory behaviour in the presence of vehicles and other pedestrians in shared spaces, and trained AV behaviour in this realistic simulation environment.
7. The studies demonstrated the effectiveness of the overall framework in avoiding collisions, respecting pedestrian personal space, and complying with expected behaviours identified in the online survey (Addressing **RQ1**).

8.3 Limitations and future work

In this thesis, we presented a human-aware navigation algorithm for operating among pedestrians in an unstructured environment and evaluated its performance in simulation. However, our proposed algorithm and the research discussed in this thesis have certain

limitations, some of which were addressed in the relevant chapters. This section summarizes key general limitations and aspects that were not covered in this thesis but have the potential to be explored in future research. These directions for further expanding this line of work, along with their relevance to the limitations of this thesis, are outlined below.

1. **Real-world field experiment in an unstructured crowded environment:** The algorithm developed in this thesis was evaluated in a simulation environment. Simulations are useful for accelerating testing, and we made efforts to ensure the simulation environments closely resembled real-world interaction scenarios. For this purpose we build our simulation based on dataset from real-world scenarios and pedestrian trajectory behaviours. However, discrepancies between simulation and real-world environments can still exist. These differences limit the direct and safe transfer of our algorithm to real-world platforms. Therefore, further experiments are required to assess the algorithm's performance in real-world settings. One challenge in transferring the learned policy to an actual golf cart is the inaccuracies that can arise from integrating the algorithm with other modules of the autonomy navigation stack. Since this work focused on the planner module of a navigation pipeline, we assumed access to accurate positional data of pedestrians, which, in real experiments, must come from the perception module that could contain errors. The propagation of these errors and how our uncertainty-aware navigation algorithm handles them should be further investigated in real-world experiments in future work.

   In-person experiments or field studies also allow for a more accurate evaluation of people's perceptions of the vehicle's behaviour by observing it in action. In this thesis, we primarily relied on simple metrics to assess the social aspects of vehicle behaviour in simulation, such as minimum distance to pedestrians or speed at the minimum distance as measures of respecting pedestrians' personal space, and the visual evaluation of the compliance of the vehicle's behaviour with people's expectations derived from an online survey. However, true perceived safety and acceptability can only be effectively assessed through real-world interactions with the physical vehicle in user studies or field experiments, which should be addressed in future work.

2. **Realistic and reactive date-driven pedestrian simulators:** The challenge of how well a trained DRL policy will perform after transferring to the real world is highly dependent on how well the simulation environment the DRL was trained in represents the real-world environment. Being cautious about this aspect from the pedestrian trajectory behaviour perspective, we went beyond the commonly used ORCA simulation environment (although we later tested our algorithm's performance in ORCA as well) and trained our model in our proposed simulation environment.

Our simulation contained realistic scenarios of pedestrian-vehicle interaction using pedestrian trajectories from a real dataset. However, a limitation of our simulation environment is that if the DRL-generated AV behaviour in the simulation deviates from the human-driven vehicle trajectory in that scenario, the pedestrian's reaction will not be realistic, as the replayed pedestrian trajectory does not react online to the AV's behaviour.

Considering this limitation, a promising future direction for research in this area is the development of reactive and realistic pedestrian simulation environments in shared spaces that could be used as a benchmark for AV crowd navigation algorithms. One potential solution is to create data-driven simulation environments using prediction models in a closed-loop format. However, this would require addressing the challenge of prediction error growth over time and access to more data. As more data from shared spaces becomes available, creating such simulation environments that represent realistic pedestrian behaviour in the presence of vehicles in an unstructured environment will become possible.

3. **Data collection from shared unstructured environments:** The focus of this work was to develop a navigation algorithm for the operation of an AV in an unstructured pedestrian-rich environment, considering the complexity of interactions happening in these dynamic spaces. Therefore, for the purpose of training and testing our data-driven trajectory prediction model and also providing realistic scenarios for training the DRL in simulation, we relied on datasets that had 1) representative behaviours in unstructured environments, 2) contained both pedestrians and vehicles and 3) contained close interactions between these two agent types. However, only a few datasets exist with such characteristics, and in some of them, either the number of data points and trajectories are limited, or few interactions happen between pedestrians and vehicles. Thus, we only trained and tested our algorithm on one dataset (HBS) which may present a limitation to the trained model's generalizability to other unstructured environments. Nevertheless, this limitation could be mitigated by providing more data to the model, as the overall design of the algorithm has proven effective on a highly interactive dataset. While the HBS dataset contains various complex interactive scenarios between pedestrians and vehicles in a shared space, the reliable real-world implementation of these learning-based algorithms will requires further training and testing across a wider range of scenarios. This can be achieved once more datasets from such environments become available. Therefore, the community in this area is encouraged to collect more data from low-speed vehicles in pedestrian-rich unstructured environments, such as vehicles operating in airport terminals or off-road campus environments, before making such learning-based algo-

rithms available for real practice.

4. **Rich features for improved pedestrian trajectory prediction:** Our model focused solely on using trajectory coordinates to predict pedestrian movements. This approach allowed us to utilize existing trajectory datasets. However, it is important to acknowledge that human drivers—whom AV navigation algorithms are modelled after—rely on more than just trajectory data when making predictions. Human drivers also consider body gestures, gaze, and contextual cues to anticipate pedestrians' intentions and future trajectories. Although this is a research area actively explored in road-crossing intention prediction literature [170], more work could be done on incorporating these richer behavioural features into pedestrian trajectory prediction models, especially in interactions beyond road-crossing scenarios and within more complex, unstructured environments. In unstructured environments, gestures, gaze, and emotional states can provide valuable information, such as identifying distracted pedestrians or predicting a pedestrian's willingness to cooperate. These factors could reduce uncertainty in predictions and result in more accurate forecasts, ultimately improving the decision-making algorithms for AVs. This is particularly relevant in shared, unstructured environments where navigation relies heavily on negotiation and cooperation, considering the mutual interactions between pedestrians and vehicles.

5. **Formal safety guarantees for learning based navigation methods:** To overcome scalability issues and the computational time burden of formal methods in interaction modelling during navigation, we proposed our model based on DRL. However, this method, by its nature, has limitations concerning formally guaranteed safety. While we showed through simulation that we can achieve zero collisions with our engineered reward function and uncertainty consideration, there is no formal safety guarantee for collision avoidance of these learning methods. On the other hand, we showed that MPC methods, which can provide guarantees through hard constraint definition in the formulation, can encounter infeasible solution in crowded environments. Therefore, providing such guarantees for RL methods could be pursued in future work through using safe-RL techniques, such as providing safety shields after the DRL's action output [73]. Other approaches could involve combining DRL with MPC, where DRL provides high-level suggestions for MPC and helps MPC avoid infeasibility. These directions could be investigated in future work.

References

[1] M Zamani Abnili and Nasser L Azad. Short term predictions of preceding vehicle speeds for connected and automated vehicles. *Controls, Dynamic Systems and Robotics (CDSR)*, 2019.

[2] Timur Akhtyamov, Aleksandr Kashirin, Aleksey Postnikov, and Gonzalo Ferrer. Social robot navigation through constrained optimization: a comparative study of uncertainty-based objectives and constraints. In *2023 European Conference on Mobile Robots (ECMR)*, pages 1–8. IEEE, 2023.

[3] 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.

[4] Rachid Alami et al. Hateb-2: Reactive planning and decision making in human-robot co-navigation. In *2020 29th IEEE international conference on robot and human interactive communication (RO-MAN)*, pages 179–186. IEEE, 2020.

[5] Daniel Althoff, Matthias Althoff, Dirk Wollherr, and Martin Buss. Probabilistic collision state checker for crowded environments. In *2010 IEEE international conference on robotics and automation*, pages 1492–1498. IEEE, 2010.

[6] Javad Amirian, Bingqing Zhang, Francisco Valente Castro, Juan Jose Baldelomar, Jean-Bernard Hayet, and Julien Pettré. Opentraj: Assessing prediction complexity in human trajectories datasets. In *Proceedings of the Asian Conference on Computer Vision*, 2020.

[7] Cyrus Anderson, Ram Vasudevan, and Matthew Johnson-Roberson. Off the beaten sidewalk: Pedestrian prediction in shared spaces for autonomous vehicles. *IEEE Robotics and Automation Letters*, 5(4):6892–6899, 2020.

[8] Bani Anvari, Michael GH Bell, Panagiotis Angeloudis, and Washington Y Ochieng. Long-range collision avoidance for shared space simulation based on social forces. *Transportation Research Procedia*, 2:318–326, 2014.

[9] Bani Anvari, Michael GH Bell, Aruna Sivakumar, and Washington Y Ochieng. Modelling shared space users via rule-based social force model. *Transportation Research Part C: Emerging Technologies*, 51:83–103, 2015.

[10] Georges S Aoude, Brandon D Luders, Joshua M Joseph, Nicholas Roy, and Jonathan P How. Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns. *Autonomous Robots*, 35:51–76, 2013.

[11] Andrea Bajcsy, Anand Siththaranjan, Claire J Tomlin, and Anca D Dragan. Analyzing human models that adapt online. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2754–2760. IEEE, 2021.

[12] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting linear mixed-effects models using lme4. *arXiv preprint arXiv:1406.5823*, 2014.

[13] Mariana Batista and Bernhard Friedrich. Investigating spatial behaviour in different types of shared space. *Transportation Research Procedia*, 60:44–51, 2022.

[14] Maren Bennewitz, Wolfram Burgard, Grzegorz Cielniak, and Sebastian Thrun. Learning motion patterns of people for compliant robot motion. *The International Journal of Robotics Research*, 24(1):31–48, 2005.

[15] Aniket Bera, Tanmay Randhavane, Rohan Prinja, Kyra Kapsaskis, Austin Wang, Kurt Gray, and Dinesh Manocha. The emotionally intelligent robot: Improving social navigation in crowded environments. *arXiv preprint arXiv:1903.03217*, 2019.

[16] Apratim Bhattacharyya, Daniel Olmeda Reino, Mario Fritz, and Bernt Schiele. Europvi: Pedestrian vehicle interactions in dense urban centers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6408–6417, 2021.

[17] Huikun Bi, Zhong Fang, Tianlu Mao, Zhaoqi Wang, and Zhigang Deng. Joint prediction for kinematic trajectories in vehicle-pedestrian-mixed scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10383–10392, 2019.

[18] 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)*, 11(3):1–24, 2022.

[19] Julian Bock, Robert Krajewski, Tobias Moers, Steffen Runde, Lennart Vater, and Lutz Eckstein. The ind dataset: A drone dataset of naturalistic road user trajectories at german intersections. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 1929–1934. IEEE, 2020.

[20] Raul Borsche and Anne Meurer. Microscopic and macroscopic models for coupled car traffic and pedestrian flow. *Journal of Computational and Applied Mathematics*, 348:356–382, 2019.

[21] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.

[22] 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.

[23] Christopher G Burns, Luis Oliveira, Vivien Hung, Peter Thomas, and Stewart Birrell. Pedestrian attitudes to shared-space interactions with autonomous vehicles–a virtual reality study. In *International Conference on Applied Human Factors and Ergonomics*, pages 307–316. Springer, 2019.

[24] Jiannan Cai, Yuxi Zhang, Liu Yang, Hubo Cai, and Shuai Li. A context-augmented deep learning approach for worker trajectory prediction on unstructured and dynamic construction sites. *Advanced Engineering Informatics*, 46:101173, 2020.

[25] Yingfeng Cai, Lei Dai, Hai Wang, Long Chen, Yicheng Li, Miguel Angel Sotelo, and Zhixiong Li. Pedestrian motion trajectory prediction in intelligent driving from far shot first-person perspective video. *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[26] Fanta Camara and Charles Fox. Space invaders: Pedestrian proxemic utility functions and trust zones for autonomous vehicle interactions. *International Journal of Social Robotics*, 13(8):1929–1949, 2021.

[27] Sandra Carrasco, D Fernández Llorca, and MA Sotelo. Scout: Socially-consistent and understandable graph attention network for trajectory prediction of vehicles and

vrus. In *2021 IEEE Intelligent Vehicles Symposium (IV)*, pages 1501–1508. IEEE, 2021.

[28] Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Traphic: Trajectory prediction in dense and heterogeneous traffic using weighted interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8483–8492, 2019.

[29] Rohan Chandra, Uttaran Bhattacharya, Christian Roncal, Aniket Bera, and Dinesh Manocha. Robusttp: End-to-end trajectory prediction for heterogeneous road-agents in dense traffic with noisy sensor inputs. In *Proceedings of the 3rd ACM Computer Science in Cars Symposium*, pages 1–9, 2019.

[30] 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.

[31] 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.

[32] Junan Chen, Yan Wang, Ruihan Wu, and Mark Campbell. Spatial-temporal graph neural network for interaction-aware vehicle trajectory prediction. In *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, pages 2119–2125. IEEE, 2021.

[33] Kai Chen, Xiao Song, and Xiaoxiang Ren. Pedestrian trajectory prediction in heterogeneous traffic using pose keypoints-based convolutional encoder-decoder network. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(5):1764–1775, 2020.

[34] Kai Chen, Xiao Song, and Xiaoxiang Ren. Modeling social interaction and intention for pedestrian trajectory prediction. *Physica A: Statistical Mechanics and its Applications*, 570:125790, 2021.

[35] 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.

[36] 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.

[37] Yuxiao Chen, Ugo Rosolia, Wyatt Ubellacker, Noel Csomay-Shanklin, and Aaron D Ames. Interactive multi-modal motion planning with branch model predictive control. *IEEE Robotics and Automation Letters*, 7(2):5365–5372, 2022.

[38] Hao Cheng, Yao Li, and Monika Sester. Pedestrian group detection in shared space. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1707–1714. IEEE, 2019.

[39] Hao Cheng, Wentong Liao, Michael Ying Yang, Monika Sester, and Bodo Rosenhahn. Mcenet: multi-context encoder network for homogeneous agent trajectory prediction in mixed traffic. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8. IEEE, 2020.

[40] Hao Cheng and Monika Sester. Mixed traffic trajectory prediction using lstm–based models in shared space. In *Geospatial Technologies for All: Selected Papers of the 21st AGILE Conference on Geographic Information Science 21*, pages 309–325. Springer, 2018.

[41] Hao Cheng and Monika Sester. Modeling mixed traffic in shared space using lstm with probability density mapping. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3898–3904. IEEE, 2018.

[42] Jianhui Cheng and Dewei Li. A study of interference between pedestrians and vehicles in drop-off area at railway station based on celluar automata. In *2019 5th International Conference on Transportation Information and Safety (ICTIS)*, pages 913–919. IEEE, 2019.

[43] Amit Chougule, Vinay Chamola, Aishwarya Sam, Fei Richard Yu, and Biplab Sikdar. A comprehensive review on limitations of autonomous driving and its impact on accidents and collisions. *IEEE Open Journal of Vehicular Technology*, 2023.

[44] Daisuke Chugo, Zhaoyu Liu, Satoshi Muramatsu, Yuki Sakaida, Sho Yokota, and Hiroshi Hashimoto. A dynamic personal space of a pedestrian for a personal mobility vehicle. In *2013 IEEE 4th International Conference on Cognitive Infocommunications (CogInfoCom)*, pages 223–228. IEEE, 2013.

[45] Emma Clarke. Shared space-the alternative approach to calming traffic. *Traffic engineering & control*, 47(8), 2006.

[46] Alexander Cui, Sergio Casas, Abbas Sadat, Renjie Liao, and Raquel Urtasun. Lookout: Diverse multi-future prediction and planning for self-driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16107–16116, 2021.

[47] Daniel Dauner, Marcel Hallgarten, Tianyu Li, Xinshuo Weng, Zhiyu Huang, Zetong Yang, Hongyang Li, Igor Gilitschenski, Boris Ivanovic, Marco Pavone, et al. Navsim: Data-driven non-reactive autonomous vehicle simulation and benchmarking. *arXiv preprint arXiv:2406.15349*, 2024.

[48] Juan Pablo de Vicente. Deepsocnav: Social navigation by imitating human behaviors. Master's thesis, Pontificia Universidad Catolica de Chile (Chile), 2021.

[49] Patrick Dendorfer, Aljosa Osep, and Laura Leal-Taixé. Goal-gan: Multimodal trajectory prediction based on goal position estimation. In *Proceedings of the Asian Conference on Computer Vision*, 2020.

[50] Debargha Dey, Azra Habibovic, Bastian Pfleging, Marieke Martens, and Jacques Terken. Color and animation preferences for a light band ehmi in interactions between automated vehicles and pedestrians. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13. ACM, 2020.

[51] Debargha Dey, Marieke Martens, Berry Eggen, and Jacques Terken. The impact of vehicle appearance and vehicle behavior on pedestrian interaction with autonomous vehicles. In *Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications Adjunct*, pages 158–162. ACM, 2017.

[52] Anderson Carlos Dos Santos and Valdir Grassi. Pedestrian trajectory prediction with pose representation and latent space variables. In *2021 Latin American Robotics Symposium (LARS), 2021 Brazilian Symposium on Robotics (SBR), and 2021 Workshop on Robotics in Education (WRE)*, pages 192–197. IEEE, 2021.

[53] Noel E Du Toit and Joel W Burdick. Probabilistic collision checking with chance constraints. *IEEE Transactions on Robotics*, 27(4):809–815, 2011.

[54] Noel E Du Toit and Joel W Burdick. Robot motion planning in dynamic, uncertain environments. *IEEE Transactions on Robotics*, 28(1):101–115, 2011.

[55] Stuart Eiffert, He Kong, Navid Pirmarzdashti, and Salah Sukkarieh. Path planning in dynamic environments using generative rnns and monte carlo tree search. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10263–10269. IEEE, 2020.

[56] Stuart Eiffert, Kunming Li, Mao Shan, Stewart Worrall, Salah Sukkarieh, and Eduardo Nebot. Probabilistic crowd gan: Multimodal pedestrian trajectory prediction using a graph vehicle-pedestrian attention network. *IEEE Robotics and Automation Letters*, 5(4):5026–5033, 2020.

[57] 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.

[58] Michael Everett, Yu Fan Chen, and Jonathan P How. Motion planning among dynamic, decision-making agents with deep reinforcement learning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3052–3059. IEEE, 2018.

[59] Muhammad Fahad, Zhuo Chen, and Yi Guo. Learning how pedestrians navigate: A deep inverse reinforcement learning approach. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 819–826. IEEE, 2018.

[60] 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.

[61] 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.

[62] NHTSA's National Center for Statistics and Analysis. Critical reasons for crashes investigated in the national motor vehicle crash causation survey. Technical Report DOT HS 812 115, National Highway Traffic Safety Administration, February 2015.

[63] Anthony Francis, Claudia Pérez-d'Arpino, Chengshu Li, Fei Xia, Alexandre Alahi, Rachid Alami, Aniket Bera, Abhijat Biswas, Joydeep Biswas, Rohan Chandra, et al. Principles and guidelines for evaluating social robot navigation algorithms. *arXiv preprint arXiv:2306.16740*, 2023.

[64] David Fridovich-Keil, Ellis Ratner, Lasse Peters, Anca D Dragan, and Claire J Tomlin. Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games. In *2020 IEEE international conference on robotics and automation (ICRA)*, pages 1475–1481. IEEE, 2020.

[65] Óscar Gil, Anaís Garrell, and Alberto Sanfeliu. Social robot navigation tasks: Combining machine learning techniques and social force model. *Sensors*, 21(21):7087, 2021.

[66] Harshayu Girase, Haiming Gang, Srikanth Malla, Jiachen Li, Akira Kanehara, Karttikeya Mangalam, and Chiho Choi. Loki: Long term and key intentions for trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9803–9812, 2021.

[67] Mahsa Golchoubian, Moojan Ghafurian, Nasser Lashgarian Azad, and Kerstin Dautenhahn. What are social norms for low-speed autonomous vehicle navigation in crowded environments? an online survey. In *Proceedings of the 9th International Conference on Human-Agent Interaction*, pages 148–156, 2021.

[68] Mahsa Golchoubian, Moojan Ghafurian, Nasser Lashgarian Azad, and Kerstin Dautenhahn. Characterizing structured versus unstructured environments based on pedestrians' and vehicles' motion trajectories. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pages 2888–2895. IEEE, 2022.

[69] Mahsa Golchoubian, Moojan Ghafurian, Kerstin Dautenhahn, and Nasser Lashgarian Azad. Pedestrian trajectory prediction in pedestrian-vehicle mixed environments: A systematic review. *IEEE Transactions on Intelligent Transportation Systems*, 2023.

[70] Mahsa Golchoubian, Moojan Ghafurian, Kerstin Dautenhahn, and Nasser Lashgarian Azad. Polar collision grids: Effective interaction modelling for pedestrian trajectory prediction in shared space using collision checks. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pages 791–798. IEEE, 2023.

[71] Mahsa Golchoubian, Moojan Ghafurian, Kerstin Dautenhahn, and Nasser Lashgarian Azad. Uncertainty-aware drl for autonomous vehicle crowd navigation in shared space. *IEEE Transactions on Intelligent Vehicles*, 2024.

[72] Andrea Gorrini, Kenichiro Shimura, Stefania Bandini, Kazumichi Ohtsuka, and Katsuhiro Nishinari. Experimental investigation of pedestrian personal space: toward

modeling and simulation of pedestrian crowd dynamics. *Transportation research record*, 2421(1):57–63, 2014.

[73] Shangding Gu, Long Yang, Yali Du, Guang Chen, Florian Walter, Jun Wang, Yaodong Yang, and Alois Knoll. A review of safe reinforcement learning: Methods, theory and applications. *arXiv preprint arXiv:2205.10330*, 2022.

[74] Dario Calogero Guastella and Giovanni Muscato. Learning-based methods of perception and navigation for ground vehicles in unstructured environments: A review. *Sensors*, 21(1):73, 2021.

[75] 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.

[76] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2255–2264, 2018.

[77] Jason Hardy and Mark Campbell. Contingency planning over probabilistic obstacle predictions for autonomous road vehicles. *IEEE Transactions on Robotics*, 29(4):913–929, 2013.

[78] Muhammad Ahmed Hassan, Muhammad Usman Ghani Khan, Razi Iqbal, Omer Riaz, Ali Kashif Bashir, and Usman Tariq. Predicting humans future motion trajectories in video streams using generative adversarial network. *Multimedia Tools and Applications*, pages 1–23, 2021.

[79] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.

[80] Michael Herman, Jörg Wagner, Vishnu Prabhakaran, Nicolas Möser, Hanna Ziesche, Waleed Ahmed, Lutz Bürkle, Ernst Kloppenburg, and Claudius Gläser. Pedestrian behavior prediction for automated driving: Requirements, metrics, and relevant features. *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[81] Omar Hesham and Gabriel Wainer. Advanced models for centroidal particle dynamics: short-range collision avoidance in dense crowds. *Simulation*, 97(8):529–543, 2021.

[82] Sakif Hossain, Fatema T. Johora, Jörg P. Müller, and Sven Hartmann. A conceptual model of conflicts in shared spaces. In *2020 The 6th International Conference on Industrial and Business Engineerin*, pages 228–235, 2020.

[83] Yue Hu, Siheng Chen, Ya Zhang, and Xiao Gu. Collaborative motion prediction via neural motion message passing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6319–6328, 2020.

[84] Markus Huber, Yi-Huang Su, Melanie Krüger, Katrin Faschian, Stefan Glasauer, and Joachim Hermsdörfer. Adjustments of speed and path when avoiding collisions with another pedestrian. *PloS one*, 9(2):e89589, 2014.

[85] Constantin Hubmann, Jens Schulz, Marvin Becker, Daniel Althoff, and Christoph Stiller. Automated driving in uncertain environments: Planning with interaction and uncertain maneuver prediction. *IEEE transactions on intelligent vehicles*, 3(1):5–17, 2018.

[86] Abdi Hadi Ibrahim, Claude Pegard, and El-Mustapha Mouaddib. Collision prediction between vehicles in an unstructured environment. In *14th International Conference on Sciences and Techniques of Automatic Control & Computer Engineering-STA'2013*, pages 509–514. IEEE, 2013.

[87] Boris Ivanovic, Yifeng Lin, Shubham Shrivastava, Punarjay Chakravarty, and Marco Pavone. Propagating state uncertainty through trajectory forecasting. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2351–2358. IEEE, 2022.

[88] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2375–2384, 2019.

[89] Zhuozhu Jian, Zihong Yan, Xuanang Lei, Zihong Lu, Bin Lan, Xueqian Wang, and Bin Liang. Dynamic control barrier function-based model predictive control to safety-critical obstacle-avoidance of mobile robot. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3679–3685. IEEE, 2023.

[90] Rui Jiang, Shuzhi Sam Ge, Nagacharan Teja Tangirala, and Tong Heng Lee. Interactive navigation of mobile robots based on human's emotion. In *Social Robotics: 8th International Conference, ICSR 2016, Kansas City, MO, USA, November 1-3, 2016 Proceedings 8*, pages 243–252. Springer, 2016.

[91] Andrés C Jiménez, John Anzola, and Alexander Jimenez-Triana. Pedestrian counting estimation based on fractal dimension. *Heliyon*, 5(4):e01449, 2019.

[92] Collin Johnson and Benjamin Kuipers. Socially-aware navigation using topological maps and social norm learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 151–157. ACM, 2018.

[93] Fatema T Johora, Hao Cheng, Jörg P Müller, and Monika Sester. An agent-based model for trajectory modelling in shared spaces: a combination of expert-based and deep learning approaches. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1878–1880, 2020.

[94] Fatema T Johora and Jörg P Müller. Modeling interactions of multimodal road users in shared spaces. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3568–3574. IEEE, 2018.

[95] Fatema T Johora and Jörg P Müller. Zone-specific interaction modeling of pedestrians and cars in shared spaces. *Transportation research procedia*, 47:251–258, 2020.

[96] Fatema T Johora and Jorg P Muller. On transferability and calibration of pedestrian and car motion models in shared spaces. *Transportation letters*, 13(3):172–182, 2021.

[97] R Jyothi, K Mahalakshmi, CK Vaishnavi, U Apoorva, and S Nitya. Driver assistance for safe navigation under unstructured traffic environment. In *2019 Global Conference for Advancement in Technology (GCAT)*, pages 1–5. IEEE, 2019.

[98] Maria Kabtoul, Anne Spalanzani, and Philippe Martinet. Towards proactive navigation: A pedestrian-vehicle cooperation based behavioral model. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6958–6964. IEEE, 2020.

[99] Maria Kabtoul, Anne Spalanzani, and Philippe Martinet. Proactive and smooth maneuvering for navigation around pedestrians. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 4723–4729. IEEE, 2022.

[100] Kapil D Katyal, Gregory D Hager, and Chien-Ming Huang. Intent-aware pedestrian prediction for adaptive crowd navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3277–3283. IEEE, 2020.

[101] Stefan Kerscher, Norbert Balbierer, Sebastian Kraust, Andreas Hartmannsgruber, Nikolaus Müller, and Bernd Ludwig. Intention-based prediction for pedestrians and vehicles in unstructured environments. In *VEHITS*, pages 307–314, 2018.

[102] Harmish Khambhaita and Rachid Alami. Viewing robot navigation in human environment as a cooperative activity. In *Robotics Research: The 18th International Symposium ISRR*, pages 285–300. Springer, 2020.

[103] Kyungdo Kim, Yoon Kyung Lee, Hyemin Ahn, Sowon Hahn, and Songhwai Oh. Pedestrian intention prediction for autonomous driving using a multiple stakeholder perspective model. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7957–7962. IEEE, 2020.

[104] 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.

[105] Sascha Kolski, Dave Ferguson, Mario Bellino, and Roland Siegwart. Autonomous driving in structured and unstructured environments. In *2006 IEEE Intelligent Vehicles Symposium*, pages 558–563. IEEE, 2006.

[106] Abhisek Konar, Bobak H Baghi, and Gregory Dudek. Learning goal conditioned socially compliant navigation from demonstration using risk-based features. *IEEE Robotics and Automation Letters*, 6(2):651–658, 2021.

[107] Ioannis Kostavelis, Andreas Kargakos, Dimitrios Giakoumis, and Dimitrios Tzovaras. Robot's workspace enhancement with dynamic human presence for socially-aware navigation. In *Computer Vision Systems: 11th International Conference, ICVS 2017, Shenzhen, China, July 10-13, 2017, Revised Selected Papers 11*, pages 279–288. Springer, 2017.

[108] Henrik Kretzschmar, Markus Spies, Christoph Sprunk, and Wolfram Burgard. Socially compliant mobile robot navigation via inverse reinforcement learning. *The International Journal of Robotics Research*, 35(11):1289–1307, 2016.

[109] Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726–1743, 2013.

[110] Markus Kuderer, Henrik Kretzschmar, Christoph Sprunk, and Wolfram Burgard. Feature-based prediction of trajectories for socially compliant navigation. In *Robotics: science and systems*, volume 8, pages 193–200, 2012.

[111] Wei-Cheng Lai, Zi-Xiang Xia, Hao-Siang Lin, Lien-Feng Hsu, Hong-Han Shuai, I-Hong Jhuo, and Wen-Huang Cheng. Trajectory prediction in heterogeneous environment via attended ecology embedding. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 202–210, 2020.

[112] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 336–345, 2017.

[113] Florin Leon and Marius Gavrilescu. A review of tracking and trajectory prediction methods for autonomous driving. *Mathematics*, 9(6):660, 2021.

[114] Jiachen Li, Hengbo Ma, and Masayoshi Tomizuka. Conditional generative neural system for probabilistic trajectory prediction. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6150–6156. IEEE, 2019.

[115] Jiachen Li, Hengbo Ma, Zhihao Zhang, Jinning Li, and Masayoshi Tomizuka. Spatio-temporal graph dual-attention network for multi-agent prediction and tracking. *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[116] Jiachen Li, Fan Yang, Hengbo Ma, Srikanth Malla, Masayoshi Tomizuka, and Chiho Choi. Rain: Reinforced hybrid attention inference network for motion forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16096–16106, 2021.

[117] Kunming Li, Mao Shan, Karan Narula, Stewart Worrall, and Eduardo Nebot. Socially aware crowd navigation with multimodal pedestrian trajectory prediction for autonomous vehicles. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8. IEEE, 2020.

[118] Lingyun Luke Li, Bin Yang, Ming Liang, Wenyuan Zeng, Mengye Ren, Sean Segal, and Raquel Urtasun. End-to-end contextual perception and prediction with interaction transformer. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5784–5791. IEEE, 2020.

[119] Xiang Li and Jian-Qiao Sun. Studies of vehicle lane-changing to avoid pedestrians with cellular automata. *Physica A: Statistical Mechanics and its Applications*, 438:251–271, 2015.

[120] Zirui Li, Jianwei Gong, Chao Lu, and Yangtian Yi. Interactive behavior prediction for heterogeneous traffic participants in the urban road: A graph-neural-network-based multitask learning framework. *IEEE/ASME Transactions on Mechatronics*, 26(3):1339–1349, 2021.

[121] Zirui Li, Chao Lu, Yangtian Yi, and Jianwei Gong. A hierarchical framework for interactive behaviour prediction of heterogeneous traffic participants based on graph neural network. *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[122] Shuijing Liu, Peixin Chang, Zhe Huang, Neeloy Chakraborty, Kaiwen Hong, Weihang Liang, D Livingston McPherson, Junyi Geng, and Katherine Driggs-Campbell. Intention aware robot crowd navigation with attention-based interaction graph. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12015–12021. IEEE, 2023.

[123] Shuijing Liu, Peixin Chang, Weihang Liang, Neeloy Chakraborty, and Katherine Driggs-Campbell. Decentralized structural-rnn for robot crowd navigation with deep reinforcement learning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3517–3524. IEEE, 2021.

[124] Wei Liu, Seong-Woo Kim, Scott Pendleton, and Marcelo H Ang. Situation-aware decision making for autonomous driving on urban road using online pomdp. In *2015 IEEE Intelligent Vehicles Symposium (IV)*, pages 1126–1133. IEEE, 2015.

[125] Yuexin Ma, Xinge Zhu, Sibo Zhang, Ruigang Yang, Wenping Wang, and Dinesh Manocha. Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. *33rd AAAI Conference on Artificial Intelligence, AAAI 2019, 31st Innovative Applications of Artificial Intelligence Conference, IAAI 2019 and the 9th AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019*, pages 6120–6127, 2019.

[126] Andre Mateus, David Ribeiro, Pedro Miraldo, and Jacinto C Nascimento. Efficient and robust pedestrian detection using deep learning for human-aware navigation. *Robotics and Autonomous Systems*, 113:23–37, 2019.

[127] Kohei Matsumoto, Akihiro Kawamura, Qi An, and Ryo Kurazume. Mobile robot navigation using learning-based method based on predictive state representation in a dynamic environment. In *2022 IEEE/SICE International Symposium on System Integration (SII)*, pages 499–504. IEEE, 2022.

[128] Christoforos Mavrogiannis, Francesca Baldini, Allan Wang, Dapeng Zhao, Pete Trautman, Aaron Steinfeld, and Jean Oh. Core challenges of social robot navigation: A survey. *ACM Transactions on Human-Robot Interaction*, 12(3):1–39, 2023.

[129] 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.

[130] Qianhui Men and Hubert PH Shum. Pytorch-based implementation of label-aware graph representation for multi-class trajectory prediction. *Software Impacts*, 11:100201, 2022.

[131] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.

[132] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[133] Xiaoyu Mo, Zhiyu Huang, Yang Xing, and Chen Lv. Multi-agent trajectory prediction with heterogeneous edge-enhanced graph attention network. *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[134] Yoichi Morales, Takahiro Miyashita, and Norihiro Hagita. Social robotic wheelchair centered on passenger and pedestrian comfort. *Robotics and Autonomous Systems*, 87:355–362, 2017.

[135] Khaled A Mustafa, Daniel Jarne Ornia, Jens Kober, and Javier Alonso-Mora. Racp: Risk-aware contingency planning with multi-modal predictions. *arXiv preprint arXiv:2402.17387*, 2024.

[136] 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.

[137] Junhyuk Oh, Satinder Singh, and Honglak Lee. Value prediction network. *Advances in neural information processing systems*, 30, 2017.

[138] Kotaro Okamoto, Akira Utsumi, Hirotake Yamazoe, Takahiro Miyashita, Shinji Abe, Kazuhiro Takahashi, and Norihiro Hagita. Classification of pedestrian behavior in a shopping mall based on lrf and camera observations. In *MVA*, pages 1–5, 2011.

[139] Ovgu Ozturk, Toshihiko Yamasaki, and Kiyoharu Aizawa. Detecting dominant motion flows in unstructured/structured crowd scenes. In *2010 20th International Conference on Pattern Recognition*, pages 3533–3536. IEEE, 2010.

[140] F Pascucci, N Rinke, C Schiermeyer, V Berkhahn, and B Friedrich. A discrete choice model for solving conflict situations between pedestrians and vehicles in shared space. *arXiv preprint arXiv:1709.09412*, 2017. Presented at the 97th annual meeting of the Transportation Research Board (January 2018 in Washington).

[141] Francesco Pascucci, Niels Rinke, Clemens Timmermann, Volker Berkhahn, and Benjamin Friedrich. Dataset used for discrete choice modeling of conflicts in shared spaces. https://doi.org/10.24355/dbbs.084-202111081809-0, 2021.

[142] 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.

[143] Björn Petrak, Gundula Sopper, Katharina Weitz, and Elisabeth André. Do you mind if i pass through? studying the appropriate robot behavior when traversing two conversing people in a hallway setting. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*, pages 369–375. IEEE, 2021.

[144] Mike Phillips and Maxim Likhachev. Sipp: Safe interval path planning for dynamic environments. In *2011 IEEE international conference on robotics and automation*, pages 5628–5635. IEEE, 2011.

[145] Manon Predhumeau, Lyuba Mancheva, Julie Dugdale, and Anne Spalanzani. An agent-based model to predict pedestrians trajectories with an autonomous vehicle in shared spaces. In *AAMAS 2021-20th International Conference on Autonomous Agents and Multiagent Systems*, pages 1–9, 2021.

[146] Manon Prédhumeau, Lyuba Mancheva, Julie Dugdale, and Anne Spalanzani. Agent-based modeling for predicting pedestrian trajectories around an autonomous vehicle. *Journal of Artificial Intelligence Research*, 73:1385–1433, 2022.

[147] Manon Prédhumeau, Anne Spalanzani, and Julie Dugdale. Pedestrian behavior in shared spaces with autonomous vehicles: An integrated framework and review. *IEEE Transactions on Intelligent Vehicles*, 8(1):438–457, 2021.

[148] Michael J Procopio, Jane Mulligan, and Greg Grudic. Learning terrain segmentation with classifier ensembles for autonomous robot navigation in unstructured environments. *Journal of Field Robotics*, 26(2):145–175, 2009.

[149] Jianmin Qin, Jiahu Qin, Jiaxin Qiu, Qingchen Liu, Man Li, and Qichao Ma. Srl-orca: A socially aware multi-agent mapless navigation algorithm in complex dynamic scenes. *IEEE Robotics and Automation Letters*, 9(1):143–150, 2023.

[150] Quecheng Qiu, Shunyi Yao, Jing Wang, Jun Ma, Guangda Chen, and Jianmin Ji. Learning to socially navigate in pedestrian-rich environments with interaction capacity. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 279–285. IEEE, 2022.

[151] Ben A Rainbow, Qianhui Men, and Hubert PH Shum. Semantics-stgcnn: A semantics-guided spatial-temporal graph convolutional network for multi-class trajectory prediction. In *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 2959–2966. IEEE, 2021.

[152] Tanmay Randhavane, Aniket Bera, Emily Kubin, Austin Wang, Kurt Gray, and Dinesh Manocha. Pedestrian dominance modeling for socially-aware robot navigation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5621–5628. IEEE, 2019.

[153] Amir Rasouli, Iuliia Kotseruba, Toni Kunic, and John K Tsotsos. Pie: A large-scale dataset and models for pedestrian intention estimation and trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6262–6271, 2019.

[154] Amir Rasouli and John K Tsotsos. Autonomous vehicles that interact with pedestrians: A survey of theory and practice. *IEEE transactions on intelligent transportation systems*, 21(3):900–918, 2019.

[155] Arun Kumar Reddy, Vaibhav Malviya, and Rahul Kala. Social cues in the autonomous navigation of indoor mobile robots. *International Journal of Social Robotics*, pages 1–24, 2020.

[156] Daniela A Ridel, Nachiket Deo, Denis Wolf, and Mohan Trivedi. Understanding pedestrian-vehicle interactions with vehicle mounted vision: An lstm model and empirical analysis. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 913–918. IEEE, 2019.

[157] Nils Rinke, Chris Schiermeyer, Federico Pascucci, Volker Berkhahn, and Bernhard Friedrich. A multi-layer social force approach to model interactions in shared spaces using collision prediction. *Transportation research procedia*, 25:1249–1267, 2017.

[158] Jorge Rios-Martinez, Anne Spalanzani, and Christian Laugier. From proxemics theory to socially-aware navigation: A survey. *International Journal of Social Robotics*, 7(2):137–153, 2015.

[159] Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *European conference on computer vision*, pages 549–565. Springer, 2016.

[160] Mikel Rodriguez, Saad Ali, and Takeo Kanade. Tracking in unstructured crowded scenes. In *2009 IEEE 12th International Conference on Computer Vision*, pages 1389–1396. IEEE, 2009.

[161] Alessandra Rossi, Fernando Garcia, Arturo Cruz Maya, Kerstin Dautenhahn, Kheng Lee Koay, Michael L Walters, and Amit K Pandey. Investigating the effects of social interactive behaviours of a robot on people's trust during a navigation task. In *Annual Conference Towards Autonomous Robotic Systems*, pages 349–361. Springer, 2019.

[162] Kanghyun Ryu and Negar Mehr. Integrating predictive motion uncertainties with distributionally robust risk-aware control for safe robot navigation in crowds. *arXiv preprint arXiv:2403.05081*, 2024.

[163] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 683–700. Springer, 2020.

[164] Adarsh Jagan Sathyamoorthy, Jing Liang, Utsav Patel, Tianrui Guan, Rohan Chandra, and Dinesh Manocha. Densecavoid: Real-time navigation in dense crowds using anticipatory behaviors. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11345–11352. IEEE, 2020.

[165] Friederike Schneemann and Irene Gohl. Analyzing driver-pedestrian interaction at crosswalks: A contribution to autonomous driving in urban environments. In *2016 IEEE intelligent vehicles symposium (IV)*, pages 38–43. IEEE, 2016.

[166] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[167] Wilko Schwarting, Alyssa Pierson, Javier Alonso-Mora, Sertac Karaman, and Daniela Rus. Social behavior for autonomous vehicles. *Proceedings of the National Academy of Sciences*, 116(50):24972–24978, 2019.

[168] Emmanuel Senft, Satoru Satake, and Takayuki Kanda. Would you mind me if i pass by you? socially-appropriate behaviour for an omni-based social robot in narrow environment. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pages 539–547, 2020.

[169] Roya Salek Shahrezaie, Bethany N Manalo, Aaron G Brantley, Casey R Lynch, and David Feil-Seifer. Advancing socially-aware navigation for public spaces. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1015–1022. IEEE, 2022.

[170] Neha Sharma, Chhavi Dhiman, and S Indu. Pedestrian intention prediction for autonomous vehicles: A comprehensive survey. *Neurocomputing*, 508:120–152, 2022.

[171] Phani Teja Singamaneni, Pilar Bachiller-Burgos, Luis J Manso, Anaís Garrell, Alberto Sanfeliu, Anne Spalanzani, and Rachid Alami. A survey on socially aware robot navigation: Taxonomy and future challenges. *The International Journal of Robotics Research*, page 02783649241230562, 2024.

[172] 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.

[173] Jack Stewart. Why people keep rear-ending self-driving cars, 2018. https://www.wired.com/story/self-driving-car-crashes-rear-endings-why-charts-statistics/.

[174] T. Stewart. Overview of motor vehicle traffic crashes in 2021. Technical Report DOT HS 813 435, National Highway Traffic Safety Administration, April 2023.

[175] Shan Su, Cheng Peng, Jianbo Shi, and Chiho Choi. Potential field: Interpretable and unified representation for trajectory prediction. *arXiv preprint arXiv:1911.07414*, 2019.

[176] Yuchao Su, Jie Du, Yuanman Li, Xia Li, Rongqin Liang, Zhongyun Hua, and Jiantao Zhou. Trajectory forecasting based on prior-aware directed graph convolutional neural network. *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[177] Lingfeng Sun, Wei Zhan, Di Wang, and Masayoshi Tomizuka. Interactive prediction for multiple, heterogeneous traffic participants with multi-agent hybrid dynamic bayesian network. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 1025–1031. IEEE, 2019.

[178] Muchen Sun, Francesca Baldini, Peter Trautman, and Todd Murphey. Move beyond trajectories: Distribution space coupling for crowd navigation. *Robotics: Science and Systems*, 2021.

[179] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[180] Lei Tai, Jingwei Zhang, Ming Liu, and Wolfram Burgard. Socially compliant navigation through raw depth inputs with generative adversarial imitation learning. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 1111–1117. IEEE, 2018.

[181] Behrad Toghi, Rodolfo Valiente, Dorsa Sadigh, Ramtin Pedarsani, and Yaser P Fallah. Altruistic maneuver planning for cooperative autonomous vehicles using multi-agent advantage actor-critic. *arXiv preprint arXiv:2107.05664*, 2021.

[182] Anh-Tuan Tran, Arun Muraleedharan, Hiroyuki Okuda, and Tatsuya Suzuki. Scenario-based stochastic mpc for vehicle speed control considering the interaction with pedestrians. *IFAC-PapersOnLine*, 53(2):15325–15331, 2020.

[183] 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*, 34(3):335–356, 2015.

[184] 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.

[185] Peter Trautman, Jeremy Ma, Richard M Murray, and Andreas Krause. Robot navigation in dense human crowds: the case for cooperation. In *2013 IEEE international conference on robotics and automation*, pages 2153–2160. IEEE, 2013.

[186] 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*, 14(4):1743–1760, 2017.

[187] Chieh-En Tsai and Jean Oh. A generative approach for socially compliant navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2160–2166. IEEE, 2020.

[188] Vaibhav V Unhelkar, Claudia Pérez-D'Arpino, Leia Stirling, and Julie A Shah. Human-robot co-navigation using anticipatory indicators of human walking motion. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6183–6190. IEEE, 2015.

[189] Rodolfo Valiente, Behrad Toghi, Ramtin Pedarsani, and Yaser P Fallah. Robustness and adaptability of reinforcement learning-based cooperative autonomous driving in mixed-autonomy traffic. *IEEE Open Journal of Intelligent Transportation Systems*, 3:397–410, 2022.

[190] 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.

[191] 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.

[192] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[193] Araceli Vega, Luis J Manso, Douglas G Macharet, Pablo Bustos, and Pedro Núñez. Socially aware robot navigation system in human-populated and interactive environments based on an adaptive spatial density function and space affordances. *Pattern Recognition Letters*, 118:72–84, 2019.

[194] Araceli Vega-Magro, Luis J Manso, Pablo Bustos, and Pedro Núnez. 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.

[195] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *2018 IEEE international Conference on Robotics and Automation (ICRA)*, pages 4601–4607. IEEE, 2018.

[196] Heinrich Von Stackelberg. *Market structure and equilibrium*. Springer Science & Business Media, 2010.

[197] YingQiao Wang. Ltn: Long-term network for long-term motion prediction. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 1845–1852. IEEE, 2021.

[198] Yu Wang and Shiwei Chen. Multi-agent trajectory prediction with spatio-temporal sequence fusion. *IEEE Transactions on Multimedia*, 2021.

[199] Benjamin W Wargo and Norman W Garrick. Shared space: Could less formal streets be better for both pedestrians and vehicles? Technical report, University of Connecticut, 2016.

[200] Lang Wu. *Mixed effects models for complex data*. Chapman and Hall/CRC, 2009.

[201] Zhanteng Xie and Philip Dames. Drl-vo: Learning to navigate through crowded dynamic scenes using velocity obstacles. *IEEE Transactions on Robotics*, 2023.

[202] Lu Xiong, Yixuan Zhang, Yuanzhi Liu, Hongyu Xiao, and Chen Tang. Integrated decision making and planning based on feasible region construction for autonomous vehicles considering prediction uncertainty. *IEEE Transactions on Intelligent Vehicles*, 2023.

[203] Wenda Xu, Jia Pan, Junqing Wei, and John M Dolan. Motion planning under uncertainty for on-road autonomous driving. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2507–2512. IEEE, 2014.

[204] Hao Xue, Du Q Huynh, and Mark Reynolds. Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1186–1194. IEEE, 2018.

[205] Dongfang Yang, Arda Kurt, Keith Redmill, and Ümit Özgüner. Agent-based microscopic pedestrian interaction with intelligent vehicles in shared space. In *Proceedings*

of the 2nd International Workshop on Science of Smart City Operations and Platforms Engineering, pages 69–74, 2017.

[206] Dongfang Yang, Linhui Li, Keith Redmill, and Ümit Özgüner. Top-view trajectories: A pedestrian dataset of vehicle-crowd interaction from controlled experiments and crowded campus. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 899–904. IEEE, 2019.

[207] Dongfang Yang, John M Maroli, Linhui Li, Menna El-Shaer, Bander A Jabr, Keith Redmill, Füsun Özguner, and Ümit Özguner. Crowd motion detection and prediction for transportation efficiency in shared spaces. In *2018 IEEE International Science of Smart City Operations and Platforms Engineering in Partnership with Global City Teams Challenge (SCOPE-GCTC)*, pages 1–6. IEEE, 2018.

[208] Dongfang Yang, Ümit Özgüner, and Keith Redmill. Social force based microscopic modeling of vehicle-crowd interaction. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1537–1542. IEEE, 2018.

[209] Dongfang Yang, Ümit Özgüner, and Keith Redmill. A social force based pedestrian motion model considering multi-pedestrian interaction with a vehicle. *ACM Transactions on Spatial Algorithms and Systems (TSAS)*, 6(2):1–27, 2020.

[210] Steven Yang. VCI Dataset (CITR), 2024. Accessed: 2024-06-27. URL: https://github.com/dongfang-steven-yang/vci-dataset-citr.

[211] Yu Yao, Ella Atkins, Matthew Johnson-Roberson, Ram Vasudevan, and Xiaoxiao Du. Bitrap: Bi-directional pedestrian trajectory prediction with multi-modal goal estimation. *IEEE Robotics and Automation Letters*, 6(2):1463–1470, 2021.

[212] Ziyi Yin, Ruijin Liu, Zhiliang Xiong, and Zejian Yuan. Multimodal transformer networks for pedestrian trajectory prediction. In *IJCAI*, pages 1259–1265, 2021.

[213] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 507–523. Springer, 2020.

[214] Ye Yuan, Xinshuo Weng, Yanglan Ou, and Kris M Kitani. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9813–9823, 2021.

[215] Chi Zhang and Christian Berger. Learning the pedestrian-vehicle interaction for pedestrian trajectory prediction. In *2022 8th International Conference on Control, Automation and Robotics (ICCAR)*, pages 230–236. IEEE, 2022.

[216] Lin Zhang, Kang Yuan, Hongqing Chu, Yanjun Huang, Haitao Ding, Jiawei Yuan, and Hong Chen. Pedestrian collision risk assessment based on state estimation and motion prediction. *IEEE Transactions on Vehicular Technology*, 71(1):98–111, 2021.

[217] Ronghui Zhang, Fuliang Li, Jiali Zhou, Feng You, Wenxiao Zeng, and Tonghai Jiang. A review on pedestrian crossing detection and behavior analysis: In intelligent transportation system. *Recent Patents on Computer Science*, 8(1):2–12, 2015.

[218] Xuexiang Zhang, Weiwei Zhang, Xuncheng Wu, and Wenguan Cao. Probabilistic trajectory prediction of heterogeneous traffic agents based on layered spatio-temporal graph. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 235(9):2413–2424, 2021.

[219] Zhao Zhang and Daocheng Fu. Modeling pedestrian–vehicle mixed-flow in a complex evacuation scenario. *Physica A: Statistical Mechanics and its Applications*, 599:127468, 2022.

[220] Zhou Zhou, Gang Huang, Zhaoxin Su, Yongfu Li, and Wei Hua. Dynamic attention-based cvae-gan for pedestrian trajectory prediction. *IEEE Robotics and Automation Letters*, 2022.

[221] Zhiqiang Zuo, Xiao Zhang, and Yijing Wang. Map feature based trajectory prediction with multi-class traffic participants. In *2021 40th Chinese Control Conference (CCC)*, pages 7312–7317. IEEE, 2021.

APPENDICES

Appendix A

Online survey

A.1 Ethics clearance certificate

UNIVERSITY OF WATERLOO

Notification of Ethics Clearance to Conduct Research with Human Participants

---

Principal Investigator: Kerstin Dautenhahn (Electrical and Computer Engineering)

Student investigator: Mahsa Golchoubian (Systems Design Engineering)

Co-Investigator: Nasser Lashgarian Azad (Systems Design Engineering)

Collaborator: Moojan Ghafurian (Computer Science Computing Facility)

File #: 43081

Title: Investigating humans' preferences for the motion of autonomous passenger-carrying golf carts operating close to pedestrians

---

The Human Research Ethics Committee is pleased to inform you this study has been reviewed and given ethics clearance.

Initial Approval Date: 05/04/21 (m/d/y)

University of Waterloo Research Ethics Committees are composed in accordance with, and carry out their functions and operate in a manner consistent with, the institution's guidelines for research with human participants, the Tri-Council Policy Statement for the Ethical Conduct for Research Involving Humans (TCPS, 2nd edition), International Conference on Harmonization: Good Clinical Practice (ICH-GCP), the Ontario Personal Health Information Protection Act (PHIPA), the applicable laws and regulations of the province of Ontario. Both Committees are registered with the U.S. Department of Health and Human Services under the Federal Wide Assurance, FWA00021410, and IRB registration number IRB00002419 (HREC) and IRB00007409 (CREC).

This study is to be conducted in accordance with the submitted application and the most recently approved versions of all supporting materials.

Expiry Date: 05/04/22 (m/d/y)

Multi-year research must be renewed at least once every 12 months unless a more frequent review has otherwise been specified. Studies will only be renewed if the renewal report is received and approved before the expiry date. Failure to submit renewal reports will result in the investigators being notified ethics clearance has been suspended and Research Finance being notified the ethics clearance is no longer valid.

Level of review: Delegated Review

Signed on behalf of the Human Research Ethics Committee

A.2 Online survey - Demographic questions

Mechanical Turk ID

Demographic questions

Gender :
(write NA if you prefer not to share)

Age :

- 18 - 24
- 25 - 34
- 35 - 44
- 45 - 54
- 55 - 64
- 65 - 74
- 75 - 84
- 85 or older
- Prefer not to say

Country of residence :
(write NA if you prefer not to share)

A.3 Online survey - Instructions for pedestrian role

Instructions:
Please imagine yourself in an indoor environment such as a big shopping mall or an airport terminal. There are a couple of golf carts around that are operating in the same indoor environment close to pedestrians. These golf carts are for transferring people from one place to another. Therefore, they could have passengers onboard while they move in the crowd. Below you can see some illustrations of the envisaged operation of these golf carts, e.g. in malls and airport terminals.
Please pay careful attention to the following features of the cart and its environment:

- These golf carts do not have any specified separate path in the mall/airport terminal for operation and pedestrians are sharing the same space with these golf carts.
- Throughout this survey, it is important for you to imagine that these **golf carts are operating autonomously without any driver**. Therefore, they should be considered **self-driving golf carts.**
- The maximum speed of the cart in absence of any obstacle is approximately 9.5 mph (4.2 m/s) which is around **3 times the human's walking speed**.
- As these golf carts are electric, they do not produce much sound.

In this study we will ask about your opinions on how such a driverless cart, e.g. in a shopping mall or airport terminal, should move among the other pedestrians in different situations. Please read the questions carefully and pay attention to the provided illustrations.

Please try to **imagine yourself in the scene** in **the place of a pedestrian** as being described or indicated in each question.

Images

At the end of the instruction step for the survey, participants were shown realistic images of a golf cart operating in a shopping mall or airport terminal alongside pedestrians. The images included both onboard and third-person perspectives, with **pedestrians** highlighted and marked with arrows labeled "You are here" to help participants visualize their role as **pedestrians** for the upcoming questions. These images have been removed here for copyright reasons.

A.4 Online survey - Instructions for passenger role

Instructions:
Please imagine yourself in an indoor environment such as a big shopping mall or an airport terminal. There are a couple of golf carts around that are operating in the same indoor environment close to pedestrians. These golf carts are for transferring people from one place to another. Therefore, they could have passengers onboard while they move in the crowd. Below you can see some illustrations of the envisaged operation of these golf carts, e.g. in malls and airport terminals.
Please pay careful attention to the following features of the cart and its environment:

- These golf carts do not have any specified separate path in the mall/airport terminal for operation and pedestrians are sharing the same space with these golf carts.

- Throughout this survey, it is important for you to imagine that these **golf carts are operating autonomously without any driver**. Therefore, they should be considered **self-driving golf carts.**
- The maximum speed of the cart in absence of any obstacle is approximately 9.5 mph (4.2 m/s) which is around **3 times the human's walking speed**.
- As these golf carts are electric, they do not produce much sound.

In this study we will ask about your opinions on how such a driverless cart, e.g. in a shopping mall or airport terminal, should move among the other pedestrians in different situations. Please read the questions carefully and pay attention to the provided illustrations.

Assume that **you have got into one of these carts as a passenger to take a ride**. Therefore, when answering the questions, please try to **imagine yourself in the scene** in **the place of a passenger onboarad the cart**.

Images

At the end of the instruction step for the survey, participants were shown realistic images of a golf cart operating in a shopping mall or airport terminal alongside pedestrians. The images included both onboard and third-person perspectives, with **passenger** seats highlighted and marked with arrows labeled "You are here" to help participants visualize their role as a **passenger** for the upcoming questions. These images have been removed here for copyright reasons.

A.5 Online survey - Pre-survey questions

You are walking straight ahead when you approach a person walking in the opposite direction as shown in the illustration. **For passing this person**, **which direction do you choose**? Assume there is enough space on both side of the person for you to pass (You are the person indicated by the green circle)

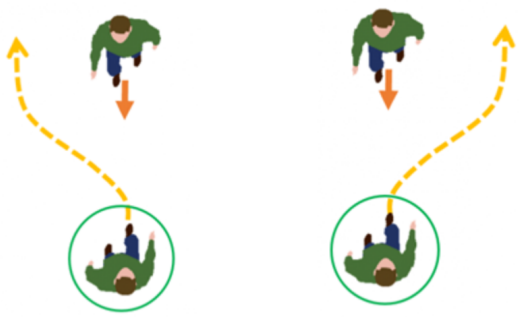

a) b) c) I have no preferred direction

You are walking straight ahead when you get behind a person walking in the same direction in front of you as shown in the illustration. **For overtaking this person**, **which direction do you choose**? Assume there is enough space on both side of the person for you to overtake. (You are the one indicated by the circle)

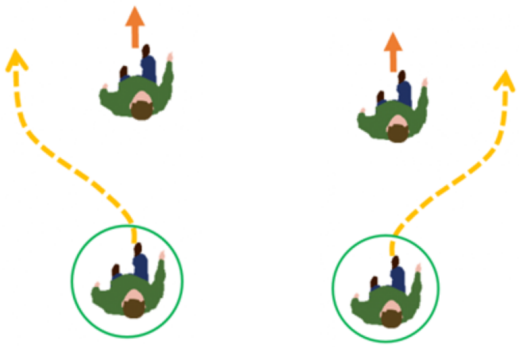

a) b) c) I have no preferred direction

A.6 Online survey - Questions for pedestrian role

Suppose **you are walking** straight ahead in a mall/airport terminal and a passenger-carrying golf cart is **approaching you from behind**. There is enough space on both of your sides for the cart to pass. **How do you feel if the cart behaves in each of the following ways** in this situation? The cart is moving autonomously without any driver.

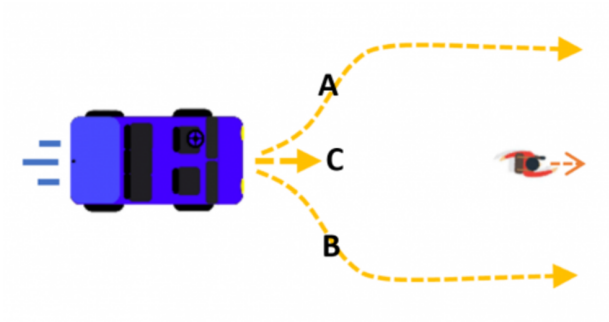

A. The cart overtakes you from **left**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

B. The cart overtakes you from **right**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

C. The cart **gets behind you** with a safe distance from you and follows your speed until you clear the path.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Suppose **you are walking** straight ahead in a shopping mall/airport terminal and a passenger-carrying golf cart is **approaching you from front**. There is enough space on both of your sides for the cart to pass. **How do you feel if the car behaves in each of the following ways** in this situation? The cart is moving autonomously without any driver.

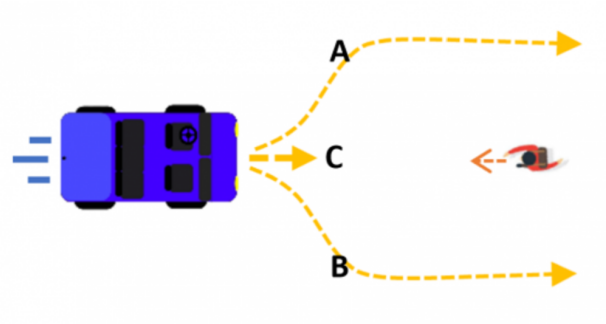

A. The cart passes you from your **right**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

B. The cart passes you from your **left**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

C. The cart decelerates in its straight path **until you clear its path** by going to your left or right.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Suppose **you are walking straight ahead** in a shopping mall or airport terminal and suddenly you recognize a **golf cart approaching you from behind** in the situation below. **What is your reaction**? You are the person circled in green. (The cart is moving autonomously without any driver)

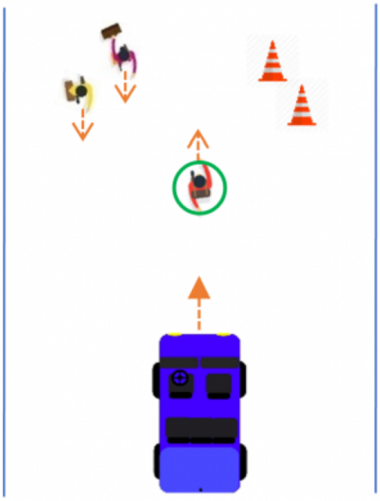

- a) I clear the cart's path by moving to the right
- b) I clear the cart's path by moving to the left
- c) I will keep walking straight ahead until the cart behind me reaches a location where it could overtake me.
- d) I pay no attention
- e) Others

Please explain the **reasons** behind your answer and the exact factors you paid attention to in the above illustration for making this choice.

You are walking along a corridor with a group of people in an airport terminal. On **which side of the corridor**, **you prefer the golf cart to move**? You can choose all the options that you prefer. (The cart is moving autonomously without any driver)

- [ ] a)

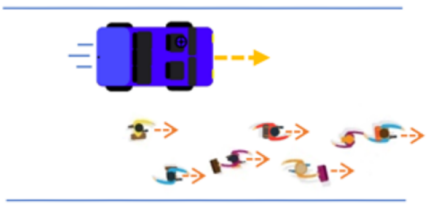

- [ ] b)

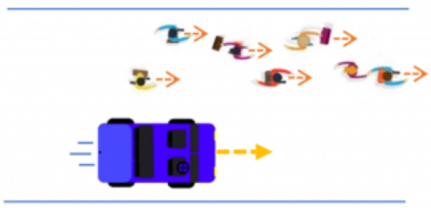

- [ ] c)

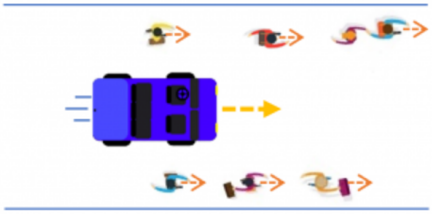

- [ ] d) I have no specific preference and I will follow the flow of pedestrian
- [ ] e) Others

Please explain your reason:

You are walking in a shopping mall/airport terminal and a couple of other pedestrians are walking in the same direction as you. **A golf cart arrives behind you** and its way is blocked by the flow of pedestrians. **How do you feel about each of the below actions taken by the cart** in this situation? You are the person circled in green. The cart is moving autonomously without any driver.
(the maximum speed of the cart in absence of any obstacle is approximately 9.5 mph (4.2 m/s) which is around 3 times the human's walking speed)

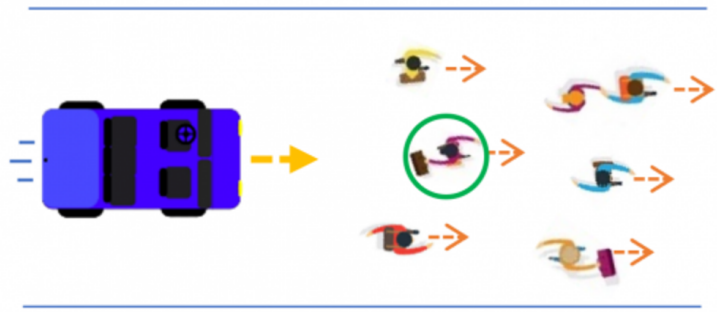

A. If the cart slows down its speed and <u>follows the flow</u> of pedestrian with a safe distance from them.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Please also answer the following questions about the same behavior of the cart

I feel comfortable with this behavior of the cart

Strongly Disagree — Strongly Agree

I think this behavior is dangerous

Strongly Disagree — Strongly Agree

I think this behavior is appropriate

Strongly Disagree — Strongly Agree

I would not have expected this behavior

Strongly Disagree — Strongly Agree

I would trust a cart that shows this behavior

Strongly Disagree — Strongly Agree

The cart did not seem to be aware of me when it showed this behavior

Strongly Disagree — Strongly Agree

B. If the cart requests a way through by using visual and auditory signals.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Please also answer the following questions about the same behavior of the cart

I feel comfortable with this behavior of the cart

Strongly Disagree — Strongly Agree

I think this behavior is dangerous

Strongly Disagree — Strongly Agree

I think this behavior is appropriate

Strongly Disagree — Strongly Agree

I would not have expected this behavior

Strongly Disagree — Strongly Agree

I would trust a cart that shows this behavior

Strongly Disagree — Strongly Agree

The cart did not seem to be aware of me when it showed this behavior

Strongly Disagree — Strongly Agree

Consider a situation in a shopping mall or an airport terminal where **the golf cart and the pedestrian are approaching each other** at a 90 degree angle. How do you feel in each of the following situations? **Consider yourself the person circled in green**. Both the pedestrian and the cart are aware of each other. (The cart is moving autonomously without any driver)

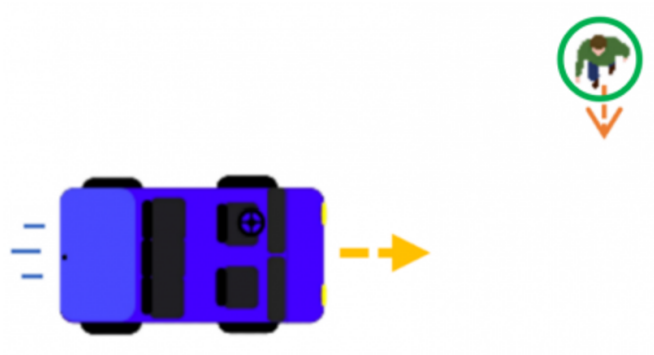

A. The cart stops and, in this way, gives you the priority to cross.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

B. The cart slows down in its straight path without a complete stop, to let you go first.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

C. The cart adjusts its path to the left or right to avoid you while keeping the same speed.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

D. The cart continues its path without decelerating to take the priority to cross.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

E. The cart shows its intention to cross first by visual and auditory signals and continues to cross your path with the same speed.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Consider the situation where the golf cart gets into the flow of pedestrian in a shopping mall/airport terminal as shown below. The intended direction of the cart is shown in the picture. **Which of the following behaviors do you expec**t from the driverless cart? **Consider yourself as one of the pedestrians in the flow**. You can select multiple choices.
(The cart is moving autonomously without any driver. The maximum speed of the cart in absence of any obstacle is approximately 9.5 mph (4.2 m/s) which is around 3 times the human's walking speed)

☐ a) The cart slows down and follows the flow of group A

☐ b) The cart slows down and tries to open up a passage in middle of the two flows to overtake the pedestrians and continue at its own higher speed, afterwards

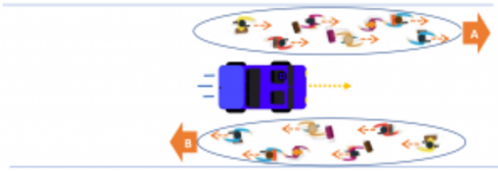

☐ c) The cart slows down and tries to overtake the pedestrians in group A from the left

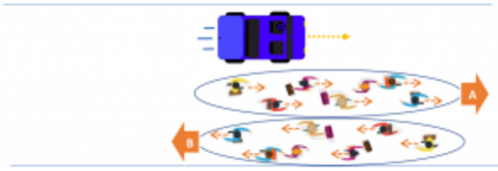

☐ d) The cart slows down and tries to overtake the pedestrian in group B from the right

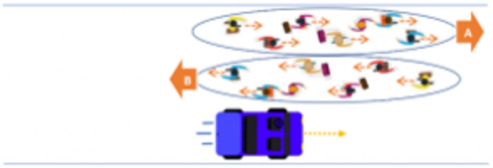

☐ e) Others

Please explain your reason:

The passenger-carrying golf cart arrives at the intersection in a shopping mall/airport terminal shown below and **its intention is to take a right turn**. Below are given three possible ways by which the cart can **indicate** its “turning right” intention. **How do you rate your feelings towards each of the following actions** taken by the cart for performing this turn? You are the person circled in green. The cart is moving autonomously without any driver.

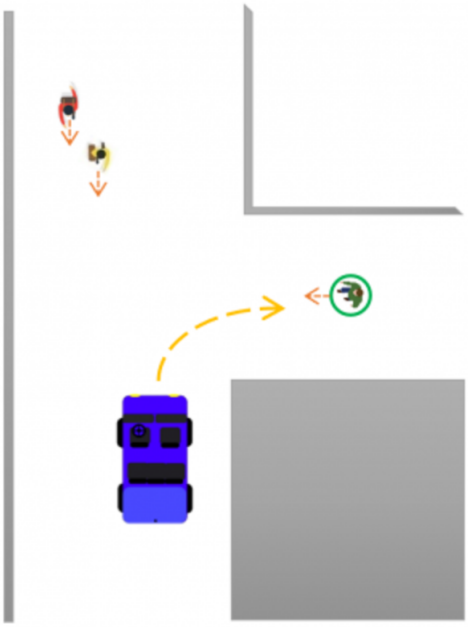

A. The cart decelerates and initiates the turn slowly, showing its intention **only through its movements** and making pedestrians to free the path. (no explicit signal)

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

B. The cart decelerates and by showing a **turn signal with flashing lights**, as in conventional cars, pauses and waits until it can turn safely.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

C. The cart decelerates and **projects a light** on the floor in front of it which **shows the exact path it is going to follow** in the next 10 seconds.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Consider the passenger-carrying **golf cart approaching a child** (under 5 years old) who is running around in a shopping mall or an airport terminal. Below are given two different behaviors the cart can show when encountering this situation. **Rate your impression of the cart** and **your emotional state** regarding each behavior. You are viewing this situation as a pedestrian close to the scene. (The cart is moving autonomously without any driver)

A. The cart <u>stops</u> within a safe distance in front of the child and waits until the child gets in a safe zone before continuing its straight path.

Please rate your **emotional state** on these scales:

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Please rate your **impression** of the cart on these scales:

Incompetent	○	○	○	○	○	Competent
Ignorant	○	○	○	○	○	Knowlegable
Irresponsible	○	○	○	○	○	Responsible
Unintelligent	○	○	○	○	○	Intelligent
Foolish	○	○	○	○	○	Sensible

B. The cart decelerates and <u>changes its path</u> to the right or left to avoid getting close to the child.

Please rate your **emotional state** on these scales:

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Please rate your **impression** of the cart on these scales:

Incompetent	○	○	○	○	○	Competent
Ignorant	○	○	○	○	○	Knowlegable
Irresponsible	○	○	○	○	○	Responsible
Unintelligent	○	○	○	○	○	Intelligent
Foolish	○	○	○	○	○	Sensible

A.7 Online survey - Questions for passenger role

Suppose **you are onboard** a driverless golf cart in a shopping mall/airport terminal. The cart **approaches a pedestrian that is walking straight ahead in the same direction as the cart**. There is enough space on both sides of the pedestrian for the cart to pass. **How do you feel if** the cart that you are in, behaves in each of the following ways in this situation? (The cart is moving autonomously without any driver)

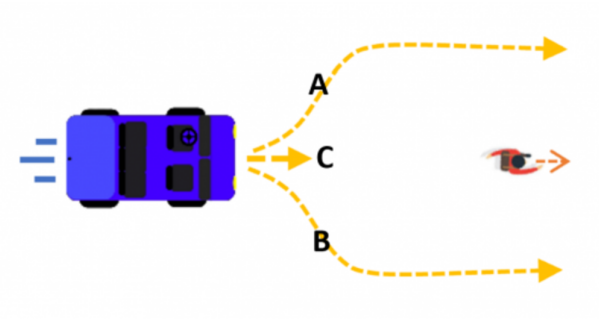

A. The cart overtakes the pedestrian from **left**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

B. The cart overtakes the pedestrian from **right**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

C. The cart gets behind the pedestrian with a safe distance from him/her and **follows the pedestrian**'s speed **until he/she clears the path**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Suppose **you are onboard** a driverless golf cart in a shopping mall/airport terminal. The cart **approaches a pedestrian that is walking straight ahead in the opposite direction** while getting close to the golf cart. There is enough space on both sides of the pedestrian for the cart to pass. **How do you feel if** the cart that you are in, behaves in each of the following ways in this situation? (The cart is moving autonomously without any driver)

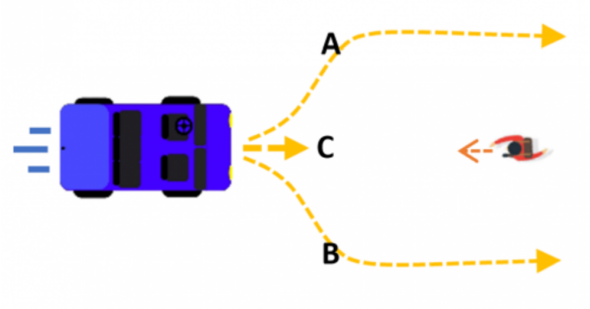

A. The cart passes the pedestrian on the **left**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

B. The cart passes the pedestrian on the **right**.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

C. The cart decelerates in its straight path **until the pedestrian clears the cart's path** by going to his/her left or right.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Suppose **you are onboard** a driverless golf cart in a shopping mall or airport terminal. The cart gets behind pedestrian "A" in the situation shown below. **The cart's path is blocked** and so, in this particular moment, the cart cannot overtake pedestrian "A" safely due to other pedestrians or obstacles present. **How should the cart behave**? (The cart is moving autonomously without any driver)

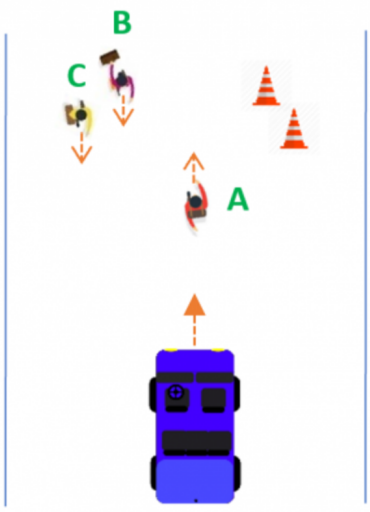

- ○ a) It should expect the pedestrian "A" to clear the cart's path so the cart can keep going.
- ○ b) It should expect pedestrians "B" and "C" to clear the cart's path so the cart can move to the left and keep going.
- ○ c) The cart should slow down and keep moving straight slowly behind pedestrian "A" until the cart reaches a location where it could safely overtake pedestrian "A".
- ○ d) Others

Please explain the **reasons** behind your answer and the exact factors you paid attention to in the above illustration for making this choice.

Suppose **you are onboard** a driverless golf cart in a shopping mall/airport terminal. As shown, the cart gets in a corridor with **pedestrians walking in the same direction as the cart**. On **which side of the corridor** should the cart move in this situation. You can choose all the options that you prefer. (The cart is moving autonomously without any driver)

- [ ] a)

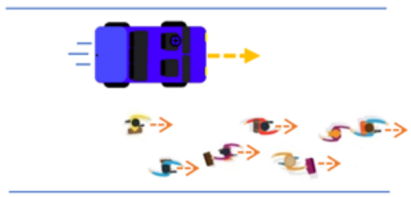

- [ ] b)

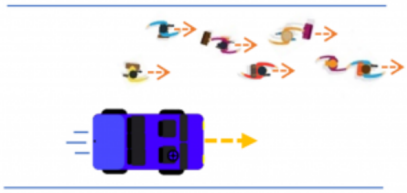

- [ ] c)

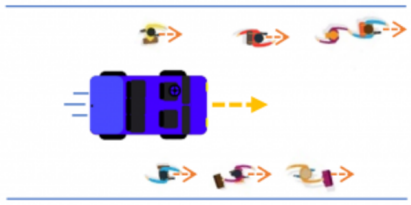

- [ ] d) I have no specific preference
- [ ] e) Others

Please explain your reason:

Suppose **you are onboard** a driverless golf cart in a shopping mall/airport terminal. **The cart gets behind** a couple of pedestrians walking in the same direction as the cart. **The cart's path is blocked** by the flow of pedestrians. **How do you feel about each of the below actions** taken by the cart in this situation? You are onboard the cart and the cart is moving autonomously without any driver.
(the maximum speed of the cart in absence of any obstacle is approximately 9.5 mph (4.2 m/s) which is around 3 times the human's walking speed)

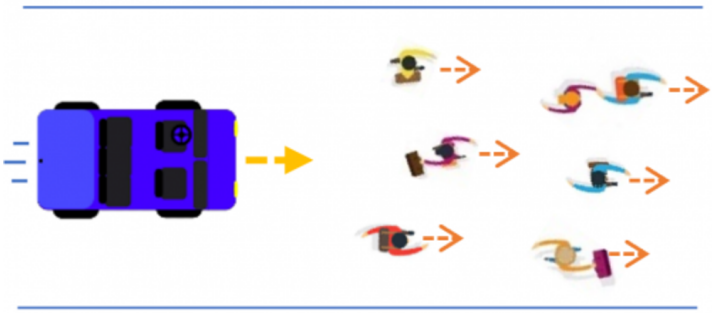

A. If the cart slows down its speed and <u>follows the flow of pedestrian</u> with a safe distance from them.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Please also answer the following questions about the same behavior of the cart

I feel comfortable with this behavior of the cart

Strongly Disagree — Strongly Agree

I think this behavior is dangerous

Strongly Disagree — Strongly Agree

I think this behavior is appropriate

Strongly Disagree — Strongly Agree

I would not have expected this behavior

Strongly Disagree — Strongly Agree

I would trust a cart that shows this behavior

Strongly Disagree — Strongly Agree

The cart did not seem to be aware of the pedestrians when it showed this behavior

Strongly Disagree — Strongly Agree

B. If the cart requests a way through by using visual and auditory signals.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Please also answer the following questions about the same behavior of the cart

I feel comfortable with this behavior of the cart

Strongly Disagree ——— Strongly Agree

I think this behavior is dangerous

Strongly Disagree ——— Strongly Agree

I think this behavior is appropriate

Strongly Disagree ——— Strongly Agree

I would not have expected this behavior

Strongly Disagree ——— Strongly Agree

I would trust a cart that shows this behavior

Strongly Disagree ——— Strongly Agree

The cart did not seem to be aware of the pedestrians when it showed this behavior

Strongly Disagree ——— Strongly Agree

Suppose you are **onboard a driverless golf cart** in a shopping mall/airport terminal. Consider a situation where **the cart and the pedestrian are approaching** each other at a 90 degree angle. **How do you feel** in each of the following situations? Answer the question from the view of a person inside the cart. Both the pedestrian and the cart are aware of each other. (The cart is moving autonomously without any driver)

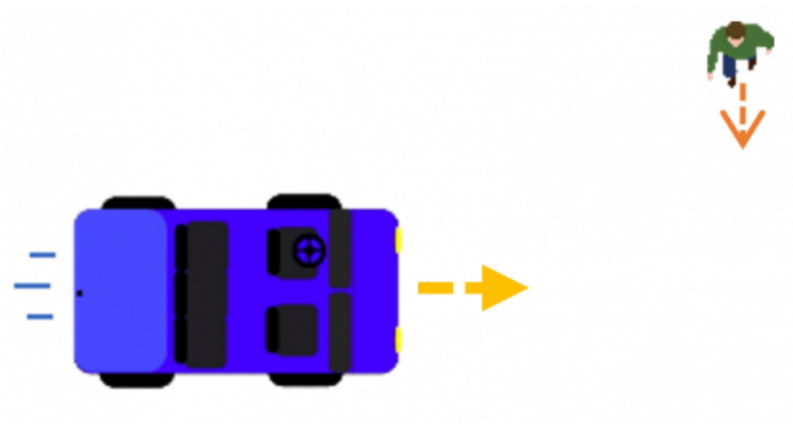

A. The cart stops and, in this way, gives the pedestrian the priority to cross.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

B. The cart slows down in its straight path without a complete stop, to let the passenger go first.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

C. The cart adjusts its path to the left or right for avoiding the pedestrian while keeping the same speed.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

D. The cart continues its path without decelerating to take the priority to cross.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

E. The cart shows its intention to cross first, by visual and auditory signals and continues to cross the pedestrian's path with the same speed.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Suppose **you are onboard** a driverless golf cart in a shopping mall/airport terminal. Consider the situation where the cart gets into the flow of pedestrian as shown below. The intended direction of the cart is shown in the picture. **Which of the following behaviors do you expect from the cart**? Answer the question **from the view of a passenger onboard the cart**. You can select multiple choices.
(The cart is moving autonomously without any driver. The maximum speed of the cart in absence of any obstacle is approximately 9.5 mph (4.2 m/s) which is around 3 times the human's walking speed)

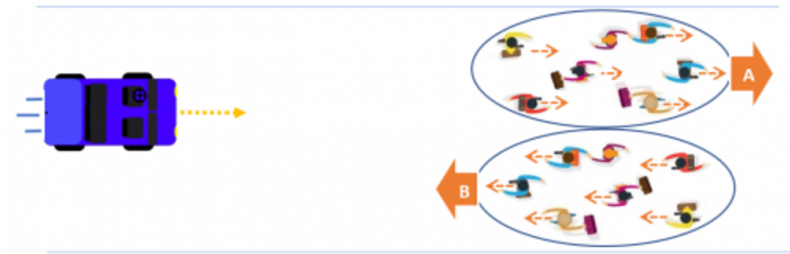

- [ ] a) The cart slows down and follows the flow of group A

- [ ] b) The cart slows down and tries to open up a passage in middle of the two flows to overtake the pedestrians and continue at its own, higher speed afterwards

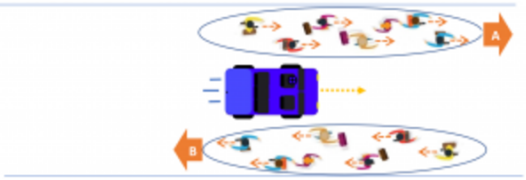

- [ ] c) The cart slows down and tries to overtake the pedestrians in group A from the left

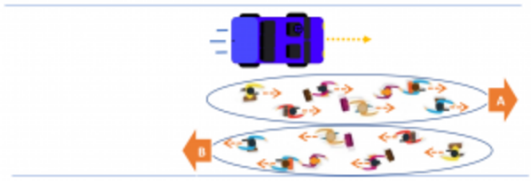

- [ ] d) The cart slows down and tries to overtake the pedestrian in group B from the right

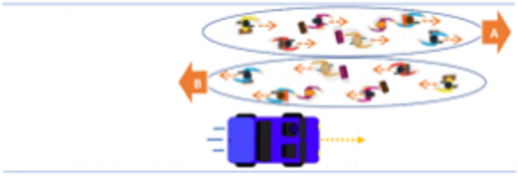

- [ ] e) Others

Please explain your reason:

Suppose **you are onboard** a driverless golf cart in a shopping mall/airport terminal. The cart arrives at the intersection shown below and its intention is to **take a right turn**. Below are given three possible ways by which the cart can **indicate** its "turning right" intention. How do you rate your feelings towards each of the following actions taken by the cart for performing this turn? Answer the questions **from the view of a passenger onboard the cart**. (The cart is moving autonomously without any driver)

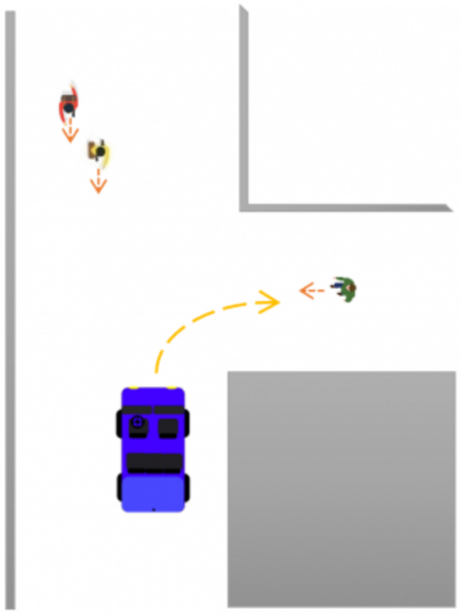

A. The cart decelerates and initiates the turn slowly, showing its intention **only through its movements** and making pedestrians to free the path. (no explicit signal)

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

B. The cart decelerates and by showing a **turn signal with flashing lights**, as in conventional cars, pauses and waits until it can turn safely.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

C. The cart decelerates and **projects a light** on the floor in front of it which **shows the exact path it is going to follow** in the next 10 seconds.

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Suppose **you are onboard** a driverless golf cart in a shopping mall/airport terminal. Consider **the cart approaching a child** (under 5 years old) who is running around. Below are given two different behaviors the cart can show when encountering this situation. Rate your impression of the cart and your emotional state regarding each behavior. **You are viewing this situation as a passenger inside the cart**. (The cart is moving autonomously without any driver)

A. The cart <u>stops</u> within a safe distance in front of the child and waits until the child gets in a safe zone before continuing its straight path.

Please rate your **emotional state** on these scales:

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Please rate your **impression** of the cart on these scales:

Incompetent	○	○	○	○	○	Competent
Ignorant	○	○	○	○	○	Knowlegable
Irresponsible	○	○	○	○	○	Responsible
Unintelligent	○	○	○	○	○	Intelligent
Foolish	○	○	○	○	○	Sensible

B. The cart decelerates and <u>changes its path</u> to the right or left to avoid getting close to the child.

Please rate your **emotional state** on these scales:

Anxious	○	○	○	○	○	Relaxed
Agitated	○	○	○	○	○	Calm
Quiescent	○	○	○	○	○	Surprised

Please rate your **impression** of the cart on these scales:

Incompetent	○	○	○	○	○	Competent
Ignorant	○	○	○	○	○	Knowlegable
Irresponsible	○	○	○	○	○	Responsible
Unintelligent	○	○	○	○	○	Intelligent
Foolish	○	○	○	○	○	Sensible

A.8 Online survey - Post-survey questions

Suppose you want to go to a shopping mall or an airport terminal where you know some driverless golf carts will be operating indoors for transporting passengers around. **Is there anything you would be concerned about in this environment**?

When responding the questions in this study and selecting your preferences about the cart's path, did you have any specific concern which you think might have influenced your responses? You can choose multiple answers and also type any other concerns in the field provided under 'others'.

- ☐ a) The presence of **children**
- ☐ b) The presence of people with **disabilities**
- ☐ c) The presence of older adults with **walking difficulties**
- ☐ d) The presence of pedestrians **carrying many suitcases and bags**
- ☐ e) The cart's approaching **sound not being noticeable** by the pedestrians
- ☐ f) Others
- ☐ g) None

Suppose you are walking in a shopping mall or an airport terminal. There are a couple of driverless golf carts operating around for transporting passengers. Which **approach direction of the cart** you are **more concerned** about when the cart is getting close to you? (you can select multiple choices)

- [ ] a) The cart getting close to you from your front
- [ ] b) The cart getting close to you from your behind
- [ ] c) The cart getting close to you from your left or right side
- [ ] d) Others

What are your thoughts about the operation of a driverless cart close to **children** in a shopping mall or an airport terminal?

In situations where the cart approaches or passes near a **child** in a shopping mall or an airport terminal, what type of behaviors from the cart can **help you feel safe** about the situation?

Appendix B

Letter of copyright permission

B.1 Copyright permission for reuse of Fig. 5.1(a) and Fig. 5.1(b)

Rightslink® by Copyright Clearance Center - Google Chrome

s100.copyright.com/AppDispatchServlet#formTop

CCC RightsLink

Sign in/Register

IEEE Requesting permission to reuse content from an IEEE publication

Top-view Trajectories: A Pedestrian Dataset of Vehicle-Crowd Interaction from Controlled Experiments and Crowded Campus

Conference Proceedings: 2019 IEEE Intelligent Vehicles Symposium (IV)

Author: Dongfang Yang

Publisher: IEEE

Date: June 2019

Copyright © 2019, IEEE

Thesis / Dissertation Reuse

The IEEE does not require individuals working on a thesis to obtain a formal reuse license, however, you may print out this statement to be used as a permission grant:

Requirements to be followed when using any portion (e.g., figure, graph, table, or textual material) of an IEEE copyrighted paper in a thesis:

1) In the case of textual material (e.g., using short quotes or referring to the work within these papers) users must give full credit to the original source (author, paper, publication) followed by the IEEE copyright line © 2011 IEEE.
2) In the case of illustrations or tabular material, we require that the copyright line © [Year of original publication] IEEE appear prominently with each reprinted figure and/or table.
3) If a substantial portion of the original paper is to be used, and if you are not the senior author, also obtain the senior author's approval.

Requirements to be followed when using an entire IEEE copyrighted paper in a thesis:

1) The following IEEE copyright/ credit notice should be placed prominently in the references: © [year of original publication] IEEE. Reprinted, with permission, from [author names, paper title, IEEE publication title, and month/year of publication]
2) Only the accepted version of an IEEE copyrighted paper can be used when posting the paper or your thesis on-line.
3) In placing the thesis on the author's university website, please display the following message in a prominent place on the website: In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of [university/educational entity's name goes here]'s products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a License from RightsLink.

If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

BACK

CLOSE WINDOW

Privacy - Terms

B.2 Copyright permission for reuse of Fig. 5.1(c) and 5.1(d)

Rightslink® by Copyright Clearance Center - Google Chrome

s100.copyright.com/AppDispatchServlet#formTop

CCC RightsLink

Sign in/Register

IEEE Requesting permission to reuse content from an IEEE publication

The inD Dataset: A Drone Dataset of Naturalistic Road User Trajectories at German Intersections

Conference Proceedings: 2020 IEEE Intelligent Vehicles Symposium (IV)

Author: Julian Bock

Publisher: IEEE

Date: 19 October 2020

Copyright © 2020, IEEE

Thesis / Dissertation Reuse

The IEEE does not require individuals working on a thesis to obtain a formal reuse license, however, you may print out this statement to be used as a permission grant:

Requirements to be followed when using any portion (e.g., figure, graph, table, or textual material) of an IEEE copyrighted paper in a thesis:

1) In the case of textual material (e.g., using short quotes or referring to the work within these papers) users must give full credit to the original source (author, paper, publication) followed by the IEEE copyright line © 2011 IEEE.
2) In the case of illustrations or tabular material, we require that the copyright line © [Year of original publication] IEEE appear prominently with each reprinted figure and/or table.
3) If a substantial portion of the original paper is to be used, and if you are not the senior author, also obtain the senior author's approval.

Requirements to be followed when using an entire IEEE copyrighted paper in a thesis:

1) The following IEEE copyright/ credit notice should be placed prominently in the references: © [year of original publication] IEEE. Reprinted, with permission, from [author names, paper title, IEEE publication title, and month/year of publication]
2) Only the accepted version of an IEEE copyrighted paper can be used when posting the paper or your thesis on-line.
3) In placing the thesis on the author's university website, please display the following message in a prominent place on the website: In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of [university/educational entity's name goes here]'s products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a License from RightsLink.

If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

BACK

CLOSE WINDOW

Privacy - Terms

B.3 Copyright permission for reuse of Fig. 5.1(e)

Rightslink® by Copyright Clearance Center - Google Chrome

s100.copyright.com/AppDispatchServlet#formTop

CCC RightsLink

Sign in/Register

IEEE Requesting permission to reuse content from an IEEE publication

Pedestrian Group Detection in Shared Space

Conference Proceedings: 2019 IEEE Intelligent Vehicles Symposium (IV)

Author: Hao Cheng

Publisher: IEEE

Date: June 2019

Copyright © 2019, IEEE

Thesis / Dissertation Reuse

The IEEE does not require individuals working on a thesis to obtain a formal reuse license, however, you may print out this statement to be used as a permission grant:

Requirements to be followed when using any portion (e.g., figure, graph, table, or textual material) of an IEEE copyrighted paper in a thesis:

1) In the case of textual material (e.g., using short quotes or referring to the work within these papers) users must give full credit to the original source (author, paper, publication) followed by the IEEE copyright line © 2011 IEEE.
2) In the case of illustrations or tabular material, we require that the copyright line © [Year of original publication] IEEE appear prominently with each reprinted figure and/or table.
3) If a substantial portion of the original paper is to be used, and if you are not the senior author, also obtain the senior author's approval.

Requirements to be followed when using an entire IEEE copyrighted paper in a thesis:

1) The following IEEE copyright/ credit notice should be placed prominently in the references: © [year of original publication] IEEE. Reprinted, with permission, from [author names, paper title, IEEE publication title, and month/year of publication]
2) Only the accepted version of an IEEE copyrighted paper can be used when posting the paper or your thesis on-line.
3) In placing the thesis on the author's university website, please display the following message in a prominent place on the website: In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of [university/educational entity's name goes here]'s products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a License from RightsLink.

If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

BACK

CLOSE WINDOW

Privacy - Terms

B.4 Copyright permission for reuse of Fig. 5.1(f) and Fig. 6.4

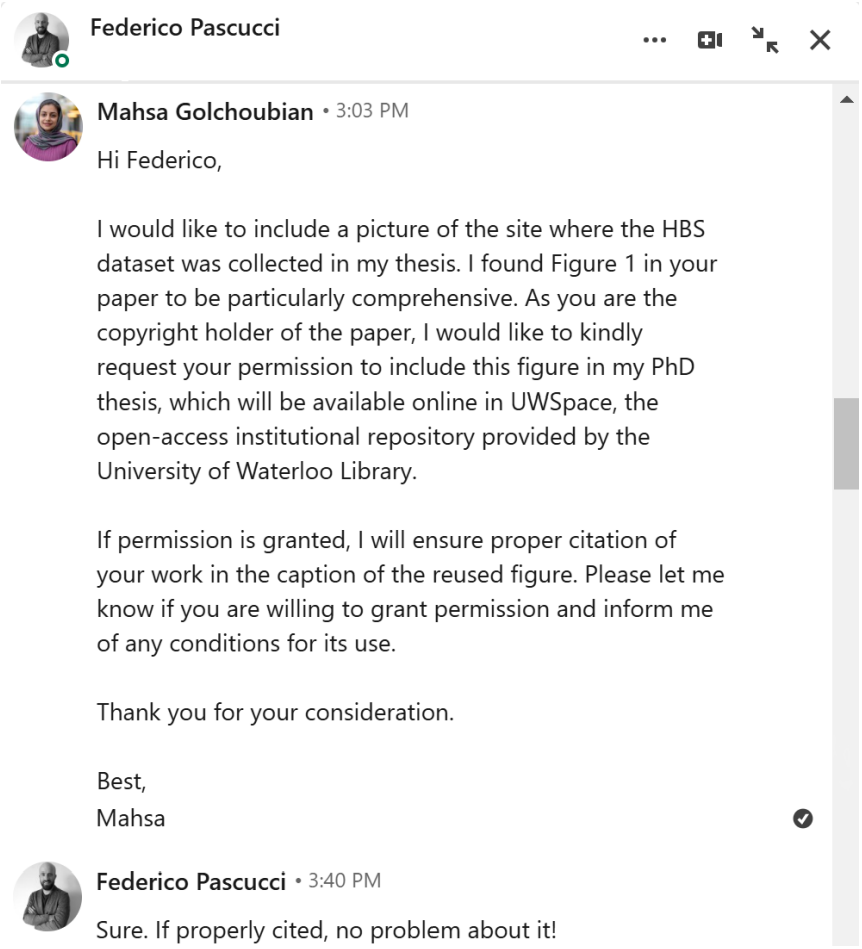

Federico Pascucci

Mahsa Golchoubian • 3:03 PM

Hi Federico,

I would like to include a picture of the site where the HBS dataset was collected in my thesis. I found Figure 1 in your paper to be particularly comprehensive. As you are the copyright holder of the paper, I would like to kindly request your permission to include this figure in my PhD thesis, which will be available online in UWSpace, the open-access institutional repository provided by the University of Waterloo Library.

If permission is granted, I will ensure proper citation of your work in the caption of the reused figure. Please let me know if you are willing to grant permission and inform me of any conditions for its use.

Thank you for your consideration.

Best,
Mahsa

Federico Pascucci • 3:40 PM

Sure. If properly cited, no problem about it!