

NONLINEAR OPINION DYNAMICS  
FOR SOCIAL ROBOT NAVIGATION:  
A DECISION-MAKING FRAMEWORK FOR SOCIALLY AWARE  
MOTION IN HUMAN-ROBOT INTERACTION

CHARLOTTE LOUISE CATHCART

A DISSERTATION  
PRESENTED TO THE FACULTY  
OF PRINCETON UNIVERSITY  
IN CANDIDACY FOR THE DEGREE  
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE  
BY THE DEPARTMENT OF  
MECHANICAL AND AEROSPACE ENGINEERING  
ADVISER: NAOMI EHRICH LEONARD

SEPTEMBER 2025

© Copyright by Charlotte Louise Cathcart, 2025.

All Rights Reserved.

## Abstract

Robots navigating spaces shared with humans must be capable of making fast, socially appropriate decisions about their movements while remaining adaptable to dynamic and uncertain environments. This dissertation presents a novel opinion-driven navigation framework for embodied, cooperative decision-making by a social robot. We focus on modeling the opinion formation such that opinions are *proactive* – swiftly prescribing robust socially aware motion (e.g., passing a person on the expected side) as soon as a robot is sensitive to a human – and *agile* – rapidly adapting in response to changing environmental cues – while remaining analytically tractable.

This work centers around an adaptation of the Nonlinear Opinion Dynamics (NOD) model by Bizyaeva et al. [1], originally formulated to describe collective decision-making in multi-agent systems, to a social navigation setting for human-robot interaction. We present this specialized NOD model, where opinions represent motion preferences (e.g., to turn left or right) and directly affect the robot’s steering control. We use bifurcation theory to analyze opinion formation behavior and guarantee deadlock-breaking of indecision in the embodied robot.

We improve the sequential decision-making capabilities of NOD by introducing excitable dynamics inspired by spiking neuron behavior to create the Spiking Nonlinear Opinion Dynamics (S-NOD) model. The spiking behavior of the opinion in S-NOD enables more responsive and flexible decision-making while maintaining the tractability of NOD. We present a geometric analysis of the S-NOD model, showing how the system’s bifurcation diagram and resultant spiking opinion behavior depend on model parameters.

We validate both NOD and S-NOD models with physical experiments of a mobile robot navigating around human movers. Qualitative results show each system’s decision-making agility and social awareness, and quantitative results demonstrate their efficiency and social intrusiveness. Comparative experiments reveal that S-NOD provides the robot with increased efficiency over NOD in more complex navigation scenarios.

This work connects social robot navigation, dynamical systems, and biologically inspired decision-making models to produce a simple yet powerful framework for human-aware robot navigation. The proposed models enable robots to move cooperatively around humans without relying on predictive tools or machine learning for motion planning, enabling more interpretable and computationally lightweight social navigation control.

## Acknowledgements

While a dissertation is a demonstration of a student's extensive research, I believe that mine is also a testament to the community that supported me through its preparation. I am grateful for this space to acknowledge all of those who made my journey possible, if not easier, over my six incredible years at Princeton.

I want to begin by expressing my utmost gratitude to my advisor, Naomi Ehrich Leonard. When I first started researching for potential graduate schools and advisors, Naomi's was truly the first name on my list, even though she was at a school that I never dreamed would have me. Fast-forwarding through my elation at admission, all the way to my welcome day as a first-year, I remember being asked what I was looking for in a PI. The most essential descriptor that I could come up with was "kind." I feel incredibly lucky to say that I found exactly what I needed by having Naomi as my mentor, teacher, and collaborator. Thank you for your faith in me, for taking the time to explain concepts and explore ideas with me, and for your patience as I found my footing throughout my degree. Your intelligence and creativity are nothing short of inspiring. I left every meeting with you feeling relieved from whatever anxieties I may have had, reinvigorated by the project plan, and excited to push forward. You have been a fantastic and inspiring mentor; I owe so much of my success and genuine enjoyment in my work to you. Thank you for everything.

The work of this dissertation would not be possible without the help of my collaborators. Thank you to Alessio Franci for proposing and spearheading the incorporation of excitability into my work - and for patiently helping me through understanding and analyzing it. It was so fun, and I have learned so much from you. Thank you to Shinkyu Park and María Santos for your advice, help, and patience as I wrote my first manuscript, and for all of your work helping set me up with the hardware for my experiments. I loved working with you both; thank you for never making me feel small or foolish as I stumbled through the early stages of my research. Finally, thank you to Ian Xul Belaustegui for all of his help shaping and explaining the theory in our work together. I am excited to see how much further you take it and have no doubt that with your immense capabilities, it will be amazing.

I thank my committee members, Jaime Fernández Fisac and Anirudha Majumdar, for all of their guidance over the years. They have sat on my committee through my general examination and every committee meeting since, and have even taught a few classes that I took while here. Thank you for being such approachable experts, for your ongoing interest in my work, and for all of the fantastic advice on various approaches and presentation methods that made my work better. Thank you for your decision to promote me to candidacy in this program – it feels full circle that it will be your decision that allows me to finish my degree.

Thank you to Shinkyu Park and Christine Allen-Blanchette for their time and effort in reading and

reviewing this entire document. It is no small ask to request that someone parse and evaluate such a dense collection of research, so I thank them both for their kindness in accepting such a task. Further, I thank them for their comments that strengthened the work contained in this dissertation and my own confidence.

I want to thank the fantastic professors and teachers who have taught me throughout my career. At Princeton, this includes Michael Littman, Luc Deike, Clarence Rowley, Alexander Glaser, Anirudha Majumdar, Naomi Leonard, Maximilian Eggl, Jaime Fernández Fisac, and Judith Swan. Princeton's scholarly value is high for good reason, and it is in no small part due to the talent of these professors. I've also had the privilege of serving as a teaching assistant for professors Andrej Košmrlj, Aimy Wissa, and Christine Allen-Blanchette. Teaching was a great joy of my Ph.D., made better by the interesting class material and these fantastic teachers I assisted. Further, thank you to all of the professors and teachers who nurtured my love of math, science, engineering, coding, and robots. That includes my undergraduate research advisor, Keith Brown, and mentor, Tommaso Ranzani, at Boston University. Thank you also to the recommendation letter writers who supported and aided my journey to graduate school, Kamil Ekinci and Paul Lipton (who also made the switch and joined Princeton when I did!). Thank you also to Dr. Tammy Scott for all of her advice and support – you made graduate school feel less overwhelming when I felt lost. And finally, to the many math and science teachers who inspired me over my adolescence. Thank you especially to my high school robotics teacher, Jay Lijoi, who noted my interest in bots and building, and was the first to gently suggest I look into studying mechanical engineering. Look at the journey that comment set me off on!

I am so grateful for my time spent as a part of the Leonard Lab, especially in the years following lockdown, during which I was lucky enough to make friends out of my colleagues. I've worked with and alongside this set of amazing, unbelievably intelligent researchers who are also fantastic, funny, lovely people. I mention here a non-exhaustive list of my Leonard Lab community, first with the postdocs that inspired me and somehow never made me feel lesser than, even though they were the picture of capability and proficiency that I hoped to become someday. Thank you to Shinkyu Park, Kenza Hamidouche, Christine Allen-Blanchette, María Santos, Juncal Arbelaitz, Himani Sinhmar, and Mallory Gaspard for any and all research discussions, presentations, conversations, and laughs we shared over my time here. To the many lab mates I've had the privilege of knowing and watching succeed, I'm so thankful to have had you as my coworkers. This includes Desmond Zhong, Yunxiu Zhou, Mari Kawakatsu, Anthony Savas, Anastasia Bizyaeva, Udari Madhushani, Justice Mason, Isla Xi Han, Justin Lidard, Giovanna Amorim, Daniel Shen, Kathryn Wantlin, Marcela Ordorica Arango, Ian Xul Belaustegui, and Daniel Levy. From presentations to puzzles, early mornings to all-nighters, Italy to Denver (or just the office to the robot room), I've had the best time with you all. Your commiseration in hard times and celebrations in the good made my time here easier, and I can't imagine how I could have gone through work without you all making my days brighter. I will dearly miss you all. So

many of you have been a massive part of my support system, so let this be a record that I promise to remain a part of yours! I am so excited to see where you all go and so proud to say that I was a part of this group.

The early stages of graduate school were marked by grueling classes, atrium meetings, and communal dining hall dinners with those in my first year. I want to acknowledge and thank my cohort members who were there with me through this transition into Princeton life, whether they remained companions through to the end of our dissertations or just for those late-night struggle sessions as we worked on problem sets. Thank you to Dennis Anthony, Tejas Dethe, Ian Gunady, Jihye Jeon, Eric Lepowsky, Marcel Louis, Yue Qiao, Sheetal Ramsurrun, Allen Ren, Wilson Ricks, James Roggeveen, Lena Sabidussi, David Snyder, Anvitha Sudhakar, and Wenbin Xu for being my first community at Princeton. As a later-stage dissertation defender, it has been an honor to watch so many of you succeed through your time at Princeton, see you become doctors, and move on to your next chapters. We got through a crazy time for graduate school together, and I'm thankful that I made it through alongside you all. I wish you all the best of luck in future endeavors as we splinter off, and I look forward to the great work I'm sure you will all accomplish.

I would like to thank my landlord, Garth Whiteside, Ph.D., for betting on my success when he took a chance on leasing his beautiful condo to me and Lena as young graduate students. He explicitly told us that he wanted to provide us a comfortable, safe place to stay as we navigated graduate school, and I had no idea how helpful it would be not to have to worry about my living situation for the years that followed. It was an incredibly generous and considerate gesture, and I thank him immensely for his kindness.

I would be remiss not to mention the friends I have made within Princeton who made my days and nights outside of campus that much better. Thank you to Gi Amorim for all of the laughter, food, gossip, and (required for conferences) sleepovers. I've had so much fun with you, and I love you dearly. You single-handedly served to improve so much of my experience at Princeton, and I'm incredibly lucky to have you as a friend. Thank you to Valeria Saro-Cortes and Hannah Wiswell for all of the Wawa snacks and dedicated-to-watching-TV nights that turned into talking-and-ignoring-the-TV hangouts with me and Lena. Thank you to Meghan Booker for all of the fun we got up to and for the nights in NYC. Thank you to Justin Lidard for dinners in the GC all the way to London, and for all the fun you, Gi, and I had as a trio in the lab. Thank you further to Justice Mason, Ian Xul Belaustegui, and Marcela Ordorica for keeping spirits high and always lending an ear in office to chat away. I'd often convince myself to get up and out of the house during the week with the hope of seeing you all in the office (with group puzzle breaks, of course). Finally, thank you to Emma Kruse and Allison Corchado for the giggles, picnics, gossip, and dog-sitting opportunities that I never knew could come out of a few stops to the admin office. I am lucky to have you all as friends.

I gratefully acknowledge the long-term support of my friends from home who kept me grounded. To Grace Fellman and Beth Hrehovcik, my best friends from kindergarten into adulthood, thank you for the

endless laughs and so much love throughout my life. I cherish the moments we have spent together since I've been in graduate school, be they socially-distanced nights on the porch or at Beth's beautiful wedding. You two have been my longest supporters and my most loving friends, and I thank you for keeping me company through so many stages of our lives. To Kim Kroupa, one of my Mechanigals, thank you for always inspiring me with your grace and success as you navigated your own doctorate program, and for the FaceTimes and visits that always brightened up my day (though truthfully, they were mostly nights). Thank you to Rachel Petherbridge for our classic five-hour phone calls and for always being someone I can rely on. I love you very much and am always in awe of your intelligence and beautiful heart, and of all of the work you do and the light you bring to everyone you meet. Finally, thank you to my friends Kris Costello and Jerami Goodwin for all of the laughs we have shared and for helping me break free from my work.

Thank you to Lena Sabidussi, my long-time best friend, for her companionship over the past decade of my life. Lena and I were classmates, then friends, and then capstone team members (long live the Mechanigals!) in the Mechanical Engineering department at Boston University during our college years, and then peers in MAE and roommates as we navigated time in Princeton together. Lena, arguably the luckiest thing that has happened to me was that I got to continue life alongside you through graduate school. I cannot picture how a single second of my time here would have gone had you not been here with me. You have been my first and favorite person to talk through just about everything with, especially when it came to grad school gripes or gossip. Your success was my motivation to keep going in my own work many times, and I'm so proud to be friends with such an intelligent, incredible person. Further, I would like to thank Lena's partner and my friend, Neil Sadhu, for all the laughs, food, support, and reality checks from the grad school bubble that he has given me over the years. Lena and Neil, you have been my family here. I can't quite picture my life without you right down the road from me, and I will try to avoid having to for as long as possible!

Thank you to my family, who have been my inspiration for academic achievements throughout my life. That includes my grandparents, Sam and Kathleen Sheber and Gerry Cathcart, who have cheered me on the whole way as I ventured into graduate school. Also, my grandfather Edgar, the first Dr. Cathcart there was, who was felt here deeply in spirit and I know would have been very proud of me. To my parents, Drs. Charles and Kathleen Cathcart, thank you for everything. For feeding my love of scholarly pursuits, for supporting me through my anxious perfectionism, and for genuinely never pushing me to be anything or anyone except for myself. You are the smartest people that I know, and I am lucky to be your daughter. Mom and Dad, you have always been my heroes, and a small silver lining to an otherwise horrific time was that the world finally caught up on seeing it too when you both braved the front lines of the pandemic – all while taking me and the rest of the kids back into the house. You are amazing. Thank you for my life, my siblings, and for being who you are. I love you both so much. Everything that I have ever done has only

been possible because I had your love and support with me at all times. This has always been for you.

I now want to acknowledge my favorite people on earth, my siblings. To Sam, Lizzie, Kadie, Fiona, and Eamon, you are the people I am most excited to talk about and to. I love you all dearly, and any time spent texting, talking, or seeing you has made my life fuller and my time away from home easier. It boggles my mind that I somehow look back fondly on the disaster of the world through the lockdown of 2020, but I know it's because I got to move back home with all of you. Returning to family dinners, movie nights, and casual cohabiting made for some of the best times of my adult life that I never would have guessed I would experience again after college. I love you all so much more than I could adequately put into words. Thank you for always believing in me. Further, thank you to my sister-in-law, Becca, for being the older sister I never got to have and a hugely important part of my finding confidence as I navigated through this Ph.D. And, of course, thank you for giving me my new favorite person – my niece, Phoebe – who reminds me that the world outside of graduate school is big, open, and fun. My favorite job is, always has been, and always will be being your sister, and now I have been lucky enough to be an aunt.

Finally, to my partner Tom Papa. Thank you for all of the laughter, snack runs, trips to New Hope, walks in the park, and for sleeping the weekends away, keeping me company after conferences, and getting into reality TV that has been the cause of so much joy in my life for the past two years. Mostly, thank you for loving me and for sticking this out with me. My life has only gotten better with you in it. I love you and could not have had a more supportive partner to vent to and center me throughout this writing process. And thank you for the 5am reminders to go to sleep when I got too caught up in work. Here's to a much more relaxed future together now that I do not have this dissertation to complain about!

I've had such an amazing time at Princeton. Thank you for having me, and thank you to everyone who helped me along the journey. I can only hope that I have made you proud.

I am grateful for the funding provided through the ONR grant N00014-19-1-2556 and Princeton University, through the generosity of Lydia and William Addy '82, for allowing me to conduct this research. Thank you for supporting my academic endeavors. This dissertation carries T#3485 in the records of the Department of Mechanical and Aerospace Engineering.

To all the Drs. Cathcart (and O'Brien-Sheber) before me.

# Contents

Abstract . . . . .	iii
Acknowledgements . . . . .	iv
List of Tables . . . . .	xiv
List of Figures . . . . .	xvi
<b>I Opinion-Driven Navigation: Nonlinear Decision-Making in Social Robots</b>	<b>1</b>
<b>1 Introduction</b>	<b>2</b>
1.1 Overview and motivation . . . . .	2
1.1.1 Social robot navigation . . . . .	2
1.1.2 Decision-making in dynamic contexts . . . . .	5
1.1.3 Excitability and spiking models . . . . .	6
1.2 Outline of contributions . . . . .	7
<b>2 Background</b>	<b>10</b>
2.1 Nonlinear system analysis . . . . .	10
2.1.1 Overview and linearization . . . . .	10
2.1.2 Phase plane geometric analysis . . . . .	11
2.1.3 Bifurcations . . . . .	14
Pitchfork bifurcation . . . . .	14
2.1.3.1 Supercritical . . . . .	15
2.1.3.2 Subcritical . . . . .	16
2.1.3.3 Subcritical - extended . . . . .	20
Hopf bifurcation . . . . .	23

2.2	Nonlinear Opinion Dynamics (NOD) model	25
2.3	Hardware for experimental demonstration	28
2.3.1	Jackal UGV research robot	28
2.3.2	Vicon motion capture system	28
<b>3</b>	<b>Nonlinear Opinion Dynamics (NOD) for proactive social robot navigation</b>	<b>30</b>
3.1	Motivations	30
3.1.1	Navigation environment and notation	31
3.2	NOD specialized for social robot navigation	31
3.2.1	Physical interpretation of opinion $z_i$	32
3.2.2	Opinion-driven motion control	33
3.2.3	Model refinements	35
3.3	Guarantee on deadlock-free navigation	36
3.4	Experimental validations of NOD	37
3.4.1	General experimental setup	39
3.4.2	Results	40
3.4.2.1	Metrics reported	40
3.4.2.2	Single-robot, single-human experiments	41
3.4.2.3	Single-robot, two-human experiments	47
<b>4</b>	<b>Spiking Nonlinear Opinion Dynamics (S-NOD) for agile decision-making</b>	<b>52</b>
	Statement of contributions	52
	Abstract	53
4.1	Introduction	53
4.2	Fast and flexible decision-making: NOD	55
4.2.1	NOD for single decision-maker and two options	55
4.2.2	Analysis of single-agent, two-option NOD	56
4.2.3	Limitation on tuning of NOD	58
4.3	Agile decision-making: S-NOD	60
4.3.1	S-NOD for a single agent and two options	60
4.3.2	Geometric analysis of single-agent, two-option S-NOD	60
4.4	Agile multi-agent decision-making: S-NOD	63
4.4.1	S-NOD for multiple agents and two options	63
4.4.2	Application to social robot navigation	64

4.5	Final remarks . . . . .	65
4.6	Appendix . . . . .	66
<b>5</b>	<b>Spiking Nonlinear Opinion Dynamics (S-NOD) for social robot navigation</b>	<b>68</b>
5.1	S-NOD specialized for social robot navigation . . . . .	68
5.2	Experimental validations of S-NOD . . . . .	70
5.2.1	General experimental setup . . . . .	71
5.2.2	Results . . . . .	71
5.2.2.1	Metrics reported . . . . .	71
5.2.2.2	Single-robot, single-human experiments . . . . .	72
	S-NOD 1 experiment set . . . . .	73
	S-NOD 2 experiment set . . . . .	79
	Compared results of NOD, S-NOD 1, and S-NOD 2 experiment sets . . . . .	84
5.2.2.3	Single-robot, two-human experiments . . . . .	90
	Comparisons of NOD and S-NOD performance . . . . .	95
5.3	Discussion of effect of $K_u$ and $U$ on navigation . . . . .	97
5.4	Simulated multi-human navigation: NOD vs. S-NOD performance . . . . .	105
<b>6</b>	<b>Conclusions</b>	<b>109</b>
6.1	Final remarks . . . . .	109
6.2	Future work . . . . .	111
<b>II</b>	<b>Published work</b>	<b>114</b>
<b>7</b>	<b>Proactive opinion-driven robot navigation around human movers</b>	<b>115</b>
	Statement of contributions . . . . .	115
	Abstract . . . . .	116
7.1	Introduction . . . . .	116
7.2	Nonlinear Opinion Dynamics in social navigation . . . . .	119
7.2.1	Nonlinear Opinion Dynamics model . . . . .	120
7.2.2	Dynamic model for opinion-driven robot navigation . . . . .	120
7.3	Guarantee on deadlock-free navigation . . . . .	122
7.4	Experiments . . . . .	124
7.4.1	Validation of flexibility of the approach . . . . .	125

7.4.1.1	Experimental setup	125
7.4.1.2	Results	125
7.4.2	Validation of the deadlock breaking	126
7.4.2.1	Experimental setup	126
7.4.2.2	Results	126
7.5	Discussion and final remarks	128
	<b>Bibliography</b>	<b>129</b>

# List of Tables

3.1	Geometry of the human-robot navigation setting. . . . .	32
3.2	Performance metrics and their definitions. . . . .	40
3.3	Performance metrics of the robot navigating in trials 1-4 (reported in Figures 3.10- 3.13). . .	51
5.1	Updated performance metrics and their definitions. . . . .	72
5.2	Comparison of average navigation time (NT) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/4$ . . . . .	84
5.3	Comparison of average navigation time (NT) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/6$ . . . . .	85
5.4	Comparison of average path length (PL) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/4$ . . . . .	85
5.5	Comparison of average path length (PL) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/6$ . . . . .	85
5.6	Comparison of average centerline deviation (CD) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/4$ . . . . .	86
5.7	Comparison of average centerline deviation (CD) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/6$ . . . . .	86
5.8	Comparison of average minimum distance (MD) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/4$ . . . . .	87
5.9	Comparison of average minimum distance (MD) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/6$ . . . . .	87
5.10	Comparison of average invasion distance (ID) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/4$ . . . . .	88
5.11	Comparison of average invasion distance (ID) across each experimental case in NOD and S-NOD experiments where $\beta = \pi/6$ . . . . .	88

5.12 Comparison of average control input magnitude $\ z\ $ across each experimental case in NOD and S-NOD experiments where $\beta = \pi/4$ . . . . .	89
5.13 Comparison of average control input magnitude $\ z\ $ across each experimental case in NOD and S-NOD experiments where $\beta = \pi/6$ . . . . .	89
5.14 Performance metrics of the single-robot, two-human trials 1-4 of NOD and S-NOD robots, and additionally of trial 5 for the S-NOD robots. . . . .	96
5.15 Performance metrics of the simulated navigation paths illustrated in Fig. 5.19 taken by NOD and S-NOD controlled robots across various values of $U$ and $K_u$ as the robot interacts with one human. . . . .	99
5.16 Performance metrics of the simulated navigation paths illustrated in Fig. 5.20 taken by NOD and S-NOD controlled robots across various values of $U$ and $K_u$ as the robot interacts with two humans. . . . .	103
5.17 Performance metrics for the NOD and S-NOD robots along the trajectories illustrated in Fig. 5.21. . . . .	106

# List of Figures

2.1	Phase planes and sample trajectories of the FitzHugh-Nagumo model (2.4). . . . .	13
2.2	Supercritical pitchfork bifurcation. . . . .	17
2.3	Unfolded supercritical pitchfork bifurcation diagram. . . . .	17
2.4	Subcritical pitchfork bifurcation. . . . .	19
2.5	Unfolded subcritical pitchfork bifurcation diagram. . . . .	19
2.6	Subcritical pitchfork extended with a stabilizing fifth-order term. . . . .	22
2.7	Unfolded subcritical pitchfork bifurcation extended with a stabilizing fifth-order term. . . .	22
2.8	Phase planes and sample trajectories with a supercritical Hopf bifurcation. . . . .	24
2.9	Phase planes and sample trajectories with a subcritical Hopf bifurcation. . . . .	25
3.1	Geometry and notation of the human-robot navigation setting. . . . .	32
3.2	The $u_{0_{ie}}$ function, constructed as an inverse hill function. . . . .	35
3.3	Deadlock breaking for $u_0 > u_0^*$ in a robot's opinion dynamics. . . . .	38
3.4	The trajectory data of a NOD robot with different $\beta$ values and one human for fifty-four trials. .	42
3.5	The average navigation times of a NOD robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. . . . .	43
3.6	The average path lengths of a NOD robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. . . . .	43
3.7	The average centerline deviations of a NOD robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. . . . .	44
3.8	The average minimum distances between a NOD robot and a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. . . . .	46
3.9	The average invasion distances by a NOD robot to a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. . . . .	46
3.10	NOD Trial 1 with two humans. . . . .	49

3.11 NOD Trial 2 with two humans. . . . .	49
3.12 NOD Trial 3 with two humans. . . . .	50
3.13 NOD Trial 4 with two humans. . . . .	50
4.1 Social navigation with NOD vs. S-NOD robot. . . . .	54
4.2 The effect of $K_u$ on the bifurcation diagram of (4.2) and the cubic and quintic terms of (4.3). . . . .	58
4.3 Opinion solutions of NOD and S-NOD over time and associated bifurcation diagrams. . . . .	59
4.4 The system solutions and $(u_s, z)$ phase portrait as the basal sensitivity $u_0$ increases. . . . .	62
4.5 S-NOD solutions over time of three agents. . . . .	64
4.6 Trajectories of social robots using S-NOD to navigate around approaching human movers. . . . .	64
4.7 An extended presentation of Fig. 4.3. . . . .	67
5.1 The trajectory data of a S-NOD 1 robot with different $\beta$ values and one human for fifty-four trials. . . . .	74
5.2 The average navigation times of a S-NOD 1 robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. . . . .	75
5.3 The average path lengths of a S-NOD 1 robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. . . . .	76
5.4 The average centerline deviations of a S-NOD 1 robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. . . . .	77
5.5 The average minimum distances between a S-NOD 1 robot and a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. . . . .	78
5.6 The average invasion distances by a S-NOD 1 robot to a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. . . . .	78
5.7 The trajectory data of a S-NOD 2 robot with different $\beta$ values and one human for fifty-four trials. . . . .	80
5.8 The average navigation times of a S-NOD 2 robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. . . . .	81
5.9 The average path lengths of a S-NOD 2 robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. . . . .	81
5.10 The average centerline deviations of a S-NOD 2 robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. . . . .	82
5.11 The average minimum distances between a S-NOD 2 robot and a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. . . . .	83

5.12 The average invasion distances between a S-NOD 2 robot and a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. . . . .	84
5.13 S-NOD Trial 1 with two humans. . . . .	91
5.14 S-NOD Trial 2 with two humans. . . . .	92
5.15 S-NOD Trial 3 with two humans. . . . .	93
5.16 S-NOD Trial 4 with two humans. . . . .	94
5.17 S-NOD Trial 5 with two humans. . . . .	95
5.18 Trajectory comparisons of NOD and S-NOD Trials 1-4 with two humans. . . . .	96
5.19 Trajectory and opinion comparisons between NOD and S-NOD social robots with various $K_u$ and $U$ values around a single stationary human. . . . .	100
5.20 Trajectory and opinion comparisons between NOD and S-NOD social robots with various $K_u$ and $U$ values around two humans. . . . .	102
5.21 Simulations of NOD and S-NOD social robots navigating around five humans. . . . .	107
7.1 A robot using opinion-driven navigation to pass two humans. . . . .	117
7.2 Multiple experimental trials with two humans and a robot using NOD. . . . .	119
7.3 An illustration of notation for human-robot passing. . . . .	121
7.4 Analysis of deadlock breaking in the robot's opinion dynamics. . . . .	123
7.5 The trajectory data of NOD robot with different $\beta$ values and one human for forty-five trials. . . . .	127
7.6 Percent increase of the robot's path length for various $\beta$ values for each of the nine configurations. . . . .	128
7.7 Average minimum distance between the robot and human for each of the nine configurations. .	129

## Part I

# Opinion-Driven Navigation: Nonlinear Decision-Making in Social Robots

# Chapter 1

## Introduction

### 1.1 Overview and motivation

Robots operating and navigating in environments with humans must be capable of making rapid and reliable decisions about how to act, even in the presence of noise or rapidly changing contexts. There exists an entire field of robotics research on these physical human-robot interaction (HRI) scenarios, aiming to identify the most effective, efficient, and reliable methods for producing successful and safe robot actions. Many turn to biology as an inspiration for its naturally occurring agile behavior. In this work, we develop navigation algorithms that utilize mathematical models of two key characteristics of animal behavior: their capacity to form consensus in decision-making within a group, and their ability to generate rapid, robust, yet flexible responses to environmental changes. Our algorithms provide reliable, adaptable, and natural navigation behavior around human movers with low levels of computational power. This section provides an overview of relevant works on each of these topics. Subsequent chapters present our models of Nonlinear Opinion Dynamics (NOD) and Spiking Nonlinear Opinion Dynamics (S-NOD), along with their detailed analysis and validation through simulation and physical experiments.

#### 1.1.1 Social robot navigation

A significant amount of research exists on the development of navigation algorithms to facilitate successful interaction between humans and robots, balancing motion planning that is efficient, reliable, safe, and legible. In this work, we utilize *opinion dynamics* to quantify and compute the collaborative turning preference (be it towards the left or right, and by how much) of a robot that passes a human. This process is highly reactive to its environment, robust to uncertainty, relatively lightweight in terms of sensor requirements,

computationally efficient, and proactive in its decision-making, anticipating and avoiding a collision when it becomes imminent. Our algorithms provide a compact and powerful framework that is novel in its ability to generate this analytically tractable reactive, robust, and proactive behavior. We contextualize our approach and its features within recent and relevant works in the field.

Numerous surveys of the field provide comprehensive overviews of the behavioral goals and navigation approaches of various social robot algorithms [2–4]. These works emphasize the challenge of achieving legible, safe, and socially compliant motion in real-world environments with changing uncertainty. They also highlight the diverse human behavior modeling strategies in use, ranging from minimally sensor-dependent methods that rely on observing a human’s orientation and walking speed, to more sensor and computationally complex approaches that consider human gaze and attention cues. Some methods take a more data-driven approach by training robots on datasets of human-in-the-loop navigation. These reviews categorize and compare methods within the core paradigms of robot navigation behavior, including reactive, predictive, learned, and hybrid approaches. We adopt this framing in this discussion, beginning with the reactive robot approaches, which are most relevant to the work of this dissertation. We then discuss predictive motion planning strategies, which generally introduce greater computational demands to output various paths forward, especially when they incorporate learning-based methods that model human behavior.

Several reactive navigation methods are relevant to our work. Most comparable to our minimally sensor-dependent method is that of the Social Momentum framework of Mavrogiannis et al. [5]. They propose a reactive robot that uses only the geometry of an observed human to gauge and align with their angular momentum, turning in a cooperative fashion through a passing scenario. While this approach shares a similar spirit with the opinion-driven method we present, it organizes decision-making within if/else statements in its algorithm and does not guarantee the resolution of decision-making in the face of low information. Our work utilizes a continuous set of equations that enable predictable and intuitive decision-making, while also providing analytic guarantees on the system’s behavior and ability to form decisions even in environments where no decision is clearly best. More broadly, the Social Force Model (SFM) [6] sets the foundational framework for many other reactive social robot algorithms. SFM represents human motion as that of particles drawn towards their goals by an attractive force and deterred from colliding with other pedestrians by repellent forces. While this provides a simple physics approach and an intuitive lens through which to approach the problem, it struggles to find success in modeling human movement in high-density crowds. Many works incorporate this premise into more sophisticated models capable of more complex planning. Reddy et al. [7] present a hybrid model that combines a potential field into its computation, adding another social force to the SFM framework and increasing the robot’s proactivity. While this is primarily reactive, elements of predictive path planning are still present, which increase its complexity. Kamezaki et al. [8] build

upon the SFM by inducing reactivity only when humans are observed to engage with the robot. Interestingly, this paper equates proactivity in robot movement with the prediction of human motion goals. In contrast, we define proactivity through the lens of opinion dynamics: robust, rapid decision-making that incorporates the observed preferred passing side of a human yields cooperative behavior without requiring explicit forecasting or probabilistic modeling.

Predictive methods of social robot navigation use environmental cues to generate a set of prospective forward motion plans, weighted by the likelihood of human cooperation. These approaches build on historical data of the navigation interaction, predicting future states of the human, and planning the best course forward to interact with the most likely next steps. However, their performance is only as sound as the approach’s prediction or the human behavior it uses. Many rely on machine learning methods to model human behavior. While powerful, these approaches are computationally expensive and often sacrifice interpretability in their decision-making. Given the prevalence of learning-based methods in the field of HRI navigation, we focus our discussion on a few representative examples. Kretzschmar et al. [9] present an algorithm that utilizes machine learning to fit model parameters to best match the navigation behavior in a given set of human trajectories in various navigation scenarios. This results in a trained model that considers how a human indicates that they will pass on the robot’s left or right and incorporates the variance of many other likelihoods of its next steps, allowing it to generalize to new settings. However, this process requires many computational steps to analyze, fit, and model human behavior, predict multiple forward paths the robot may take, and optimize these predicted paths against various design constraints. To reduce the costly computation, the authors assume a well-initialized trajectory to reduce the convergence time of the algorithm. Chen et al. [10] admit that the modeling of human behavior for direct replication by a social robot is challenging and instead defer to having the robot learn based on a reward function that penalizes breaking social conventions. They specifically utilize deep reinforcement learning to develop their navigation policy, avoiding explicit human modeling. Neural networks can already sacrifice interpretability of their output due to their hidden layers; the deep neural networks used within this work further obscure how specific observations lead to particular navigation decisions. Okal and Arras [11] similarly present an algorithm to prompt the robot to follow social norms, but instead utilize Bayesian inverse reinforcement learning (BIRL) to model its behavior accordingly. The robot’s reward function is similarly built upon demonstrations of socially appropriate trajectories, but is more interpretable as its features (such as maintaining a distance from humans and approximating heading directions) are hand-crafted. Still, its results depend on the quality and level of representation of relevant navigation scenarios within the demonstration data. While this line of physical HRI navigation produces exciting results, our work prides itself on its intuitive, tractable, and low-cost models for social robot navigation, which still result in legible, cooperative, and efficient motion.

A problem widely considered in social robot navigation, included within each of the works mentioned above, is that of *decision deadlock* by the robot, also known as the *freezing robot* problem. As defined in [12], this describes the navigation scenario in which all paths forward become indistinguishable because they are all equally unsafe, leaving the robot unable to commit to a forward direction and effectively paralyzed. Many works, some mentioned above, utilize stops or if/else statements to prompt the robot out of its indecision, thereby removing it from a deadlock state. However, the results of [13] showed that a robot specifically designed to signal and seek cooperation with interacting humans was able to avoid freezing. Their findings, supported by surveys in [14], indicated that humans feel safer and prefer when robots move in a human-like fashion, moving cooperatively and at a constant speed. The opinion dynamics model used in this dissertation achieves precisely that: it enables proactive, socially cooperative navigation with analytically guaranteed deadlock-breaking to resolve motion indecision.

### 1.1.2 Decision-making in dynamic contexts

Models of decision-making aim to capture the mechanisms of information exchange, the influence of neighbor interactions, and the evolution of opinions prevalent across the social world through mathematical rules. They require rapid formation and adaptive adjustment of decisions in response to new information. Opinion dynamics address the same challenge as seen in the countermeasures for the freezing robot problem, aiming to avoid indecision or delays in decision-making when options must be chosen quickly. There is an extensive history of opinion dynamic formulations, with the work of [15] providing an overview of previous opinion dynamics modeling within the field. This survey predates the publication of the Bizyaeva et al. NOD model [1], but highlights a similar formulation called the Baumann model [16] in which an agent's opinion state develops according to negative and saturated, nonlinear positive feedback. However, this model and others from its time lack the tunable sensitivity that is crucial to the flexibility of the NOD model.

We highlight some recent works that incorporate opinion dynamics into robotic decision-making, demonstrating the value of such models for producing collective behavior. In Montes de Oca et al. [17], an opinion dynamics model is used for task allocation decisions within a decentralized swarm of robots, showing that a preference for the fastest executable task consistently and reliably spreads through the team. Salem et al. [18] propose a “gossip”-based mechanism of opinion dynamics for a robot team, in which a single random robot at each timestep shares information to influence the opinion of a single nearby robot within a group. Even with this slow incorporation of neighbor opinion, the opinion dynamics ultimately create predictable levels of consensus within the group across various neighbor graph structures. In Li et al. [19], opinion dynamics are ruled by a model that incorporates higher-order derivatives to predict complex dynamics in

multi-robot systems. While this method reports higher precision in its model of interaction between robots, it comes with more expensive mathematical computation. Collectively, these approaches demonstrate that opinion dynamics provide a valuable framework for modeling coordination and decision-making in robots.

The work of this dissertation builds on the Nonlinear Opinion Dynamics (NOD) model presented in [1], which draws inspiration from biological sources such as honeybee cooperation [20,21]. In subsequent chapters, we propose an approach in which a robot uses environmental context and responsive sensitivity to rapidly form an *opinion* that represents the strength of its preferred passing side — left or right, and by how much — for navigating around an oncoming human. This opinion is then used to control the robot’s motion directly. The use of a NOD model creates analytical tractability of the system. It reveals that a neutral opinion, indicating indecision in motion, becomes unstable when a parameter of sensitivity exceeds a predictable implicit threshold. Further work [22, 23] revealed the NOD model’s ability to prompt and manage opinion cascades and opinion switching in a multi-agent system, generalizing to decision-making over multiple options [24]. The generality and benefits of the model have made it attractive for various applications in recent years. Many works have adopted this formulation to describe robot motion in human-robot interaction [25–27], learn parameterization links between passing behaviors in an autonomous racing scenario [28], and govern the behavior of real unmanned surface vessels between options to explore, exploit, or migrate in their environment [29].

This dissertation adapts the NOD model to enable proactive social navigation and then introduces increased responsiveness and flexibility by incorporating excitable behavior into that model. This increases the system’s agility while preserving NOD’s soundness in decision modeling and control.

### 1.1.3 Excitability and spiking models

The fast, flexible, and adaptive behavior observed in biology can be credited to the excitable nature of cellular signaling [30–33]. Mathematical models of excitability seek to represent the spiking response to stimuli seen in analog biophysical processes. The foundational FitzHugh-Nagumo model [34] provides a simplified form of the Hodgkin-Huxley equations [35] that describe the spiking activation of neurons to current changes in their membrane. Spiking control systems [36] and neuromorphic engineering [37] build control policies from these excitable models that combine the adaptability of continuous-time analog systems with the reliability of discrete-event signaling. These are characterized by on-demand, event-based spikes that occur even in the presence of small-scale inputs so long as the system is sufficiently close to its excitability threshold.

Despite their advantages and applications to control systems, we note that these existing models describe only single-input/single-output [38] spike-based signal processing. This binary activity encodes a binary

decision—to spike or not to spike—based on the input signals that push the system toward its excitability threshold [39]. This restricts current spiking models from being applied to decision-making involving multiple options. In scenarios with many options, such as neurons tuned to specific visual orientations in the cortex [40], or decision-making in sensorimotor control tasks [41], excitatory behaviors must be superimposed with other mechanisms to produce more complex outputs.

The primary contribution of this dissertation concerning excitability lies in the development and validation of *two*-option spiking in a two-dimensional, two-timescale model to create an agile system that is fast, flexible, and adaptive. We call this model Spiking Nonlinear Opinion Dynamics (S-NOD) as it derives from the aforementioned Nonlinear Opinion Dynamics (NOD) framework in [1, 42]. With the introduction of a single additional term inspired by the dynamics of excitable systems, S-NOD equips the NOD model with self-regulating, spiking decision-making that can occur in as many directions as there are options. We apply this extended framework to the social robot navigation problem and leverage S-NOD to enable efficient and agile movement while preserving the proactive motion generation of NOD.

## 1.2 Outline of contributions

We detail the development of NOD and S-NOD models as applied to social robot navigation in subsequent chapters. Motivated by the open questions raised in the discussion above, we present the contributions listed below in this dissertation.

**Part I** presents the core contributions of this thesis and contains the following sections:

1. In **Chapter 2**, we present key theoretical tools and context for the mathematical concepts discussed within this dissertation. We begin with a review of bifurcations, equilibrium solutions, and stability within a nonlinear dynamical system. We present and discuss the normal forms of fundamental bifurcations that are leveraged within this work: the supercritical pitchfork, the subcritical pitchfork, the extended normal form of the subcritical pitchfork, and the Hopf. We then introduce the general Nonlinear Opinion Dynamics (NOD) model presented by Bizyaeva et al. [1] that the work of this dissertation adapts for social robot navigation. We conclude the chapter with information about the hardware used for the experimental demonstrations presented in later chapters.
2. In **Chapter 3**, we motivate and present a specialized form of NOD for social robot navigation. We begin by defining the working environment of the robot and then present the adapted form of the NOD equations to implement the robot’s opinion as an element of its navigation control. We clarify

the refinements made to this model, which were updated from previously published iterations of the work. We present an analysis that proves the guarantee of deadlock-free navigation when controlled with NOD, supplemented with illustrations and simulations of the robot navigation. We conclude this chapter with new experimental validations of this deadlock-free guarantee and the flexibility of this navigation approach in various passing scenarios. We report the detailed experimental setups to be discussed in later chapters and define the metrics to be used for performance reviews of the robot’s behavior.

3. In **Chapter 4**, we build on the NOD model within the previous chapter and introduce the Spiking Nonlinear Opinion Dynamics (S-NOD) model for increased agility in decision making. We motivate this work with the biologically inspired excitable behavior in spiking systems and apply a slow-negative feedback to the NOD equations to restore the system to a state of sensitivity, enabling spiking decision-making. We present a formal analysis of a system with a single agent and two options, uncovering and discussing how the shape of its pitchfork bifurcation depends on the model parameters. To the best of our knowledge, this is the first such model to generalize spiking to more than one option. The limitations of NOD are discussed, which motivates the development of S-NOD. We present a geometric analysis of the phase plane of the S-NOD system, showcasing the spiking behavior and potential limit cycles that exist within the system. We extend these results to a multi-agent setting and discuss the synchronization and anti-synchronization behaviors possible. We conclude with an application of S-NOD to a simulated social robot navigation setting, following the work of NOD. This Chapter was based on the peer-reviewed publication of [26] in IEEE Control Systems Letters (L-CSS) and has been lightly revised for presentation in this thesis.
4. In **Chapter 5**, we extend the S-NOD application to physical experiments with humans and more than two option settings. We begin by introducing the specialized equations for S-NOD in the context of social robot navigation. We then validate the deadlock-breaking guarantee and flexibility of the approach, following the same experimental protocol as with NOD. The behaviors of the robot in the S-NOD and NOD experiments are compared, and the differences in control efficiency between them are discussed. We discuss the effect of key parameters with the NOD and S-NOD navigation behaviors, highlighting the perceived limitations of the sequential decision-making capabilities in the NOD performance. These are further revealed in a comparative simulation with NOD and S-NOD as they navigate a larger, more crowded navigation setting.
5. In **Chapter 6**, we conclude this part of the dissertation with an overview of the results of this work, identify its limitations, and suggest future pathways and extensions.

**Part II** presents an unmodified published work to provide a complete record of the original research that contextualizes the contributions of Part I.

6. In **Chapter 7**, we present the peer-reviewed publication of [25] in the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2023). In this work, we present the original formulation of the NOD model for social robot navigation. We present an analysis of this model to demonstrate the guarantee of deadlock avoidance in the robot's decision-making. We conclude the paper with experiments that validate this deadlock-breaking guarantee and the flexibility of the original formulation of the NOD approach for social robot navigation. The experimental results are analyzed and reveal the trade-off between navigation efficiency and spatial intrusiveness present within the original formulation of NOD when varying a single parameter. This work strongly motivated the subsequent reformulation of NOD and its experimental validations in Chapters 3 and 5.

# Chapter 2

## Background

The work detailed in this thesis focuses on the modeling of a nonlinear dynamical system for social robot navigation and its application to real-world experiments. We provide an analysis of behavioral guarantees offered with the model using bifurcation analysis. Therefore, we provide background and context here, including bifurcation theory, phase plane visualizations, and fundamental forms of relevant bifurcations in nonlinear models, as well as an overview of the opinion dynamics model employed in this thesis. We conclude with a description of the equipment used for social robot navigation experiments.

### 2.1 Nonlinear system analysis

This section is prepared based on definitions and analysis methods documented in various textbooks [43–46].

#### 2.1.1 Overview and linearization

A *dynamical system* is a system of equations describing the evolution of a set of state variables over time. Time can be considered as either a discrete or continuous variable in the system. Dynamical systems that evolve continuously in time are described by either ordinary differential equations (ODEs) or partial differential equations (PDEs), the latter of which being generally more difficult to analyze. The work of this thesis considers the development, analysis, and application of a continuous-in-time dynamical system of the form of a general ODE:

$$\frac{dz}{dt} = \dot{z} = f(z), \quad (2.1)$$

where  $z(t) \in \mathbb{R}^n$  depends on time  $t$ . When the vector field  $f(z)$  has no explicit time  $t$  dependency, we call this an *autonomous* system. If the state variable  $z$  has only one dimension, the ODE is a *scalar* differential

equation. We are interested in understanding how a dynamical system evolves from an initial state  $z(t_0) = z_0$  at  $t = t_0$ .

An *equilibrium solution*  $z^{eq}$  of the system satisfies  $\dot{z} = 0 = f(z^{eq})$  and corresponds to a steady-state. Identifying equilibrium solutions and performing a *stability analysis* of the equilibrium solutions allows us to determine the system's long-term behavior. For example, it can help us predict conditions under which the system will converge to an equilibrium, oscillate around a solution, or exhibit other, more complex behavior.

In our work, we analyze *nonlinear* ODEs that describe a dynamical system. Nonlinearity is intrinsic to most models of complex natural behavior, but its analysis can be substantially more challenging than its linear counterpart. However, we gain insight into local system behavior and utilize tools from local bifurcation theory, which involves a Taylor series expansion of the dynamics about a singular equilibrium solution.

To study the stability of equilibria away from the singular equilibrium solution, we can use linear stability analysis by linearizing the nonlinear ODE about a nonsingular equilibrium point. The linearized dynamics are  $\dot{y} = J(z^{eq})y$  where  $J(z^{eq})$  is the Jacobian evaluated at the equilibrium point  $z^{eq}$ . The Jacobian at an equilibrium point  $z^{eq}$  is defined by

$$J(z^{eq}) = \left. \frac{\partial \dot{z}}{\partial z} \right|_{z=z^{eq}}. \quad (2.2)$$

The eigenvalues  $\lambda$  of this Jacobian matrix determine the local system stability at the equilibrium point  $z = z^{eq}$ . If  $\text{Re}(\lambda_i) < 0 \ \forall i = 1, \dots, N$  the equilibrium is *locally asymptotically stable*. If any eigenvalue has a positive real part, that is, if  $\exists \text{Re}(\lambda_i) > 0$ , the equilibrium is *unstable*. We next introduce *phase plane* methods, which offer an intuitive and geometric visualization of the evolution of solution trajectories over time, and can qualitatively capture the system's global behavior.

### 2.1.2 Phase plane geometric analysis

A *phase portrait* can be used to visualize the global behavior in a system of two or more dynamic equations. They are constructed by considering information about two equations of the ODE. For convenience, we note these as dynamic equations  $\dot{z}_1$  and  $\dot{z}_2$ . The phase portrait includes lines that represent the zero-level set of an  $\dot{z}_1 = 0$  and  $\dot{z}_2 = 0$ , vector fields representing the values of  $\dot{z}_1(z_1, z_2)$  and  $\dot{z}_2(z_1, z_2)$  throughout the plane, and solution trajectories of  $z_1, z_2$  following these trajectory fields. The resultant diagram can be used to assess the equilibria and stability within the system and identify the presence of limit cycles.

We step through the elements of a phase portrait below with an illustrative example of the FitzHugh-Nagumo model for excitable behavior discussed in Chapter 1.1.3. The equations for this dynamical system are

$$\frac{dz}{dt} = \frac{d}{dt} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} \dot{z}_1 \\ \dot{z}_2 \end{bmatrix} = \begin{bmatrix} f_1(z_1, z_2) \\ f_2(z_1, z_2) \end{bmatrix}, \quad (2.3)$$

where  $z_1(t), z_2(t) \in \mathbb{R}$  and the vector fields of each are described by  $f_1(z_1, z_2)$ , and  $f_2(z_1, z_2)$  respectively. Note that this is an explicitly two-dimensional form of the general ODE (2.1).

The equilibrium points of the system describe where all state derivatives are zero, i.e.,  $\dot{z}_1 = \dot{z}_2 = 0$ . With a phase plane, we can locate these points geometrically while simultaneously observing the system's global behavior by examining its *nullclines*, which are the curves in the phase plane where one component of the vector field is a steady-state. That is, the  $z_1$ -nullcline is defined by solutions to  $\dot{z}_1 = 0 = f_1(z_1, z_2)$ , and the  $z_2$ -nullcline is defined by solutions to  $\dot{z}_2 = 0 = f_2(z_1, z_2)$ . These nullclines can be sketched on the  $z_1 - z_2$  plane, and the system equilibria can be found where the curves intersect.

To complete the phase plane, we consider the direction of the vector field at various points within the plane. That is, we evaluate the velocity vector  $(\dot{z}_1, \dot{z}_2) = (f_1(z_1, z_2), f_2(z_1, z_2))$  throughout the  $z_1 - z_2$  plane. Arrows are drawn to indicate the direction and magnitude at each point in the plane.

From the completed sketch of the phase plane, the behavior of the system can be observed with sketches of trajectories, which reveal the stability of equilibria. Trajectories represent the evolution of the states  $z_1, z_2$  over time. Consider an initial solution of the system at  $t = t_0$  defined by  $(z_1, z_2)_0 = (z_1(t_0), z_2(t_0))$ . This can be plotted as a point on the phase plane. From this point, we can follow the direction of the vector field to sketch the shape of how the solution evolves from this initial point. This completes the phase portrait. If we observe that the system's trajectories diverge from an equilibrium point, we can classify its behavior as unstable. Conversely, if the trajectories are attracted and trapped in an equilibrium solution, we can classify its behavior as stable.

We present a representative phase portrait. To illustrate its structure and its application in stability interpretation. This phase plane represents that of the excitable FitzHugh-Nagumo model [34], a key motivation for the work within this dissertation. The following equations describe this model:

$$\frac{dz}{dt} = \frac{d}{dt} \begin{bmatrix} z_1 = v \\ z_2 = w \end{bmatrix} = \begin{bmatrix} \dot{v} \\ \dot{w} \end{bmatrix} = \begin{bmatrix} v - v^3/3 - w + I \\ c(v + a - bw) \end{bmatrix}, \quad (2.4)$$

where the states are  $z = [z_1, z_2]^T = [v, w]^T \in \mathbb{R}^2$  and parameters are  $a, b, c$ , and  $I$ . See Fig. 2.1 for three

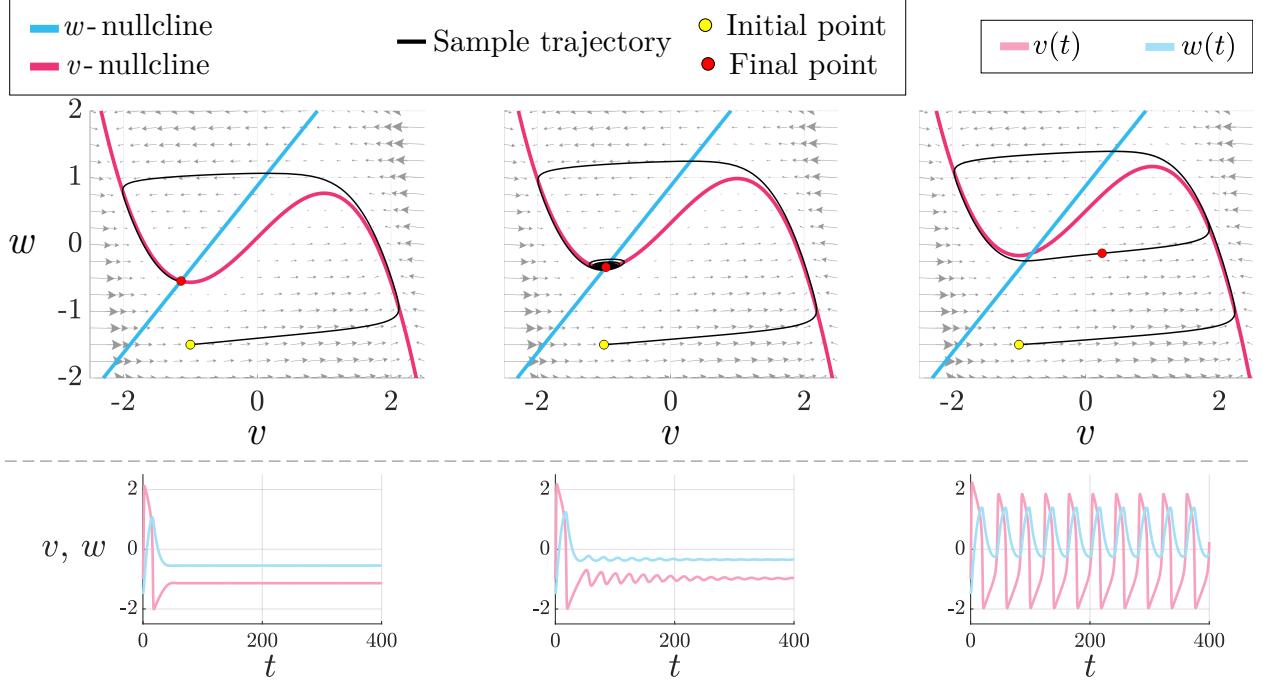


Figure 2.1: Phase planes and sample trajectories of the FitzHugh-Nagumo model (2.4). (Top) The  $v - w$  phase plane. The  $v$ - and  $w$ -nullclines are drawn in pink and blue, respectively. A sample trajectory is superimposed with a black trace that moves from the same initial point (yellow dot) to its final point (red dot) in a simulation. (Bottom) The  $v$  and  $w$  solutions over time that correspond with the sample trajectory shown in the top graphs. Parameters are:  $a = 0.7$ ,  $b = 0.8$ ,  $c = 12.5$  and  $I = \{0.1, 0.32, 0.5\}$  from left to right graphs.

phase planes shown for the same parameters  $a, b$ , and  $c$ , but different current input value  $I$ .

As stipulated, the phase planes feature the nullclines of each state variable, with a vector field shown with gray arrows indicating the flow about the system. An equilibrium point exists at the single point where the nullclines intersect, but the behavior of the solution trajectory varies depending on the changing  $I$  value. In the first, leftmost phase plane, we observe that a sample trajectory moves through the system's flow but is directly attracted to the equilibrium point; therefore, we can conclude that the equilibrium is stable. In the middle image, we highlight that the vector field has changed due to the increased value of  $I$ . The solution trajectory follows a similar initial path to the first example, but the solution oscillates about the equilibrium, is quickly damped, and then settles. Therefore, we can conclude that the equilibrium is stable, although the system exhibits oscillatory behavior. In the last, rightmost phase plane, we observe that a sample trajectory moves through the system's flow near and past the equilibrium point, becoming locked in a steady cycle; therefore, we can conclude that the system exhibits a limit cycle.

A phase plane with solution tracing (i.e., a phase portrait) provides a geometric analysis that qualitatively explains the behavior of a dynamical system. However, it only captures the behavior of a system at fixed parameter values. In the above example illustration, changing the input  $I$  caused drastically different

behaviors within the system's solutions. However, a single phase plane could not have revealed the varied potential behaviors of the ODE. This motivates the subsequent discussion on theoretical tools to uncover and study the parameter values that cause these qualitative changes, known as *bifurcations*.

### 2.1.3 Bifurcations

Some dynamical systems can behave in qualitatively distinct ways depending on a system parameter  $p$ . To study these qualitative changes in behavior, we can consider what happens when we smoothly vary a parameter  $p$  in a dynamical system (2.1). Loosely speaking, if a stark qualitative behavior change occurs as  $p$  is varied, the system is said to have undergone a *bifurcation*. This can happen when equilibrium solutions are created or destroyed, equilibria gain or lose stability, oscillatory responses emerge or disappear, or any combination of these behaviors. Bifurcation theory focuses on identifying these points in the parameter space where even a small perturbation of  $p$  can significantly alter the system's response.

Bifurcation points can be identified by analyzing how the eigenvalues  $\lambda_i$  of the Jacobian matrix at an equilibrium solution  $z = z^{eq}$  depend on a parameter  $p$ . That is, a bifurcation may occur at  $z = z^{eq}$  if there exists a parameter value  $p = p^*$  such that  $Re(\lambda_i(p^*)) = 0$  for at least one eigenvalue. The point  $(z, p) = (z^{eq}, p^*)$  is then referred to as a *singular point*. In the case that the singular point is a *pitchfork bifurcation*, to be defined below, the equilibrium point  $z^{eq}$  transitions between stability and instability as  $p$  passes through  $p^*$ . We can visualize these and other transitions with a *bifurcation diagram*, which graphs the solutions  $z = z^{eq}$  of the dynamical system against the varied system parameter  $p$  on the  $(z, p)$ -plane. Stable solutions are typically shown with solid lines, while unstable solutions are shown with dotted lines. These provide a valuable tool for identifying regions of changing numbers of solutions and stability within the system.

The work in this thesis is primarily focused on the existence and study of two fundamental types of bifurcations: pitchfork and Hopf bifurcations. Each of these types of bifurcation is accompanied by characteristic variations, which are described in detail below.

## Pitchfork bifurcation

The pitchfork bifurcation describes a symmetric system that undergoes a transition from a single equilibrium point to three equilibria as a system parameter  $p$  is varied, producing a distinctive pitchfork shape in its bifurcation diagram. The general normal form of a pitchfork bifurcation is

$$\dot{z} = p z - \epsilon q z^3, \quad (2.5)$$

where  $z$  is the system state,  $p \in \mathbb{R}$  is the bifurcation parameter of note, and  $\epsilon \in [-1, 1]$  assigns the sign to the positive scaling factor  $q > 0$ . Depending on the sign of  $\epsilon$ , the pitchfork bifurcation's characteristic behavior of the stability of its one and three equilibria changes.

The point at which the system transitions between one and three equilibria is called the *bifurcation point*, also noted as the critical (singular) point of the system. Consider the behavior of the system as the parameter  $p$  increases through this critical point. Two characteristic behaviors may occur at this bifurcation, depending on the sign of  $\epsilon$  in (2.5). A *supercritical* pitchfork bifurcation occurs with  $\epsilon = 1$  and describes the creation and existence of a pair of stable equilibria beyond (super) the bifurcation (critical) point. A *subcritical* pitchfork bifurcation occurs with  $\epsilon = -1$  and describes the destruction and existence of a pair of stable equilibria before (sub) the bifurcation (critical) point.

There is another variable that can be introduced that perturbs the pitchfork bifurcation, which we call a *bias*  $b \in \mathbb{R}$ . Assume that  $b$  is small. Adding  $b$  to the normal form (2.5), we get a more general form:

$$\dot{z} = p z - \epsilon q z^3 + b. \quad (2.6)$$

When  $b = 0$ , this recaptures the special case of the normal form (2.5). However, when  $b \neq 0$ , this term acts to break the symmetry of the system and skews the solution branches toward the sign of this parameter. This peels the solution branches apart from the bifurcation point, and the pitchfork is said to *unfold*.

We present the characteristic forms and solutions of the normal pitchfork bifurcation types and conclude each with an illustration of its associated unfolded diagrams.

### 2.1.3.1 Supercritical

The supercritical pitchfork can be described by (2.5) with  $\epsilon = 1$ , thus

$$\dot{z} = p z - q z^3. \quad (2.7)$$

We now consider the solutions for the zero-level set of the equations, where  $\dot{z} = 0$ , that describe the system's equilibria. This is described by  $0 = p z - q z^3 = z(p - q z^2)$ . Thus, we see that  $z_0^{eq} = 0$  and  $z_{1,2}^{eq} = \pm\sqrt{\frac{p}{q}}$  are these solutions. Note that the equilibrium  $z_0^{eq} = 0$  exists for all parameter values  $p$ ; thus, it is always an equilibrium point. However,  $z_{1,2}^{eq} \in \mathbb{R}$  only when  $p \geq 0$ , and grow in magnitude as  $p$  increases. At  $p = 0$ , a single solution  $z_0^{eq} = 0$  exists.

To analyze the stability of the equilibria, we first compute the Jacobian  $J(z, p) = \partial\dot{z}/\partial z = p - 3qz^2$ . Evaluating the Jacobian at the equilibrium  $z^{eq} = 0$  we get  $J(z = 0, p) = p$ , indicating that the stability of

the origin changes as  $p$  crosses zero. The equilibrium  $z^{eq} = 0$  is singular when  $p = p^* = 0$ , and this point  $(z, p) = (z^{eq}, p^*) = (0, 0)$  marks the pitchfork bifurcation. We can use the Jacobian to study the stability of  $z = 0$  for  $p \neq 0$  and the other pair of equilibria points for  $p > 0$ . However, because  $J(0, 0) = 0$ , the linearization fails to determine the stability of the bifurcation point. In this case, we can use a Lyapunov function — a scalar function that decreases along the system's trajectories — to assess stability without relying on linearization. Note that when  $p = 0$  the dynamics reduce to  $\dot{z} = -qz^3$ , therefore a Lyapunov function like  $V(z) = z^2$  yields  $\dot{V} = 2z\dot{z} = -2qz^4 \leq 0$ . This indicates that the singular point  $z = 0$  is asymptotically stable, and all trajectories will converge to this point.

When  $p < 0$ , only  $z_0^{eq} = 0$  exists. At this point,  $J(z_0^{eq}, p) = J(0, p) = p$ , and as  $p < 0$ , we see  $J(z_0^{eq}, p) < 0$ . Thus,  $z = 0$  is globally exponentially stable for  $p < 0$ .

When  $p > 0$ , three equilibria  $z_{0,1,2}^{eq}$  exist. We begin by considering the same equilibrium as above, at  $z_0^{eq} = 0$ . At this point,  $J(z_0^{eq}, p) = J(0, p) = p$ , and as  $p > 0$ , we see  $J(z_0^{eq}, p) > 0$ . Thus,  $z = 0$  is unstable for  $p > 0$ . We now consider the two other equilibria,  $z_1^{eq}$  and  $z_2^{eq}$ . These points describe solutions at equal and opposite values  $|z| = \sqrt{p/q}$ , so we refer to them together with a slight abuse of notation as  $z_{1,2}^{eq}$ . At these points,  $J(z_{1,2}^{eq}, p) = J(\pm\sqrt{p/q}, p) = p - 3q(\pm\sqrt{p/q})^2 = -2p$ , and as  $p > 0$ , we see  $J(z_{1,2}^{eq}, p) < 0$ . Thus,  $z = \pm\sqrt{p/q}$  are locally exponentially stable for  $p > 0$ .

Putting this information all together, the characteristic bifurcation diagram of a supercritical pitchfork can be seen in Fig. 2.2, which visually captures the information described above.

We also present the representative bifurcation diagrams of an *unfolded* supercritical pitchfork bifurcation in Fig. 2.3. This illustrates the bifurcation of the biased supercritical pitchfork described by

$$\dot{z} = p z - q z^3 + b, \quad (2.8)$$

with  $b \in \mathbb{R}$ . When  $b > 0$ , the upper stable solution branch is favored. When  $b < 0$ , the lower stable solution branch is favored. Note that for higher magnitudes of bias  $b$ , the previously discussed critical point (at  $p = 0$ ) appears to shift to the left and completely disappear with sufficiently high bias.

### 2.1.3.2 Subcritical

The subcritical pitchfork can be described by (2.5) with  $\epsilon = -1$ , thus

$$\dot{z} = p z + q z^3. \quad (2.9)$$

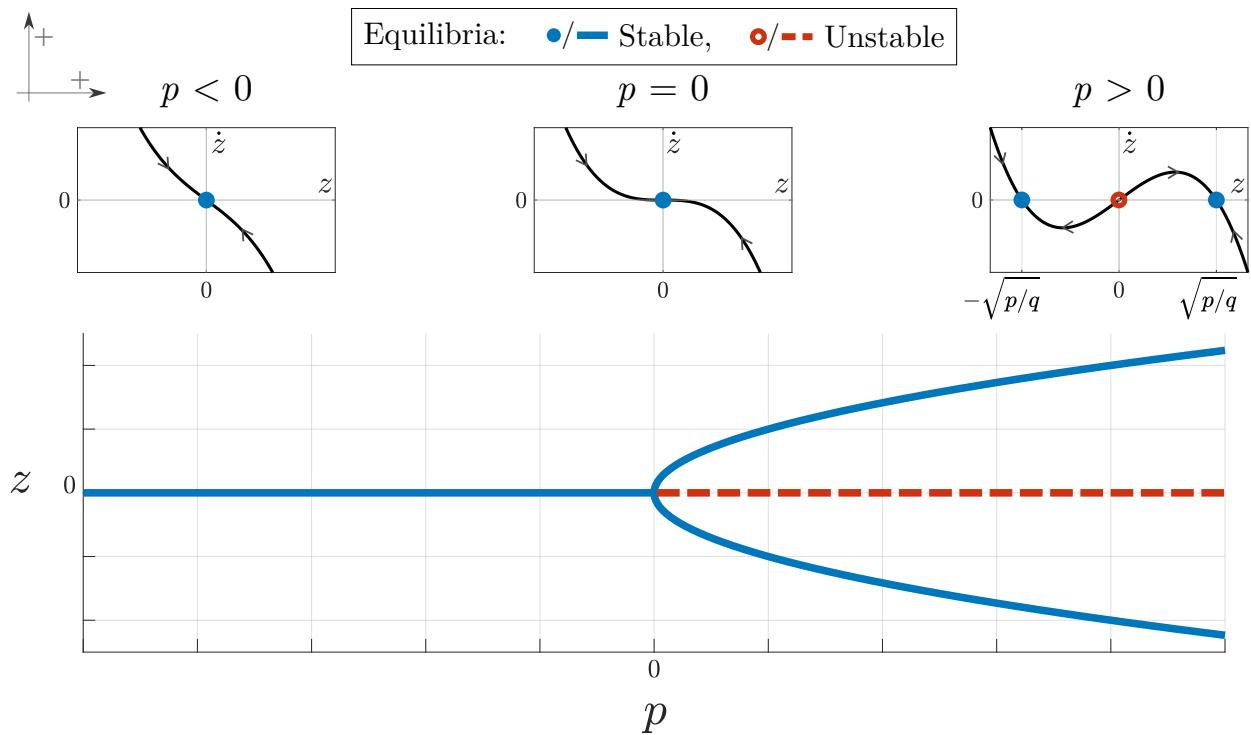


Figure 2.2: Supercritical pitchfork bifurcation. Top: Function  $\dot{z}$ . Bottom: Bifurcation diagram of  $z$  solutions against parameter  $p$ .

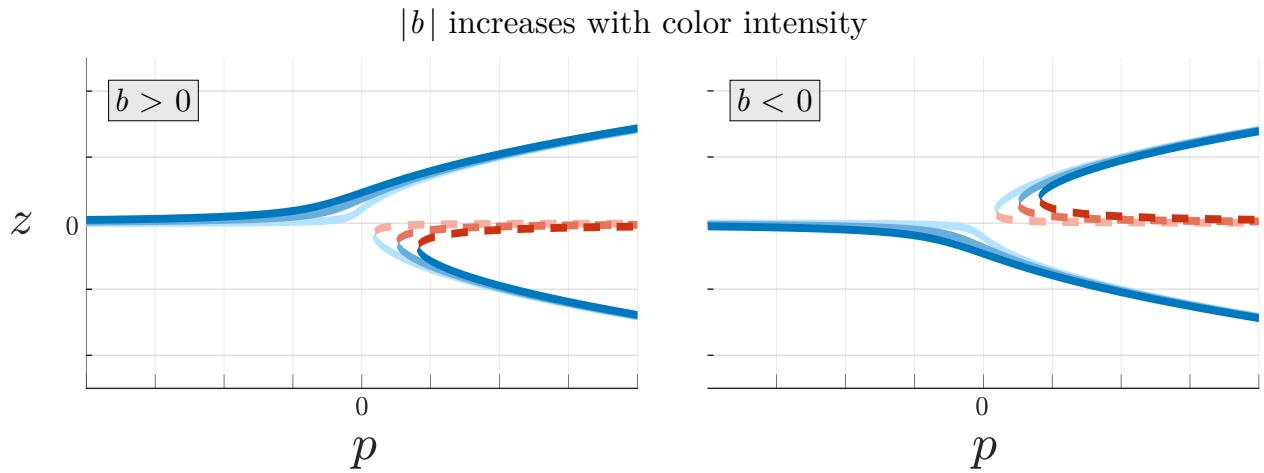


Figure 2.3: Unfolded supercritical pitchfork bifurcation diagram by bias  $b$ . (Left) Positive bias,  $b > 0$ . (Right) Negative bias,  $b < 0$ . Line features to illustrate stability are as in Fig. 2.2, with colors of higher magnitude  $b$ .

We now consider the solutions for the zero-level set of the equations, where  $\dot{z} = 0$ , that describe the system's equilibria. This is described by  $0 = p z + q z^3 = z(p + q z^2)$ . Thus, we see that  $z_0^{eq} = 0$  and  $z_{1,2}^{eq} = \pm i \sqrt{\frac{p}{q}}$  are these solutions. Note that the equilibrium  $z_0^{eq} = 0$  exists for all parameter values  $p$ ; thus, it is always an equilibrium point. However,  $z_{1,2}^{eq} \in \mathbb{R}$  only when  $p < 0$ , and grow in magnitude as  $|p|$  increases. At  $p = 0$ , a single solution  $z_0^{eq} = 0$  exists.

To analyze the stability of the equilibria, we first compute the Jacobian  $J(z, p) = \partial \dot{z} / \partial z = p + 3qz^2$ . Evaluating the Jacobian at the equilibrium  $z^{eq} = 0$  we get  $J(z = 0, p) = p$ , indicating that the stability of the origin changes as  $p$  crosses zero. The equilibrium  $z^{eq} = 0$  is singular when  $p = p^* = 0$ , and this point  $(z, p) = (z^{eq}, p^*) = (0, 0)$  marks the pitchfork bifurcation. We can use the Jacobian to study the stability of  $z = 0$  for  $p \neq 0$  and the other pair of equilibria points for  $p < 0$ . However, because  $J(0, 0) = 0$ , the linearization fails to determine the stability of the bifurcation point. As before, we can use a Lyapunov function to assess stability without relying on linearization. Note that when  $p = 0$  the dynamics reduce to  $\dot{z} = qz^3$ , therefore a Lyapunov function like  $V(z) = z^2$  yields  $\dot{V} = 2z\dot{z} = 2qz^4 \geq 0$ . This indicates that the singular point  $z = 0$  is unstable, and all trajectories will diverge from this point.

When  $p > 0$ , only  $z_0^{eq} = 0$  exists. At this point,  $J(z_0^{eq}, p) = J(0, p) = p$ , and as  $p > 0$ , we see  $J(z_0^{eq}, p) > 0$ . Thus,  $z = 0$  is unstable for  $p > 0$ .

When  $p < 0$ , three equilibria  $z_{0,1,2}^{eq}$  exist. We begin by considering the same equilibrium as above, at  $z_0^{eq} = 0$ . At this point,  $J(z_0^{eq}, p) = J(0, p) = p$ , and as  $p < 0$ , we see  $J(z_0^{eq}, p) < 0$ . Thus,  $z = 0$  is locally stable for  $p < 0$ . We now consider the two other equilibria,  $z_1^{eq}$  and  $z_2^{eq}$ . These points describe solutions at equal and opposite values  $|z| = \sqrt{|p/q|}$ , so we refer to them together with a slight abuse of notation as  $z_{1,2}^{eq}$ . At these points,  $J(z_{1,2}^{eq}, p) = J(\pm i \sqrt{|p/q|}, p) = p - 3q(\pm i \sqrt{|p/q|})^2 = -2p$ , and as  $p < 0$ , we see  $J(z_{1,2}^{eq}, p) > 0$ . Thus,  $z = \pm i \sqrt{|p/q|}$  are unstable for  $p < 0$ .

Putting this information all together, the characteristic bifurcation diagram of a subcritical pitchfork can be seen in Fig. 2.4, which visually captures the information described above.

We also show the subcritical pitchfork unfolded by bias  $b \in \mathbb{R}$  with

$$\dot{z} = p z + q z^3 + b. \quad (2.10)$$

Fig. 2.5 shows the unfolding effect of varying values of  $b$  on the subcritical pitchfork. The single stable solution branch moves farther from the previously discussed critical point (at  $p = 0$ ), and disappears as the magnitude of the bias increases.

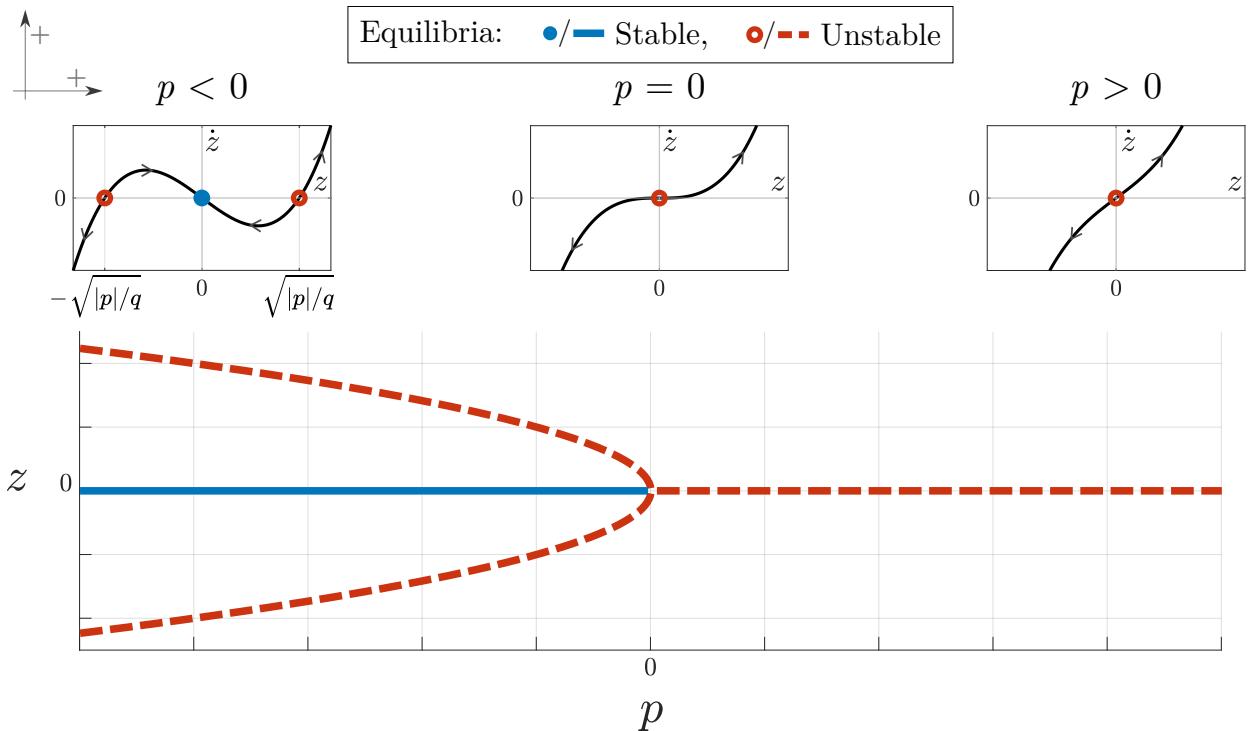


Figure 2.4: Subcritical pitchfork bifurcation. Top: Function  $\dot{z}$ . Bottom: Bifurcation diagram of  $z$  solutions against parameter  $p$ .

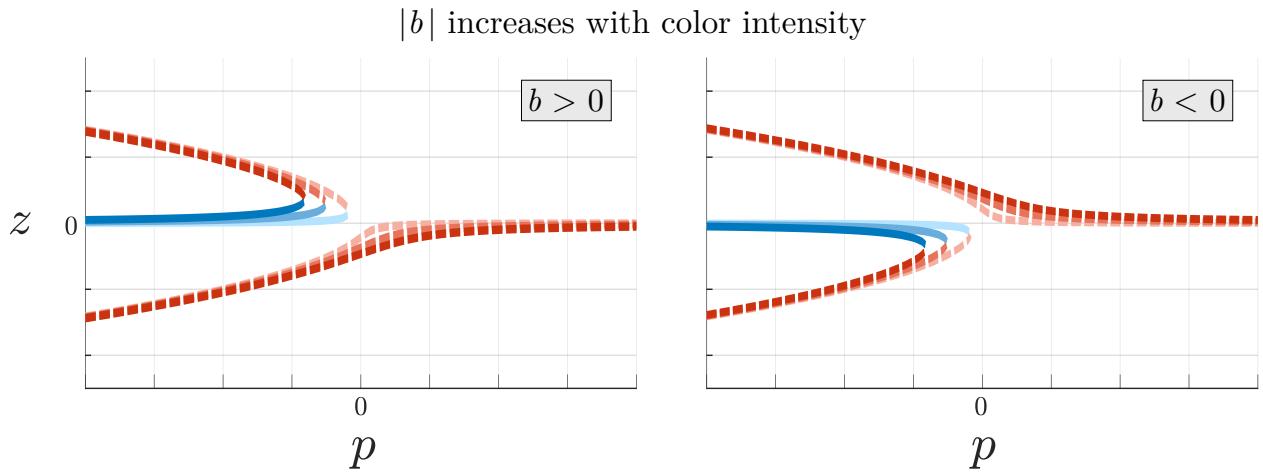


Figure 2.5: Unfolded subcritical pitchfork bifurcation diagram by bias  $b$ . (Left) Positive bias,  $b > 0$ . (Right) Negative bias,  $b < 0$ . Line features to illustrate stability are as in Fig. 2.4, with colors of higher intensity associated with higher magnitude  $b$ .

### 2.1.3.3 Subcritical - extended

The subcritical pitchfork defined by equation (2.9) defines a dynamical system in which solutions can grow unbounded due to the unrestrained  $z^3$  term. While this is a useful normal form for contrasting the supercritical pitchfork, it does not accurately reflect the behavior of most physical systems, where typically some form of saturation or other constraint limits growth. To model this bounded behavior, a higher-order stabilizing term can be added to the normal form. Preserving the symmetry structure in the pitchfork normal form, the next appropriate term is of order  $z^5$ . Therefore, the *extended* normal form of the subcritical pitchfork with a fifth-order stabilizing term is described by

$$\dot{z} = p z + q z^3 - r z^5, \quad (2.11)$$

with parameter  $r > 0$ .

We now consider the solutions for the zero-level set of the equations, where  $\dot{z} = 0$ , that describe the system's equilibria. This is described by  $0 = p z + q z^3 - r z^5 = z(p + q z^2 - r z^4)$ . Thus, we see that  $z_0^{eq} = 0$  and solutions to  $(p + q z^2 - r z^4) = 0$  are the equilibria. We solve the function  $(p + q z^2 - r z^4) = 0$  using the quadratic equation and the substitution method. This yields four further solutions, written here in shorthand pairs as  $z_1^*$  and  $z_2^*$ . Real-valued solutions  $z_1^* = \pm \sqrt{(q + \sqrt{q^2 + 4rp}) / (2r)}$  exist only when  $p > -q^2/4r$  and grow in magnitude as  $p$  increases. Real-valued solutions  $z_2^* = \pm \sqrt{(q - \sqrt{q^2 + 4rp}) / (2r)}$  exist only when  $-q^2/4r < p < 0$  and grow in magnitude as  $|p|$  increases. Note that the equilibrium  $z_0^{eq} = 0$  exists for all parameter values  $p$ ; thus, it is always an equilibrium point. At  $p = 0$ , a single solution  $z_0^{eq} = 0$  exists.

To analyze the stability of the equilibria, we first compute the Jacobian  $J(z, p) = \partial \dot{z} / \partial z = p + 3qz^2 - 5rz^4$ . Evaluating the Jacobian at the equilibrium  $z^{eq} = 0$  we get  $J(z = 0, p) = p$ , indicating that the stability of the origin changes as  $p$  crosses zero. The equilibrium  $z^{eq} = 0$  is singular when  $p = p^* = 0$ , and this point  $(z, p) = (z^{eq}, p^*) = (0, 0)$  marks the pitchfork bifurcation. We can use the Jacobian to study the stability of  $z = 0$  for  $p \neq 0$  and the equilibria points for  $p > -q^2/4r$ . However, because  $J(0, 0) = 0$ , the linearization fails to determine the stability of the bifurcation point. In this case, we can use a Lyapunov function — a scalar function that decreases along the system's trajectories — to assess stability without relying on linearization. Note that when  $p = 0$  the dynamics reduce to  $\dot{z} = qz^3 - rz^5$ , therefore a Lyapunov function like  $V(z) = z^2$  yields  $\dot{V} = 2z\dot{z} = 2z^4(q - 2rz^2)$ . Note that near  $z = 0$ , the  $q > 0$  term dominates the smaller  $2rz^2$  term in  $\dot{V}$ , thus  $\dot{V} > 0$  near  $z = 0$ . This indicates that the singular point  $z = 0$  is unstable, and all trajectories will diverge from this point.

When  $p < -q^2/4r$ , only  $z_0^{eq} = 0$  exists. At this point,  $J(z_0^{eq}, p) = J(0, p) = p$ , and as  $p < 0$ , we see

$J(z_0^{eq}, p) < 0$ . Thus,  $z = 0$  is locally asymptotically stable for  $p < -q^2/4r$ .

When  $-q^2/4r < p < 0$ , five equilibria  $z_0^{eq}, z_1^*, z_2^*$  exist where  $z_1^*$  and  $z_2^*$  each describe equal and opposite pairs of solutions. We begin by considering the same equilibrium as above, at  $z_0^{eq} = 0$ . At this point,  $J(z_0^{eq}, p) = J(0, p) = p$ , and as  $p < 0$ , we see  $J(z_0^{eq}, p) < 0$ . Thus,  $z = 0$  is locally asymptotically stable for all  $p < 0$ . At the equilibrium points described by  $z_1^*$ , we evaluate the Jacobian  $J(z_1^*, p)$ . We save the reader from the complicated resultant algebraic function and announce its solution, that  $J(z_1^*, p) < 0$  with  $-q^2/4r < p < 0$ . Thus, both of the equilibria described by  $z_1^*$  are locally asymptotically stable. Similarly, at the equilibrium points described by  $z_2^*$ , we evaluate the Jacobian  $J(z_2^*, p)$ . We again save the reader from the complicated resultant algebraic function and announce its solution, that  $J(z_2^*, p) > 0$  with  $-q^2/4r < p < 0$ . Thus, both of the equilibria described by  $z_2^*$  are unstable.

When  $p > 0$ , three equilibria  $z_0^{eq}$  and  $z_1^*$  exist. We begin by considering the equilibrium at  $z_0^{eq} = 0$ . At this point,  $J(z_0^{eq}, p) = J(0, p) = p$ , and as  $p > 0$ , we see  $J(z_0^{eq}, p) > 0$ . Thus,  $z = 0$  is unstable for  $p > 0$ . We now consider the two other equilibria,  $z_1^*$ . At the equilibrium points described by  $z_1^*$ , we evaluate the Jacobian  $J(z_1^*, p)$ . We again save the reader from the complicated resultant algebraic function and announce its solution, that  $J(z_1^*, p) < 0$  with  $p > 0$ . Thus, both of the equilibria described by  $z_1^*$  are locally asymptotically stable for  $p > 0$ .

Putting this information all together, the characteristic bifurcation diagram of a subcritical pitchfork, extended with an additional stabilizing term, can be seen in Fig. 2.6, which visually captures the information described above. Notice that the bifurcation diagram at the bottom of Fig. 2.6 appears to be a merging of the supercritical and subcritical bifurcation diagrams for the pitchfork normal forms.

This diagram can also be unfolded by bias  $b \in \mathbb{R}$  with

$$\dot{z} = p z + q z^3 - r z^5 + b. \quad (2.12)$$

Fig. 2.7 shows the unfolding effect of varying values of  $b$  on the subcritical pitchfork extended with a stabilizing term. Its influence echoes that of the supercritical pitchfork. When  $b > 0$ , the upper stable solution branch is favored. When  $b < 0$ , the lower stable solution branch is favored. The differences lie in the fact that these solution branches feature a small, unstable region along the favored solution branch, which was nonexistent in the supercritical case. Notably, the critical point (at  $p = 0$ ) disappears quickly with the nonzero bias.

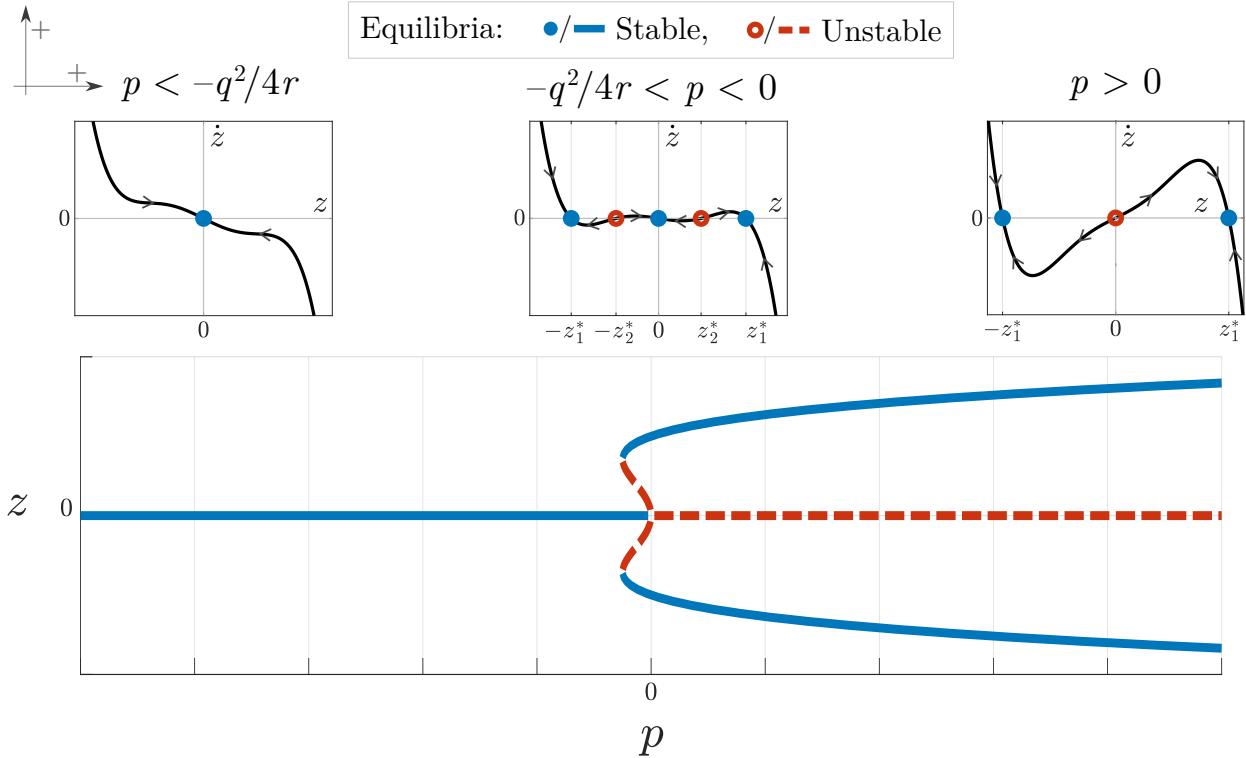


Figure 2.6: Subcritical pitchfork extended with a stabilizing fifth-order term. Top: Function  $\dot{z}$ . Bottom: Bifurcation diagram of  $z$  solutions against parameter  $p$ .

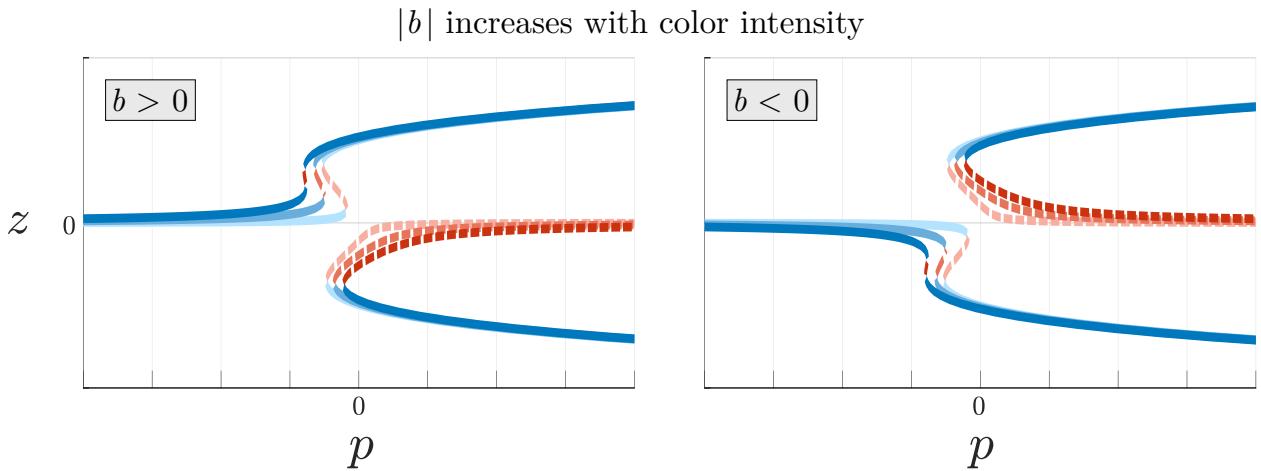


Figure 2.7: Unfolded subcritical pitchfork bifurcation extended with a stabilizing fifth-order term by bias  $b$ . (Left) Positive bias,  $b > 0$ . (Right) Negative bias,  $b < 0$ . Line features to illustrate stability are as in Fig. 2.6, with colors of higher intensity associated with higher magnitude  $b$ .

## Hopf bifurcation

The Hopf bifurcation describes a system (which is necessarily not one-dimensional) that undergoes a change in stability of an equilibrium point, generating or losing a limit cycle and causing periodic orbits in its solutions. In this case, we define the system behavior that classifies the bifurcation. This can be done by analyzing the Jacobian matrix of a general dynamical system:  $J(z, p) = \frac{\partial \dot{z}}{\partial z} = \frac{\partial}{\partial z}(f(z, p))$ .

All Hopf bifurcations have this general form: the Jacobian matrix has a pair of complex-conjugate eigenvalues that cross the imaginary axis as a bifurcation parameter is varied.

We turn to Theorem 11.12 in [45] to guide the definition of this bifurcation. Assume that  $(z^{eq}, p^*) = (0, 0)$  is an equilibrium of the system. This point is a Hopf bifurcation if it satisfies the following criteria:

1. There is a complex conjugate pair of eigenvalues  $(\lambda(p))$  of the Jacobian matrix  $J(z^{eq}, p^*)$  that are purely imaginary. Thus,  $\lambda(p^*) = \lambda(0) = \pm i\omega(0)$  where  $\omega(0) > 0$ .
2. All other eigenvalues of the Jacobian matrix satisfy  $Re(\lambda(p^*)) < 0$  and would otherwise indicate equilibrium stability.
3. The real part of the conjugate pair of eigenvalues crosses the imaginary axis with nonzero speed. That is,  $\frac{\partial Re(\lambda(p))}{\partial p} \Big|_{p=0} \neq 0$ .

This would indicate that as  $p$  is varied, the pair of complex conjugate eigenvalues crosses the imaginary axis. As  $p > 0$ , a limit cycle forms. The stability of this limit cycle is determined by the stability of the function  $f(z, p)$ . A Hopf bifurcation is considered *supercritical* when its created limit cycle is stable, attracting nearby solution trajectories into its orbit; and a Hopf bifurcation is considered *subcritical* when its created limit cycle is unstable, repelling nearby solution trajectories away from its orbit.

While a bifurcation diagram can reveal the *location* of a Hopf bifurcation, the distinction in its supercritical and subcritical behaviors is best observed with a phase portrait.

Fig. 2.8 showcases a supercritical Hopf bifurcation. Its system is defined by the ODE below, following an example in [46, Ch. 8.2]:

$$\frac{dz}{dt} = \frac{d}{dt} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} \mu x - \omega y - x(x^2 + y^2) \\ \mu y + \omega x - y(x^2 + y^2) \end{bmatrix}, \quad (2.13)$$

where the states are  $z = [x, y]^T \in \mathbb{R}^2$  and parameters are  $\mu$  and  $\omega$ . We observe that as the parameter  $\mu$  transitions from a negative to a positive value, the system shifts from having a single stable solution at the origin to exhibiting a stable limit cycle.

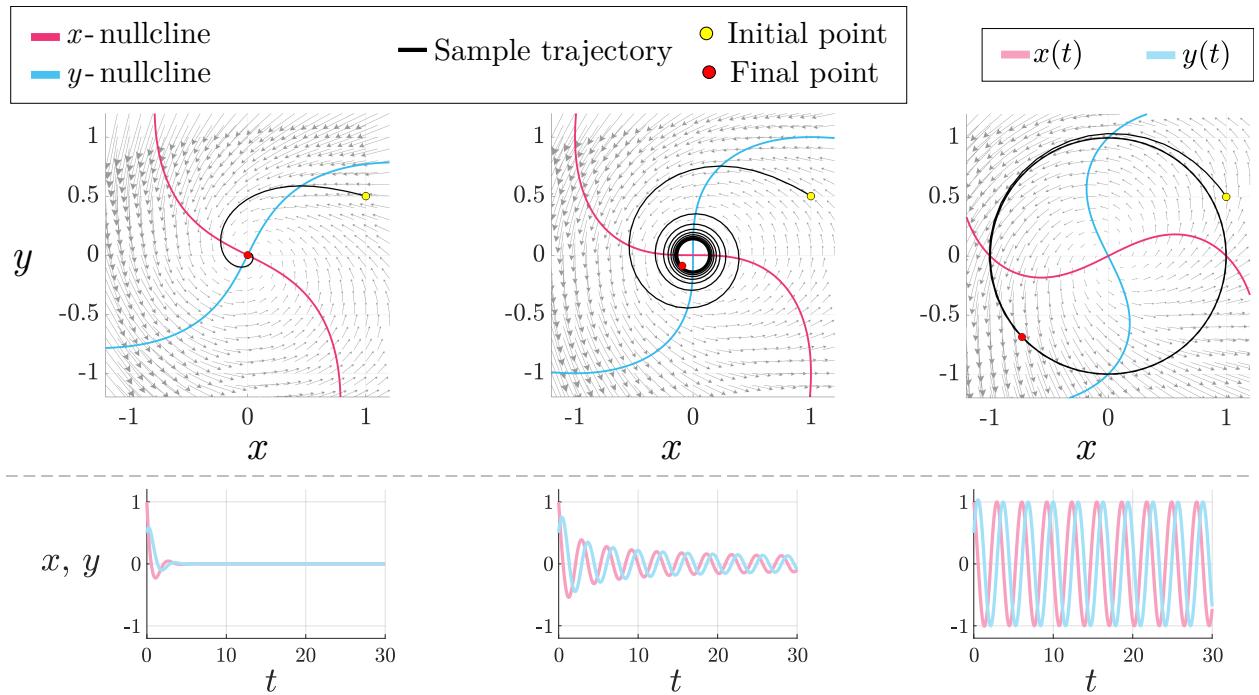


Figure 2.8: Phase planes and sample trajectories of the model (2.13) with a supercritical Hopf bifurcation. (Top) The  $x - y$  phase plane. The  $x$ - and  $y$ -nullclines are drawn in pink and blue, respectively. A sample trajectory is superimposed with a black trace that moves from the same initial point (yellow dot) to its final point (red dot) in a simulation. (Bottom) The  $x$  and  $y$  solutions over time that correspond with the sample trajectory shown in the top graphs. Parameters are:  $\omega = 2$  and  $\mu = -1, 0, 1$  from left to right graphs.

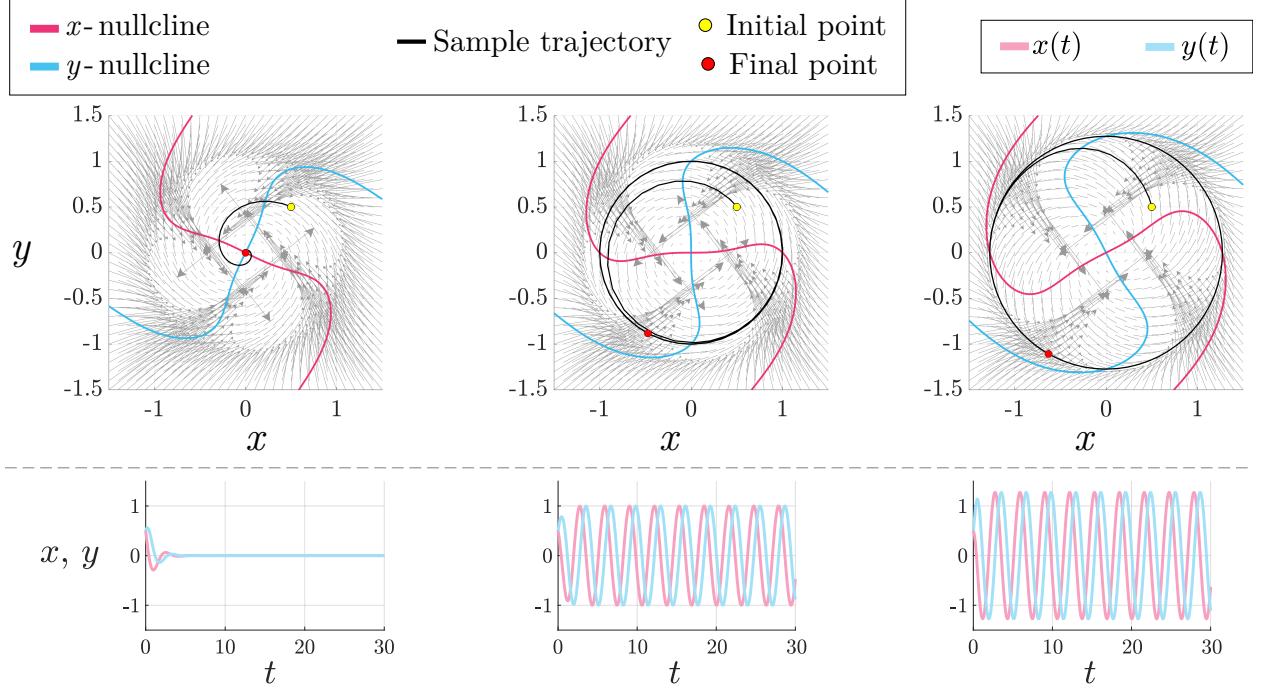


Figure 2.9: Phase planes and sample trajectories of the model (2.14) with a subcritical Hopf bifurcation. (Top) The  $x - y$  phase plane. The  $x$ - and  $y$ -nullclines are drawn in pink and blue, respectively. A sample trajectory is superimposed with a black trace that moves from the same initial point (yellow dot) to its final point (red dot) in a simulation. (Bottom) The  $x$  and  $y$  solutions over time that correspond with the sample trajectory shown in the top graphs. Parameters are:  $\omega = 2$  and  $\mu = -1, 0, 1$  from left to right graphs.

Fig. 2.9 showcases a subcritical Hopf bifurcation. Its system is defined by the ODE below, following an example in [46, Ch. 8.2]:

$$\frac{dz}{dt} = \frac{d}{dt} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} \mu x - \omega y + x((x^2 + y^2) - (x^2 + y^2)^2) \\ \mu y + \omega x + y((x^2 + y^2) - (x^2 + y^2)^2) \end{bmatrix}, \quad (2.14)$$

where the states are  $z = [x, y]^T \in \mathbb{R}^2$  and parameters are  $\mu$  and  $\omega$ , as above. We observe that as the parameter  $\mu$  transitions from a negative to a positive value, a solution with an initial condition near the origin shifts from moving toward the stable origin to moving away from the unstable origin into a larger-amplitude limit cycle.

## 2.2 Nonlinear Opinion Dynamics (NOD) model

Virtually all actions made in life are the result of a decision-making process that considers various options and influences and can adapt to any changing context surrounding the decision. Nearly everything can be categorized into a decision-making process, such as determining which task to spend time on, what

food to eat, or where to move within a room. Due to its universal relevance, decision-making has long attracted research seeking to formulate and generalize models that encapsulate its nuanced dynamics. The work presented within this dissertation builds upon the general Nonlinear Opinion Dynamics (NOD) model developed by Bizyaeva et al., presented in [1] and detailed further in [42]. These papers introduce a framework for a nonlinear, saturated opinion exchange between agents (a term that can describe either an individual or a population) about a multitude of options as a decision-making process. The NOD model captures the building of a collective opinion within a group or cooperative team due to social and environmental influences, and maintains flexibility in changing contexts due to its nonlinearity that saturates extreme opinions.

The Bizyaeva et al. model of Nonlinear Opinion Dynamics is presented here in its most general form, and its parameters are described. This serves as a reference point in this thesis before the presentation and discussion of a specialized form of NOD for social robot navigation.

Consider the formation of an *opinion*, represented by variable  $z$ . The NOD model provides a framework for a nonlinear and saturated opinion exchange between agents (a term that can describe either an individual or a population) about a multitude of options. Consider agent  $i = 1, \dots, N_a$  within a group of  $N_a$  agents as it forms an opinion about option  $j = 1, \dots, N_o$  within a group of  $N_o$  options. It follows that the opinion about this option is noted by the variable  $z_{ij}$ . The magnitude of  $z_{ij}$  represents the strength of the agent's opinion about this option, and its sign (positive, negative) represents its preference (favor, disfavor) for the option. When the opinion is zero,  $z_{ij} = 0$ , agent  $i$  is *neutral* about option  $j$ . If we consider when all agents are neutral about all options,  $\mathbf{z} = \mathbf{0}$ , then the system is *undecided*. These opinions are modulated within the model by various parameters that encourage intra- and inter-agent opinion exchanges about each option. The opinion is tuned by each agent's *sensitivity*, represented by variable  $u$ . According to this model, an agent  $i$ 's sensitivity  $u_i$  can change dynamically, dependent on some basal level of sensitivity  $u_0$  that it may have and/or the increasing opinions of the agent itself and/or its neighbors. With these most relevant variables now defined, the general form of the Nonlinear Opinion Dynamics of Bizyaeva et al. is presented below:

$$\dot{z}_{ij} = -d_i z_{ij} + S \left( u_i \left( \alpha_i^j z_{ij} + \sum_{\substack{k=1 \\ k \neq i}}^{N_a} a_{ik}^a z_{kj} + \sum_{\substack{l=1 \\ l \neq j}}^{N_o} a_{jl}^o z_{il} + \sum_{\substack{k=1 \\ k \neq i}}^{N_a} \sum_{\substack{l=1 \\ l \neq j}}^{N_o} a_{ik}^a a_{jl}^o z_{kl} \right) \right) + b_{ij}, \quad (2.15)$$

$$\tau_u \dot{u}_i = -u_i + u_0 + K_u \sum_{j=1}^{N_o} \sum_{k=1}^{N_a} a_{ik}^u (z_{kj})^2. \quad (2.16)$$

The parameters for the opinion  $z_{ij}$  in equation (2.15) are as follows:

- $d_i \geq 0$  describes the *resistance* of the agent to forming an opinion about option  $i$  and serves to damp

the linear impact of the agent's current opinion state.

- $S : \mathbb{R} \rightarrow \mathbb{R}$  is a bounded saturation function that satisfies the conditions  $S(0) = 0, S'(0) = 1$ , and  $S'''(0) = 0$ . For the duration of this thesis and much of other work involving NOD,  $S(\cdot) = \tanh(\cdot)$ .
- $\alpha_i^j \geq 0$  describes the *self-reinforcement* of an agent's opinion about an option.
  - Note that the signs of  $d_i, \alpha_i^j \geq 0$  create no redundancy, the first serving to only dampen and the second to only strengthen an opinion when the parameters are nonzero, otherwise serving to have no effect if zero.
- $a_{ik}^a \in \mathbb{R}$  describes the *inter-agent, same-option* gain on the influence from different agents opinions  $z_{kj}, k \neq i$  on agent  $i$  about the same option  $j$ .
  - Note the signed influence of this term. When positive,  $a_{ik}^a > 0$  reflects an excitatory influence of the communicated opinion of agent  $k$  on agent  $i$ . When negative,  $a_{ik}^a < 0$  reflects an inhibitory influence of the communicated opinion of agent  $k$  on agent  $i$ . When zero,  $a_{ik}^a = 0$  indicates that there is no influence communicated between agents  $i$  and  $k$ .
- $a_{jl}^o \in \mathbb{R}$  describes the *intra-agent, different-option* gain on the influence from agent  $i$ 's opinions  $z_{il}$  about all other options than  $j, l \neq j$ .
  - Note the signed influence of this term. When positive,  $a_{jl}^o > 0$  reflects an excitatory influence of agent  $i$ 's opinion of option  $j$  on its opinion of option  $l$ . When negative,  $a_{jl}^o < 0$  reflects an inhibitory influence of agent  $i$ 's opinion of option  $j$  on its opinion of option  $l$ . When zero,  $a_{jl}^o = 0$  indicates that there is no link between options  $j$  and  $l$ .

The parameters for the sensitivity  $u_i$  in equation (2.16) are as follows:

- $\tau_u > 0$  is a time constant on the sensitivity,
- $u_0 \geq 0$  is the *basal sensitivity* present in the agent,
- $K_u \geq 0$  is a gain on the influence of agents' opinions on the increase of its sensitivity,
- $a_{ik}^u \geq 0$  is the gain on the influence from agent  $k$ 's opinions  $z_{kj}$  on the sensitivity of agent  $i$ .

As a final note, for much of the work presented within this dissertation we consider that  $\tau_u \rightarrow 0$  and thus the dynamic sensitivity of (2.16) becomes algebraic with  $u_i = u_0 + K_u \sum_{j=1}^{N_o} \sum_{k=1}^{N_a} a_{ik}^u (z_{kj})^2$ .

## 2.3 Hardware for experimental demonstration

The experimental portion of this work was conducted using two primary hardware sources. The Jackal UGV was utilized for navigation in all human-robot interaction experiments within this body of work, and all motion-related data were acquired using a Vicon motion capture system. The relevant specifications of each are described in detail below.

### 2.3.1 Jackal UGV research robot

Navigation experiments utilized the Jackal UGV (Unmanned Ground Vehicle) made by Clearpath Robotics [47]. The Jackal is a small, four-wheeled mobile robot equipped with GPS, IMU, and upgraded with a Jetson TX2 computing device [48] for increased memory and bandwidth. Its internal computer runs the Linux Ubuntu 16.04 version with the ROS (Robot Operating System) Kinetic Kame distribution.

The robot uses skid-steer drive, a multi-wheeled form of differential drive, to move and turn as it navigates. This entails turning multiple wheels per side of the robot (rather than one wheel per side) to increase the rotational velocity higher than that on the other side, creating a turn towards the side with lower rotational velocity. For example, both right wheels of the Jackal will turn faster than its two left wheels, causing the robot to turn towards the left.

For additional sensing, the Jackal is also fitted with a Velodyne PUCK LiDAR sensor [49] at its front. This equipment can measure depths of up to 100 meters in a 360° horizontal and 30° vertical range. In experiments, this 3D visual field information is measured to verify that the robot is a safe operating distance away from nearby obstacles. Suppose the robot observes an obstacle nearby but is unable to navigate away from a collision safely. In that case, the LiDAR information will trigger a backup stop protocol as a fail-safe. This ensures that the Jackal will not collide with any obstacles, whether they are static elements of the environment or humans.

For detailed, real-time data on the Jackal's position and orientation, specialized markers are attached to the robot's top and tracked by the Vicon motion capture cameras. Data is broadcast through the `vicon_bridge` driver on ROS [50] such that this reliable pose information can be collected and used within the robot's algorithm.

### 2.3.2 Vicon motion capture system

The position and orientation data of humans and robots within experiments are tracked using the Vicon motion capture system in the H-121 laboratory space. Unique marker plates are used to identify relevant objects in each study. For example, each human participant in a study wears a unique plate attached to a

hat during a navigation experiment, and each robot has a unique plate attached to its top. Twelve Vicon Vero v2.2 cameras [51] line the perimeter of the room’s ceiling, are calibrated before each batch of trials, and provide real-time pose information of tracked objects at 330 Hz. The precise position and orientation data for tracked objects is then shared through the ROS `vicon_bridge` software as a quaternion.

## Chapter 3

# Nonlinear Opinion Dynamics (NOD) for proactive social robot navigation

In this chapter, we introduce the setting and design of a Nonlinear Opinion Dynamics (NOD) model to control the navigation behavior of robots interacting with humans. This chapter draws heavily from the published work of *Proactive Opinion-Driven Robot Navigation around Human Movers* of [25] presented at the 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems, which can be seen in Chapter 7. The application to a navigation problem for a social robot is expanded upon from the published work. We leverage bifurcation theory to analyze the model and provide a guarantee on its deadlock-breaking ability. We conclude with experimental validations of NOD to create flexible and reliable navigation between a single robot interacting with either one or two humans.

### 3.1 Motivations

Robots are increasingly desired to be, and some are already becoming, commodities in human-centric environments, from hospitals to homes. The deployment of robots in these human-populated environments requires robotic navigation systems that are not only collision-free but also socially appropriate and aware. Many state-of-the-art navigation strategies leverage data-driven, computationally intensive tools, such as reinforcement learning and machine learning, to construct models of human behavior, predict navigation path options, or make more complex decisions. However, these methods are often difficult to analyze and do not guarantee that indecision (deadlock) will not occur when the robot faces the vast navigation options and constraints that exist in human interaction. We instead turn to considering the analytically tractable

Nonlinear Opinion Dynamics (NOD) of [1] to create a novel opinion-driven model for socially competent robot motion control. This work builds on the preliminary efforts of [52] to design legible robot motion for navigation with human interaction using a simplified version of the NOD model. We construct a set of continuous dynamical equations for a robot’s evolving opinion on navigating around humans and leverage bifurcation theory to establish a guarantee of deadlock-free decision-making in robot-human navigation. We aim to showcase how this computationally light model can proactively respond to human movement in real time, enabling robots to make real-time decisions that are responsive to social cues and provide the ability to comfortably adapt to others in a shared space.

### 3.1.1 Navigation environment and notation

We define the setup of the navigation environment and the notation describing the spatial and angular relationships in human-robot interaction scenarios. Let robot  $i = 1, \dots, N_r$  describe a mobile robot within a team of  $N_r$  robots. Let human  $\ell = 1, \dots, N_h$  describe a human detected in the environment by robot  $i$  within a set of  $N_h$  humans. Each robot  $i$  moves with constant velocity  $V_i$  and each human  $\ell$  moves with velocity  $V_{h_\ell}$ .

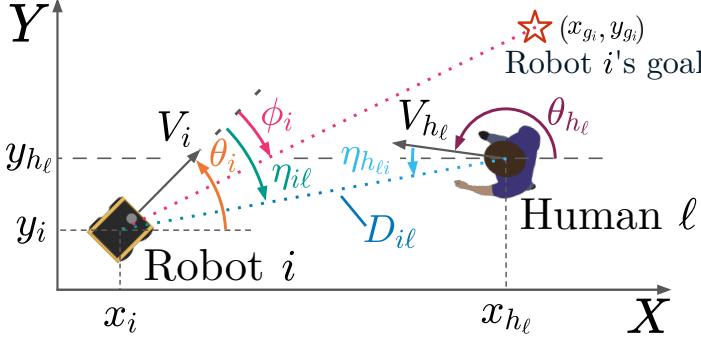
We define the relevant spatial and angular geometry for a navigation scenario that involves human-robot interaction. The two-dimensional position of robot  $i$  is given by  $\mathbf{x}_i = (x_i, y_i)$  and its goal is located at  $\mathbf{x}_{g_i} = (x_{g_i}, y_{g_i})$ . The heading (or turning) angle of robot  $i$  relative to the horizontal axis is given by  $\theta_i$ . Similarly, human  $\ell$ ’s position and heading angle are given by  $\mathbf{x}_{h_\ell} = (x_{h_\ell}, y_{h_\ell})$  and  $\theta_{h_\ell}$ , respectively. The Euclidean distance between robot  $i$  and human  $\ell$  is denoted  $D_{i\ell}$ . We consider three relative angles within the environment:

1.  $\phi_i$ , the angle between robot  $i$ ’s heading and the direction of its goal,
2.  $\eta_{i\ell}$ , the angle between robot  $i$ ’s heading and the direction to human  $\ell$ ,
3.  $\eta_{h_\ell i}$ , the angle between human  $\ell$ ’s heading and the direction to robot  $i$ .

These angles, along with the robot-human distance, are illustrated in Fig. 3.1, and their functions are described in Table 3.1. The function  $\text{atan2}(y, x)$  computes a real solution of the four-quadrant inverse tangent of  $y$  and  $x$  following the east counter-clockwise convention.

## 3.2 NOD specialized for social robot navigation

We build upon the NOD model by Bizyaeva et al., simplified to a case with two mutually exclusive options, as discussed in Chapter 2.2. Thus,  $N_o = 2$  and  $z_{i1} = -z_{i2}$ . The opinion state is therefore captured by a single



Variable	Function
$\phi_i$	$\text{atan2}(y_{g_i} - y_i, x_{g_i} - x_i) - \theta_i$
$\eta_{i\ell}$	$\text{atan2}(y_{h_\ell} - y_i, x_{h_\ell} - x_i) - \theta_i$
$\eta_{h_\ell i}$	$\text{atan2}(y_i - y_{h_\ell}, x_i - x_{h_\ell}) - \theta_{h_\ell}$
$D_{i\ell}$	$\sqrt{(x_{h_\ell} - x_i)^2 + (y_{h_\ell} - y_i)^2}$

Figure 3.1 & Table 3.1: Geometry and notation of the human-robot navigation setting.

scalar variable  $z_i$ , whose sign indicates the option preference and whose magnitude indicates the strength of the preference. We assume  $\tau_u \rightarrow 0$  in (2.16), which simplifies the dynamic sensitivity to an algebraic expression for the social robot navigation setting. We therefore consider the simplified set of equations below as the general two-option NOD model to adapt for navigation:

$$\dot{z}_i = -d_i z_i + S \left( u_i \left( \alpha_i z_i + \sum_{k=1}^{N_a} a_{ik} z_k \right) \right) + b_i, \quad (3.1a)$$

$$u_i = u_0 + K_u \sum_{j=1}^{N_o} \sum_{k=1}^{N_a} a_{ik}^u (z_{kj})^2. \quad (3.1b)$$

We propose a robot navigation model that uses NOD to direct the navigation of a robot so that it moves comfortably around nearby human movers. We begin by outlining the desired behavior of a robot that navigates with NOD and then introduce the ruling equations to satisfy these design requirements.

### 3.2.1 Physical interpretation of opinion $z_i$

We first detail the meaning of the opinion variable  $z_i$  in the context of human-robot interaction. We define the opinion  $z_i$  to represent robot  $i$ 's opinion for passing a human on the left or the right and by how much. That is, for robot  $i$ :  $z_i = 0$  represents no preferred passing side (i.e., move straight),  $z_i > 0$  represents a left-preferred passing side, and  $z_i < 0$  represents a right-preferred passing side. Robot  $i$  should form this preferred passing side based on social cues relevant to comfortable and cooperative navigation. These cues include the opinions of nearby robots on a team ( $z_k$ , where  $k \neq i$ ) and an estimate of the preference of passing humans, which we represent as a proxy opinion ( $\hat{z}_{h_\ell}$ ) for each human  $\ell$ . These cues should be weighted more heavily as collisions become more likely. This implies that the sensitivity  $u_i$  that acts as a gain on the social cues should increase as collisions become more likely. Finally, we consider that robot  $i$  may have an internal

bias  $b_i$  on which side to pass someone, e.g., associated with a social convention like passing on the right in the United States. These behavioral requirements for social robot navigation guide our specialization of the  $z_i$  and  $u_i$  variables within the general NOD equations (3.1).

### 3.2.2 Opinion-driven motion control

With the behavioral interpretation of the preferred passing side captured by opinion  $z_i$ , we can now consider how its value should affect robot motion. This extends beyond the general Bizyaeva et al. model, requiring us to construct an angular velocity command  $\dot{\theta}_i$  that incorporates the opinion's influence and the goal direction  $\phi_i$ . The robot's sensitivity to humans can be used to control the balance between these two turning influences. These required turning behaviors are incorporated into the motion control formulation for the robot.

The specialized NOD equations for social robot navigation are

$$\tau_{z_i} \dot{z}_i = -d_i z_i + \tanh \left( u_i \left( a_i z_i + \sum_{k=1}^{N_a} a_{ik} z_k + \sum_{\ell=1}^{N_h} a_{i\ell} \hat{z}_{h_\ell} \right) + b_i \right), \quad (3.2a)$$

$$u_i = \max(u_{0_{i\ell}}) + K_{u_i} z_i^2, \quad (3.2b)$$

$$\dot{\theta}_i = k_{\theta_i} \sin(\widetilde{u_{0_{i\ell}}} \cdot \beta_i \cdot \tanh(z_i) + (1 - \widetilde{u_{0_{i\ell}}}) \phi_i). \quad (3.2c)$$

We begin by defining the parameters and variables within the opinion update in equation (3.2a). Recall that  $S(\cdot) = \tanh(\cdot)$  acts as the saturating function within NOD. The parameters are as follows:  $\tau_{z_i} > 0$  is the opinion time constant,  $d_i \geq 0$  is the damping coefficient,  $a_i \geq 0$  is the self-reinforcement gain,  $a_{ik} \in \mathbb{R}$  is the weight of robot  $k$ 's opinion,  $a_{i\ell} \in \mathbb{R}$  is the weight of human  $\ell$ 's proxy opinion,  $\hat{z}_{h_\ell}$  is the proxy opinion of human  $\ell$ , and  $b_i \in \mathbb{R}$  is an internal bias. Note that this bias  $b_i$  is moved within the saturating function such that the robot will only act upon this internal turning preference when a human is observed (i.e.,  $u_i \neq 0$ ) and will otherwise move directly toward its goal. We define  $\hat{z}_{h_\ell} = \sin(-\eta_{h_\ell})$  as the proxy opinion of human  $\ell$ . The weight  $a_{i\ell}$  is defined as  $a_{i\ell} = \gamma_i \cdot \max(0, u_{0_{i\ell}} / \sum_{\ell=1}^{N_h} u_{0_{i\ell}})$  where  $\gamma_i > 0$  is a gain and  $u_{0_{i\ell}}$  is the basal level of sensitivity of robot  $i$  to human  $\ell$ , where with a slight abuse of language we allow this “basal” sensitivity to change value. Note that the max function in  $a_{i\ell}$  serves to avoid divisions by 0 when there are no observed humans (i.e., when  $\sum_{\ell=1}^{N_h} u_{0_{i\ell}} = 0$ ). In this chapter, we focus on a single robot case ( $i = N_r = 1$ ) and therefore defer the definition of  $a_{ik}$  to later chapters.

We now describe the sensitivity parameters in equation (3.2b), along with the geometric calculations

specific to the navigation setting. Each robot  $i$  calculates a basal level of sensitivity  $u_{0,i\ell}$  to human  $\ell$  based on their proximity, which represents how strongly the robot should respond to the human's movements. The basal sensitivity is computed as a function of the distance,  $\chi_{i\ell} = D_{i\ell} = \|\mathbf{x}_i - \mathbf{x}_{h\ell}\|$  and direction,  $\kappa_{i\ell} = D_{crit,i} \cdot \exp(1 - \max(\cos \eta_{i\ell}, 0)^{-1})$ , between robot  $i$  and human  $\ell$ . Here,  $D_{crit,i}$  is the critical distance at which robot  $i$  begins to meaningfully respond to human  $\ell$ . The function to describe the basal sensitivity is based on the inverse Hill equation, a bounded sigmoidal curve, and is  $u_{0,i\ell} = U_i(\kappa_{i\ell}^n / (\kappa_{i\ell}^n + \chi_{i\ell}^n))$ . The parameter  $U_i > 0$  describes the upper bound of the basal level of sensitivity, and  $n > 0$  describes the Hill coefficient of the function that determines the curve's steepness. Fig. 3.2 illustrates the sensitivity function across distance and various interaction angles  $\eta_{h\ell}$ . Equation (3.2b) describes the total sensitivity experienced by the robot. Its parameters are as follows:  $u_{0,i\ell}$  is the basal level of sensitivity robot  $i$  shows to human  $\ell$ , and  $K_{u,i} \geq 0$  is the gain on the robot's own opinion. The maximum basal level of sensitivity  $\max(u_{0,i\ell})$  represents the single largest value of  $u_{0,i\ell}$  for all humans  $\ell \in [1, N_h]$ .

Finally, we combine the influence of both the opinion and sensitivity on the turning command sent to robot  $i$  by its angular velocity,  $\dot{\theta}_i$ , as defined in equation (3.2c). The parameters are as follows:  $k_{\theta,i} > 0$  is a gain,  $\widetilde{u_{0,i\ell}}$  is the normalized maximum basal sensitivity experienced by robot  $i$ , and  $\beta_i \in (0, \frac{\pi}{2}]$  is the maximum dodging angle used when interacting with a human. To interpret the role of design parameter  $\beta_i$ , note that when  $z_i$  is sufficiently large such that  $\tanh(z_i) \approx 1$  (resp. -1), (3.2c) steers the robot  $i$ 's heading angle by an additional  $\beta_i$  radians in the counterclockwise (resp. clockwise) direction from the orientation toward the goal location  $\phi_i$ . Thus, we can tune  $\beta_i$  to prescribe how much robot  $i$ 's heading angle should deviate from its direct path toward its goal when it detects a human and forms a strong preference for a passing direction. We calculate  $\widetilde{u_{0,i\ell}} = \max_{\ell}(u_{0,i\ell})/U_i$ . This value represents the largest normalized basal level of sensitivity of robot  $i$ . Accordingly, when robot  $i$ 's sensitivity is high, even if only to one human, it turns based on its formed opinion  $z_i$ . Otherwise, robot  $i$  will proceed directly toward its goal direction  $\phi_i$ .

For the duration of this chapter, only a single robot ( $i = N_r = 1$ ) is considered. Therefore, the subscript  $i$  can be dropped from variable and parameter notation. The equations for a single robot interacting with  $N_h$  humans then simplify to

$$\tau_z \dot{z} = -d z + \tanh \left( u \left( az + \sum_{\ell=1}^{N_h} a_{\ell} \hat{z}_{h\ell} \right) + b \right), \quad (3.3a)$$

$$u = \max(u_{0,\ell}) + K_u z^2, \quad (3.3b)$$

$$\dot{\theta} = k_{\theta} \sin(\widetilde{u_{0,\ell}} \cdot \beta \cdot \tanh(z) + (1 - \widetilde{u_{0,\ell}}) \phi). \quad (3.3c)$$

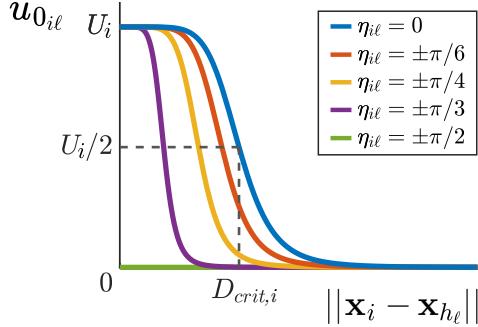


Figure 3.2: The  $u_{0_{i\ell}}$  function, constructed as an inverse hill function. Various values of  $\eta_{i\ell}$  are shown.

We use these equations to analyze a navigation interaction scenario with a single robot and a single human to guarantee the deadlock-breaking properties of the NOD model.

### 3.2.3 Model refinements

There are two significant differences between the model presented and studied in this chapter and those in the published work of [25, 26] and in Chapters 4 and 7. First, we consider a different function for the proxy opinion of a human,  $\hat{z}_h$ , within NOD. Second, we consider a different function for the sensitivity  $u$  of the robot. Both changes ensure consistency throughout the work presented in this dissertation without sacrificing the analysis of NOD.

In previous iterations of this work, the proxy opinion for a human was defined as  $\hat{z}_h = \tan(-\eta_h)$ . The new definition of  $\hat{z}_h = \sin(-\eta_h)$  maintains the sign and opinion information conveyed by this previous proxy, but improves subsequent work by maintaining continuity between proxy opinions formed as  $\eta_h$  passes through 0 and  $2\pi$  radians. This change does not affect any reported results in published works, but its update in simulations and experiments conducted for this dissertation improves the performance of robot navigation. Previously, experiments could be compromised by the discontinuity caused by a human crossing from one side to the other of a robot's heading direction (thus, a crossing between the asymptote at 0 and  $2\pi$  radians in tangent).

Finally, the sensitivity variable  $u$  was previously defined as a dynamic variable for consistency with the general NOD model of Bizyaeva et al. in [1], transcribed in Chapter 2. The definition was  $\tau_u \dot{u} = -u + g(\cdot)$ , with  $g(\cdot)$  a function of the environment's geometry. However, it was noted that experimental trials benefited from assuming  $\tau_u \rightarrow 0$ , yielding an algebraic function of sensitivity  $u = g(\cdot)$ . This change provided more control over the desired behavior of the robot. In this work, we treat sensitivity solely as an algebraic function and do not consider  $u$  in a dynamic state. This unifies the algorithms used in simulations and experiments, enabling a more accurate comparison of the robot's observed behaviors.

### 3.3 Guarantee on deadlock-free navigation

A key contribution of this work is the guarantee of deadlock-free navigation for the robot. We establish this performance guarantee by analyzing the single robot navigation model (3.3) with a single human ( $\ell = N_h = 1$ ) as a proof of concept. Using tools from nonlinear dynamical systems theory [43], we can show how the robot can always rapidly and reliably form a strong opinion to select one of the two options - either move left ( $z > 0$ ) or right ( $z < 0$ ) - to avoid colliding with a human. We achieve this by considering the challenging case in which a human maintains a path straight for the robot, such that  $\dot{z}_h = 0$ , which provides no information on which passing side to prefer, and the robot has no bias ( $b = 0$ ) to internally prompt a preferred passing side. This reduces the navigation NOD model to

$$\tau_z \dot{z} = -d z + \tanh(u a z), \quad (3.4a)$$

$$u = u_0 + K_u z^2, \quad (3.4b)$$

$$\dot{\theta} = k_\theta \sin\left(\frac{u_0}{U} \cdot \beta \cdot \tanh(z) + \left(1 - \frac{u_0}{U}\right) \phi\right). \quad (3.4c)$$

We focus on equations (3.4a)-(3.4b). Substituting in the function of  $u$  into (3.4a), we get

$$\tau_z \dot{z} = -d z + \tanh((u_0 + K_u z^2) a z). \quad (3.5)$$

We use the Jacobian  $J(z)$  of the system to linearly approximate the local behavior of the system (3.5) around point  $z$ :

$$J(z) = \frac{\partial \dot{z}}{\partial z} = -d + (a(u_0 + K_u z^2) + 2K_u a z^2) \cdot \text{sech}^2((u_0 + K_u z^2) a z).$$

We analyze the equilibrium of equation (3.5) to gain insight into the dynamics and stability of the system. The neutral opinion  $z = 0$  is always an equilibrium solution of (3.5). If the robot remains in this unopinionated state, it will experience deadlock in its decision-making, ultimately leading to a collision with the oncoming human directly ahead. Therefore, we linearize<sup>1</sup> the system (3.5) about this equilibrium to characterize its stability. Deadlock-breaking is guaranteed if this  $z = 0$  equilibrium is unstable.

The Jacobian at the neutral equilibrium is  $J(z = 0) = -d + a u_0$ , which is a function of only  $u_0$ . As it is a single value, this is the unique eigenvalue  $\lambda$  of the linearization at equilibrium  $z = 0$ . The neutral equilibrium is stable if its eigenvalue is negative ( $\lambda < 0$ ) and unstable if its eigenvalue is positive ( $\lambda > 0$ ). Here, a bifurcation occurs when the neutral equilibrium changes stability. Setting  $\lambda = 0 = -d + a u_0$  yields

---

<sup>1</sup>We will explore further analysis methods in subsequent chapters to further investigate the effects of various parameters on system behavior. Currently, linearization serves as a tool to identify and characterize the bifurcation parameter of interest.

the bifurcation point  $u_0^* = d/a$ . As stated, deadlock-breaking is guaranteed if there exists some parameter regime where the neutral  $z = 0$  equilibrium is unstable. We can now assess: when  $u_0 < u_0^*$ ,  $\lambda < 0$ , and the neutral equilibrium is stable; when  $u_0 > u_0^*$ ,  $\lambda > 0$ , and the neutral equilibrium is unstable. Thus, when the basal level of attention  $u_0$  of the robot exceeds this  $u_0^*$  value, deadlock-breaking is guaranteed as the neutral solution becomes unstable.

Furthermore, it can be demonstrated that in the parameter regime where  $u_0 > u_0^*$ , two symmetric bistable solutions emerge, one corresponding to moving to the left and the other to moving to the right. These equilibria are equal and opposite, such that  $z_1^{\text{eq}} = -z_2^{\text{eq}}$ , and both are locally exponentially stable. This meets the criteria of the *pitchfork bifurcation* described in Chapter 2. For this chapter, we consider the parameters that create a *supercritical* pitchfork bifurcation in the system. A discussion of the parameters that cause a subcritical pitchfork bifurcation is provided in the following chapter.

The supercritical pitchfork bifurcation diagram for this social navigation scenario is seen at the top of Fig. 3.3a. The bifurcation diagram shows the equilibrium values of opinion  $z$  plotted as a function of the bifurcation parameter  $u_0$ . The curves in blue represent stable solutions, and those in dashed red represent unstable solutions (saddle points). Note that the strength of preferences (i.e., the magnitude of the opinion  $z$ ) increases with increasing  $u_0 > u_0^*$ . In this regime, deadlock is unstable, and thus the robot's opinion will necessarily converge on one of the two non-neutral, opinionated solutions. The selected solution branch will depend on the initial conditions of the system and measurement noise. When the robot is biased ( $b \neq 0$ ) or the human is approaching the robot obliquely ( $\eta_h \neq 0$ ), the pitchfork bifurcation *unfolds*, as illustrated at the bottom of Fig. 3.3a.

We illustrate the deadlock-breaking behavior with simulations in Fig. 3.3b. The human (trajectory in black) heads straight for the robot. In the unbiased case ( $b = 0$ ), the robot (trajectory in pink) moves straight briefly before arbitrarily choosing to turn right to pass around the human. This corresponds to behavior indicated by the negative blue curve at the top of Fig. 3.3a. In the biased case ( $b > 0$ ), the robot (trajectory in wine) follows its bias and moves left, departing sooner from the centerline than in the unbiased case. This corresponds to the positive blue curve at the bottom of Fig. 3.3a.

### 3.4 Experimental validations of NOD

The NOD algorithm was validated through both simulated and experimental trials with NOD-embodied robots navigating past a single oncoming moving human. Simulations were rendered using MATLAB, and the behaviors observed in these idealized, noise-free simulations provided evidence for the general range of parameter values to be used in real experiments.

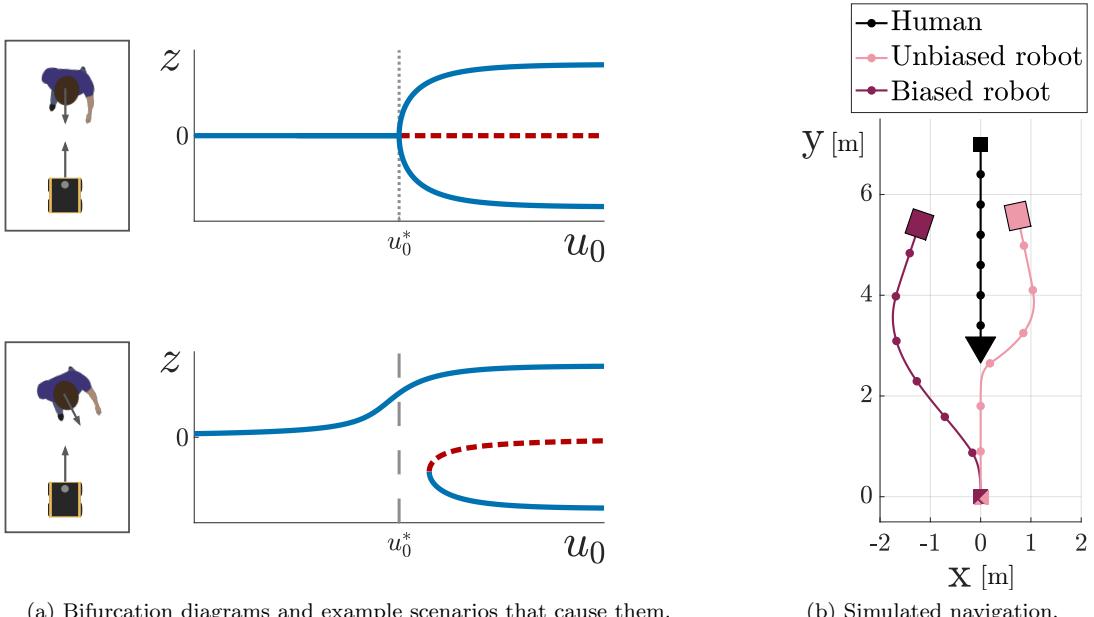


Figure 3.3: Deadlock breaking for  $u_0 > u_0^*$  in a robot's opinion dynamics. (a) Example human-robot interaction scenarios and their associated pitchfork bifurcation diagram. Top: Symmetric, Bottom: Unfolded. (b) Example trajectories of unbiased ( $b = 0$ ) and biased ( $b > 0$ ) robots facing a directly oncoming human. Dots along trajectories are temporal markers.

Previous work by [25], as seen in Chapter 7, expanded on a NOD robot's behavior in an environment with two humans. This work employed an algorithm with minor differences from the NOD model presented in this chapter. Therefore, although these experiments were not fully recreated for this dissertation, the movement data from the previous experiments were used to simulate a robot's response using NOD, consistent with the findings of this dissertation.

Our experimental objectives are threefold:

1. To demonstrate the flexibility of our approach by showing that the robot can reliably navigate environments and interact with multiple human movers across a variety of scenarios.
2. To validate our algorithm's analysis, confirming that the robot can always resolve deadlock situations—gracefully maneuvering around an oncoming human, even if the human is unaware of or ignores the robot, and even when the robot's preferred passing side conflicts with that of the human.
3. To test our hypothesis that the trade-off between efficiency and spatial intrusiveness in navigation behavior can be tuned using a single algorithm parameter,  $\beta$ .

### 3.4.1 General experimental setup

Experiments were conducted in an uncluttered 8 m  $\times$  8 m laboratory space equipped with 12 Vicon motion capture cameras to track the real-time positions and orientations of marked objects in the space. A single UGV Jackal robot, fitted with a Vicon tracking marker, was placed at an initial location and provided a goal point (set with another Vicon tracking marker) to navigate towards at a constant speed. Further specifications of the cameras and robot are described in Section 2.3. Each human participant in the space wore a hat fitted with a Vicon tracking marker. The robot utilized real-time position and orientation Vicon data from nearby humans to react to their movements.

In all trials, the robot's objective was to move through the room toward its goal. If no humans were present in the robot's forward field of view, the robot proceeded directly toward its goal, mimicking the classic "go-to-goal" behavior seen in PID control. If humans were present, the robot's behavior depended on whether an internal bias  $b \neq 0$  was included in the NOD model to predispose the robot's preference for moving around nearby humans, either to the left or to the right. If no internal bias was set ( $b = 0$ ), the robot develops its preferred passing side around a human based solely on its observations of the human's position and orientation. The NOD algorithm guided the robot away from approaching or nearby humans cooperatively and collaboratively (e.g., moving towards its left if a human appeared to be moving on the robot's right side).

Two forms of experiments were conducted, wherein

1. a single robot navigates around a single human, each moving toward the other's initial location,
2. a single robot navigates around two humans moving within the space.

The instructions and objectives given to the human participants varied depending on the experiment format.

For the single robot/single human experiments, the human was initially placed opposite the robot in the lab space, facing the robot directly. The participant's goal location was the robot's initial position. The human was instructed to move toward this goal in one of three ways: walk directly toward the goal as if unaware of the robot's presence, or walk toward the goal with a curving trajectory toward either their left or right. The human was instructed to walk at a comfortable pace and to stop their movement only once they had arrived at their goal location. Further details are described in Section 3.4.2.2.

For the single robot/two humans experiments, the humans were allowed to move freely through the laboratory space. The experiment naturally produced trials in which humans directly moved from an initial location to a goal location. Participants could start or stop their movement at will throughout a trial. Further details are described in Section 3.4.2.3.

In each experiment, the positions and orientations of all tracked objects were recorded, along with their timestamps. The relative angles and distances between objects relevant to the robot (e.g., its goal, any nearby humans) were also logged. The robot’s opinion  $z$ , opinion rate of change  $\dot{z}$ , basal level of sensitivity  $u_0$ , sensitivity  $u$ , and angular velocity  $\dot{\theta}$  were calculated and recorded. This information, along with all parameters used within the NOD model, was compiled into a single .csv file for each experiment. All data analysis was performed in MATLAB. Each experiment was recorded on video from the start of movement until the robot reached its goal.

Participants in the study were recruited through word of mouth, and any identifying features in both the data and video were anonymized. Study #14788 has been approved by the International Review Board (IRB) of Princeton University.

### 3.4.2 Results

#### 3.4.2.1 Metrics reported

We first define the metrics used to analyze the trajectories taken by the robot in experimental trials. They are as follows:

Efficiency Metrics		Spatial Intrusiveness Metrics*	
Definition	Unit	Definition	Unit
<b>Navigation time:</b> The time it took for the robot to arrive at its goal location from its initial position.	s	<b>Minimum distance**:</b> The distance between the robot and a passing human at the closest point in their paths.	m
<b>Path length:</b> The total distance the robot traveled as it moved between its initial and goal locations.	m	<b>Invasion distance:</b> The distance the robot traveled while invading the personal space (surrounding 1.2 m*** of a passing human).	m
<b>Centerline deviation:</b> The cumulative area under the robot’s curved path from the straight path to its goal.	$m^2$		

Table 3.2: Performance metrics and their definitions.

\*Note that these spatial intrusiveness metrics have been categorized as “indicators of human discomfort” in [53]. The discomfort levels of humans are of interest to this work; however, this metric was not measured as it requires a participant survey.

\*\* Recall from Section 3.2 that  $D_\ell$  represents the distance between a single robot and human  $\ell$ . Therefore, we can shorthand this metric as  $\min(D_\ell)$  for the minimum distance between a single robot and human  $\ell$ .

\*\*\* The 1.2 m value comes from the reported threshold in [54] between a “personal zone” preferred by humans for comfortable interactions with friends and a “social zone” for comfortable interactions with

strangers, according to data collected from northern Europeans (in England, France, and Germany). These results summarize the work of Hall [55] on the introduction of *proxemics*, the study of how much space is necessary and understood to be maintained throughout various social interactions. Setting a metric for the desired minimum separation between a robot and a human throughout an interaction allows us to analyze whether and how well the NOD navigation can innately adhere to social conventions without their explicit inclusion in the model.

These metrics are generally accepted as common and necessary results to gauge social robot interaction as in [5]. These are also supported in part by the findings of [56], which surveyed employees of various organizations deploying social robots and ranked personal distance violations, path length, and path efficiency among the most important success metrics.

### 3.4.2.2 Single-robot, single-human experiments

The setup of each experiment was as follows. Pairs of starting and goal locations for the robot and the human were fixed across all trials at (0, 0) m and (0, 6) m, respectively. The goal location of the robot was set at (0, 6.4)m, effectively mirroring the initial location of the human. The human was instructed to move toward the robot’s initial location. Thus, the initial and goal locations of the robot and human were swapped, causing them to move head-on toward one another.

We combined three robot bias cases with three movement prompt cases for the human participant, yielding nine unique trial configurations. We conducted three experiments within each trial configuration, resulting in a total of twenty-seven trials. For each trial, we tested two values of  $\beta$  in equation (3.3c):  $\beta = \pi/4$  and  $\beta = \pi/6$ . This resulted in a total of fifty-four trials. The resultant trajectories in each of these trials are documented in Fig. 3.4, which is modeled after Fig. 7.5 in [25] and can be seen in Chapter 7. Further, we note that the resultant metrics of these trials are illustrated in figures which are modeled after Fig. 7.6 and 7.7 in [25] and can be seen in Chapter 7.

#### Robot instructions

The robot was programmed to move at a constant speed  $V = 0.7\text{m/s}$  toward its goal location, modifying the direction of its trajectory when encountering movers according to the navigation model (3.3) with parameters  $\tau_z = 0.008, d = 1, a = 2, \gamma = 2, K_u = 0.3, U = 1, n = 2, D_{crit} = 5.5\text{m}$ , and  $k_\theta = 3$ . We designed three cases corresponding to three different values of the robot’s bias  $b$ :

1. unbiased ( $b = 0$ ),
2. biased to its left ( $b = 0.3$ ),
3. biased to its right ( $b = -0.3$ ).

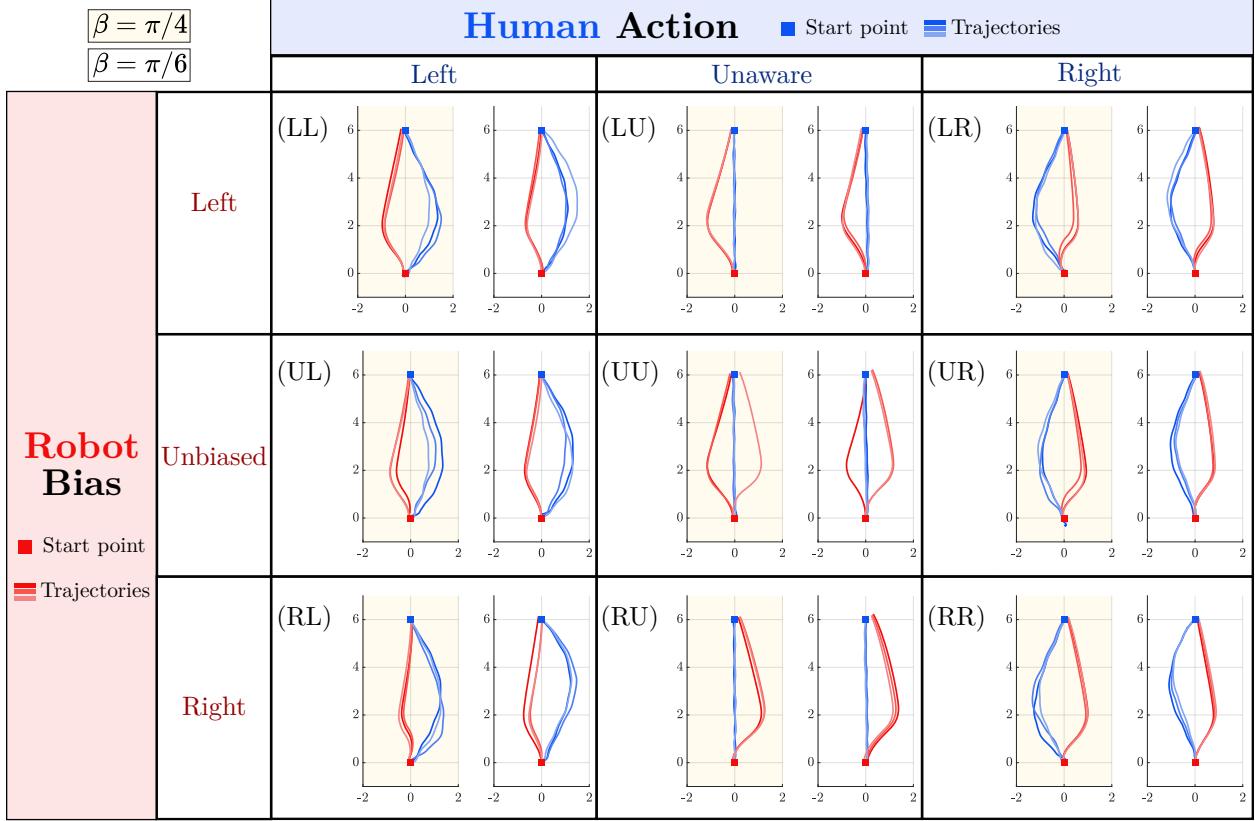


Figure 3.4: The trajectory data of a **NOD** robot for three runs each of the nine trial configurations, with both  $\beta_r = \pi/4$  (shaded yellow) and  $\beta_r = \pi/6$  (unshaded). Axes correspond to the  $xy$ -plane in meters. The robot paths are shown in red, with a red box indicating the robot’s starting position at approximately  $(0, 0)$  m. The human paths are shown in blue, with a blue box indicating the human’s starting position at approximately  $(0, 6)$  m. In trial configuration labels, L=left, U = unaware/unbiased, and R = right.

### Human participant instructions

The participant was instructed to walk at their normal pace towards their goal location according to one of three prompts: 1) go straight (labeled as the human was unaware of the robot), 2) bear to the left, and 3) bear to the right.

### Efficiency results

We first consider the efficiency metrics and their trends across each trial.

The navigation time from the beginning of the experiment to the robot’s arrival at its goal is shown in Fig. 3.5. Shorter navigation time corresponds to a more efficient movement by the robot. We observe a consistent trend of increased time to goal when the robot has the larger  $\beta$  value. This aligns with our intuition: the robot that turns more sharply to avoid a passing human will spend time traveling to its goal as it moves further off course. Thus, tuning the  $\beta$  parameter to a smaller value, thereby limiting the size of a turn that the robot can make, can improve efficiency.

A comparable trend is observed in the robot’s path length, as shown in Fig. 3.6. A shorter path length

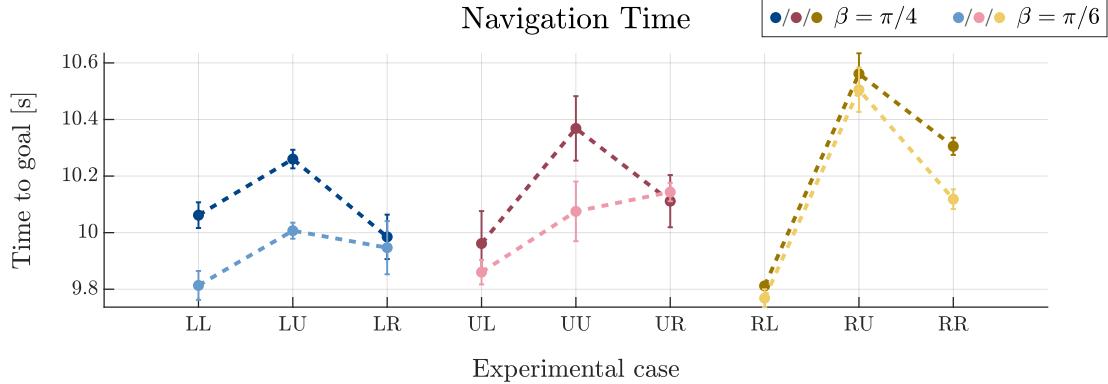


Figure 3.5: The average navigation times of a **NOD** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

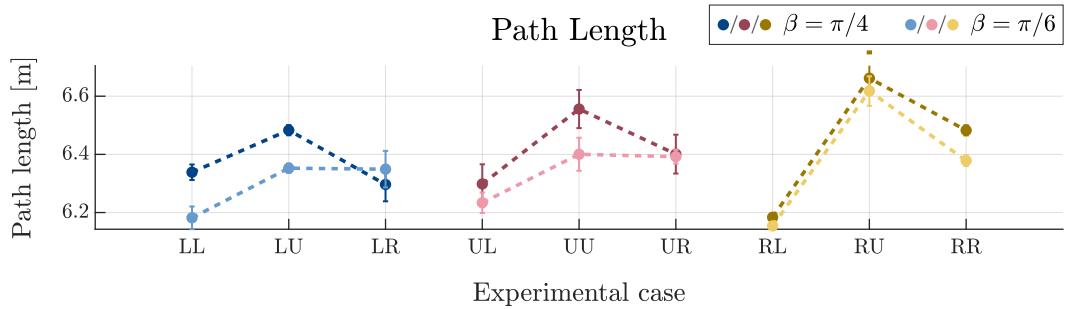


Figure 3.6: The average path lengths of a **NOD** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

corresponds to a more direct and efficient movement toward the goal. Again, a higher  $\beta$  value in NOD consistently produced longer paths towards the goal. This indicates that the robot moved less efficiently when permitted to make a larger turn away from an approaching human. We note that the LR and UR experimental cases, where the human passed on their right, did not follow this trend. The data of each trial reveals that in these LR cases and for each  $\beta$  value, there are two out of three trials with markedly similar path lengths and one outlier value: for  $\beta = \pi/4$ , the outlier is lower than the others, and for  $\beta = \pi/6$ , the outlier is higher than the others. If these outliers were removed from the dataset for the case, the trend would hold as the average path length would be higher for the larger  $\beta$  value.

Finally, we examine the centerline deviation of the paths taken by the robot in each trial. This metric quantifies how far the robot deviated from the straight go-to-goal path while navigating around the oncoming human and is shown in Fig. 3.7. This was computed by summing the perpendicular distances of all trajectory

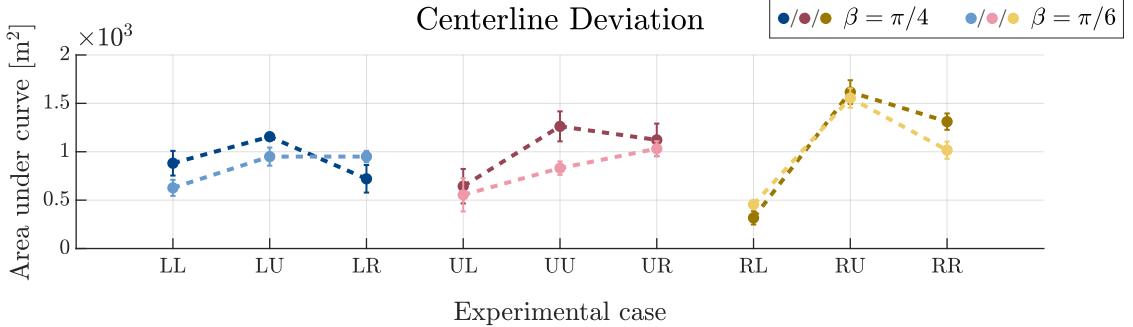


Figure 3.7: The average centerline deviations of a **NOD** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

points  $\mathbf{x}(t) = \mathbf{x}_t$  within the trial from the line created connecting the initial location  $\mathbf{x}(t = 0) = \mathbf{x}_0$  and goal location  $\mathbf{x}_g$ . This equation is as follows:

$$\text{centerline deviation}(\mathbf{x}_0, \mathbf{x}_t, \mathbf{x}_g) = \int_{\mathbf{x}_0}^{\mathbf{x}_g} \frac{|(y_g - y_0)x_t - (x_g - x_0)y_t + x_g y_0 + y_g x_0|}{\sqrt{(y_g - y_0)^2 + (x_g - x_0)^2}} d\mathbf{x}_t.$$

This follows from the distance found between a single point (here,  $\mathbf{x}_t$ ) and a line that is defined by two points (here, the line between points  $\mathbf{x}_0$  and  $\mathbf{x}_g$ ). A lower deviation value indicates that the robot remained more aligned with a direct trajectory, reflecting higher efficiency. We again observe that a larger  $\beta$  value results in a greater deviation from the centerline, consistent with the findings on path length. However, this pattern does not hold for the experimental cases of LR and RL, in which the robot initially moved with a bias that conflicted with the human's intended passing side. This created paths that initially veered one way but ultimately turned in the other direction, allowing for cooperative passing. In these cases, the robot's internal bias was overridden by a more substantial influence from the human's proxy opinion state. Still, the robot was internally influenced by its bias  $b$  that contradicted the direction taken to adjust its course to accommodate a human. Therefore, these trials exhibited lower deviations from the centerline in their trajectories than other trials with the same  $\beta$ , suggesting that this resistant yet responsive adjustment to human movement can produce more efficient interactions than those with immediate, bias-driven avoidance. In all efficiency metrics, we observe the same general trend across experimental cases – that robot performance is least efficient when paired with an unaware human (cases LU, UU, RU). This result is expected, as an unaware human provides no clear proxy opinion to the robot, leading the robot to build its preferred passing side based solely on small perturbations in the environment rather than a dedicated direction to cooperate. This causes a delay in the formation of strong opinions, which in turn delays the cooperative passing process.

An increase in efficiency was noted in the biased robot (cases LU and RU) of previous work [25] and in the left biased robot (case LU) in these experiments, compared to the unbiased robot facing an unaware human case. This follows logically, as a biased robot has an internal prompt to form an opinion and take some action when interacting with a human. In contrast, an unbiased robot must build from an unopinionated state with minimal input from the human. However, an unbiased robot is the most sensitive and responsive to the movements of a passing human. This pattern highlights the inherent trade-off between efficiency in goal-directed movement and flexibility as a cooperative passing partner in a human-robot interaction.

We note that there is a consistent anomaly in navigation duration across trials where the robot's bias  $b = -0.3$  indicates a right-turning preference. One would expect relatively symmetric results between biased trials with  $b \neq 0$ , as their biases were equal but opposite; therefore, passing behavior was anticipated to be symmetric and mirrored across the centerline. However, these right-biased trials resulted in longer paths to the goal, affecting the navigation time, path length, and centerline deviation. Upon reflection, we believe that the robot's goal position may have become slightly misaligned as the trials progressed, resulting in an uneven goal distance, particularly in right-biased experiments.

### Spatial intrusiveness results

We now consider the spatial intrusiveness metrics and their trends across each trial.

The minimum distance reached between the robot and human in each case is shown in Fig. 3.8. The lower the minimum distance, the closer the robot came to colliding with the human, which indicates more spatially intrusive passing behavior. We observe a consistent pattern across cases: the unbiased robot reaches a markedly lower minimum distance than its biased counterparts. The lowest minimum distance for both  $\beta$  values was the unbiased robot, unaware human (UU) case. As expected, trials with the lower  $\beta$  value consistently resulted in shorter minimum distances across all cases. These results indicate that a biased robot with ample turning allowance, as provided by a high  $\beta$ , creates the least spatially intrusive setting for successful passing.

Finally, we examine the invasion distance metric to analyze the robot's ability to naturally maintain socially acceptable passing conventions. These results are illustrated in Fig. 3.9. Ideally, the robot should not invade the personal space (1.2 m) surrounding the passing human, and personal space would be maintained if the invasion distance were 0. In all cases where the human's movement was directed—either left (LL, UL, RL) or right (LR, UR, RR)—the robot maintained a zero invasion distance. Only when the human moved in the unaware case directly towards the goal (LU, UU, RU) did the robot invade their personal space. The invasion distance was noticeably lower when the  $\beta$  value was higher, such that even if the robot could not maintain full personal space, it was able to move itself further away quickly.

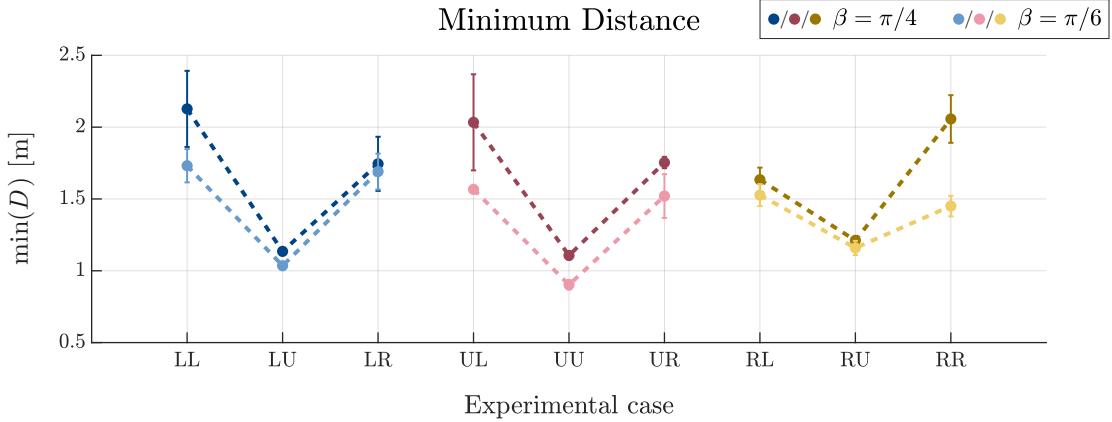


Figure 3.8: The average minimum distances between a **NOD** robot and a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

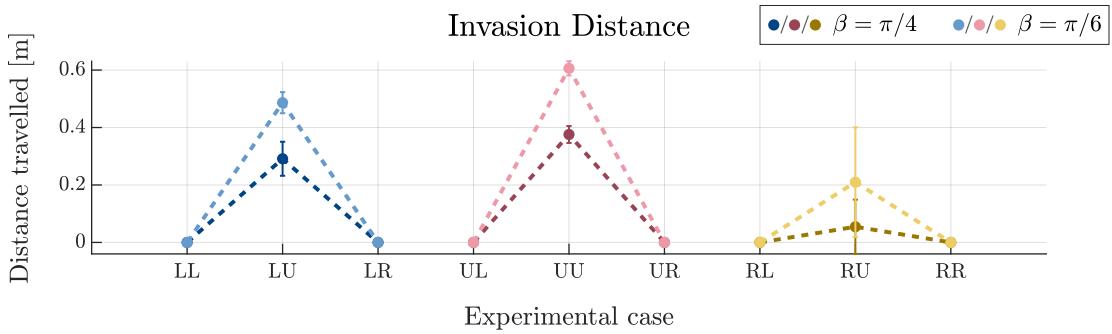


Figure 3.9: The average invasion distances by a **NOD** robot to a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 3.4. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

### 3.4.2.3 Single-robot, two-human experiments

The setup of each experiment was as follows. The robot was placed at an initial location (3.5, -3.5)m in the space, facing its goal at (-3.5,  $\sim$ 3.7)m directly. Two human participants would move around the room of their own accord. Four representative navigation trials are reported in this work to showcase the flexibility of the NOD control of the robot. Trials are each qualitatively described, and their performance metrics are subsequently compared.

#### Robot instructions

The robot was programmed to move at a constant speed  $V_r = 0.75$ m/s toward its goal location, modifying its trajectory when encountering movers according to the navigation model (3.3) with parameters  $\tau_z = 0.1, d = 0.1, a = 0.1, \gamma = 4, b = 0, K_u = 0.0027, U = 1.1, n = 7, k_\theta = 1$  and  $\beta = \pi/4$ . The robot was kept unbiased ( $b = 0$ ) for maximum flexibility in passing behavior. The critical distance  $D_{crit}$  is the only parameter that differed between trials, set such that the resultant paths qualitatively match those observed in [25].

#### Human participant instructions

Two participants were instructed to walk freely in the room at any speed, without a requirement for consistency in direction or pace, if desired. The selected trials presented here show each human moving from an initial location directly to a single goal location.

#### Trial descriptions

Trial 1: Two humans walk independently within the laboratory space as shown in Fig. 3.10. This experiment used human navigation data from Fig. 7.2a in [25], reproduced in Chapter 7. Human 1 moves perpendicularly across the robot's path from the robot's right side towards its left. We see the sensitivity rise as Human 1 approaches, the robot's opinion pitches negative, and the robot turns towards its right to quickly avoid Human 1. The sensitivity drops as Human 1 is passed, but rises again as Human 2 remains stationary along the robot's path on the robot's right side. The opinion pitches upward and becomes positive, causing the robot to turn to the left. The robot then reaches its goal.

Trial 2: Two humans walk independently but close together within the laboratory space as shown in Fig. 3.11. This experiment used human navigation data from Fig. 7.2c in [25], reproduced in Chapter 7. This interaction causes a more pronounced turn by the robot and extends the trial duration by approximately one second compared to Trial 1. Human 2 moves at an angle across the robot's path from the robot's left side towards its right side. The sensitivity is higher than that seen in Trial 1, as this human is initially closer to the robot. The robot's opinion becomes positive and turns to its left to avoid Human 2. Human 1 then appears very nearby in the robot's path on its left side. The robot's sensitivity increases further, but the

opinion reverses sharply to become negative, causing the robot to turn right. With no further humans to avoid, the robot moves towards its goal.

Trial 3: Two humans walk closely together within the laboratory space as shown in Fig. 3.13. This experiment used human navigation data from Fig. 7.2c in [25], reproduced in Chapter 7. Humans 1 and 2 walk so close to each other that a single sensitivity peak captures the influence of both. This rise in sensitivity prompts a more positive opinion to form, one of larger magnitude than the opinions seen in trials 1 and 2, and the robot turns left. This single turn avoids both humans, allowing the robot to then go towards its goal unobstructed for the duration of its path.

Trial 4: Two humans walk together within the laboratory space as shown in Fig. 3.13. This experiment used human navigation data from Fig. 7.2e in [25], reproduced in Chapter 7. Though still nearby one another, Humans 1 and 2 begin moving at different times. This allows the robot to interact with each independently rather than as a unit, as seen in Trial 3. Human 2 is seen first on the robot's right side, increasing the robot's sensitivity and pitching its opinion positively. The robot turns to the left, but then is immediately on a collision path with Human 1, who is positioned on its left side. This raises the sensitivity again, and the opinion becomes negative. The robot turns right to avoid Human 1, then moves towards its goal. Visual inspection of the resultant trajectories shows that the robot comes noticeably closer to the passing humans in this trial than in the others.

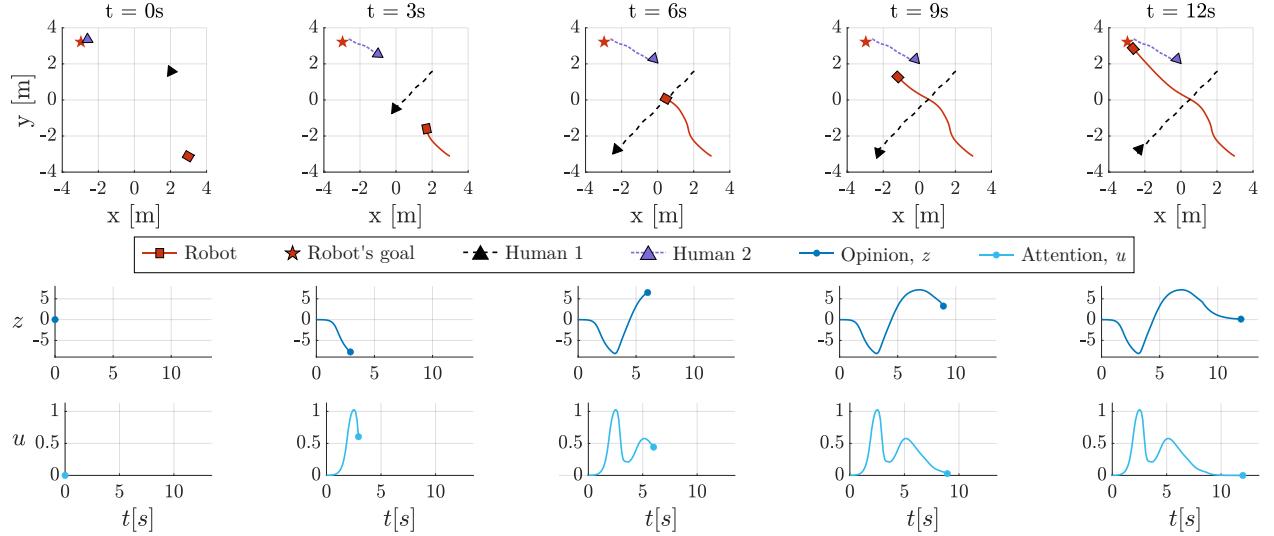


Figure 3.10: NOD Trial 1. Human 1 walks across the robot's path, causing the robot to turn right to avoid them. Meanwhile, Human 2 abruptly stops moving while in the robot's path, prompting the robot to turn left to avoid them.  $D_{crit} = 3\text{m}$ . Experimental data from [25], illustrated in Fig 7.2a of Chapter 7.

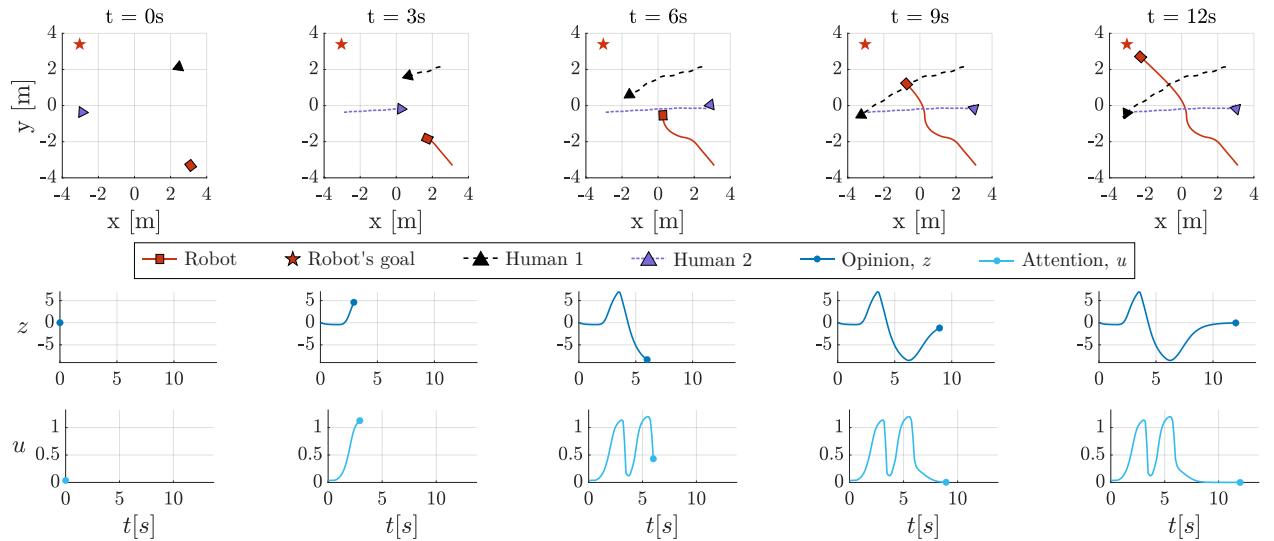


Figure 3.11: NOD Trial 2. Humans 1 and 2 walk across the robot's path. The robot first turns left to avoid Human 2, then right to avoid Human 1.  $D_{crit} = 4\text{m}$ . Experimental data from [25], illustrated in Fig 7.2c of Chapter 7.

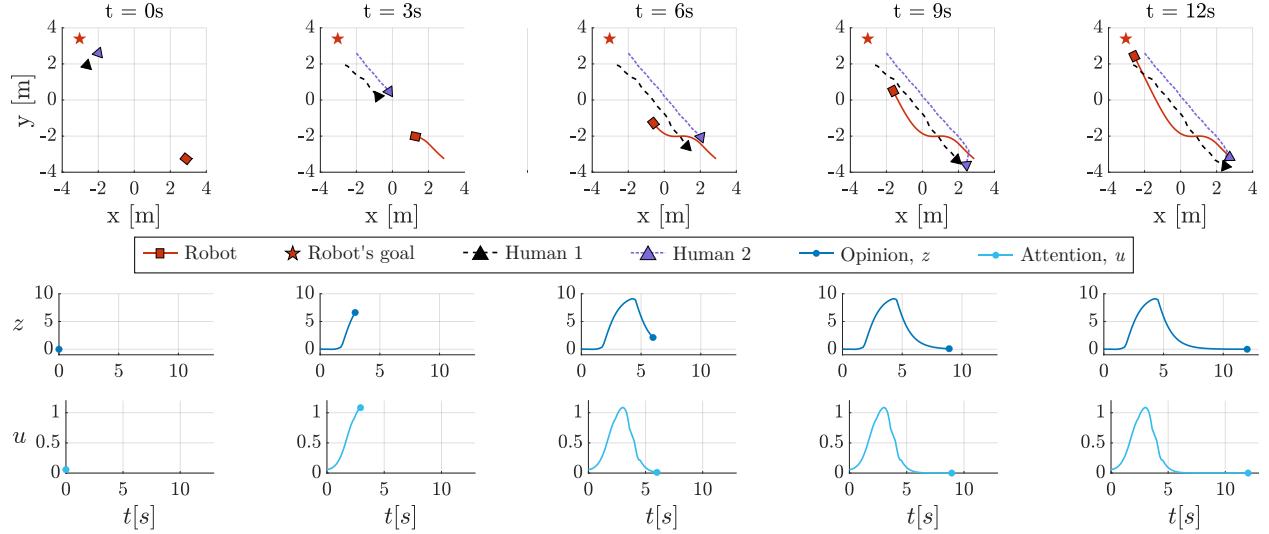


Figure 3.12: NOD Trial 3. Humans 1 and 2 walk toward the robot's initial position with no passing space in between them. The robot avoids both humans by turning left.  $D_{crit} = 5\text{m}$ . Experimental data from [25], illustrated in Fig 7.2d of Chapter 7.

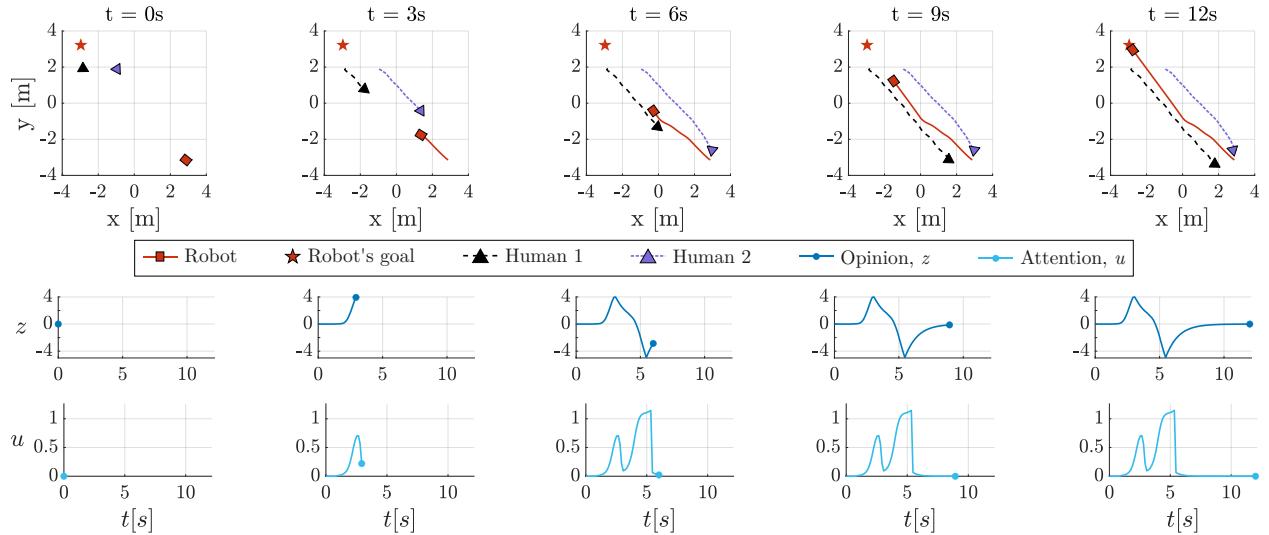


Figure 3.13: NOD Trial 4. Humans 1 and 2 walk toward the robot's initial position with passing space in between them. The robot initially dodges Human 2 with a left turn, finds itself in the path of Human 1, and dodges Human 1 with a right turn.  $D_{crit} = 2.5\text{m}$ . Experimental data from [25], illustrated in Fig 7.2e of Chapter 7.

## Trial comparisons and results

Finally, we compare the performance metrics across trials. These are listed in Table 3.3. Trial 4 exhibits the shortest navigation time and path length, but also the lowest minimum distances to both humans and the most extended invasion distances seen across all trials. Thus, it is our most efficient yet least comfortable (as it is the most spatially intrusive) case. In contrast, Trial 2 is the least efficient yet most spacious passing case with the least spatial intrusiveness.

Trials 1 and 2 are similar to one another, as are Trials 3 and 4. We compare these pairs. Between Trials 1 and 2, we observe that neither invades the personal space of either Human 1 or 2; however, Trial 2 maintains more space throughout the experiment. However, this comes at the cost of a longer path length and increased navigation time. Between Trials 3 and 4, we see that both invade the personal space of at least Human 1 along their paths. In these Trials, the robot needs to swiftly change its opinion to the opposite preferred passing side as it interacts with little room between itself and the Human it passes. We believe this is due to the inherent difficulty of managing real-time sequential decisions; however, parameters can be tuned to improve performance. We explore the trade-off between forming robust, strong opinions to produce comfortable passing behavior and the flexibility of changing these opinions quickly to better handle sequential decision-making.

Trial #	Navigation time [s]	Path Length [m]	$\min(D_1)$ [m]	$\min(D_2)$ [m]	Invasion Distance of Human 1 [m]	Invasion Distance of Human 2 [m]
1	12.53	8.77	2.03	1.38	0	0
2	13.20	9.24	2.18	1.55	0	0
3	12.93	9.05	0.54	1.29	0.28	0
4	12.13	8.49	0.36	1.01	0.47	0.28

Table 3.3: Performance metrics of the robot navigating in trials 1-4 (reported in Figures 3.10- 3.13).

## Chapter 4

# Spiking Nonlinear Opinion Dynamics (S-NOD) for agile decision-making

CHARLOTTE CATHCART, IAN XUL BELAUSTEGUI, ALESSIO FRANCI, AND NAOMI EHRICH LEONARD

### Statement of contributions

The material presented in this chapter also appears in [26] and was presented at the American Control Conference (ACC) in Denver, CO in 2025. I, Charlotte Cathcart, am the lead author of the manuscript. Ian Xul Belaustegui is the second author of the manuscript. Naomi Ehrich Leonard and Alessio Franci served as joint senior co-authors, making equal contributions to the published work. Naomi Ehrich Leonard and Alessio Franci conceived of the presented S-NOD model. All four authors developed the given theory of S-NOD and took part in discussions on its presentation in this manuscript. I, Charlotte Cathcart, developed the application and embodiment of S-NOD in a robotic navigation system, drawing on the work on NOD in a proactive social robot navigation system [25], as well as preliminary code on embodied spiking decision-making by Alessio Franci. I, Charlotte Cathcart, drafted Lemma 1. I, Charlotte Cathcart, and Alessio Franci drafted Proposition 1. Ian Xul Belaustegui drafted Proposition 2. All authors contributed to the editing of these theoretical results. Alessio Franci and Naomi Ehrich Leonard led the writing of Section 4.1. Ian Xul Belaustegui led the writing of Section 4.4.1. I, Charlotte Cathcart, led the writing of all other sections. All four authors contributed to the editing of the final manuscript. I, Charlotte Cathcart, created all figures within and animated videos for the manuscript. I, Charlotte Cathcart, and Ian Xul Belaustegui contributed to simulating various results for the work, with all presented simulations for figures done by me. Naomi

Leonard supervised the acquisition of funding for this project through ONR grant N00014-19-1-2556.

There is a minor change made in each Fig. 4.2b and 4.3 from their original printing to maintain consistent notation within the manuscript. Figure 4.4 has been reoriented from its original printing to increase its readability, with minor adjustments to its caption to match. Finally, an additional appendix is included to present an updated version of Fig. 4.3.

## Abstract

We present, analyze, and illustrate a first-of-its-kind model of two-dimensional excitable (spiking) dynamics for decision-making over two options. The model, Spiking Nonlinear Opinion Dynamics (S-NOD), provides superior agility, characterized by fast, flexible, and adaptive responses to rapid and unpredictable changes in context, environment, or information received about available options. S-NOD derives through the introduction of a single extra term to the previously presented Nonlinear Opinion Dynamics (NOD) for fast and flexible multi-agent decision-making behavior. The extra term is inspired by the fast-positive, slow-negative mixed-feedback structure of excitable systems. The agile behaviors brought about by the new excitable nature of decision-making driven by S-NOD are analyzed in a general setting and illustrated in an application to multi-robot navigation around human movers.

### 4.1 Introduction

The fast, flexible, and adaptive behavior observed in biology owes much to the excitable (spiking) nature of cellular signaling [30–33]. Models of excitability represent the analog molecular and/or biophysical processes that produce spikes in response to stimuli. These models inherit the adaptive behavior of analog (continuous) systems and the reliability of digital (discrete) systems, foundational to spiking control systems [36] and neuromorphic engineering [37]. However, existing models describe single-input/single-output [38] spike-based signal processing. This is spiking activity that can only encode binary, single-option decisions: to spike or not spike, as determined by the input signal pushing the system toward its excitability threshold [39]. This limits the use of these models in studying and designing spiking decision-making over *multiple* observed options, e.g., with neurons in cortical columns that activate at specifically oriented visual patterns [40] or in sensorimotor (sensory-based action) decision-making [41].

We present a generalized model of excitable (spiking) dynamics that allows for fast, flexible, and adaptive decision-making over multiple options. In this paper we focus on two-option spiking in a two-dimensional, two-timescale model, and we use “agile” to mean “fast, flexible, and adaptive.” To the best of our knowledge,

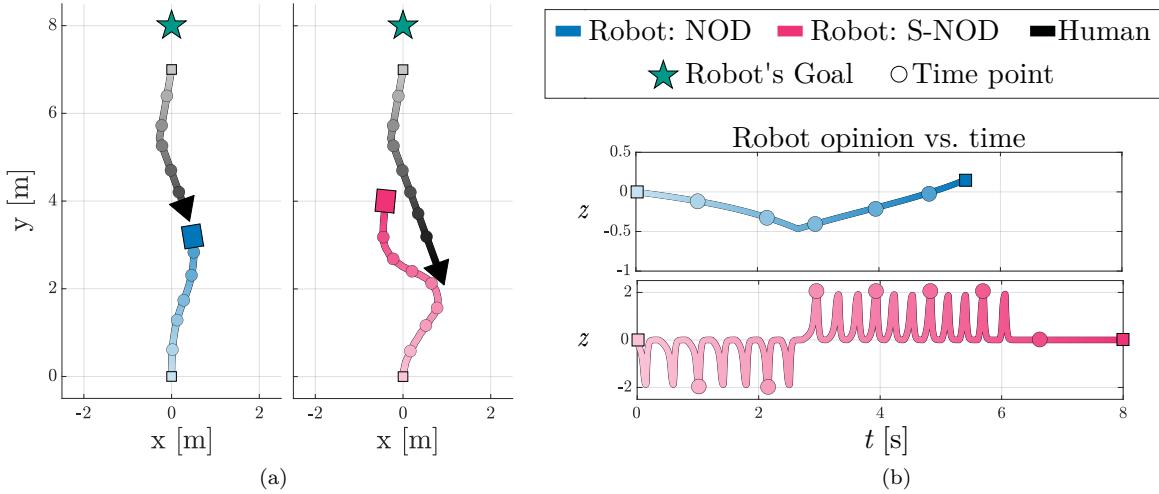


Figure 4.1: (a): Trajectory of a robot controlled with NOD (S-NOD) is shown with a blue (pink) line as it navigates towards a goal (star) in the presence of an oncoming human mover (black). The NOD robot experiences a collision while the S-NOD robot does not. (b): Opinion  $z$  of the robot over time  $t$ . Circles mark matching points in time along trajectories and time-evolution of opinions. [Figures are animated at `https://spikingNOD.github.io`.](https://spikingNOD.github.io)

this is the first such model to generalize spiking to more than one option, i.e., spiking that can occur in any of the multiple directions corresponding to the multiple options available to the excitable decision-maker. We call our model *Spiking Nonlinear Opinion Dynamics* (S-NOD) as it derives from the Nonlinear Opinion Dynamics (NOD) model of [1,42] through the introduction of an extra term that makes NOD excitable, i.e., spiking.

NOD models the time evolution of opinions of a group of agents engaged in a collective decision-making process over a set of options. The derivation of NOD was tailored to model and study the principles of fast and flexible decision-making in biological collectives [20, 21] and to use these principles to design fast and flexible decision-making in built collectives [1, 22, 24]. Decision-making driven by NOD [1, 42] is fast because it can diverge quickly from indecision even in the absence of informative inputs about the options. It is flexible because the sensitivity (versus robustness) of opinion formation to informative inputs is tunable. NOD exhibits a mixed-feedback structure [39, 57]: opinion formation arises from the balance of a negative feedback loop that regulates agent opinions to a neutral solution and positive feedback loops (at single-agent and network levels) that destabilize the neutral solution and trigger nonlinear opinion formation. Both speed and flexibility are determined by a tunable threshold for opinion formation where negative and positive feedback are perfectly balanced and the dynamics become singular.

To derive S-NOD, we introduce in NOD a *slow* regulation term inspired by the dynamics of excitable (spiking) signal processing systems. The resulting excitable dynamics give S-NOD its superior agility in

decision-making. Where NOD allows for a fast decision, S-NOD allows for autonomous fast sequential decision-making, not requiring any ad-hoc reset of the model state once a decision is made. Where NOD is flexible, S-NOD is flexible and capable of fast “changes of mind” and adaptive responses when the information about the options changes rapidly and unexpectedly. Further, S-NOD provides on-demand (event-based) opinion formation in the sense that large opinions are formed sparsely in time as “decision events” and only when context requires it. This makes S-NOD efficient. The agility of S-NOD is illustrated in Fig. 4.1 in the context of a robot navigating around a human mover as studied in [25].

Our major contributions in this paper are the presentation, analysis, and illustration of a first-of-its-kind model of two-dimensional excitable (spiking) dynamics for decision-making over two options, which provides superior agility (fast, flexible, and adaptive behavior), especially important in changing contexts. Also, in Section 4.2, we present a new analysis of the singularity in NOD for a single decision-making agent and two options. We prove how a single feedback gain  $K_u$  tunes opinion formation. We show that when  $K_u$  gets too large, an opinion can become so robust that it will not change quickly enough if a new input arrives in favor of the alternative option. In Section 4.3, we present S-NOD for a single agent and two options. S-NOD provides fast opinion changes with the arrival of new informative input even for large  $K_u$ . We use geometric analysis to show the existence of the spiking limit cycles associated with the excitable behavior of S-NOD. We show further how this provides agility in decision-making. We generalize S-NOD to a network of multiple agents and apply it to a social robot navigation problem in Section 4.4.

## 4.2 Fast and flexible decision-making: NOD

We recall NOD [1, 42] in Section 4.2.1 for a single agent evolving continuously over time its real-valued opinion about two mutually exclusive options with possible input present. In Section 4.2.2 we analyze stability of the neutral opinion solution and prove conditions on feedback gain  $K_u$  that determine the type of singularity (type of pitchfork bifurcation) in the dynamics. We show that, by shaping bifurcation branches,  $K_u$  tunes opinion formation. In Section 4.2.3, we show limits on tunability of NOD that sacrifice agility in decision-making, motivating S-NOD, which is introduced in Section 4.3.

### 4.2.1 NOD for single decision-maker and two options

We let an agent represent a single decision-maker. Let  $z(t) \in \mathbb{R}$  define the agent’s opinion at time  $t$  about two mutually exclusive options. The more positive (negative) is  $z$ , the more the agent favors (disfavors) option 1 and disfavors (favors) option 2. When  $z = 0$ , the agent is neutral about the two options, i.e., in a state of indecision. Let  $u(t) \geq 0$  define the attention of the agent at time  $t$  to its observations;  $u$  is implemented as a

gain in the dynamics. Let  $b(t) \in \mathbb{R}$  define an input signal at time  $t$  that represents external stimulus and/or internal bias. When  $b(t) > 0$  ( $b(t) < 0$ ), it provides information (evidence) in favor of option 1 (option 2).

Decision-making variables  $z$ ,  $u$  evolve in continuous time  $t$  according to the following NOD, adapted from [1, 42]:

$$\tau_z \dot{z} = -d z + \tanh(u(az) + b), \quad (4.1a)$$

$$u = u_0 + K_u z^2, \quad (4.1b)$$

where  $\dot{z} := dz/dt$ .  $\tau_z > 0$  is a time constant, and damping coefficient  $d > 0$  weights the negative feedback on  $z$  that regulates to the neutral solution  $z = 0$ . The second term in (4.1a) provides a nonlinear positive feedback on  $z$  with weight given by the product of  $u$  and amplification coefficient  $a > 0$ , plus the effects of  $b$ . The saturation nonlinearity given by the  $\tanh$  function enables fast-and-flexible decision-making through opinion-forming bifurcations [1, 42]. The positive feedback gain, attention  $u$ , is state-dependent according to (4.1b) and grows with  $z^2$ . Hence, small deviations from the neutral solution ( $z = 0$ ) in response to small inputs  $b$  leave attention  $u$  small and do not trigger large, nonlinear opinion formation. Large enough deviations from the neutral solution in response to large enough inputs  $b$  cause a sharp increase in attention  $u$  and trigger large, nonlinear opinion formation  $z$ . The resulting implicit threshold distinguishing small and large inputs is tuned by  $u_0$  and  $K_u$ .

#### 4.2.2 Analysis of single-agent, two-option NOD

We study the dynamics and stability of solutions of system (4.1a)-(4.1b) using bifurcation theory. A local bifurcation refers to a change in number and/or stability of equilibrium solutions of a nonlinear dynamical system as a (bifurcation) parameter is changed. The state and parameter values at which this change occurs is the *bifurcation point*. At a bifurcation point, one or more eigenvalues of the Jacobian of the model must have zero real part [44, 46], i.e., a bifurcation point is a singularity of the model vector field.

Our main interest is in the *pitchfork* bifurcations. There are two generic types of pitchforks. A *super-critical pitchfork* bifurcation describes how one stable solution becomes unstable and two stable solutions emerge as the bifurcation parameter increases. A *subcritical pitchfork* bifurcation describes how two unstable solutions disappear and one stable solution becomes unstable as the bifurcation parameter increases.

Our objective is to understand how thresholds of fast-and-flexible decision-making are controlled by the model parameters with the goal of designing feedback control laws for those parameters that can make

decision thresholds adaptive to context. Substituting (4.1b) into (4.1a) yields

$$\tau_z \dot{z} = -d z + \tanh \left( (u_0 + K_u z^2) \cdot (az) + b \right). \quad (4.2)$$

We first study (4.2) in the case  $b = 0$ , i.e., when there is no evidence to distinguish the options, and bifurcations are symmetric. Then, we introduce  $b \neq 0$  and use unfolding theory [44] to understand the effects of inputs.

**Lemma 1.** (*NOD Taylor expansion and singularity*): *Consider (4.2) and let  $b = 0$ . Then the solution  $z = 0$  is always an equilibrium, and the Taylor expansion of (4.2) about  $z = 0$  is*

$$\dot{z} = \frac{1}{\tau_z} \left( (a u_0 - d) z + a \left( K_u - \frac{a^2 u_0^3}{3} \right) z^3 + a^3 u_0^2 \left( \frac{2a^2 u_0^3}{15} - K_u \right) z^5 \right) + \mathcal{O}(z^7). \quad (4.3)$$

A singularity exists at  $(u_0, z) = (u_0^*, 0)$ , with  $u_0^* = \frac{d}{a}$ . The solution  $z = 0$  is stable (unstable) when  $u_0 < u_0^*$  ( $u_0 > u_0^*$ ).

*Proof.* When  $z = 0$  and  $b = 0$ , the right hand side of (4.2) is zero, thus  $z = 0$  is always an equilibrium. We expand (4.2) with  $b = 0$  about  $z = 0$ . The Taylor expansion of the hyperbolic tangent is  $\tanh(w) = w - w^3/3 + 2w^5/15 + \mathcal{O}(w^7)$ . Using this in (4.2) yields (4.3). The Jacobian  $J(z) = \frac{d\dot{z}}{dz}$  of (4.3) evaluated at  $z = 0$  is  $J(0) = (a u_0 - d)/\tau_z$ , which is singular when  $u_0 = u_0^* = \frac{d}{a}$ . When  $u < u_0^*$ ,  $J(0) < 0$  thus  $z = 0$  is exponentially stable. When  $u > u_0^*$ ,  $J(0) > 0$  thus  $z = 0$  is unstable. ■

We next explore in Proposition 1 and Fig. 4.2 the effect of parameter  $K_u$  on the cubic and quintic terms of (4.3) and its role in determining the type of singularity at  $(u, z) = (u_0^*, 0)$ .

**Proposition 1.** ( *$K_u$  determines singularity type*): *Let  $b = 0$ ,  $u_0^* = \frac{d}{a}$ . The singularity of dynamics (4.2) at  $(u, z) = (u_0^*, 0)$  as proved in Lemma 1 corresponds to a supercritical pitchfork bifurcation for  $K_u < \frac{d^3}{3a}$ , a quintic pitchfork bifurcation for  $K_u = \frac{d^3}{3a}$ , and a subcritical pitchfork bifurcation for  $K_u > \frac{d^3}{3a}$ .*

*Proof.* Denote  $p = (K_u - \frac{d^3}{3a})$  and  $q = (\frac{2d^3}{15a} - K_u)$  as the coefficients of  $az^3/\tau_z$  and  $ad^2 z^5/\tau_z$  in (4.3) resp. at  $u_0 = u_0^* = \frac{d}{a}$ . When  $K_u < \frac{d^3}{3a}$  ( $K_u > \frac{d^3}{3a}$ ), then  $p < 0$  ( $p > 0$ ) and (4.3) is the normal form of the supercritical (subcritical) pitchfork bifurcation [44]. When  $K_u = \frac{d^3}{3a}$ , then  $p = 0$ ,  $q < 0$  and (4.3) is the normal form of the quintic pitchfork, by recognition problem [44, Prop. VI.2.14] and its  $Z_2$ -symmetric universal unfolding [44, Prop. VI.3.4; Fig. VI.3.3]. ■

Proposition 1 uncovers the key role of  $K_u$  in tuning opinion formation (Fig. 4.2): (i)  $K_u$  controls the supercritical vs. subcritical nature of opinion formation, and (ii) increasing  $K_u$  increases opinion strength

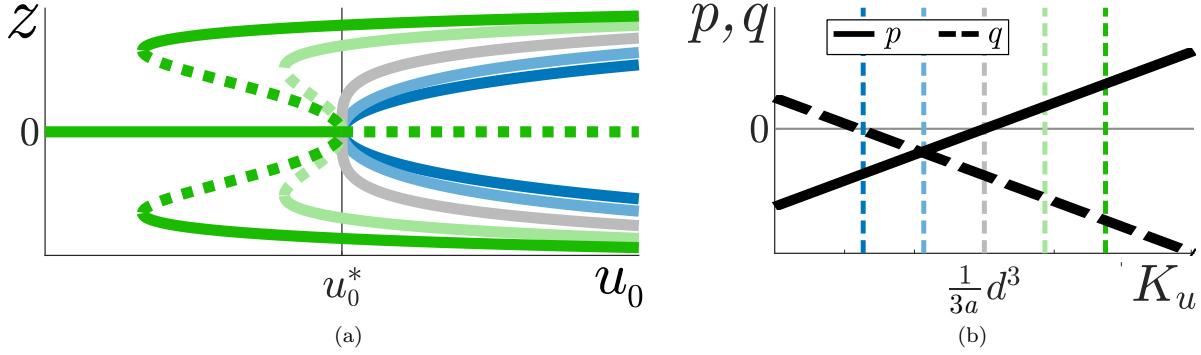


Figure 4.2: The effect of  $K_u$  on the bifurcation diagram of (4.2) and the cubic and quintic terms of (4.3). (a): Bifurcation diagrams of NOD (4.2) with  $K_u$  values corresponding to the vertical dashed lines in (b). Stable (unstable) solutions are shown with solid (dotted) lines. The bifurcation point is  $(u_0, z) = (u_0^*, 0)$ . (b): Coefficient  $p$  ( $q$ ) as a function of  $K_u$  shown as a solid (dashed) black line.

of non-neutral solutions. Bifurcation diagrams in Fig. 4.2a plot equilibrium solutions of (4.2), i.e., solutions of  $\dot{z} = 0$ , as a function of bifurcation parameter  $u_0$  for different values of  $K_u$  (Fig. 4.2b). The singularity at  $(u_0, z) = (u_0^*, 0)$  is a pitchfork bifurcation: blue, gray, green lines show supercritical, quintic, subcritical solutions, respectively. For all  $K_u$ : when  $u_0 < u_0^*$ ,  $z = 0$  is stable; when  $u_0 > u_0^*$ ,  $z = 0$  is unstable and there is a bistable symmetric pair of solutions. When  $K_u < \frac{d^3}{3a}$ , the pitchfork is supercritical: there are no other solutions. When  $K_u > \frac{d^3}{3a}$ , the pitchfork is subcritical: two stable non-neutral solutions appear for  $u_0 < u_0^*$  through saddle-node bifurcations. As  $K_u$  grows more positive, these solutions emerge for smaller values of  $u_0$  and increase in magnitude, reflecting increasing opinion strength.

#### 4.2.3 Limitation on tuning of NOD

We prove that the region of multi-stability in the subcritical bifurcation of NOD grows as  $K_u$  gets large.

**Proposition 2.** ( $K_u$  determines region of multi-stability): Let  $b = 0$  and  $(u_0^\dagger, z^\dagger)$  be either of the two saddle-node bifurcation points of the subcritical pitchfork of NOD (4.2) for  $K_u > \frac{d^3}{3a}$ . Then  $u_0^\dagger$  is a monotonically decreasing function of  $K_u$ , i.e.,  $\frac{\partial u_0^\dagger}{\partial K_u} < 0$ .

*Proof.* Let  $K_u^\dagger > \frac{d^3}{3a}$  and  $f(z, K_u, u_0) := -dz + \tanh(az(u_0 + K_u z^2))$ . By hypothesis,  $f(z^\dagger, K_u^\dagger, u_0^\dagger) = 0$ . We have  $\frac{\partial f}{\partial u_0}(z^\dagger, K_u^\dagger, u_0^\dagger) = az^\dagger \tanh'(az^\dagger(u_0^\dagger + K_u^\dagger(z^\dagger)^2)) \neq 0$ , since  $z^\dagger \neq 0$ . Following [58], the implicit function theorem shows the existence of  $g : \mathbb{R}^2 \rightarrow \mathbb{R}$  such that for some neighborhood of  $(z^\dagger, K_u^\dagger)$ ,  $f(z, K_u, g(z, K_u)) = 0$ . We get

$$\frac{\partial g}{\partial K_u} = - \left( \frac{\partial f}{\partial u_0} \right)^{-1} \left( \frac{\partial f}{\partial K_u} \right) = - \frac{a(z^\dagger)^3 \tanh'(az^\dagger(u_0^\dagger + K_u^\dagger(z^\dagger)^2))}{az^\dagger \tanh'(az^\dagger(u_0^\dagger + K_u^\dagger(z^\dagger)^2))} = -(z^\dagger)^2.$$

Since  $z^\dagger \neq 0$ ,  $u_0^\dagger$  is monotonically decreasing in  $K_u$ . ■

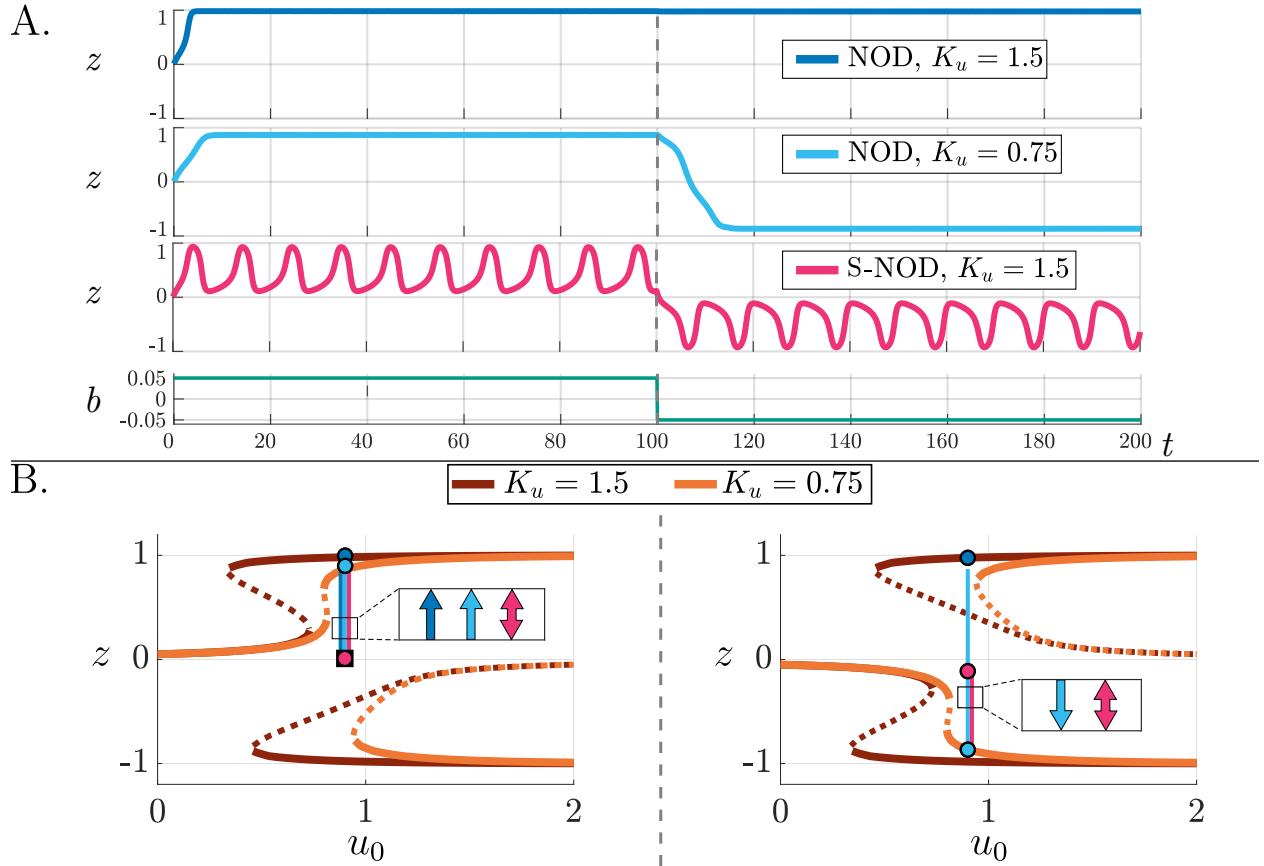


Figure 4.3: Opinion solutions of NOD and S-NOD over time and associated bifurcation diagrams. (A): Opinion solutions  $z$  of NOD (4.2) with larger and smaller values of  $K_u$  and of S-NOD (4.4) with larger value of  $K_u$ , all with initial condition  $z(0) = 0.01$  and  $u_0 = 0.9$ , and  $u_s(0) = 0$  for S-NOD. Input signal  $b$  is also shown over time. (B): Bifurcation diagrams of (4.2) for the two values of  $K_u$ , with the solutions  $z$  in A superimposed: (left) from  $t = 0$  to  $t = 100$  and (right) from  $t = 100$  to  $t = 200$ .

Proposition 2 implies that one limitation of NOD is that large  $K_u$  can make the region of multi-stability of NOD (4.2) so large and robust that solutions can get “stuck” in one of the decision attractors unless very large inputs in favor of another decision state are applied. This is illustrated in Fig. 4.3A, where the first (dark blue) and second (light blue) NOD differ only in their  $K_u$  parameters ( $K_{u_1} > K_{u_2} > \frac{d^3}{3a}$ ) but their solutions are distinctively different. At the stimulus onset ( $b > 0$  for  $0 \leq t < 100$ ), the solution of the first NOD converges to  $z > 0$  much more rapidly than that of the second NOD. However, when the input switches values ( $b < 0$  for  $t \geq 100$ ), the solution of the first NOD gets stuck at a positive value, whereas the solution of the second NOD is able to track the change in input sign. This example reveals a fundamental trade-off between speed/robustness (first, dark blue NOD) and flexibility (second, light blue NOD).

Instead of aiming to fine-tune the gain  $K_u$  around a hard-to-define fast/robust enough yet flexible enough decision-making behavior, we use mixed-feedback principles to make the system *excitable*, inheriting both

the speed of NOD with large  $K_u$  and the flexibility of NOD with small  $K_u$  and imparting system agility. The behavior of the resulting S-NOD is shown in pink in Fig. 4.3A. By generating “decision spikes” the S-NOD is as fast as the high  $K_u$ -NOD and as flexible as the low- $K_u$  NOD. In what follows, we present the S-NOD model, its analysis, and its multi-agent generalization.

## 4.3 Agile decision-making: S-NOD

We present and analyze the Spiking Nonlinear Opinion Dynamics (S-NOD) model.

### 4.3.1 S-NOD for a single agent and two options

We define S-NOD by introducing a slow regulation variable  $u_s$  to NOD (4.2), as in the fast-positive, slow-negative mixed-feedback structure of excitable systems [30–33]:

$$\tau_z \dot{z} = -d z + \tanh \left( (u_0 - u_s + K_u z^2) \cdot (az) + b \right), \quad (4.4a)$$

$$\tau_{u_s} \dot{u}_s = K_{u_s} z^4 - u_s, \quad (4.4b)$$

where  $\tau_{u_s} \gg \tau_z$  is larger by at least an order of magnitude such that  $u_s$  responds more slowly than  $z$ . S-NOD defined by (4.4) describes dynamics with excitability: a fast positive feedback (mediated by  $z$ ) acts to excite the system, while a slow negative feedback (mediated by  $u_s$ ) regulates it back to near the ultrasensitive pitchfork singularity (as seen in Fig. 4.2).

### 4.3.2 Geometric analysis of single-agent, two-option S-NOD

We use phase-plane analysis to study and illustrate the spiking and decision-making behavior of S-NOD (4.4). To construct the phase-plane, we first compute the nullclines.

The  $z$ -nullcline is defined as the solution pairs  $(u_s, z)$  that satisfy  $\dot{z} = 0$  for (4.4a). This is equivalent to solving for the equilibrium solutions of (4.2) as a function of  $u_0$  as in Section 4.2.2. Thus, the  $z$ -nullcline (pink in Fig. 4.4) is analogous to the bifurcation diagram of (4.2), mirrored about the vertical axis and shifted right by  $u_0$ , with  $u_s^* = u_0 - d/a = u_0 - u_0^*$ . When  $b = 0$ , the neutral solution  $z = 0$  is a stable (unstable) equilibrium of (4.4a) for  $u_s > u_s^*$  ( $u_s < u_s^*$ ). The  $u_s$ -nullcline (blue in Fig. 4.4) is defined as the solution pairs  $(u_s, z)$  that satisfy  $\dot{u}_s = 0$  in (4.4b), which gives the quartic parabola  $u_s = K_{u_s} z^4$ . The larger  $K_{u_s}$ , the more narrow this parabola is.

The intersections of the nullclines determine equilibrium solutions of S-NOD (4.4) and, as shown in Fig. 4.4, depend on the value of  $u_0$ . If  $b = 0$ , the *neutral solution*  $(u_s, z) = (0, 0)$  is always an equilibrium.

Two more equilibria, symmetric about  $z = 0$ , may be present for high enough  $u_0$ .

Fig. 4.4a depicts the phase-plane when  $b = 0$  and  $u_0 < u_0^*$ . The nullclines have one point of intersection at the neutral solution. The neutral solution is stable. Trajectories will converge to and settle at this point and no excitable behavior in the decision-making will take place.

Fig. 4.4b depicts the phase-plane when  $b = 0$  and  $u_0 = u_0^*$ . The nullclines have three points of intersection, the neutral solution and two unstable equilibria symmetric about  $z = 0$ . The neutral solution is a saddle-node bifurcation with one exponentially stable eigendirection (along  $z = 0$ ) and one marginally unstable eigendirection (along  $u_s = 0$ ). There are two saddle-node-homoclinic (infinite period) cycles, diverging upward and downward from the saddle-node. In the absence of noise/exogenous perturbations, all trajectories asymptotically converge to the neutral solution. The presence of any arbitrarily small noise makes the trajectory escape from the neutral solution at random time instants along either the upward or downward saddle-node-homoclinic cycle, leading to large prototypical excursions in the  $(u_s, z)$  plane.

These large prototypical excursions resemble the “spiking” trajectories of excitable neuronal system. By analogy, we call them “decision spikes” or “excitable decisions”. In contrast to neuronal spikes which happen in only one direction, decision spikes can happen in as many directions as there are options. For the one-dimensional two-option dynamics studied here, both upward (in favor of option 1) and downward (in favor of option 2) decision spikes are possible.

Fig. 4.4c depicts the phase-plane when  $b = 0$  and  $u_0 > u_0^*$ . Three fixed points are unstable and it is possible to prove, along the same lines as [59], the existence of two limit cycles, symmetric about the horizontal axis  $z = 0$ . These limit cycles are made of repetitive decision spikes, i.e., spiking decision limit cycles. Geometric singular perturbation analysis [60, 61] provides the tools to rigorously prove the existence of these spiking decision limit cycles. Such an analysis goes beyond the scope of this paper. Instead, we leverage Fig. 4.4c to describe qualitatively a typical oscillatory spiking decision behavior in the presence of small noisy perturbations.

Consider Fig. 4.4c and a trajectory with a large initial condition in  $u_s$ . Initially, the trajectory rapidly converges to the  $z = 0$  axis, then slowly slides leftward and approaches the neutral solution. As soon as the trajectory nears this equilibrium, noisy perturbations push it either upward or downward, generating an upward or downward decision spike, respectively. The decision spike trajectory brings the trajectory back to the pitchfork singularity, the next decision spike is generated, and the spiking decision cycle continues.

When input  $b \neq 0$  and  $|b|$  is sufficiently large, the  $z$ -nullcline *unfolds* accordingly to the universal unfolding of the pitchfork [44, Ch. III]. Due to the nullcline unfolding, the phase-plane geometry changes qualitatively as shown in Fig. 4.4d. Similar geometric singular perturbation analysis methods as those employed for the analysis of Fig. 4.4b and 4.4c reveal the existence of a *unique* spiking decision limit cycle associated to

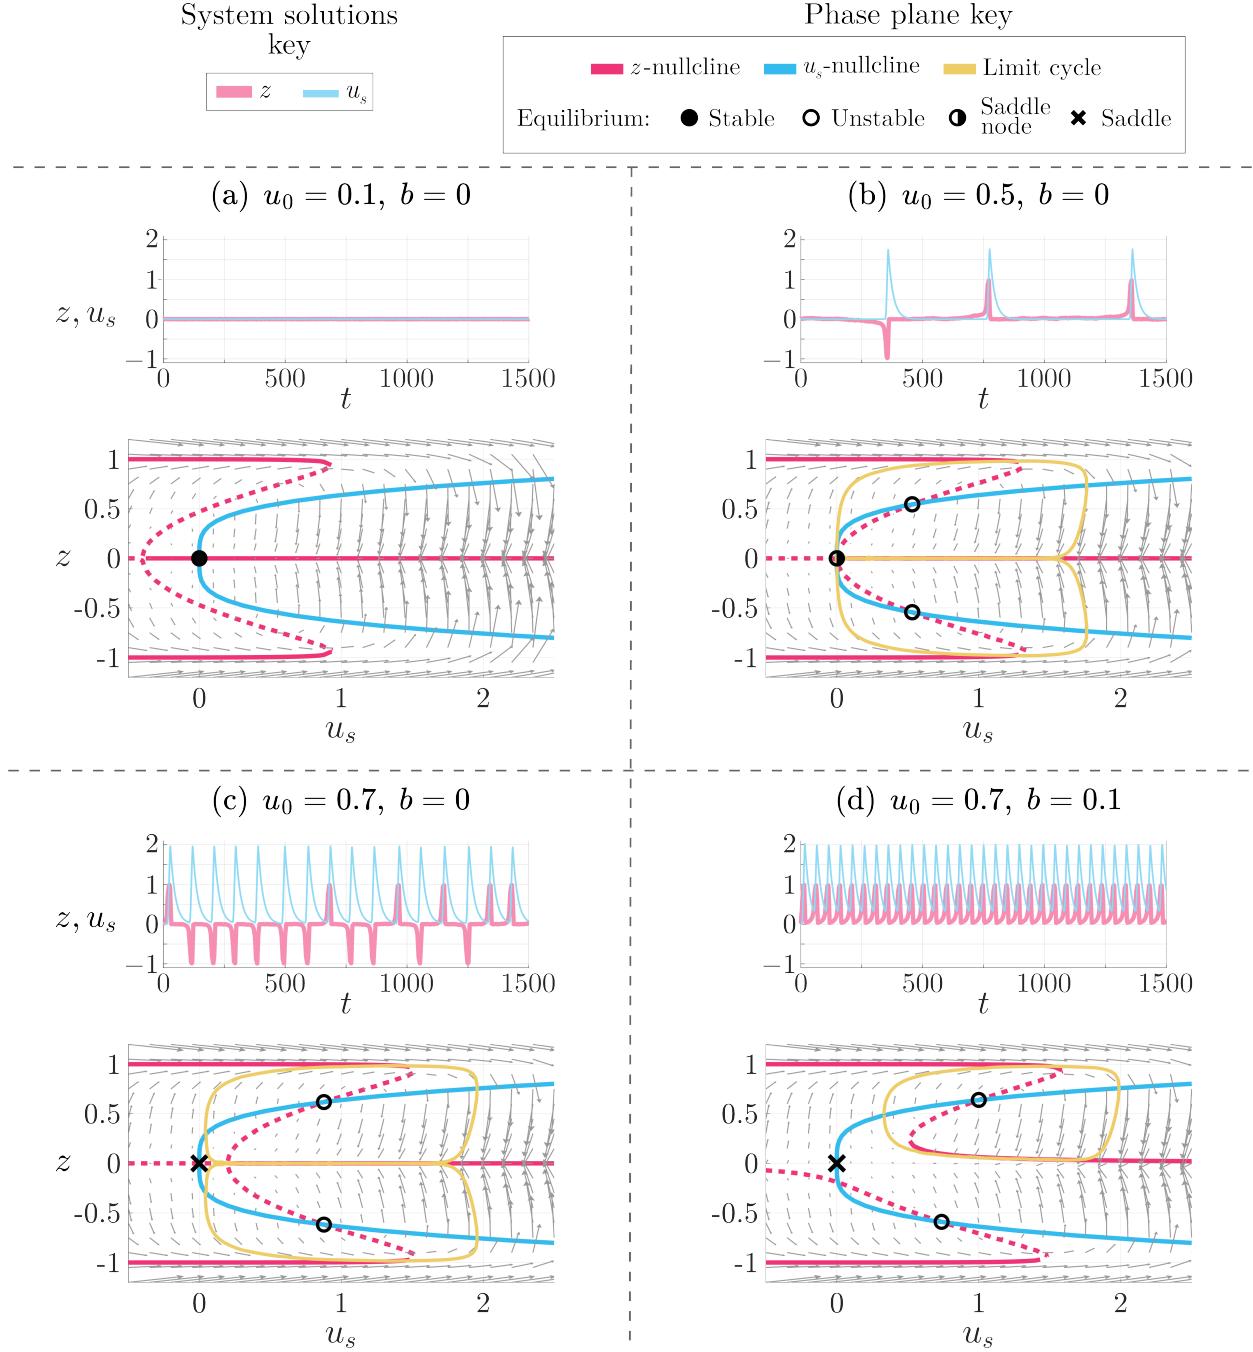


Figure 4.4: The system solutions and  $(u_s, z)$  phase portrait as the basal sensitivity  $u_0$  increases. For all:  $d = 1, a = 2$ , (thus  $u_0^* = 0.5$ ),  $K_u = 2, K_{u_s} = 6, \tau_{u_s}/\tau_z = 10$ . (Top of each panel): Solutions of  $u_s$  and  $z$  over time, with initial condition  $(u_s, z)|_{t=0} = (0.01, 0.01)$  and additive Gaussian distributed white noise. (Bottom of each panel): The  $u_s$ -nullcline ( $z$ -nullcline) is shown in blue (pink). Solid (dotted) lines indicate stable (unstable) branches of the  $z$ -nullcline with respect to (4.4a). Gray arrows denote the vector field. Black circles show equilibria: filled are stable, unfilled are unstable, partially filled are saddle-node bifurcations. Crosses show saddle equilibria. Saddle-node-homoclinic cycles in (b), and limit cycles in (c) and (d) are in yellow.

spiking decisions in the direction of the option favored by the inputs (e.g., upward decision spikes in the case of Fig. 4.4d where  $b > 0$  provides evidence in favor of option 1).

Observe that in the presence of informative inputs (Fig. 4.4d), the decision spiking frequency is higher than in the case of endogenous decision spiking oscillations (Fig. 4.4c). This feature is similar to spike frequency indicating input intensity in neural systems. In applications like robot navigation,  $u_0$  can be controlled to avoid endogenous spiking.

## 4.4 Agile multi-agent decision-making: S-NOD

### 4.4.1 S-NOD for multiple agents and two options

We can generalize the single-agent S-NOD equations (4.4) to the case of  $N_a$  agents in the same way that NOD generalizes to  $N_a$  agents [1]. The multi-agent NOD models the decision-making process of multiple agents sharing and influencing one another's opinions over a communication network. Examples include agents choosing how they distribute their time over two resource patches or deciding whether to move to the left or right when navigating a cluttered space, all while integrating information from other agents' opinions. In the multi-agent S-NOD, each agent  $i$  has two state variables  $z_i$  and  $u_{s,i}$  with dynamics given by:

$$\tau_z \dot{z}_i = -d z_i + \tanh \left( (u_0 - u_{s,i} + K_u z_i^2) \left( \sum_{k=1}^{N_a} a_{ik} z_k \right) + b_i \right), \quad (4.5a)$$

$$\tau_{u_s} \dot{u}_{s,i} = -u_{s,i} + K_{u_s} z_i^4, \quad (4.5b)$$

where  $A = [a_{ik}] \in \mathbb{R}^{N_a \times N_a}$  is the S-NOD network adjacency matrix, capturing both the strength ( $a_{ii} \geq 0$ ) of a self-reinforcing term and the strength ( $|a_{ik}|$ ) of the influence of the opinion of agent  $k$  on the opinion of agent  $i$ . If  $a_{ik}$  is positive (negative), then an opinion of agent  $k$  in favor of one of the options influences an opinion of agent  $i$  in favor of the same (other) option. We assume homogeneous agents, i.e., all agents have the same  $d$ ,  $u_0$ ,  $K_u$  and  $K_{u_s}$ . S-NOD as presented in (4.5) is the networked, distributed version of (4.4).

In Fig. 4.5 we simulate the opinion dynamics of (4.5) for three agents in a network when agent 1 receives a constant input  $b_1 > 0$  for  $t \in [10, 40]$ . For the loop-free networks of Fig. 4.5, we see that when the weights  $a_{ik}$  are positive (negative) the spiking of agent 1 for option 1 triggers synchronized (anti-synchronized) spiking of agents 2 and 3. With  $a_{ik} > 0$  (Fig. 4.5a), all agents form spikes of positive opinion due to the positive input to agent 1. With  $a_{ik} < 0$  (Fig. 4.5b), agent 1's behavior is unchanged but agents 2 and 3 instead form spikes of negative opinion. Future work will characterize different possible behaviors, e.g., opinion spike (anti)synchronization, for classes of network structures.

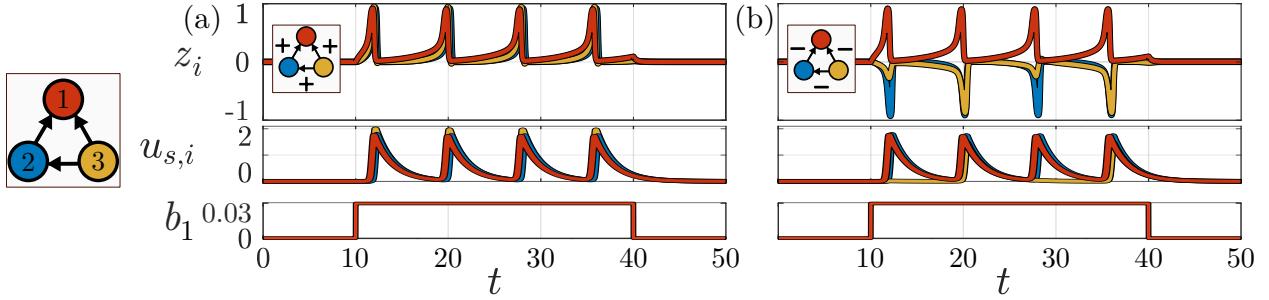


Figure 4.5: S-NOD solutions over time of three agents with the directed communication network shown and an input  $b_1$  applied only to agent 1. For  $i \neq k$ , (a)  $a_{ik} = +0.1$  and (b)  $a_{ik} = -0.1$ . For all:  $a_{ii} = 1$ ,  $d = 1$ ,  $K_u = 2.3$ ,  $K_{u_s} = 16$ ,  $u_0 = 0.9$ ,  $\tau_{u_s}/\tau_z = 20$ .

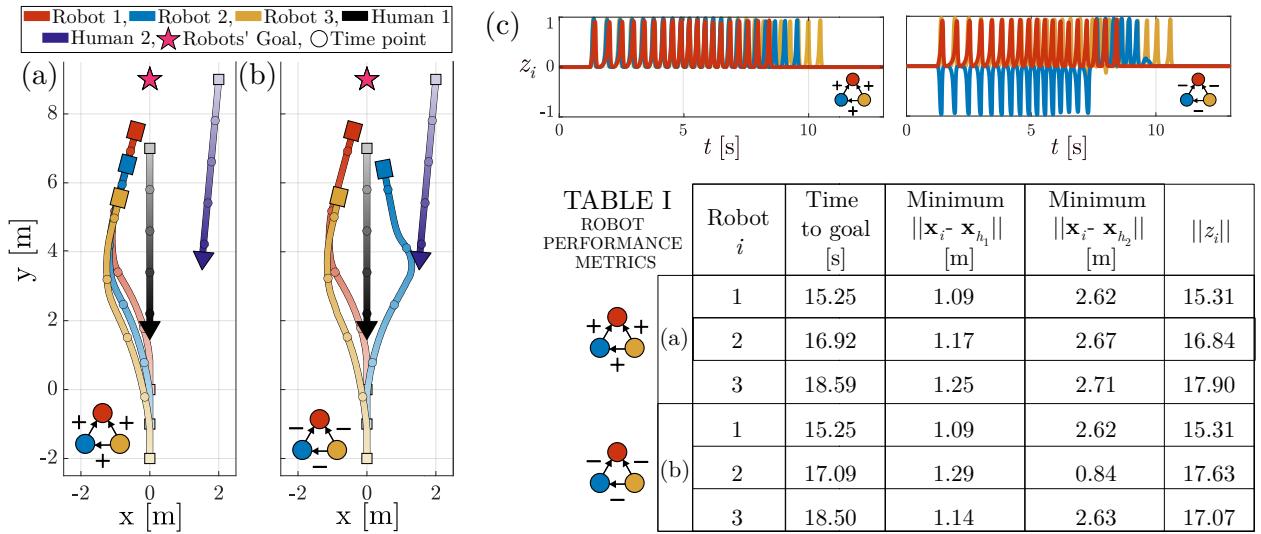


Figure 4.6: Trajectories of social robots using S-NOD to navigate around approaching human movers. Communication network and parameters as in Fig. 4.5 with (a)  $a_{ik} = +0.1$  and (b)  $a_{ik} = -0.1$ . (c) Plots of  $z$  over time  $t$  for the robots. Table I lists performance metrics of the robots.

#### 4.4.2 Application to social robot navigation

We use S-NOD (4.5) to design a decentralized, agile controller for social robots navigating around oncoming human movers in 2D. Each robot has a nominal control that steers it toward its goal by regulating its heading direction through proportional negative feedback. The S-NOD state  $z_i$  defines the strength of robot  $i$ 's preference for turning left ( $z_i > 0$ ) or right ( $z_i < 0$ ). A term mediated by  $z_i$  is added as positive feedback to the nominal steering controller. This overcomes negative feedback regulation and promotes fast reactive steering when possible collisions with oncoming human movers are imminent. Simulations of the resulting navigation behavior are in Fig. 4.1 and Fig. 4.6 and animated at <https://spikingNOD.github.io>.

To anticipate collisions, robot  $i$  can estimate its distance ( $\rho_i$ ) to a human, bearing angle ( $\eta_i$ ) on the human, and angle ( $\eta_{h,i}$ ) between the human's heading direction and the robot-human vector. Robots can

exchange steering opinions  $z_k$  over a communication network as in (4.5a). We let the robot's attention grow above its basal level  $u_0$  as collision risk grows with decreasing  $\rho_i$  and  $\eta_i$ . This increases the strength of (i) the positive feedback loop of the steering controller, and (ii) the interactions with other robots to achieve coordinated obstacle avoidance. Thus, each robot's steering opinion deviates from navigating toward its goal dependent on  $\rho_i$ ,  $\eta_i$ ,  $\eta_{h_i}$ , and on other robot opinions. We let  $\hat{z}_h(t) = \tan(\eta_h(t))$  be a proxy for the human's opinion at time  $t$  and add  $\hat{z}_h$  to the term  $\sum_{k=1}^{N_a} a_{ik} z_k$  in (4.5a). Coordination among robots derives from the sign of  $a_{ik}$ : when  $a_{ik} > 0$  ( $a_{ik} < 0$ ), robot  $i$  is influenced to make a similar (opposite) steering choice as robot  $k$ . We let  $|a_{ik}|$  decay with growing distance between robots  $i$ ,  $k$ .

Fig. 4.1 compares the trajectory and opinion of a robot using NOD (adapted from [25]) and S-NOD (4.4) to navigate around a human mover. The S-NOD robot passes the human with a minimum distance of 0.96m and arrives at its goal in 14.4 seconds. Without the return to the sensitive bifurcation point that S-NOD provides, the NOD robot's opinion change lags as the human makes a sharp turn and the robot experiences a collision (moving closer than 0.3m to the human) less than 6 seconds into the simulation.

Fig. 4.6 showcases three robots navigating towards a common goal and around two approaching humans using multi-agent S-NOD (4.5). The communication networks are those in Fig. 4.5 and we observe similar spiking opinion formation behavior. Robots in Fig. 4.6a exhibit the same synchronized opinion formation behavior seen in Fig. 4.5a. In Fig. 4.6b, due to the distance dependent  $|a_{ik}|$ , robot 3 disagrees more with robot 2 than robot 1, thus preferring the same opinion as robot 1, unlike in Fig. 4.5b, with constant  $|a_{ik}|$ , where agent 3 joins agent 2. Notably, in Fig. 4.6b, robot 2 initially turns right (opposite to robots 1 and 3) to avoid human 1, but later switches to a left turn to avoid human 2. S-NOD gracefully navigates robot 2 out of consecutive potential collisions. This embodies the agility and sequential decision-making features of S-NOD.

## 4.5 Final remarks

We presented and analyzed Spiking Nonlinear Opinion Dynamics (S-NOD) for a single agent and two options. We showcased the ability of S-NOD to swiftly form opinions and regulate back to ultrasensitivity. S-NOD provides first-of-its-kind *two-dimensional* excitable (spiking) dynamics for agile decision-making over two options. We showed how NOD can become too robust, but the self-regulation of S-NOD recovers flexibility. We analyzed existence of limit cycles for certain parameter regimes in S-NOD. We presented S-NOD for multiple agents that communicate opinions over a network and highlighted potential for agent (anti)synchronization. We illustrated S-NOD's agility in a social robot navigation application and plan to implement on physical robots. We aim to provide analytical guarantees on the onset of periodic spiking in

limit cycles, to analyze synchronization patterns for multiple agents, and to generalize to multiple options.

## 4.6 Appendix

We present an extended form of Fig. 4.3 to visualize the additional  $u_s$  dimension on the bifurcation diagram of  $z$  and  $u_0$ . The top and middle panels are unchanged from the published figure, aside from colors within the bifurcation diagrams. A final panel shows the loop in the  $u_s$  dimension created with the S-NOD model that creates the spiking behavior seen in (A) and represented with the pink bi-directional arrow in (B).

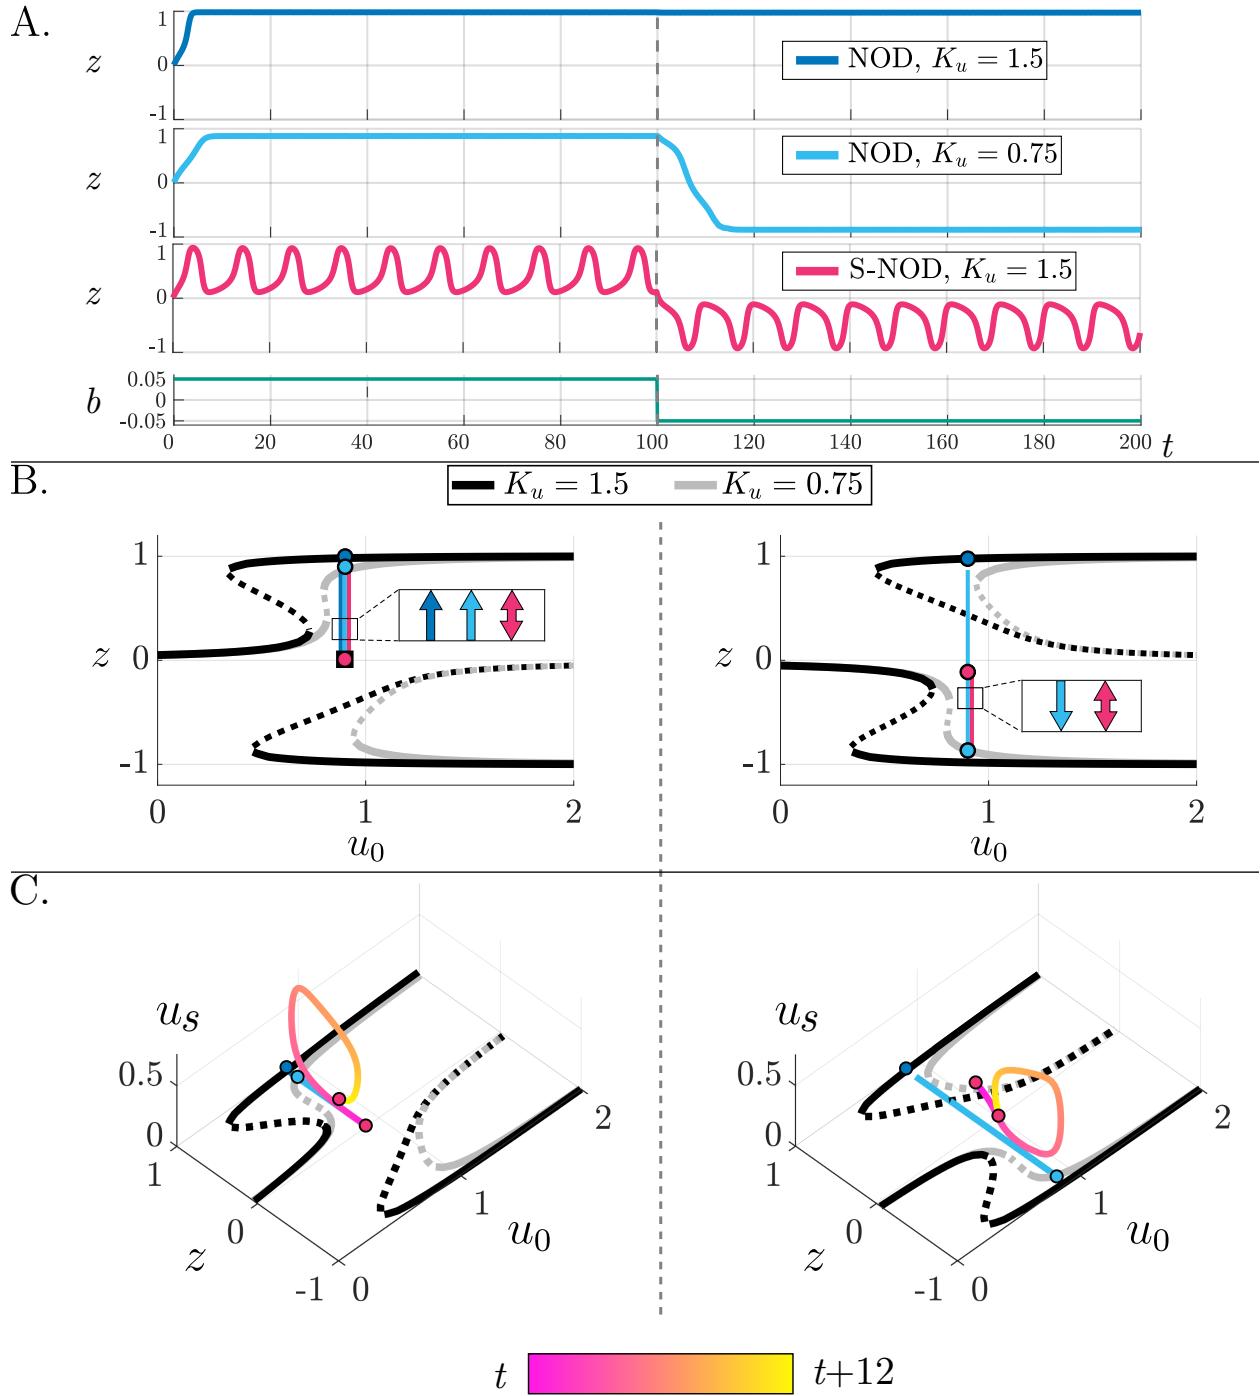


Figure 4.7: An extended presentation of Fig. 4.3. Middle section (B) has modified colors used to distinguish the bifurcation diagrams with different  $K_u$  values. Bottom section (C) illustrates the excitation loop in the  $u_s$  direction previously unseen in the 2D bifurcation diagram of (B). The color bar shows the 12-timestep shape of the loop that is followed by the trajectory before (in the left plot) and after (in the right plot) the bias input  $b$  changes sign. Dots and streaks of the NOD trajectories are also shown, but move so quickly that a color bar was unable to capture temporal information meaningfully and was therefore omitted.

## Chapter 5

# Spiking Nonlinear Opinion Dynamics (S-NOD) for social robot navigation

This chapter serves to validate and expand on the S-NOD model when applied to social robot navigation. We begin this chapter with the explicit definition of the specialized S-NOD equations for social robot navigation. We then replicate the experimental validation of S-NOD on a navigation social robot, following the one-human and two-human experimental protocols outlined in Chapter 3. Finally, we explore the parameter sensitivity of navigation behavior by a robot navigating with S-NOD.

### 5.1 S-NOD specialized for social robot navigation

We begin by introducing the equations of S-NOD for social robot navigation. These are presented in line with those behind the multi-agent, multi-human simulations of Fig. 4.6 in the previous chapter. The parameter definitions follow in line from those reported in equation (3.2) for NOD specialized for social robot navigation defined in Chapter 3.2. For brevity, only new variables and/or parameters are described here. The equations for S-NOD on a navigating social robot are as follows:

$$\tau_{z_i} \dot{z}_i = -d_i z_i + \tanh \left( \max(u_i - u_{s_i}, 0) \cdot \left( \sum_{k=1}^{N_a} a_{ik} z_k + \sum_{\ell=1}^{N_h} a_{i\ell} \hat{z}_{h_\ell} \right) + b_i \right), \quad (5.1a)$$

$$u_i = \max(u_{0_{i\ell}}) + K_{u_i} z_i^2, \quad (5.1b)$$

$$\tau_{u_{s_i}} \dot{u}_{s_i} = K_{u_{s_i}} z_i^4 - u_{s_i}, \quad (5.1c)$$

$$\dot{\theta}_i = k_{\theta_i} \sin(\widetilde{u_{0_{i\ell}}} \cdot \beta_i \cdot \tanh(z_i) + (1 - \widetilde{u_{0_{i\ell}}}) \phi_i). \quad (5.1d)$$

All parameter and variable definitions for opinion  $z_i$  update (5.1a), sensitivity  $u_i$  update (5.1b), and heading  $\theta_i$  update 5.1d can be found in Chapter 3.2. All parameter and variable definitions for the slow recovery state  $u_{s_i}$  update (5.1c) can be found in Chapter 4.4.

A new parameter not previously discussed exists only within the opinion  $z_i$  update in equation (5.1a). This is  $a_{ik} \in \mathbb{R}$ , which defines the weight of robot  $k$ 's opinion on robot  $i$ 's opinion update.  $A = [a_{ik}] \in \mathbb{R}^{N_a \times N_a}$  is the S-NOD network adjacency matrix, which captures both the strength ( $a_{ii} \geq 0$ ) of a self-reinforcing term and the strength ( $|a_{ik}|$ ) of the influence of the opinion of robot  $k$  on the opinion of robot  $i$ . Note that  $a_{ii} = a_i$  in the NOD model, combining the  $a_i z_i$  term into the sum of all robots' influence on robot  $i$ 's opinion. The term for all other robot influences where  $i \neq k$  is defined as  $a_{ik} = \Gamma \cdot \left(1 - \left(\frac{D_{ik}}{P - D_{ik}}\right)^m\right)^{-1}$ , with gain  $\Gamma \in \mathbb{R}$ , the distance between robots  $i$  and  $k$  described with  $D_{ik} = \|\mathbf{x}_i - \mathbf{x}_k\|$  and letting  $P > D_{ik}$  be the neighbor proximity threshold for communication. This function is a saturated curve between a value of 0 when  $D_{ik} > P/2$  and 1 when  $D_{ik} < P/2$  with a slope of  $m > 0$ , where  $P$  defines the distance at which  $a_{ik} = 0.5$ . Further, we let each robot  $i$  observe its robot teammate  $k$  with the same  $\eta_{ik}$  formulation as that for the relative angle  $\eta_{i\ell}$  for human movers. Therefore, we let  $a_{ik} = 0$  when  $\cos(\eta_{ik}) < 0$  to stop communication between robot  $i$  and robot  $k$  if robot  $k$  is physically behind and therefore unobservable to robot  $i$ . The sign of  $\Gamma$  determines the kind of inter-robot influence: if  $\Gamma$  is positive (negative), then  $a_{ik} \geq 0$  ( $a_{ik} \leq 0$ ) and an opinion of robot  $k$  in favor of one passing side influences an opinion of robot  $i$  in favor of the same (other) passing side.

For the duration of this chapter, only a single robot ( $i = N_r = 1$ ) is considered. Therefore, the subscript  $i$  can be dropped from variable and parameter notation. The equations for a single robot interacting with  $N_h$  humans then simplify to:

$$\tau_z \dot{z} = -d z + \tanh \left( \max(u - u_s, 0) \cdot \left( az + \sum_{\ell=1}^{N_h} a \hat{z}_{h_\ell} \right) + b \right), \quad (5.2a)$$

$$u = \max(u_{0_\ell}) + K_u z^2, \quad (5.2b)$$

$$\tau_{u_s} \dot{u}_s = K_{u_s} z^4 - u_s, \quad (5.2c)$$

$$\dot{\theta} = k_\theta \sin(\widetilde{u_{0_\ell}} \cdot \beta \cdot \tanh(z) + (1 - \widetilde{u_{0_\ell}}) \phi). \quad (5.2d)$$

We use these equations to validate the efficacy and agility of the S-NOD model in navigating social robots in both real experiments and simulations.

## 5.2 Experimental validations of S-NOD

As discussed in Chapter 4, S-NOD offers several key advantages over the original NOD framework, particularly in the context of social robot navigation. We utilize experiments to showcase these benefits. With its spike-based dynamics, S-NOD enables the robot to respond rapidly and effectively to dynamic human behavior. This makes S-NOD well-suited for human-robot interactions, as the decision-making spikes allow the robot to form an opinion on its preferred passing side, implement a turn, and then return to an unopinionated and ultrasensitive state. That is, the robot can more responsively adapt to changes in the human's motion or intent than seen with NOD. Additionally, we demonstrate that the system's behavior can be tuned by its parameters, recovering the navigation behavior of NOD.

In this section, we present the results of experiments in human-robot navigation with a social robot using the S-NOD algorithm, aiming to both validate these advantages and provide a direct comparison of robot behavior under NOD and S-NOD control.

The S-NOD algorithm was validated through both simulated and experimental trials, with S-NOD-embodied robots navigating past either a single oncoming moving human or around multiple oncoming moving humans. Simulations were rendered using MATLAB, and the behaviors observed in these idealized, noise-free simulations provide evidence for the general range of parameter values to be used in real experiments.

Experiments using the S-NOD algorithm on a social robot were conducted to directly parallel those used to validate the NOD algorithm on a social robot in Chapter 3.4. Specifically, we tested the S-NOD algorithm on a social robot under the same conditions as the single-robot, single-human and single-robot, two-human scenarios that had been used to validate the NOD algorithm on a social robot. This allowed for a direct

comparison between the behaviors of S-NOD and NOD in social robot navigation.

Our experimental objectives are fourfold, some of them the same as those for the experimental validations of NOD:

1. To demonstrate the flexibility of our approach by showing that the S-NOD robot can reliably navigate environments and interact with multiple human movers across a variety of scenarios.
2. To validate our algorithm's analysis, confirming that the robot can always resolve deadlock situations - gracefully maneuvering around an oncoming human, even if the human is unaware of or ignores the robot, and even when the robot's preferred passing side conflicts with that of the human.
3. To replicate and compare the navigation results of the S-NOD algorithm to the NOD algorithm.
4. To identify behavioral differences between S-NOD and NOD in a robot's physical motion.

### 5.2.1 General experimental setup

The setup for both the single-robot, single-human set of experiments and the single-robot, two-humans set of experiments directly mirrors that described in Chapter 3.4.1. We direct readers to this section for a detailed description of the methodology for both sets of experiments.

Two sets of experiments were conducted for the single-robot, single-human framework. The first set of experiments, designated as **S-NOD 1** trials, was performed using the same parameters as those in the NOD trials within the single-robot, single-human framework. This set of experiments serves to *compare* the behavior of NOD and S-NOD when embodied and tuned in the same manner. The second set of experiments, noted as **S-NOD 2** trials, were done with adjusted parameters to *replicate* the behavior of the NOD embodied robot. This enables us to evaluate the NOD and S-NOD algorithms through a direct comparison under identical conditions (S-NOD 1) and demonstrate S-NOD's versatility in replicating NOD-like behavior (S-NOD 2).

### 5.2.2 Results

#### 5.2.2.1 Metrics reported

The metrics used to analyze the trajectories taken by the robot in experimental trials include the same set of efficiency and spatial intrusiveness metrics described in Chapter 3.4.2.1.

Additionally, we consider the total opinion magnitude  $\|z\|$  exerted throughout a navigation trial. We interpret the total opinion control  $\|z\|$  as a measure of the NOD or S-NOD algorithm's exercised authority

in controlling the robot's turning when interacting with a human. At any time  $t$  along the trial,  $\|z(t)\|$  reflects the magnitude of the robot's opinion, therefore its cumulative value  $\|z\| = \int_0^T \|z(t)\| dt$  captures how actively the robot formed its preferences in response to a human over time. A high  $\|z\|$  value indicates that the robot was frequently—or persistently—prompted to turn away from its goal direction in response to an interaction with a human. A low  $\|z\|$  value indicates that the opinion dynamics minimally influenced the robot and thus the robot moved more directly in its goal direction. This provides us with a new efficiency metric to compare alongside the results of the similar NOD experiments being replicated. The complete set of metrics reported within this chapter is contained in the updated table below.

Efficiency Metrics		Spatial Intrusiveness Metrics	
Definition	Unit	Definition	Unit
<b>Navigation time:</b> The time it took for the robot to arrive at its goal location from its initial position.	s	<b>Minimum distance*</b> : The distance between the robot and a passing human at the closest point in their paths.	m
<b>Path length:</b> The total distance the robot traveled as it moved between its initial and goal locations.	m	<b>Invasion distance:</b> The distance the robot traveled while invading the personal space (surrounding 1.2 m) of a passing human.	m
<b>Centerline deviation:</b> The cumulative area under the robot's curved path from the straight path to its goal.	$m^2$		
<b>Opinion control <math>\ z\ </math>:</b> The total opinion magnitude exerted during the robot's full navigation.	1		

Table 5.1: Updated performance metrics and their definitions.

\* Recall from Section 3.2 that  $D_\ell$  represents the distance between a single robot and human  $\ell$ . Therefore, we can shorthand this metric as  $\min(D_\ell)$  for the minimum distance between a single robot and human  $\ell$ .

### 5.2.2.2 Single-robot, single-human experiments

The setup of these sets of experiments matches that detailed in Chapter 3.4.2.2. Pairs of starting and goal locations for the robot and human were fixed across all trials at (0, 0) m and (0, 6) m. The robot's goal location was set at (0, 6.4)m, effectively mirroring the initial location of the human. The human was instructed to move toward the robot's initial location. Thus, the initial and goal locations of the robot and human were swapped. This causes the robot and human to begin their movement by facing each other head-on.

We combined three robot bias cases with three movement prompt cases for the human participant, yielding nine unique trial configurations. We conducted three experiments within each trial configuration, resulting in a total of twenty-seven trials. As before, in each trial, we tested two values of  $\beta$  in equation

(5.2d):  $\beta = \pi/4$  and  $\beta = \pi/6$ . This resulted in a total of fifty-four trials per experiment set.

### Robot instructions

Consistent with these sets of experiments and the NOD single-robot, single-human experiments, the robot was programmed to move at a constant speed,  $V = 0.7$  m/s, toward its goal location. The robot modifies its trajectory when encountering movers according to the navigation model (5.2). Two experimental sets were completed, each with different parameters, which are detailed in their respective sections. We designed three cases corresponding to three different values of the robot's bias  $b$ :

1. unbiased ( $b = 0$ ),
2. biased to its left ( $b = 0.3$ ),
3. biased to its right ( $b = -0.3$ ).

### Human participant instructions

The participant was instructed to walk at their normal pace towards their goal location according to one of three prompts: 1) go straight (labeled as the human was unaware of the robot), 2) bear to the left, and 3) bear to the right.

### Experiment sets

We conducted two sets of experiments within this chapter with the explicit purpose of comparing them among themselves and to the NOD experiments of Chapter 3.4.2.2. One set of experiments, **S-NOD 1**, used the same base parameters as those in the NOD single-robot, single-human experiments. The other set of experiments, **S-NOD 2**, features adjusted parameters from those in the NOD single-robot, single-human experiments to produce similar qualitative navigation behavior.

#### **S-NOD 1: Parameters comparable to NOD experiments of Section 3.4.2.2**

The parameters used in the S-NOD equation (5.2) are as follows:  $\tau_z = 0.008$ ,  $d = 1$ ,  $a_{11} = a = 2$ ,  $\gamma = 2$ ,  $K_u = 0.3$ ,  $U = 1$ ,  $n = 2$ ,  $D_{crit} = 5.5m$ ,  $\tau_{u_s} = 0.08$ ,  $K_{u_s} = 10$ , and  $k_\theta = 3$ . Note that these parameters are the same as those reported in the NOD single-robot, single-human set of experiments, except for the new, S-NOD-specific variables  $\tau_{u_s}$  and  $K_{u_s}$ .

The resultant trajectories in each of these trials are documented in Fig. 5.1.

### Efficiency results

We first consider the efficiency metrics and their trends across each trial.

The navigation time from the beginning of the experiment to the robot's arrival at its goal is shown in Fig. 5.2. Shorter navigation time corresponds to a more efficient movement by the robot. We observe a

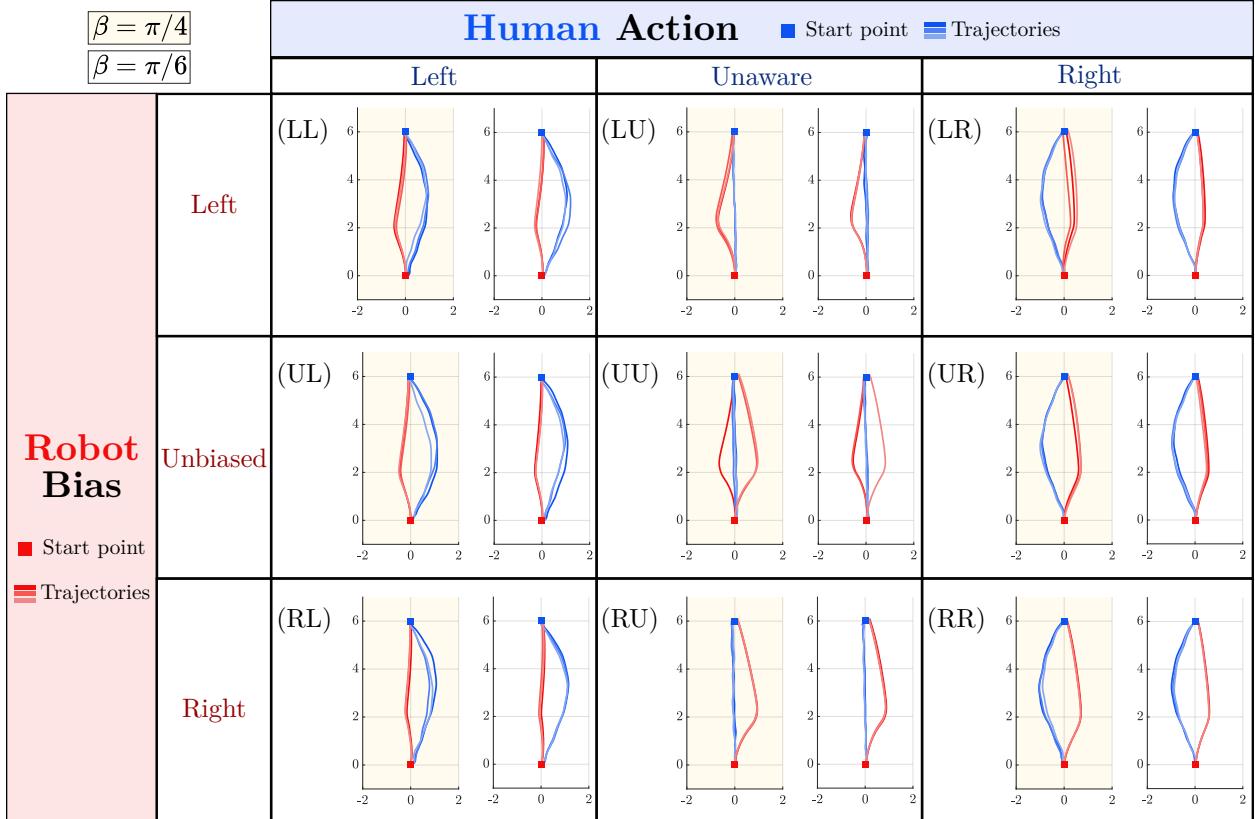


Figure 5.1: The trajectory data of a **S-NOD 1** robot for three runs each of the nine trial configurations, with both  $\beta_r = \pi/4$  (shaded yellow) and  $\beta_r = \pi/6$  (unshaded). This robot uses *parameters* comparable to the NOD robot experiments of Section 3.4.2.2. Axes correspond to the  $xy$ -plane in meters. The robot paths are shown in red, with a red box indicating the robot's starting position at approximately  $(0, 0)$ m. The human paths are shown in blue, with a blue box indicating the human's starting position at approximately  $(0, 6)$  m. In trial configuration labels, L=left, U = unaware/unbiased, and R = right.

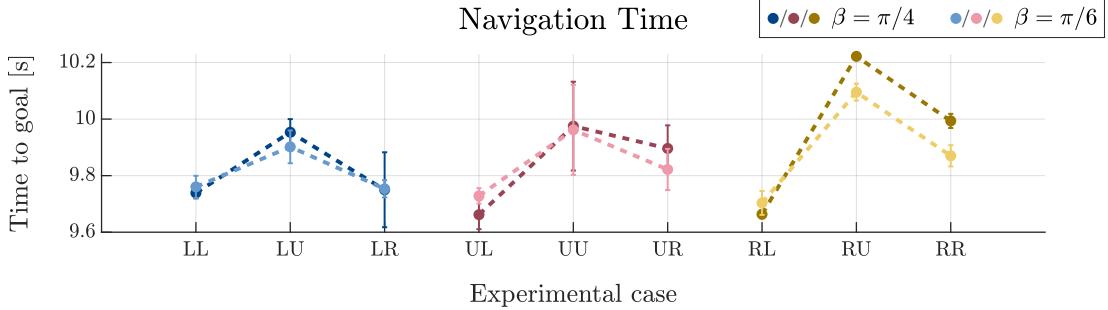


Figure 5.2: The average navigation times of a **S-NOD 1** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

similar trend to that seen in the NOD experiments, where an unaware human produces the longest navigation times in each subset of robot biases. We also observe that the trend of increased navigation time for the larger  $\beta$  value remains for all cases except for that of UL. Thus, we observe the same result as in NOD: tuning the  $\beta$  parameter to a smaller value can improve efficiency. The notable difference between these results and those from NOD is the reduction in navigation time for all cases. The S-NOD  $\beta = \pi/4$  experiments exhibit an average 2.9% decrease of navigation time from their NOD counterparts. The S-NOD  $\beta = \pi/4$  experiments exhibit an average 2.8% decrease in navigation time from their NOD counterparts. The increase of efficiency is less stark in the  $\beta = \pi/6$  experiments, with the S-NOD experiments exhibiting an average 1.8% decrease of navigation time from their NOD counterparts. Thus, we can declare that the S-NOD algorithm produces more efficient navigation times than a comparable NOD algorithm in this setting.

A comparable trend is observed in the robot's path length, as shown in Fig. 5.3. A shorter path length corresponds to a more direct and efficient movement toward the goal. Again, we observe the same result as in NOD: a higher  $\beta$  value in S-NOD consistently produced longer paths towards the goal. Thus, the robot moved less efficiently when permitted to make a larger turn away from an approaching human. This trend is maintained in all cases except for that of RL; however, we note the similar and symmetric case of LR in the NOD experiments, which displayed the same behavior. While this lack of adherence to the general trend is more minimal than that observed in NOD, we note that it exists for the same reason as discovered in NOD: in the RL case and for each  $\beta$  value, there are two out of three trials with markedly similar path lengths and one outlier value: for  $\beta = \pi/4$ , the outlier is lower than the others, and for  $\beta = \pi/6$ , the outlier is higher than the others. Removal of these outliers from the dataset would recover the trend of the increase of path length for the larger  $\beta$  value. Furthermore, we observe that the path lengths of trials conducted with  $\beta = \pi/4$  exhibit a 2.87% decrease compared to their NOD counterparts. The path lengths of trials

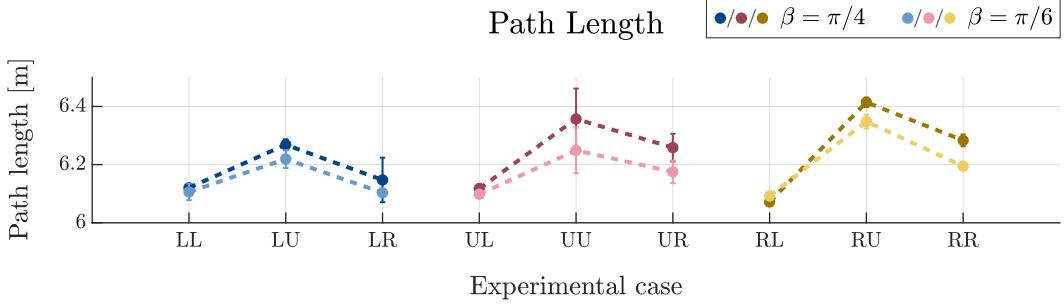


Figure 5.3: The average path lengths of a **S-NOD 1** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

conducted with  $\beta = \pi/6$  show a 2.56% decrease from their NOD counterparts. Thus, we can declare that the S-NOD algorithm produces shorter and more efficient path lengths than a comparable NOD algorithm in this setting.

Finally, we examine the centerline deviation of the paths taken by the robot in each trial. This metric quantifies how far the robot deviated from the straight go-to-goal path while navigating around the oncoming human and is shown in Fig. 5.4. We note that these cases performed with consistently lower centerline deviation values than those observed in the NOD experiments. A lower deviation value indicates that the robot remained more aligned with a direct trajectory, reflecting higher efficiency. As previously observed with NOD and consistent with the preceding results in this section, we again find that a larger  $\beta$  value results in greater deviation from the centerline, which is consistent with the findings on path length. The centerline deviations of S-NOD trials conducted with  $\beta = \pi/4$  show a 37.02% decrease from their NOD counterparts. The centerline deviations of S-NOD trials conducted with  $\beta = \pi/6$  show a 41.93% decrease from their NOD counterparts. Thus, we can declare that the S-NOD algorithm produces more goal-oriented, straight, and efficient paths than a comparable NOD algorithm in this setting.

In all efficiency metrics, we observe the same general trend across experimental cases as was observed in the comparable NOD experiments. That is, robot performance is least efficient when paired with an unaware human (cases LU, UU, RU). The efficiency of both S-NOD and NOD can be improved by interacting with an aware, cooperative human.

### Spatial intrusiveness results

We now consider the spatial intrusiveness metrics and their trends across each trial.

The minimum distance reached between the robot and human in each case is shown in Fig. 5.5. The lower the minimum distance, the closer the robot came to colliding with the human, which indicates more

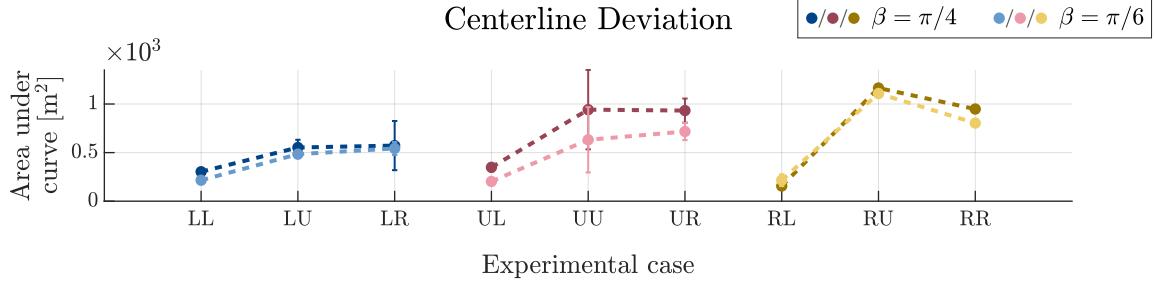


Figure 5.4: The average centerline deviations of a **S-NOD 1** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

spatially intrusive passing behavior. The unbiased robot consistently reaches a markedly lower minimum distance than its biased counterparts in all cases, as was previously observed in NOD experiments. We note that the S-NOD algorithm yields very similar minimum distances between  $\beta$  values, suggesting that this  $\beta$  tuning has a less significant impact on the tuning of spatial intrusiveness in the algorithm. As in NOD, the lowest minimum distance was observed for both  $\beta$  values in the unbiased robot, unaware human (UU) case. Furthermore, we observe that the minimum distances of trials conducted with  $\beta = \pi/4$  exhibit a 30.30% decrease compared to their NOD counterparts. The RR case with  $\beta = \pi/6$  is the most similar in quantitative result to the NOD experiments, with only an 8.36% decrease in its average value. Aside from that case, the minimum distances of trials conducted with  $\beta = \pi/6$  show an average 30.17% decrease from their NOD counterparts. Thus, we can declare that the S-NOD algorithm produces lower minimum distances between a robot and a passing human, indicating more spatially intrusive passing behavior and possibly less comfortable passing behavior than a comparable NOD algorithm. As evidenced by the earlier results, this is consistent with the S-NOD algorithm being notably more efficient. Consequently, a robot controlled by it is less likely to deviate from its path to the goal strongly.

Finally, we examine the invasion distance metric to analyze the S-NOD robot's ability to naturally maintain socially acceptable passing conventions. Ideally, the robot should not invade the personal space (1.2 m) surrounding the passing human, and personal space would be maintained if the invasion distance were 0. Differing from the NOD experiments, wherein all cases where the human's movement was directed (LL, UL, RL, LR, UR, RR) resulted in the robot maintaining a zero invasion distance, these S-NOD experiments show that only the UL, UR, and RR cases achieved this ideal result. Otherwise, the more efficient S-NOD robot invaded the personal space of the human for a non-negligible distance. We also note that the previously observed differences between the invasion distance dependent on  $\beta$  are also roughly absent. The

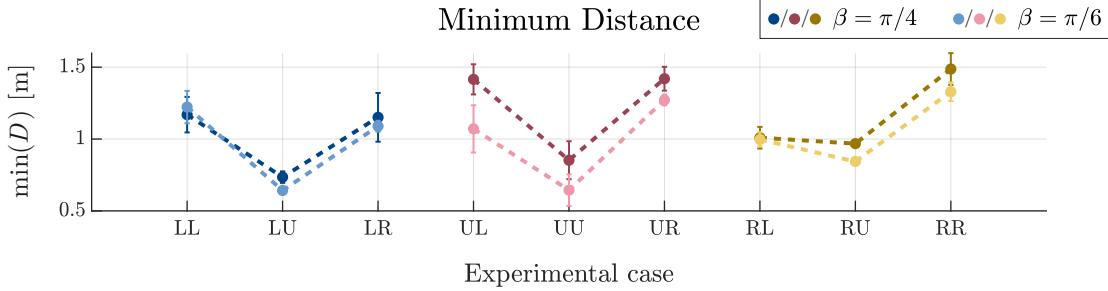


Figure 5.5: The average minimum distances between a **S-NOD 1** robot and a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

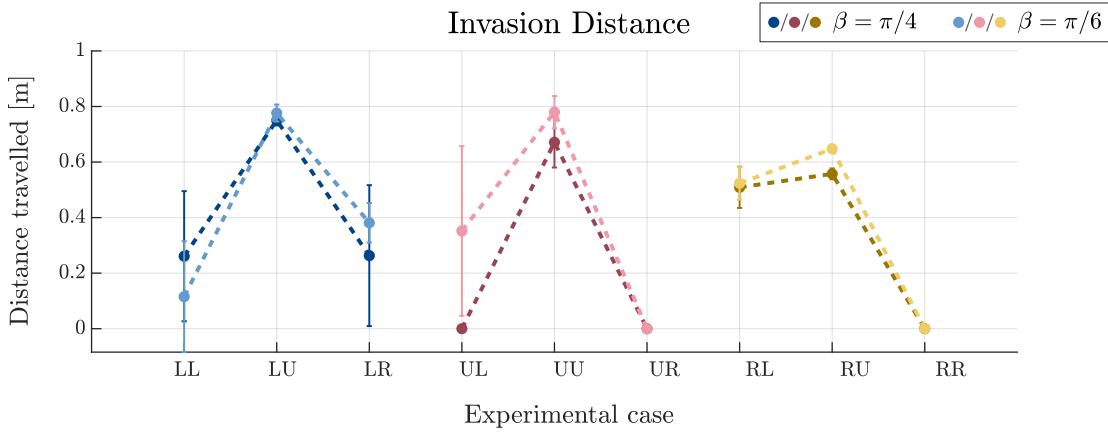


Figure 5.6: The average invasion distances by a **S-NOD 1** robot to a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.1. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

invasion distance was noticeably lower when the  $\beta$  value was higher, such that even if the robot could not maintain full personal space, it was able to move itself further away quickly. Thus, we can declare that the S-NOD algorithm produces higher invasion distances between a robot and a passing human's personal space, indicating more spatially intrusive passing behavior and possibly less comfortable passing behavior than a comparable NOD algorithm.

This consistent trend of increasing efficiency but sacrificing spatial unobtrusiveness with an S-NOD algorithm compared against a parameter-comparable NOD algorithm motivates the subsequent experiments to recover and replicate the desired *behavior* from the NOD algorithm by adjusting S-NOD's parameters. The parameters of NOD and S-NOD need not be the same, as the efficiency and agility of the S-NOD

model are its most important characteristics. These qualities were evident in these experiments and can be maintained with minor adjustments to produce less spatially intrusive, more comfortable passing with a human.

#### **S-NOD 2: *Behavior comparable to NOD experiments of Section 3.4.2.2***

The parameters used in the S-NOD equation (5.2) are as follows:  $\tau_z = 0.008$ ,  $d = 1$ ,  $a_{11} = a = 0.5$ ,  $\gamma = 4$ ,  $K_u = 2.667$ ,  $U = 4$ ,  $n = 7$ ,  $D_{crit} = 8m$ ,  $\tau_{u_s} = 0.08$ ,  $K_{u_s} = 12$ , and  $k_\theta = 2$ . Note that all parameters are changed from those reported in the NOD single-robot, single-human set of experiments, except for the values of  $\tau_z$  and  $d$ . Further, note that the S-NOD-specific variable  $\tau_{u_s}$  remains consistent with that of the S-NOD 1 set of experiments, but  $K_{u_s}$  has been increased. These parameters are set in such a way as to produce trajectories by the S-NOD robot that more closely match the trajectories by the NOD robot seen in Chapter 3.4.2.3—that is, the behavior of the S-NOD robot is comparable to the behavior of the NOD robot.

The design of each parameter adjustment can be explained intuitively in terms of the desired outcome of the behavior. For the opinion, the self-reinforcing weight  $a_{ii} = a$  was decreased, allowing it to form an opinion more quickly based on the proxy opinion of the human rather than on its own. Similarly, the weight on the human’s opinion  $\gamma$  was increased to reflect more emphasis on cooperating to form an opinion that mirrored the passing behavior of the human. The  $K_u$  value was increased to design the decision-making as more sustained, wider spikes instead of thin, bursting decision-making spikes. This, in turn, produces longer durations along the path where the robot is turning, prompting further deviations from its path to the goal. Parameters  $U$ ,  $n$ , and  $D_{crit}$  were all increased to shape the basal sensitivity  $u_0$  function to be reactive to humans that are further away, allowing the robot to more quickly respond to human behavior with collaborative decision-making. The increase of gain  $K_{u_s}$  creates more spiking behavior to maintain high spiking frequency in light of the adjustment of  $K_u$  that increased the width of the decision-making spikes. Finally,  $k_\theta$  was decreased to best suit the desired behavior of the NOD paths that these S-NOD results sought to recreate.

The resultant trajectories in each of these trials are documented in Fig. 5.7. As this set of experiments was collected expressly to replicate the behavior of NOD experiments in Chapter 3.4.2.2, the similarities and differences in their results are presented here. We direct readers to the previous chapter for a more thorough analysis of the behavior in each case, independent of comparison.

#### **Efficiency results**

We first consider the efficiency metrics and their trends across each trial.

The navigation time of the robot is illustrated in Fig. 5.8, the path length in Fig. 5.9, and the centerline deviation in Fig. 5.10. Their results, along with those of the NOD and S-NOD 1 experimental sets, are

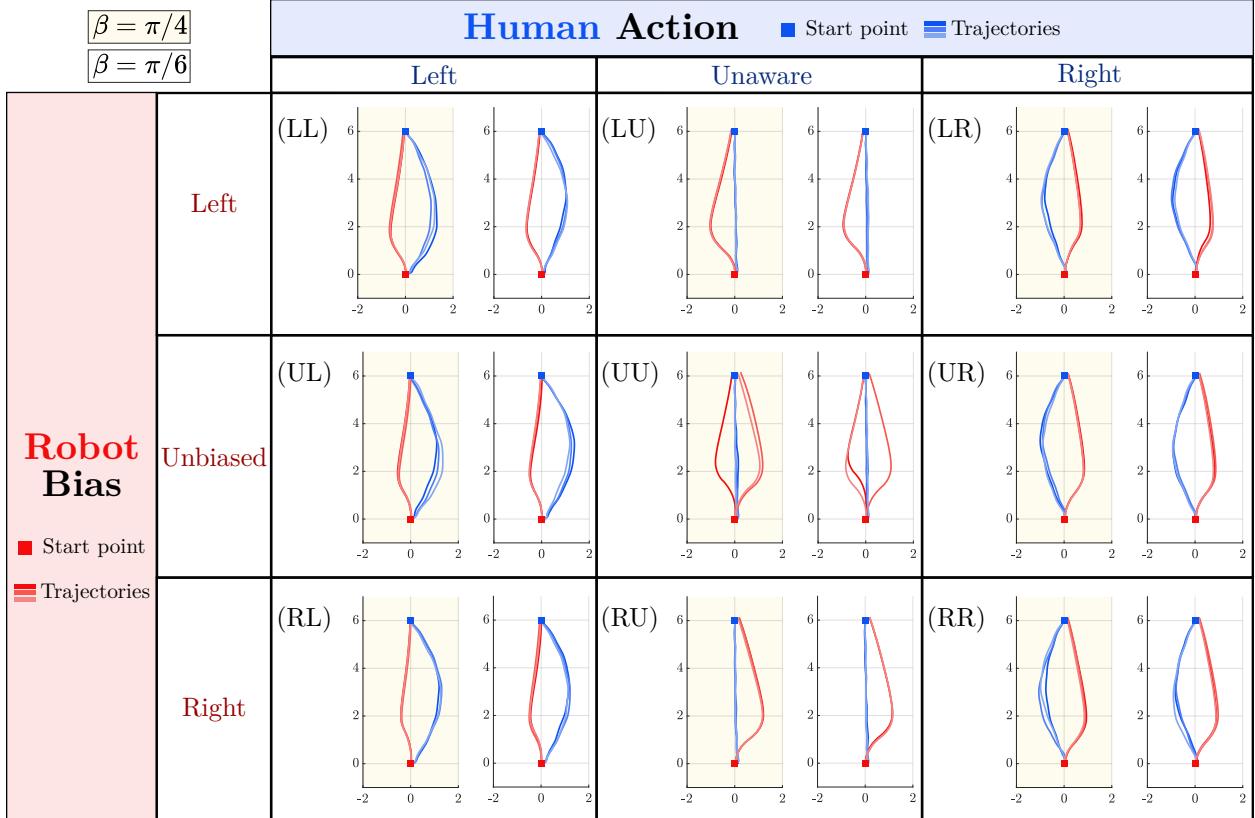


Figure 5.7: The trajectory data of a **S-NOD 2** robot for three runs each of the nine trial configurations, with both  $\beta_r = \pi/4$  (shaded yellow) and  $\beta_r = \pi/6$  (unshaded). This robot uses different parameters from the NOD experiments of Section 3.4.2.2 to create comparable *behavior* to that illustrated in Fig. 3.4. Axes correspond to the  $xy$ -plane in meters. The robot paths are shown in red, with a red box indicating the robot's starting position at approximately  $(0,0)$ m. The human paths are shown in blue, with a blue box indicating the human's starting position at approximately  $(0,6)$  m. In trial configuration labels, L=left, U = unaware/unbiased, and R = right.

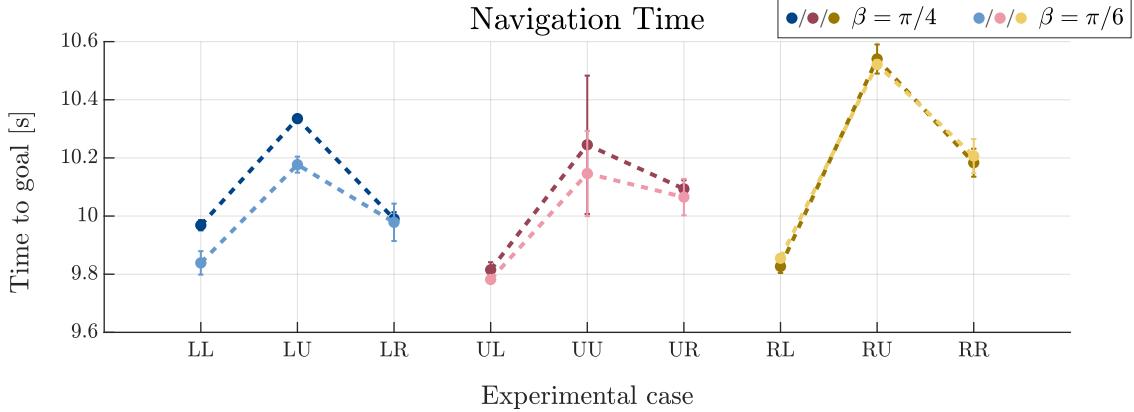


Figure 5.8: The average navigation times of a **S-NOD 2** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

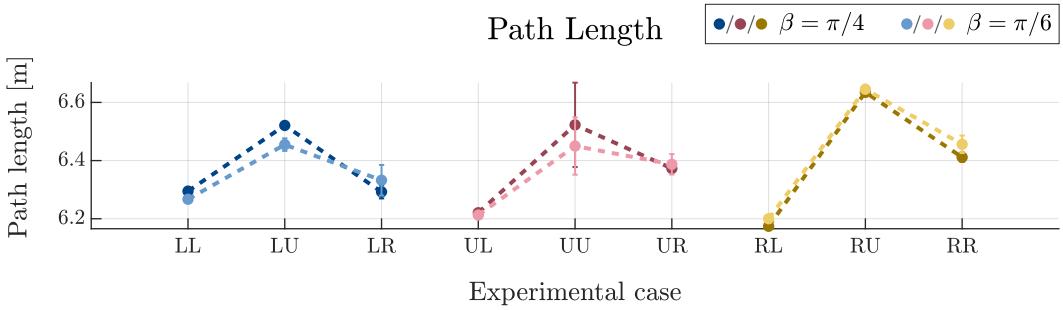


Figure 5.9: The average path lengths of a **S-NOD 2** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

described in Tables 5.2- 5.7. We highlight that the navigation times observed for both the NOD (Fig. 3.5) and S-NOD 2 (Fig. 5.8) experiment sets in the cases with a right-biased ( $b < 0$ ) robot (RL, RU, and RR) are nearly identical. When the S-NOD robot is biased left  $b > 0$ , these cases (LL, LU, LR) produce slightly higher navigation times than those observed in NOD. When the robot is unbiased ( $b = 0$ ), these cases (UL, UU, UR) experience: 1) navigation times lower than the NOD experiments for trials with  $\beta = \pi/4$ , 2) navigation times higher than the NOD experiments for trials with  $\beta = \pi/6$ . These results show us that navigation times of S-NOD experiments in each case between those with  $\beta = \pi/4$  and  $\beta = \pi/6$  are more similar than those between NOD experiments. These same trends exist in the path lengths observed for both the NOD (Fig. 3.6) and S-NOD 2 (Fig. 5.9) experiment sets. These results, along with smaller error bars in S-NOD 2's experiments, indicate that there is less variability of robot behavior between S-NOD trials than in the NOD trials.

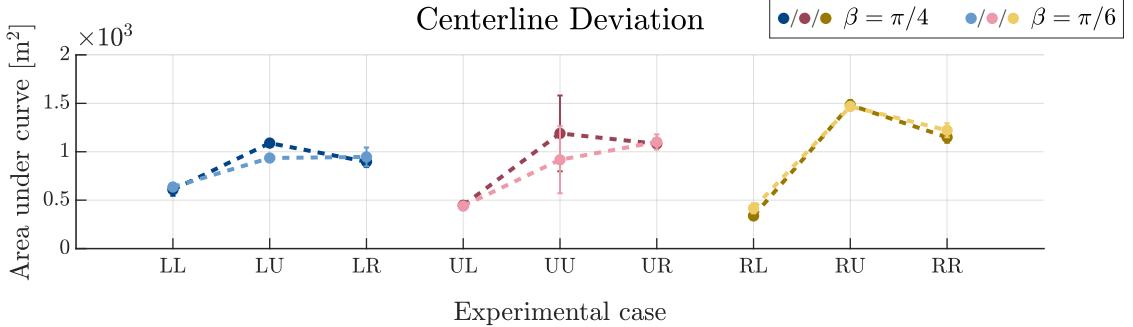


Figure 5.10: The average centerline deviations of a **S-NOD 2** robot in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

The only notable difference in centerline deviation observed between NOD (Fig. 3.7) and S-NOD 2 (Fig. 5.10) is that the cases LL and LR with  $\beta = \pi/4$  do not align. That is, the centerline deviation for case LL is higher for the NOD experiments than for the S-NOD experiments. In contrast, this metric for case LR is lower for the NOD experiments than for the S-NOD experiments. Besides this discrepancy, as well as smaller error bars observed for the NOD set of experiments, S-NOD 2 performs with similar centerline deviations to NOD for the bulk of the experiment set.

#### Spatial intrusiveness results

There are more differences observed between the NOD and S-NOD spatial intrusiveness metrics than those seen in the previously discussed efficiency metrics. These results, along with those of the NOD and S-NOD 1 experimental sets, are described in Tables 5.8- 5.11.

We first discuss the differences in minimum distances observed between NOD (Fig. 3.8) and S-NOD 2 (Fig. 5.11). The trend of the lowest minimum distances observed in unaware human cases (LU, UU, RU) compared to their aware counterparts remains, as does the smaller  $\beta = \pi/6$  producing lower minimum distances than the  $\beta = \pi/4$  case. However, S-NOD 2 produces on average a 17.79% reduction in minimum distance observed in cases LL, LR, UL, UR, and RR with  $\beta = \pi/4$  when compared to NOD. In the other cases of LU, UU, RL, and RU, this decrease was only on average 6.67%. These differences were less stark for the  $\beta = \pi/6$  set of experiments, with minimum distances by S-NOD 2 experiencing an average decrease of 2.58% from its NOD counterparts. Even still, the average minimum distance reached in the  $\beta = \pi/6$  case is only slightly within the social proximity of the passing human and not egregious.

We observe the differences in invasion distances observed between NOD (Fig. 3.9) and S-NOD 2 (Fig. 5.12). We see that the ideal result of zero invasion distance is recovered in all human-aware cases (LL, LR, UL,

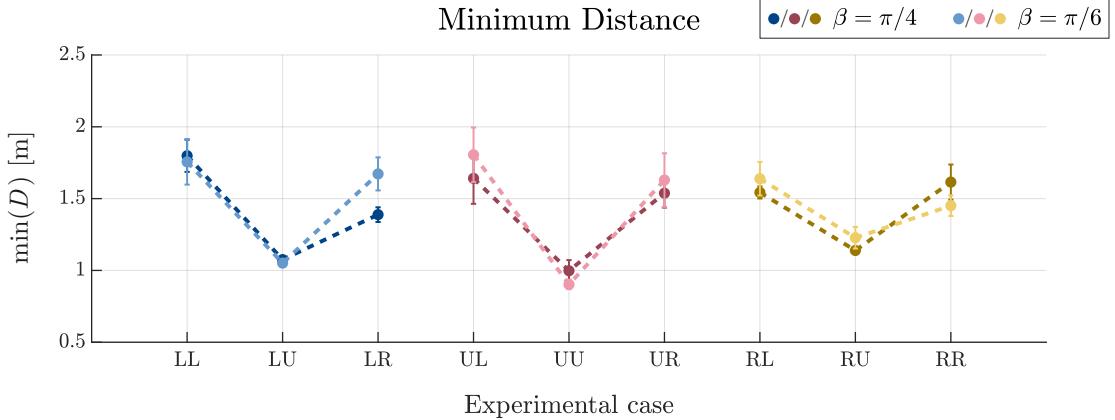


Figure 5.11: The average minimum distances between a **S-NOD 2** robot and a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

UR, RL, and RR). Therefore, we compare the results of the unaware human cases (LU, UU, and RU) in both NOD and S-NOD. We see that S-NOD 2 with  $\beta = \pi/4$  produces an average 37.98% increase in invasion distance from NOD across trials LU and UR. This is much smaller for the same trial with  $\beta = \pi/6$ , which saw an increase of only 4.72%. The case with a right-biased robot, RU, experienced exceptionally high increases in invasion distance between S-NOD 2 and NOD, with a 438.99% increase when  $\beta = \pi/4$  and an 87.12% increase when  $\beta = \pi/6$ . However, we recall that these right-biased NOD trials had peculiar paths with anomalous results due to potential shifts in goal location. We note that the S-NOD 2 (and S-NOD 1) did not produce the same abnormal behavior in the right-biased trials, which affects these comparative results. Considering these results, it appears that S-NOD, even when designed to replicate NOD's behavior, remains slightly more spatially intrusive in its passing etiquette than NOD.

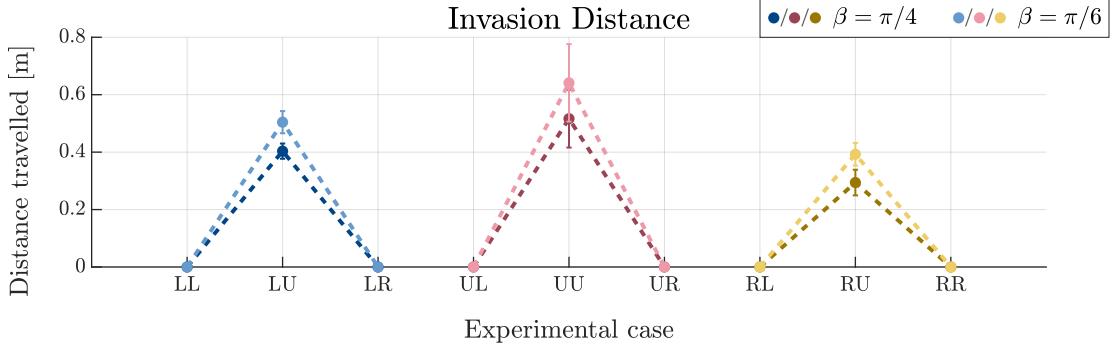


Figure 5.12: The average invasion distances between a **S-NOD 2** robot and a human in each of the nine configurations of trajectories illustrated and labeled in Fig. 5.7. Blue colors link results with the robot biased to the left ( $b = 0.3$ ), pink colors with the unbiased robot ( $b = 0$ ), and yellow colors with the robot biased to the left ( $b = -0.3$ ). Darker colors distinguish trials with  $\beta = \pi/4$ , lighter colors distinguish trials with  $\beta = \pi/6$ . Dotted lines link results associated with the same  $\beta$  value. Error bars show the standard deviation of each average.

#### Comparison tables of NOD, S-NOD 1, and S-NOD 2 metrics

The results of the average navigation times exhibited by the robot in each experimental case for NOD, S-NOD 1, and S-NOD 2 across both values of  $\beta = \pi/4$  and  $\beta = \pi/6$  are compared in Table 5.2 and Table 5.3, respectively. These tables summarize the results and observations discussed in the previous sections. The NOD results were illustrated in Fig. 3.5, the S-NOD 1 results were illustrated in Fig. 5.2, and the S-NOD 2 results were illustrated in Fig. 5.8.

Model	Experimental case with $\beta = \pi/4$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
NT [s]	NOD $\pm 0.05$	10.06 $\pm 0.05$	10.26 $\pm 0.03$	9.99 $\pm 0.08$	9.96 $\pm 0.11$	10.37 $\pm 0.11$	10.11 $\pm 0.09$	9.81 $\pm 0.01$	10.56 $\pm 0.07$	10.31 $\pm 0.03$
	S-NOD 1	9.74 $\pm 0.02$	9.95 $\pm 0.05$	9.75 $\pm 0.13$	9.66 $\pm 0.05$	9.98 $\pm 0.16$	9.90 $\pm 0.08$	9.66 $\pm 0.01$	10.22 $\pm 0.01$	9.99 $\pm 0.03$
	S-NOD 2	9.97 $\pm 0.02$	10.34 $\pm 0.00$	9.99 $\pm 0.02$	9.82 $\pm 0.03$	10.25 $\pm 0.24$	10.09 $\pm 0.03$	9.83 $\pm 0.02$	10.54 $\pm 0.05$	10.18 $\pm 0.05$

Table 5.2: Comparison of average navigation time (NT) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/4$ .

Model	Experimental case with $\beta = \pi/6$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
NT [s]	NOD	9.81 $\pm 0.05$	10.01 $\pm 0.03$	9.95 $\pm 0.09$	9.86 $\pm 0.04$	10.08 $\pm 0.11$	10.14 $\pm 0.03$	9.77 $\pm 0.03$	10.51 $\pm 0.08$	10.12 $\pm 0.04$
	S-NOD 1	9.76 $\pm 0.04$	9.90 $\pm 0.06$	9.75 $\pm 0.03$	9.73 $\pm 0.03$	9.96 $\pm 0.16$	9.82 $\pm 0.07$	9.70 $\pm 0.04$	10.10 $\pm 0.03$	9.87 $\pm 0.04$
	S-NOD 2	9.84 $\pm 0.04$	10.18 $\pm 0.03$	9.98 $\pm 0.06$	9.78 $\pm 0.00$	10.15 $\pm 0.15$	10.07 $\pm 0.06$	9.85 $\pm 0.01$	10.52 $\pm 0.02$	10.21 $\pm 0.06$

Table 5.3: Comparison of average navigation time (NT) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/6$ .

The results of the average path lengths of the robot exhibited in each experimental case for NOD, S-NOD 1, and S-NOD 2 across both values of  $\beta = \pi/4$  and  $\beta = \pi/6$  are compared in Table 5.4 and Table 5.5, respectively. These tables summarize the results and observations discussed in the previous sections. The NOD results were illustrated in Fig. 3.6, the S-NOD 1 results were illustrated in Fig. 5.3, and the S-NOD 2 results were illustrated in Fig. 5.9.

Model	Experimental case with $\beta = \pi/4$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
PL [m]	NOD	6.34 $\pm 0.03$	6.48 $\pm 0.02$	6.30 $\pm 0.06$	6.30 $\pm 0.07$	6.56 $\pm 0.07$	6.40 $\pm 0.07$	6.18 $\pm 0.01$	6.66 $\pm 0.04$	6.48 $\pm 0.02$
	S-NOD 1	6.12 $\pm 0.01$	6.27 $\pm 0.02$	6.15 $\pm 0.08$	6.12 $\pm 0.02$	6.36 $\pm 0.10$	6.26 $\pm 0.05$	6.07 $\pm 0.01$	6.41 $\pm 0.02$	6.28 $\pm 0.02$
	S-NOD 2	6.29 $\pm 0.01$	6.52 $\pm 0.01$	6.29 $\pm 0.02$	6.22 $\pm 0.01$	6.52 $\pm 0.15$	6.37 $\pm 0.00$	6.17 $\pm 0.01$	6.64 $\pm 0.01$	6.41 $\pm 0.02$

Table 5.4: Comparison of average path length (PL) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/4$ .

Model	Experimental case with $\beta = \pi/6$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
PL [m]	NOD	6.18 $\pm 0.04$	6.35 $\pm 0.02$	6.35 $\pm 0.06$	6.23 $\pm 0.04$	6.40 $\pm 0.06$	6.39 $\pm 0.02$	6.15 $\pm 0.01$	6.62 $\pm 0.05$	6.38 $\pm 0.02$
	S-NOD 1	6.11 $\pm 0.03$	6.22 $\pm 0.03$	6.10 $\pm 0.01$	6.10 $\pm 0.00$	6.25 $\pm 0.08$	6.18 $\pm 0.04$	6.09 $\pm 0.01$	6.35 $\pm 0.02$	6.19 $\pm 0.01$
	S-NOD 2	6.27 $\pm 0.01$	6.45 $\pm 0.02$	6.33 $\pm 0.05$	6.21 $\pm 0.01$	6.45 $\pm 0.10$	6.39 $\pm 0.04$	6.20 $\pm 0.01$	6.65 $\pm 0.01$	6.46 $\pm 0.03$

Table 5.5: Comparison of average path length (PL) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/6$ .

The results of the average centerline deviations exhibited by the robot in each experimental case for NOD, S-NOD 1, and S-NOD 2 across both values of  $\beta = \pi/4$  and  $\beta = \pi/6$  are compared in Table 5.6 and Table 5.7, respectively. These tables summarize the results and observations discussed in the previous sections. The NOD results were illustrated in Fig. 3.7, the S-NOD 1 results were illustrated in Fig. 5.4, and the S-NOD 2 results were illustrated in Fig. 5.10.

Model	Experimental case with $\beta = \pi/4$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
CD [m <sup>2</sup> ]	NOD	881 $\pm 127$	1156 $\pm 26$	721 $\pm 142$	644 $\pm 178$	1262 $\pm 155$	1122 $\pm 169$	317 $\pm 69$	1617 $\pm 123$	1311 $\pm 85$
	S-NOD 1	303 $\pm 38$	553 $\pm 79$	572 $\pm 253$	348 $\pm 40$	942 $\pm 408$	933 $\pm 124$	156 $\pm 2$	1165 $\pm 20$	948 $\pm 29$
	S-NOD 2	608 $\pm 63$	1089 $\pm 27$	897 $\pm 55$	447 $\pm 44$	1190 $\pm 392$	1082 $\pm 24$	337 $\pm 21$	1482 $\pm 52$	1150 $\pm 57$

Table 5.6: Comparison of average centerline deviation (CD) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/4$ .

Model	Experimental case with $\beta = \pi/6$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
CD [m <sup>2</sup> ]	NOD	627 $\pm 83$	950 $\pm 94$	950 $\pm 59$	554 $\pm 171$	831 $\pm 69$	1032 $\pm 77$	455 $\pm 32$	1555 $\pm 100$	1017 $\pm 89$
	S-NOD 1	216 $\pm 12$	484 $\pm 19$	543 $\pm 67$	203 $\pm 16$	632 $\pm 336$	717 $\pm 88$	216 $\pm 61$	1108 $\pm 39$	804 $\pm 18$
	S-NOD 2	636 $\pm 35$	935 $\pm 22$	946 $\pm 98$	438 $\pm 36$	917 $\pm 346$	1101 $\pm 79$	416 $\pm 53$	1468 $\pm 18$	1222 $\pm 72$

Table 5.7: Comparison of average centerline deviation (CD) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/6$ .

The results of the average minimum distance experienced between the robot and the human in each experimental case for NOD, S-NOD 1, and S-NOD 2 across both values of  $\beta = \pi/4$  and  $\beta = \pi/6$  are compared in Table 5.8 and Table 5.9, respectively. These tables summarize the results and observations discussed in the previous sections. The NOD results were illustrated in Fig. 3.8, the S-NOD 1 results were illustrated in Fig. 5.5, and the S-NOD 2 results were illustrated in Fig. 5.11.

Model	Experimental case with $\beta = \pi/4$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
MD [m]	NOD	2.13 $\pm 0.26$	1.14 $\pm 0.02$	1.74 $\pm 0.19$	2.03 $\pm 0.33$	1.11 $\pm 0.03$	1.75 $\pm 0.04$	1.63 $\pm 0.08$	1.21 $\pm 0.03$	2.06 $\pm 0.17$
	S-NOD 1	1.17 $\pm 0.12$	0.74 $\pm 0.04$	1.15 $\pm 0.17$	1.42 $\pm 0.10$	0.85 $\pm 0.13$	1.42 $\pm 0.08$	1.01 $\pm 0.08$	0.97 $\pm 0.02$	1.49 $\pm 0.11$
	S-NOD 2	1.80 $\pm 0.11$	1.08 $\pm 0.02$	1.39 $\pm 0.05$	1.64 $\pm 0.18$	1.00 $\pm 0.07$	1.54 $\pm 0.10$	1.54 $\pm 0.04$	1.14 $\pm 0.02$	1.62 $\pm 0.12$

Table 5.8: Comparison of average minimum distance (MD) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/4$ .

Model	Experimental case with $\beta = \pi/6$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
MD [m]	NOD	1.73 $\pm 0.12$	1.04 $\pm 0.02$	1.69 $\pm 0.12$	1.57 $\pm 0.02$	0.90 $\pm 0.03$	1.52 $\pm 0.15$	1.53 $\pm 0.08$	1.16 $\pm 0.05$	1.45 $\pm 0.07$
	S-NOD 1	1.22 $\pm 0.11$	0.64 $\pm 0.02$	1.09 $\pm 0.04$	1.07 $\pm 0.16$	0.64 $\pm 0.11$	1.27 $\pm 0.04$	1.00 $\pm 0.05$	0.84 $\pm 0.00$	1.33 $\pm 0.06$
	S-NOD 2	1.51 $\pm 0.09$	1.00 $\pm 0.04$	1.41 $\pm 0.06$	1.62 $\pm 0.09$	0.90 $\pm 0.13$	1.58 $\pm 0.06$	1.53 $\pm 0.09$	1.08 $\pm 0.03$	1.58 $\pm 0.10$

Table 5.9: Comparison of average minimum distance (MD) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/6$ .

The results of the average invasion distance experienced between the robot and the human in each experimental case for NOD, S-NOD 1, and S-NOD 2 across both values of  $\beta = \pi/4$  and  $\beta = \pi/6$  are compared in Table 5.10 and Table 5.11, respectively. These tables summarize the results and observations discussed in the previous sections. The NOD results were illustrated in Fig. 3.9, the S-NOD 1 results were illustrated in Fig. 5.6, and the S-NOD 2 results were illustrated in Fig. 5.12.

Model	Experimental case with $\beta = \pi/4$								
	LL	LU	LR	UL	UU	UR	RL	RU	RR
ID [m]	NOD	0	0.29 $\pm 0.06$	0	0	0.38 $\pm 0.03$	0	0	0.05 $\pm 0.09$
	S-NOD 1	0.26 $\pm 0.23$	0.75 $\pm 0.02$	0.26 $\pm 0.25$	0	0.67 $\pm 0.09$	0	0.51 $\pm 0.07$	0.56 $\pm 0.02$
	S-NOD 2	0	0.40 $\pm 0.03$	0	0	0.52 $\pm 0.10$	0	0	0.29 $\pm 0.04$

Table 5.10: Comparison of average invasion distance (ID) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/4$ .

Model	Experimental case with $\beta = \pi/6$								
	LL	LU	LR	UL	UU	UR	RL	RU	RR
ID [m]	NOD	0	0.49 $\pm 0.04$	0	0	0.61 $\pm 0.02$	0	0	0.21 $\pm 0.19$
	S-NOD 1	0.12 $\pm 0.20$	0.78 $\pm 0.03$	0.38 $\pm 0.07$	0.35 $\pm 0.31$	0.78 $\pm 0.06$	0	0.52 $\pm 0.06$	0.65 $\pm 0.02$
	S-NOD 2	0	0.50 $\pm 0.04$	0	0	0.64 $\pm 0.13$	0	0	0.39 $\pm 0.04$

Table 5.11: Comparison of average invasion distance (ID) across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/6$ .

Finally, we present a comparison of the total control used in each case across the NOD, S-NOD 1, and S-NOD 2 sets of experiments. The robot's control variable is its opinion  $z$ , which is used to control the robot's steering, as shown in equation (5.2d). Therefore, the total control throughout a trial is the magnitude of the opinion  $\|z\|$  for all navigation time.

Model	Experimental case with $\beta = \pi/4$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
$\ z\ $	NOD	28.47 $\pm 0.50$	28.47 $\pm 0.50$	23.96 $\pm 0.25$	22.48 $\pm 0.20$	22.70 $\pm 0.95$	22.53 $\pm 0.34$	23.18 $\pm 0.38$	27.64 $\pm 0.59$	28.58 $\pm 0.40$
	S-NOD 1	15.90 $\pm 0.11$	16.24 $\pm 0.03$	13.39 $\pm 0.10$	10.35 $\pm 0.20$	10.36 $\pm 0.47$	10.24 $\pm 0.04$	13.56 $\pm 0.10$	16.23 $\pm 0.20$	16.03 $\pm 0.03$
	S-NOD 2	25.98 $\pm 0.85$	26.61 $\pm 0.41$	21.30 $\pm 0.70$	16.83 $\pm 0.22$	17.00 $\pm 0.29$	17.49 $\pm 0.26$	23.48 $\pm 0.86$	27.08 $\pm 0.11$	26.66 $\pm 0.04$

Table 5.12: Comparison of average control input magnitude  $\|z\|$  across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/4$ .

Model	Experimental case with $\beta = \pi/6$									
	LL	LU	LR	UL	UU	UR	RL	RU	RR	
$\ z\ $	NOD	28.16 $\pm 0.27$	27.61 $\pm 0.82$	24.97 $\pm 0.28$	23.51 $\pm 0.87$	23.78 $\pm 1.33$	23.31 $\pm 0.50$	25.57 $\pm 1.28$	28.38 $\pm 0.44$	28.07 $\pm 0.26$
	S-NOD 1	15.97 $\pm 0.08$	16.29 $\pm 0.09$	13.64 $\pm 0.10$	10.42 $\pm 0.08$	10.21 $\pm 0.76$	10.51 $\pm 0.13$	13.72 $\pm 0.10$	16.14 $\pm 0.08$	16.05 $\pm 0.11$
	S-NOD 2	24.00 $\pm 1.65$	26.59 $\pm 0.33$	21.43 $\pm 0.96$	17.19 $\pm 0.13$	17.20 $\pm 0.57$	17.22 $\pm 0.14$	20.54 $\pm 0.41$	27.06 $\pm 0.32$	26.46 $\pm 0.48$

Table 5.13: Comparison of average control input magnitude  $\|z\|$  across each experimental case in NOD and S-NOD experiments where  $\beta = \pi/6$ .

The results of the total control  $\|z\|$  in each experimental case for NOD, S-NOD 1, and S-NOD 2 across both values of  $\beta = \pi/4$  and  $\beta = \pi/6$  are compared in Table 5.12 and Table 5.13, respectively.

We observe that in the S-NOD 1 experimental set, where the same parameters from NOD were applied in S-NOD, there is a drastic reduction in the total control used within the robot. For the experiments where  $\beta = \pi/4$ , we observe an average decrease in control of 46.75% in all cases. We see almost the same results in the experiments where  $\beta = \pi/6$ , where there is an average control decrease of 47.73% in all cases. We also note a massive reduction of 67.00% and 73.67% in the standard deviations for the  $\beta = \pi/4$  and  $\beta = \pi/6$  trials in S-NOD 1. This indicates that S-NOD 1 is far more efficient in control than NOD.

We now turn to the S-NOD 2 experimental set, in which the parameters of S-NOD were adjusted to produce similar navigation behavior to that of NOD. Although the behavior is similar, we still observe a reduction in total control used by the robot. For the experiments where  $\beta = \pi/4$ , we observe an average

decrease in control of 11.83% in all cases. For the experiments where  $\beta = \pi/6$ , we observe an average reduction in control of 15.93% in all cases. We observe that the standard deviations do not exhibit the same drastic reduction as seen from S-NOD 1. However, the upper deviation still maintains the total control below the lower deviation of the total control from NOD. This suggests that S-NOD 2 is more efficient in control than NOD, while maintaining similar navigation behavior in practice.

### 5.2.2.3 Single-robot, two-human experiments

The setup of each experiment was as follows, following the protocol introduced in Chapter 3.4.2.3. The robot was placed at an initial location (3, -3)m in the space, facing its goal at (-3,  $\sim 2.1$ )m directly. Two human participants would move around the room of their own accord. Five representative trials are reported in this work to showcase the flexibility of the S-NOD control of the robot. Trials are each qualitatively described, and their performance metrics are subsequently compared both to one another and to their NOD counterparts. Videos of these experiments and other animated figures throughout this thesis can be viewed at <https://cathcart-dissertation.github.io>.

#### Robot instructions

The robot was programmed to move at a constant speed  $V_r = 0.7$ m/s toward its goal location, modifying its trajectory when encountering movers according to the navigation model (5.2) with parameters. The parameters used in the S-NOD equation (5.2) are as follows:  $\tau_z = 0.008$ ,  $d = 1$ ,  $a_{11} = a = 0.5$ ,  $\gamma = 4$ ,  $K_u = 2.667$ ,  $U = 4$ ,  $n = 7$ ,  $D_{crit} = 6$ m,  $\tau_{u_s} = 0.08$ ,  $K_{u_s} = 12$ ,  $k_\theta = 2$  and  $\beta = \pi/4$ . The robot was kept unbiased ( $b = 0$ ) for maximum flexibility in passing behavior.

#### Human participant instructions

As these trials were conducted to replicate the NOD experiments in [25] and Chapter 3.4.2.3, the two participants were directed by the researcher on their desired paths. Though their paths were instructed rather than given complete independence to move throughout the space, they were still asked to move at their own pace. Once the four comparative trials were conducted, the two participants were then informed that they could walk freely in the room at any speed and without a requirement for consistency in direction or pace, if desired. This was the instruction provided for the participants in the NOD experiments. We present the mirrored trials here, as well as a selected trial of an unusual passing scenario that the S-NOD programmed robot could navigate.

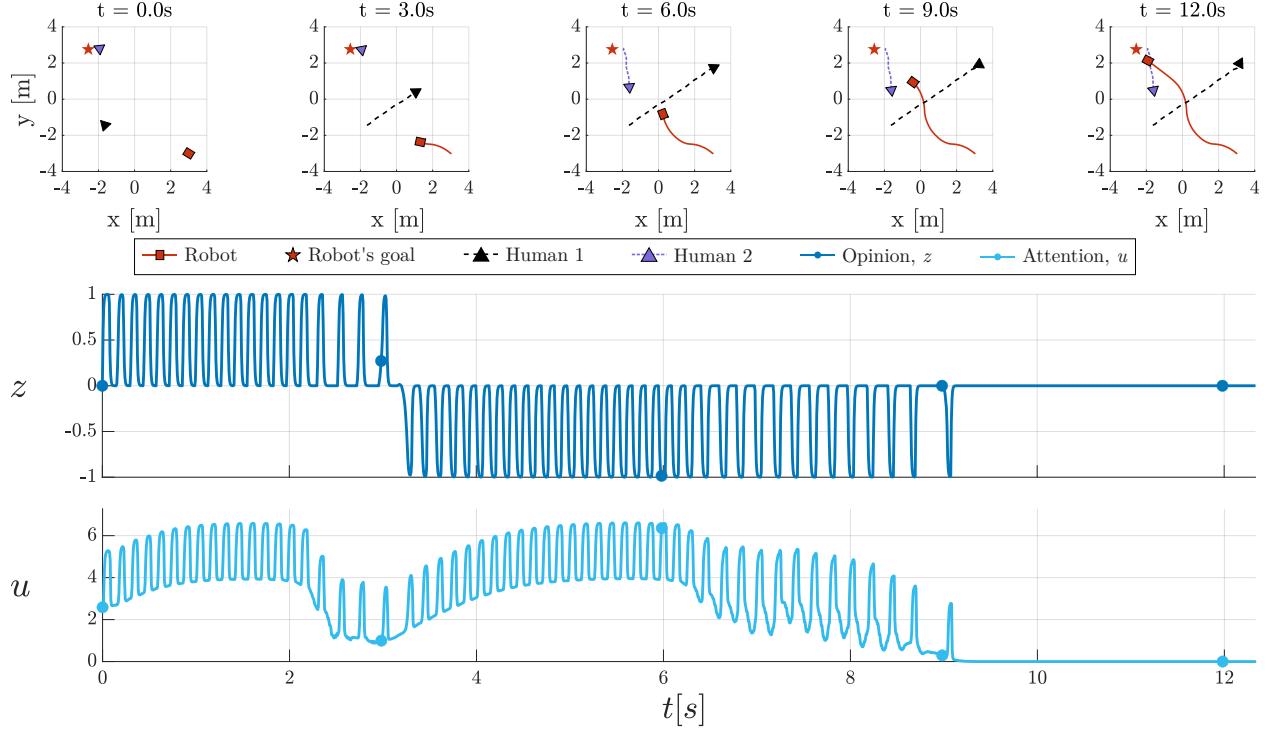


Figure 5.13: S-NOD Trial 1. (Top) The trajectories of all participants at snapshots along the duration of the trial. Human 1 walks across the robot's path, and the robot turns left to avoid them. Meanwhile, Human 2 abruptly stops moving while in the robot's path, so the robot turns right to avoid them. (Middle) Opinion  $z$  over time  $t$  of the robot. (Bottom) Sensitivity  $u$  over time  $t$  of the robot. Dots along the  $z$  and  $u$  traces correspond with the timestamps in the top trajectory snapshots.

### Trial descriptions

Trial 1: Two humans walk independently within the laboratory space as shown in Fig. 3.10. Human 1 moves perpendicularly across the robot's path from the robot's left side towards its right. We observe the sensitivity  $u$  rise in spikes as Human 1 approaches, immediately inducing high-frequency positive spiking in the robot's opinion  $z$ , which causes the robot to turn towards its left. The sensitivity drops as Human 1 is passed, thus decreasing the robot's opinion spiking frequency. The sensitivity rises again as Human 2 remains stationary along the robot's path on the robot's left side. The opinion immediately begins to exhibit high-frequency negative spiking, causing the robot to turn to the right and away from Human 2. As it turns further out along the human's path, we see the sensitivity value fall to zero, and the spikes in both opinion and sensitivity decrease in frequency. The robot then reaches its goal.

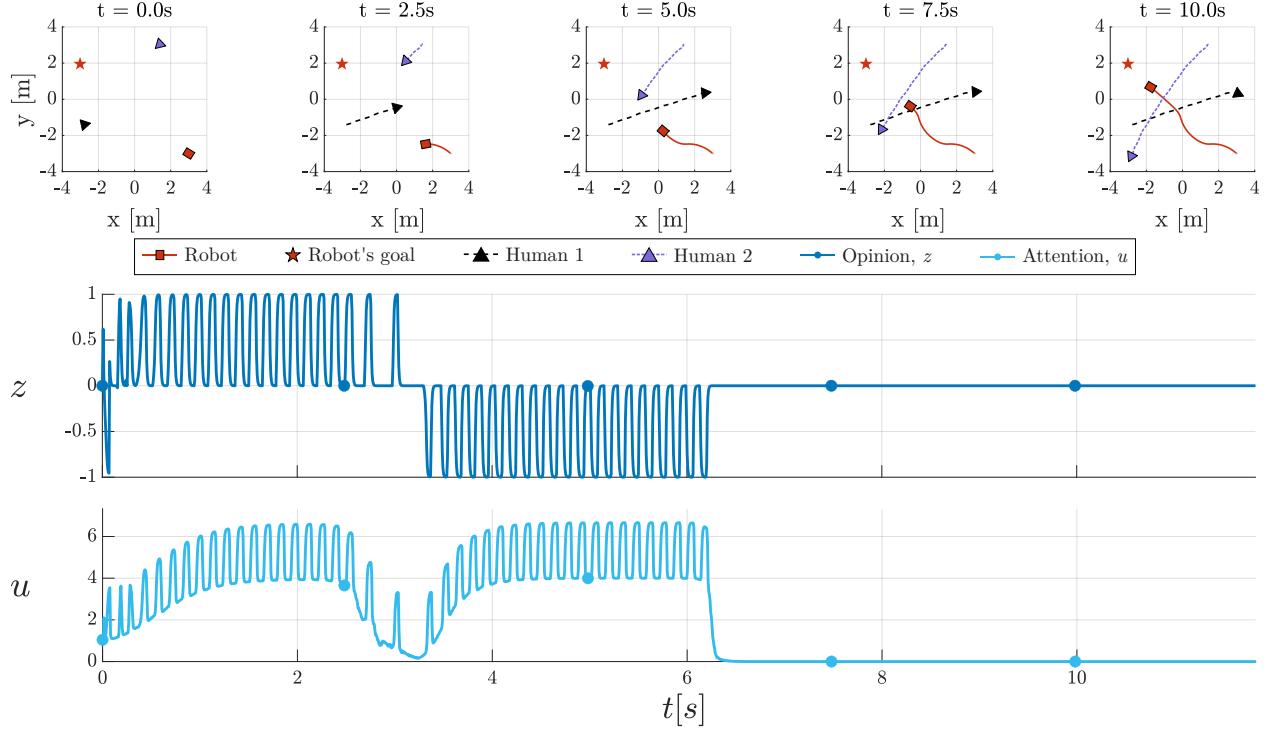


Figure 5.14: S-NOD Trial 2. (Top) Humans 1 and 2 walk across the robot's path. The robot first turns left to avoid Human 1, then right to avoid Human 2. (Middle) Opinion  $z$  over time  $t$  of the robot. (Bottom) Sensitivity  $u$  over time  $t$  of the robot. Dots along the  $z$  and  $u$  traces correspond with the timestamps in the top trajectory snapshots.

Trial 2: Two humans walk independently but close together within the laboratory space as shown in Fig. 3.11. While this form of interaction once caused a more pronounced turn by the robot and extended the trial in NOD, the robot with S-NOD gets to its goal even faster than in Trial 1. Human 1 moves at an angle across the robot's path from the robot's left side towards its right side. The sensitivity rises in spikes more slowly than that seen in Trial 1. Initially, the robot's opinion spikes negatively, but quickly corrects to spike positively and with high frequency. As Human 1 passes, Human 2 walks more slowly across the robot's path from right to left. The sensitivity begins spiking again as Human 2 appears very nearby in the robot's path on its right side, but moving towards the left. The robot's opinion starts to spike negatively, causing it to turn to the right. With no further humans to avoid, the robot moves towards its goal.

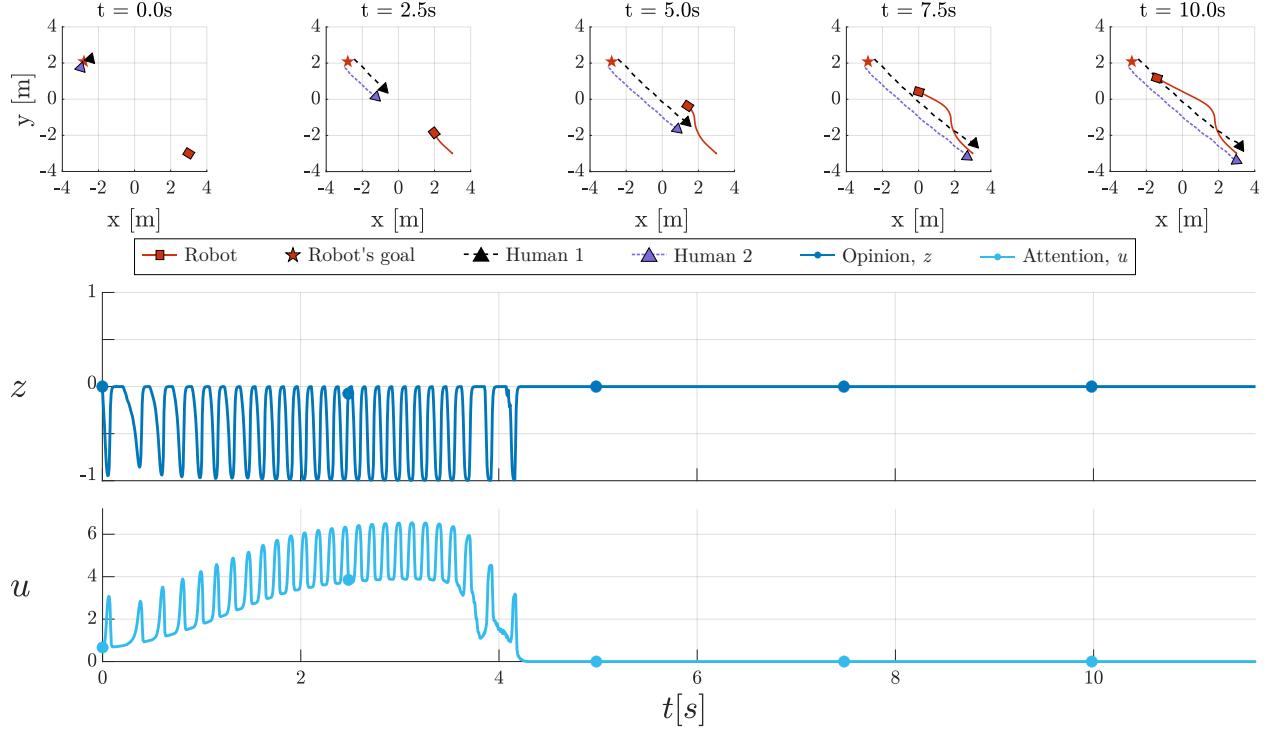


Figure 5.15: S-NOD Trial 3. (Top) Humans 1 and 2 walk toward the robot’s initial position with no passing space in between them. The robot avoids both humans by turning right. (Middle) Opinion  $z$  over time  $t$  of the robot. (Bottom) Sensitivity  $u$  over time  $t$  of the robot. Dots along the  $z$  and  $u$  traces correspond with the timestamps in the top trajectory snapshots.

Trial 3: Two humans walk closely together within the laboratory space as shown in Fig. 3.12. As the sensitivity begins to spike, the opinion begins to spike negatively, causing the robot to turn to its right. As the humans come closer to the robot, the frequency of the spikes increases. Though its control is spiking, the robot appears to make a smooth right turn that avoids both humans and then moves directly toward its goal.

Trial 4: Two humans walk together within the laboratory space as shown in Fig. 3.12. Though still nearby one another, Humans 1 and 2 begin moving at different times. This allows the robot to interact with each human independently rather than as a unit, as seen in Trial 3. We observe that at the beginning of the trial, the robot’s sensitivity is nonzero but not spiking. This causes minimal turning behavior from the robot, though its opinion is slightly nonzero as well. When Human 1 is seen sufficiently nearby and on the robot’s right side, the robot’s sensitivity begins to spike and build, causing the robot’s opinion to spike positively. The robot turns to the left, but then is immediately on a collision path with Human 2, who is positioned on its left side. This further raises the sensitivity, causing an increase in spiking frequency as the opinion spikes negatively. The robot turns right to avoid Human 2, then moves towards its goal.

Trial 5: Two humans face one another from across the room and begin walking directly across the robot’s

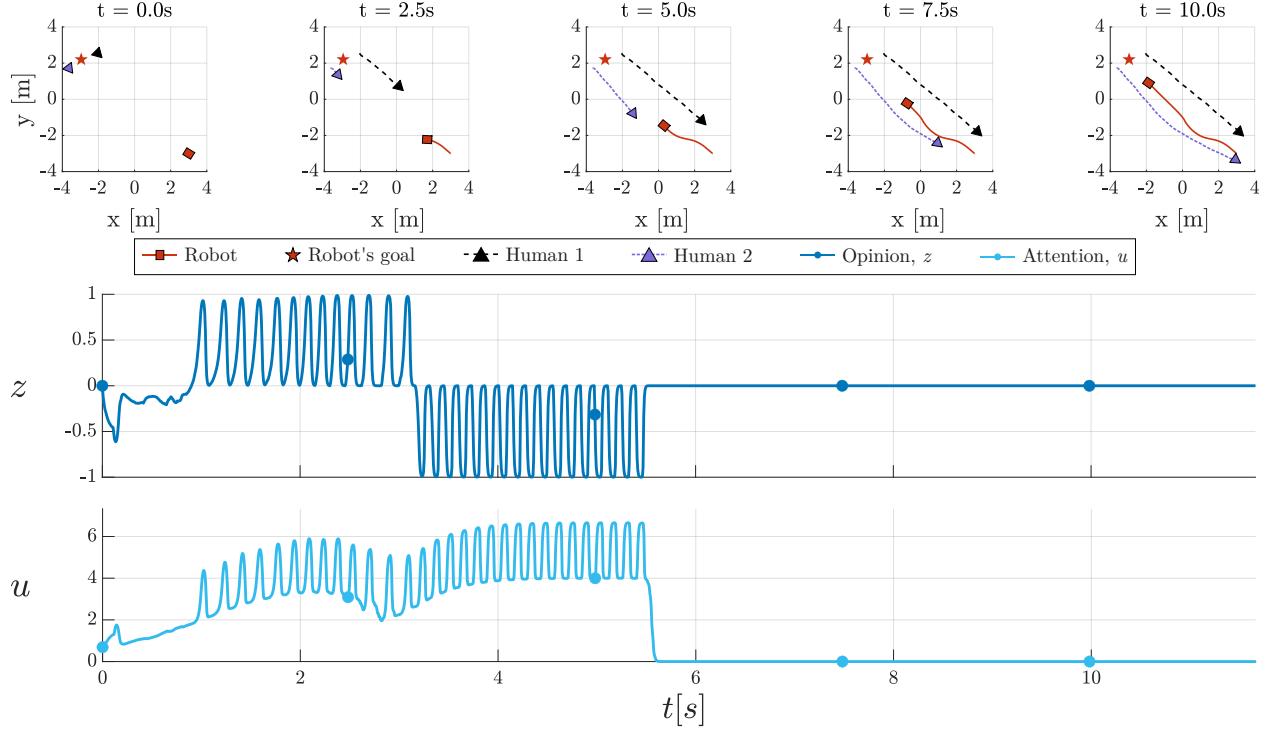


Figure 5.16: S-NOD Trial 4. (Top) Humans 1 and 2 walk toward the robot's initial position with passing space in between them. The robot initially dodges Human 1 with a left turn, finds itself in the path of Human 2, and dodges Human 2 with a right turn. (Middle) Opinion  $z$  over time  $t$  of the robot. (Bottom) Sensitivity  $u$  over time  $t$  of the robot. Dots along the  $z$  and  $u$  traces correspond with the timestamps in the top trajectory snapshots.

path. The robot's sensitivity begins to spike and build, initially causing its opinion to spike positively, prompting a left turn. As the humans meet in the center of the laboratory, they begin to turn around and look at one another. This turning of the observed relative heading direction of each human  $\eta_h$  causes the robot to move between spiking negatively and positively. This produces an almost straight-looking trajectory, although the robot's internal control is indicating oscillating decisions between left and right passing. Finally, once the humans have completed a full loop around one another, they continue along a straight path away from the robot. The robot then turns right to avoid Human 2 as it walks past the robot on its left. When both humans are away from the robot, it navigates directly towards its goal.

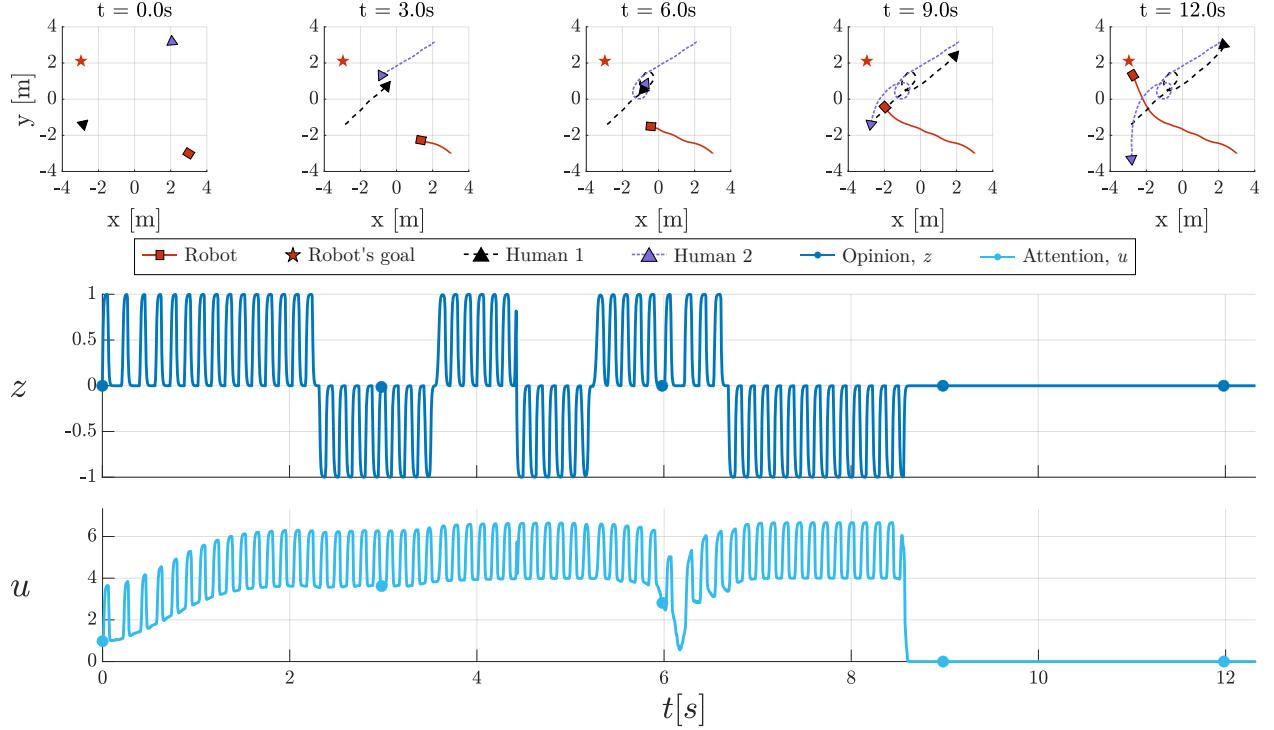


Figure 5.17: S-NOD Trial 5. (Top) Humans 1 and 2 cross the robot’s path, meeting in the middle and moving in a full loop around one another, before continuing on their paths. The robot dodges both humans on the left, although its opinion evolution shows multiple changes of opinion as the humans make their loops. (Middle) Opinion  $z$  over time  $t$  of the robot. (Bottom) Sensitivity  $u$  over time  $t$  of the robot. Dots along the  $z$  and  $u$  traces correspond with the timestamps in the top trajectory snapshots.

### Comparisons of NOD and S-NOD performance

We illustrate the comparable navigation trials 1-4 for both NOD and S-NOD social robots in Fig. 5.18 for easy reference as we compare their results. We present the performance metrics across trials 1-4 for both NOD and S-NOD robots, as well as those of trial 5 for the S-NOD robot, in Table 5.14. We then discuss the differences observed between the performance of each algorithm.

We compare the performance metrics across trials.

We first examine both sets of NOD and S-NOD experiments for each trial number. We note that *all* S-NOD robot path lengths are shorter than *any* path lengths recorded in the NOD robot experiments. We note that for each trial, *all* S-NOD robot navigation times are shorter than those in the comparable NOD experiment. Both of these results validate that S-NOD produces more efficient navigation when completing the same task as a NOD robot. Examining the metrics of spatial intrusiveness, we observe that both algorithms maintain comparable minimum distances and rarely invade the personal space of passing humans. However, occasional intrusions into personal space were observed, most notably in Trial 4 for both NOD and S-NOD robots, and in Trial 5 for the S-NOD robot. This indicates that spatial intrusiveness was

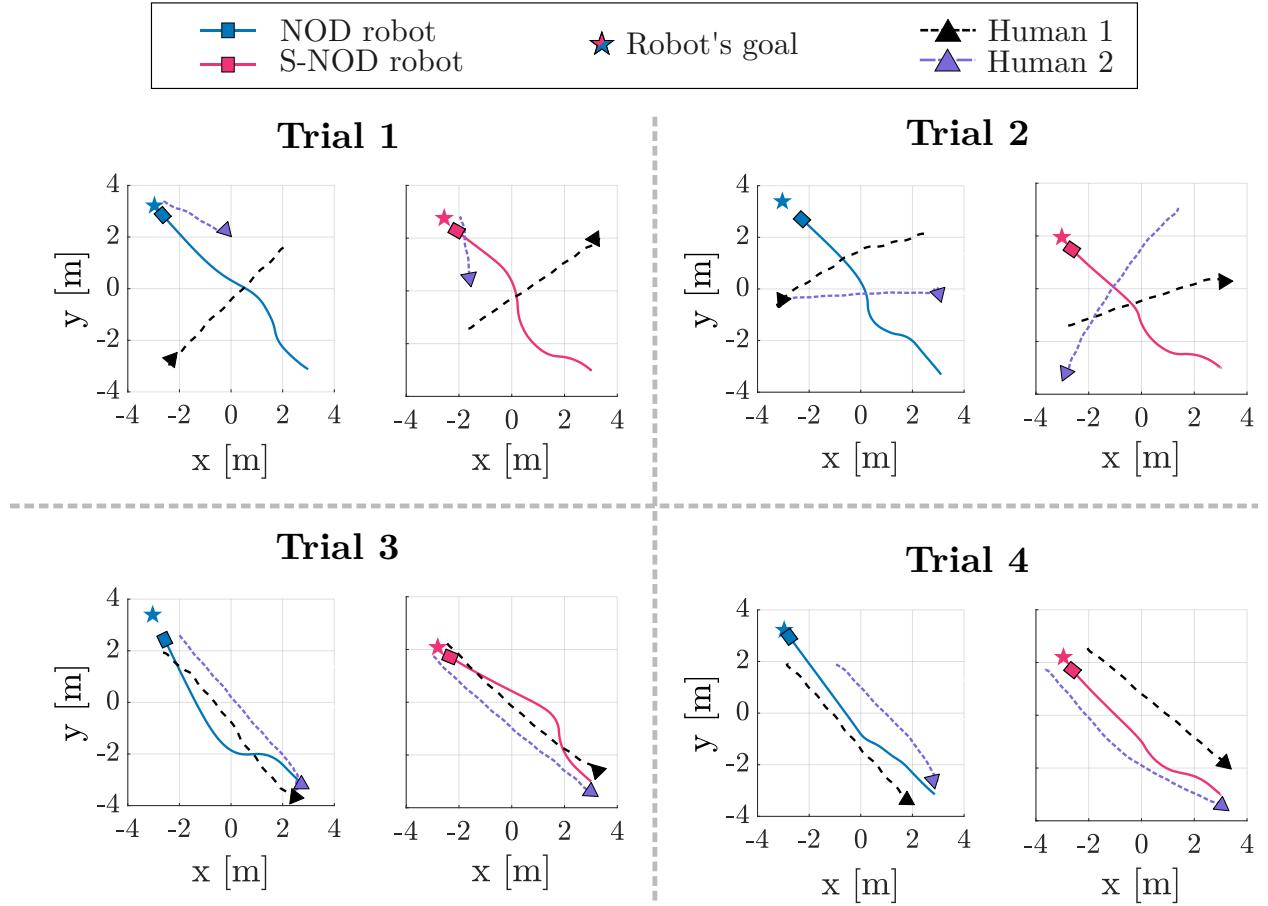


Figure 5.18: The full trajectories of Trials 1-4 as reported for the NOD robot in Chapter 3.4.2.3 and the S-NOD robot in this chapter.

Trial #	Model	NT [s]	PL [m]	MD <sub>1</sub> [m]	MD <sub>2</sub> [m]	ID <sub>1</sub> [m]	ID <sub>2</sub> [m]	$\ z\ $
1	NOD	12.53	8.77	2.03	1.38	0	0	42.32
	S-NOD	12.34	7.89	2.72	1.10	0	0	21.61
2	NOD	13.20	9.24	2.18	1.55	0	0	39.42
	S-NOD	11.84	7.57	2.25	1.50	0	0	18.87
3	NOD	12.93	9.05	0.54	1.29	0.28	0	36.22
	S-NOD	11.64	7.49	0.73	1.34	0	0	14.04
4	NOD	12.13	8.49	0.36	1.01	0.47	0.28	15.64
	S-NOD	11.66	7.51	1.60	0.80	0	0.40	16.12
5	NOD	-	-	-	-	-	-	-
	S-NOD	12.32	7.89	1.93	0.84	0	0.34	22.92

Table 5.14: Performance metrics of the single-robot, two-human trials 1-4 of NOD and S-NOD robots, and additionally of trial 5 for the S-NOD robots. The NOD data is duplicated from Table 3.3. Shorthand: NT = Navigation Time, PL = Path Length, MD <sub>$\ell$</sub>  = Minimum Distance with Human  $\ell$ , ID <sub>$\ell$</sub>  = Invasion Distance of Human  $\ell$ .

not eliminated by either approach, which follows as there was *no bound or constraint* on the distance to maintain in either algorithm. The ability of both algorithms to make decisions proactively, avoiding collisions with sufficient space to frequently avoid invading personal space, is an unexpected yet welcome feature. We believe that adding a constraint on the acceptable distance from a human could be imposed and addressed elegantly with S-NOD, thereby resolving the trade-off between efficiency and spatial intrusiveness. Finally, we consider the difference in total control  $\|z\|$  used throughout each trial. Importantly, we see that the total control  $\|z\|$  was substantially lower in S-NOD robots than NOD robots for almost all trials. Only Trial 4 shows an increase in control used, but we recall that the NOD Trial 4 caused very close interactions between the robot and the human, which moved quickly within the trial. This did not give the robot much time or space to form opinions, and therefore, less control was input into the system.

Trial 5 of an S-NOD robot, which did not have a comparable experiment with a NOD robot, further illustrates S-NOD’s robustness in complex navigation environments. A NOD robot may have been able to steer away from the interaction to move towards its goal successfully, but its opinion would have likely remained at a high value with only a single sign; that is, the robot would have formed a single opinion and not deviated from it even as the environment changed. While this robustness may be helpful, it indicates that a NOD robot is not as swiftly responsive to changing navigation interactions once it has already formed an opinion.

Overall, we observe that S-NOD can facilitate more efficient and agile decision-making in a social robot and can handle more complex navigation scenarios than its NOD counterpart. As S-NOD is more capable of sequential decision-making than NOD, a robot using the S-NOD algorithm is better equipped to be responsive to changes in the environment. As these experiments demonstrate, S-NOD enables the robot to rapidly change its opinion about its preferred passing side, even when interacting with little room between itself and humans who move across its path. The results of these experiments highlight S-NOD’s ability to navigate responsively and efficiently in various passing scenarios, albeit with some trade-offs in spatial intrusiveness under certain conditions.

### 5.3 Discussion of effect of $K_u$ and $U$ on navigation

This section utilizes simulations of a social robot navigating around one or two humans to show the progression of behavior as the tuning parameters  $K_u$  and  $U$  are increased past their theoretically set critical points. Although deadlock can be guaranteed when these thresholds are minimally passed, we explore the behavioral impacts on an embodied robot as it operates with higher values for these parameters. We recall from Chapter 4 and in [26] that a critical value  $K_u^* = d^3/3a$  marks the point where the S-NOD system’s

bifurcation diagram changes from a supercritical pitchfork to a subcritical pitchfork with hysteresis. We also recall the bifurcation point at the basal sensitivity value  $u_0^* = d/a$ . The analysis within Chapter 4 informs us that a value of  $K_u > K_u^*$  will cause a *subcritical* pitchfork bifurcation in the system, leading to a section of multistability near the ultrasensitive bifurcation point, which the S-NOD algorithm repeatedly returns the system to. It also informs us that a system that experiences  $u_0 > u_0^*$  will form a non-neutral opinion  $z$ , as the solution  $z = 0$  is unstable in this region.

In this section, we aim to explore through simulation how NOD and S-NOD robots behave under conditions where both parameters are set to force the system to exceed these critical values. We anticipate this will lead to the NOD robot becoming “stuck” in an opinion state as detailed in Chapter 4 and [26], which should appear as a robot that continues to turn in a single direction, whereas S-NOD will remain agile and responsive and will turn cooperatively throughout an interaction with humans. These simulations highlight the weakness of NOD decision-making compared to S-NOD decision-making in a social navigation scenario with increasing complexity. We extend these results in the next section with a demonstration that showcases how NOD and S-NOD robots perform in a navigation scenario involving frequent, varied human interactions.

Recall from Chapter 5.1 that in the social robot setting, we define  $u_0$  as a function of proximity to an oncoming human  $\ell$  with distance  $\chi_\ell = D_\ell = \|\mathbf{x} - \mathbf{x}_{h_\ell}\|$  and direction,  $\kappa_\ell = D_{crit} \cdot \exp(1 - \max(\cos \eta_\ell, 0)^{-1})$ . The basal sensitivity function is defined as  $u_{0,i_\ell} = U(\kappa_{i_\ell}^n / (\kappa_{i_\ell}^n + \chi_{i_\ell}^n))$ , with its upper bound set by parameter  $U$ . We can therefore set  $U$  to be a multiple of  $u_0^*$ , such that there are proximity values at which the robot will be *guaranteed* to have a  $u_0$  value that is past the system’s bifurcation point. We can also set  $K_u$  to be a multiple of  $K_u^*$  to build our intuition of the effect of large gain on the robot’s opinion, pushing the sensitivity  $u$  past the system’s bifurcation.

For all subsequent simulations, the parameters within NOD equation (3.3) and S-NOD equation (5.2) held constant are as follows:  $\tau_z = 0.02$ ,  $d = 1$ ,  $a_{11} = a = 1$ ,  $\gamma = 3$ ,  $n = 2$ ,  $D_{crit} = 3.5m$ ,  $\tau_{u_s} = 0.2$ ,  $K_{u_s} = 10$ ,  $k_\theta = 2$  and  $\beta = \pi/4$ . The robot was kept unbiased ( $b = 0$ ) for maximum flexibility in passing behavior.

We consider a scenario where a robot begins at point (0,0)m in a space, directly facing its goal at (0, 6.67)m<sup>1</sup>. A stationary human facing the robot’s starting location directly sits at (0, 3.8) m. The robot travels with velocity  $V = 0.6\text{m/s}$  as it navigates the space. We compare the resultant trajectories and opinion evolutions of NOD and S-NOD controlled social robots within this environment at two values of  $K_u \in [5K_u^*, 10K_u^*]$  and four values of  $U \in [1.25u_0^*, 1.5u_0^*, 1.75u_0^*, 1.25u_0^*]$ . The resultant trajectories are illustrated in Fig. 5.19 and each trial’s performance metrics are summarized in Table 5.15.

Several consistencies in behavior can be identified, either through visual inspection or by examining the

---

<sup>1</sup>These initial and goal locations were set to be within the bounds of the experimental space used within the single-robot, single-human setups.

$U$	Model	$K_u$	Navigation Time [s]	$\min(D_1)$	Invasion Distance of Human 1 [m]	$\ z\ $
$1.25u_0^*$	NOD	$5K_u^*$	11.18	0.60	1.37	22.106
		$10K_u^*$	11.19	0.61	1.37	45.40
	S-NOD	$5K_u^*$	10.98	0.43	1.28	8.14
		$10K_u^*$	10.98	0.43	1.28	9.21
$1.5u_0^*$	NOD	$5K_u^*$	11.34	0.78	1.21	27.03
		$10K_u^*$	11.34	0.79	1.20	47.44
	S-NOD	$5K_u^*$	11.07	0.59	1.29	10.20
		$10K_u^*$	11.06	0.57	1.28	11.26
$1.75u_0^*$	NOD	$5K_u^*$	11.43	0.90	0.89	30.34
		$10K_u^*$	11.43	0.90	0.88	49.08
	S-NOD	$5K_u^*$	11.13	0.68	1.23	11.97
		$10K_u^*$	11.11	0.67	1.24	13.07
$2u_0^*$	NOD	$5K_u^*$	11.49	0.97	0.66	32.70
		$10K_u^*$	11.49	0.98	0.65	50.36
	S-NOD	$5K_u^*$	11.16	0.74	1.16	13.55
		$10K_u^*$	11.15	0.72	1.17	14.68

Table 5.15: Performance metrics of the simulated navigation paths illustrated in Fig. 5.19 taken by NOD and S-NOD controlled robots across various values of  $U$  and  $K_u$  as the robot interacts with one human.

provided performance metrics.

We notice that all trajectories across both navigation models and all parameter options are very similar. Each trajectory chooses to dodge the human towards the robot's right. We observe that the trajectories across each case also exhibit the same pattern, determined by the  $U$  value: the lower the  $U$  value, the closer the robot approaches the human before making its turn away. This is reflected in the opinion evolutions as well, with the opinions becoming nonzero earlier in order of the highest  $U$  values.

Trajectories in the NOD social robot appear to provide more passing space than the comparable trajectories of the S-NOD social robot. This is confirmed by the consistently lower minimum distances seen in the S-NOD cases within each group of experiments with the same  $U$  value.

To understand the go-to-goal behavior exhibited by the S-NOD and NOD robots in light of the saturated opinion state of the NOD robot as in Fig. 5.19B., we recall the definition of the robot's heading update in equations (3.3c) and (5.2d):  $\dot{\theta} = k_\theta \sin(\widetilde{u_{0_\ell}} \cdot \beta \cdot \tanh(z) + (1 - \widetilde{u_{0_\ell}}) \phi)$ . The turn away from the direction of the goal  $\phi$  by some fraction of  $\beta$  only occurs when both  $z \neq 0$  and  $\widetilde{u_{0_\ell}} \neq 0$ . That means the opinion can remain non-neutral ( $z \neq 0$ ), but it will not affect the robot's path towards its goal when no humans are nearby the robot ( $\widetilde{u_{0_\ell}} = 0$ ).

It appears that the variation of  $K_u$  has no meaningful effect on the trajectories within either the NOD or S-NOD robot. Instead, its influence is evident in the evolution of opinions for each robot. The larger  $K_u$  value in the NOD social robot causes the opinion to saturate at a  $z = -1$  value, resulting in a right turn.

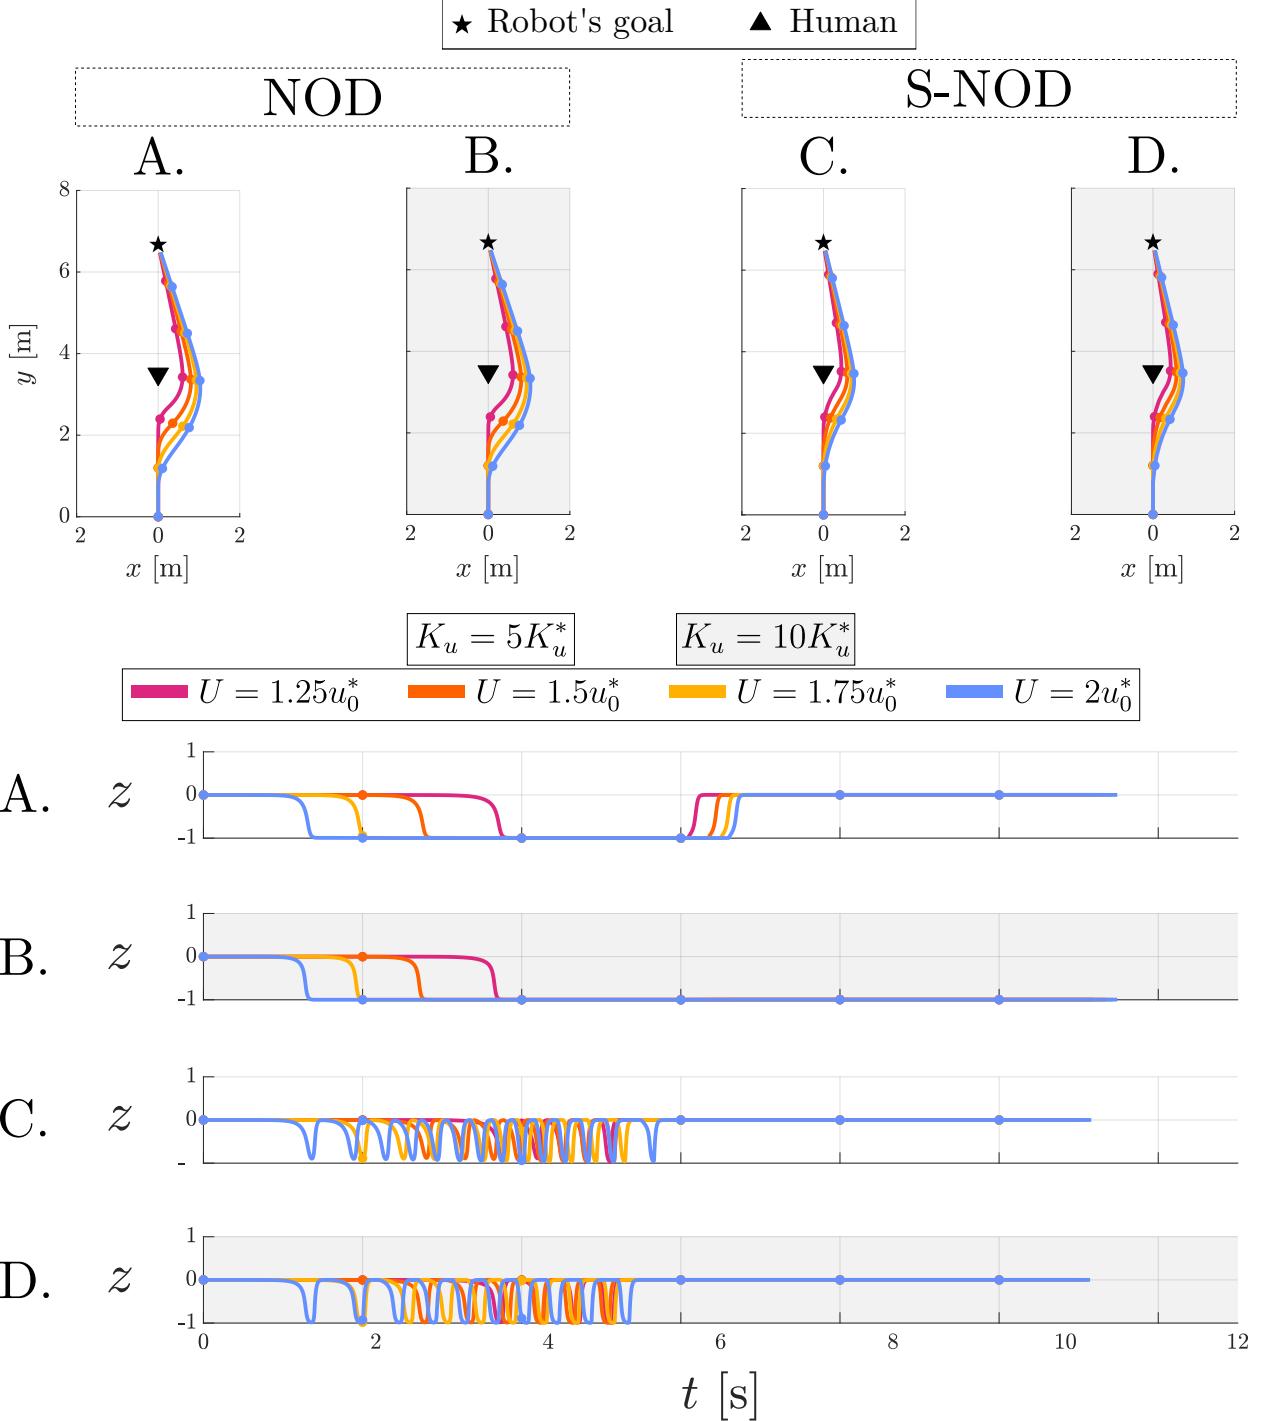


Figure 5.19: Trajectory and opinion comparisons between NOD and S-NOD social robots with various  $K_u$  and  $U$  values around a single stationary human. (Top) The trajectories of a social robot. (Bottom) The opinion of the robot  $z$  over time  $t$  within each trial. Identifiers A.-D. link each trajectory to its opinion. Dots mark even temporal steps along each trajectory. White backgrounds correspond to  $K_u = 5K_u^*$  and gray shaded backgrounds correspond to  $K_u = 10K_u^*$ .

However, this value ( $z \neq 0$ ) remains after the robot has passed the human ( $u_{0t} = 0$ ). Whereas the larger  $K_u$  value in the S-NOD social robot causes broader and deeper spikes at the  $z = -1$  value, increasing in frequency as the human approaches, but ultimately returning to the sensitive  $z = 0$  state after the robot has exited its interaction with the human. This follows directly in line with the discussions in Chapter 4.2.3 of the limitations of tuning NOD and NOD's ability to become too robust to change its opinion state.

We note the consistently shorter navigation times of the S-NOD robot compared to its comparable NOD robot, but recognize that this comes at the expense of lower minimum distances maintained and higher invasion distances for the human in all cases where  $U > 1.5u_0^*$ . However, we assert that the S-NOD algorithm is exhibiting all of the benefits that were hoped for in its formulation, in that the decision-making is extremely agile while remaining robust. When considering the results of the robot's spatial intrusiveness, we recall that there is *no explicit bound* on maintaining a certain distance from approaching humans in this social robot setting. The naturally occurring, mostly comfortable passing behavior exhibited by an S-NOD (and NOD) robot is a bonus, not a feature. The NOD robot appears to be less socially intrusive because it is less agile, forming and maintaining turning behavior past the point of necessity. We believe that adding a constraint on the acceptable distance from a human could be imposed and addressed very elegantly with S-NOD, thereby resolving the trade-off between efficiency and spatial intrusiveness.

Finally, we discuss the amount of control used within each experiment. There are apparent differences in the scale of control shown between the NOD and S-NOD robots, with the latter proving to require much less control signal to the robot. There are also inter-model differences in control: we see that a doubling of the  $K_u$  value from  $5K_u^*$  to  $10K_u^*$  results in an increase of total control within the range of 50% to 105% for the NOD robot, whereas this range is only 8.34% to 13.14% for the S-NOD robot.

With these results, we now consider increasing the complexity of the environment by introducing another human to observe how the robot responds. Let a second human begin from (0,8)m and move with a constant direction towards its left at a slow constant speed  $V_{h_2} = 0.3\text{m/s}$ . This interaction is designed to prompt the robot to interact with the second human after it has interacted with the first human. The resultant trajectories are illustrated in Fig. 5.20, and each trial's performance metrics are summarized in Table 5.16.

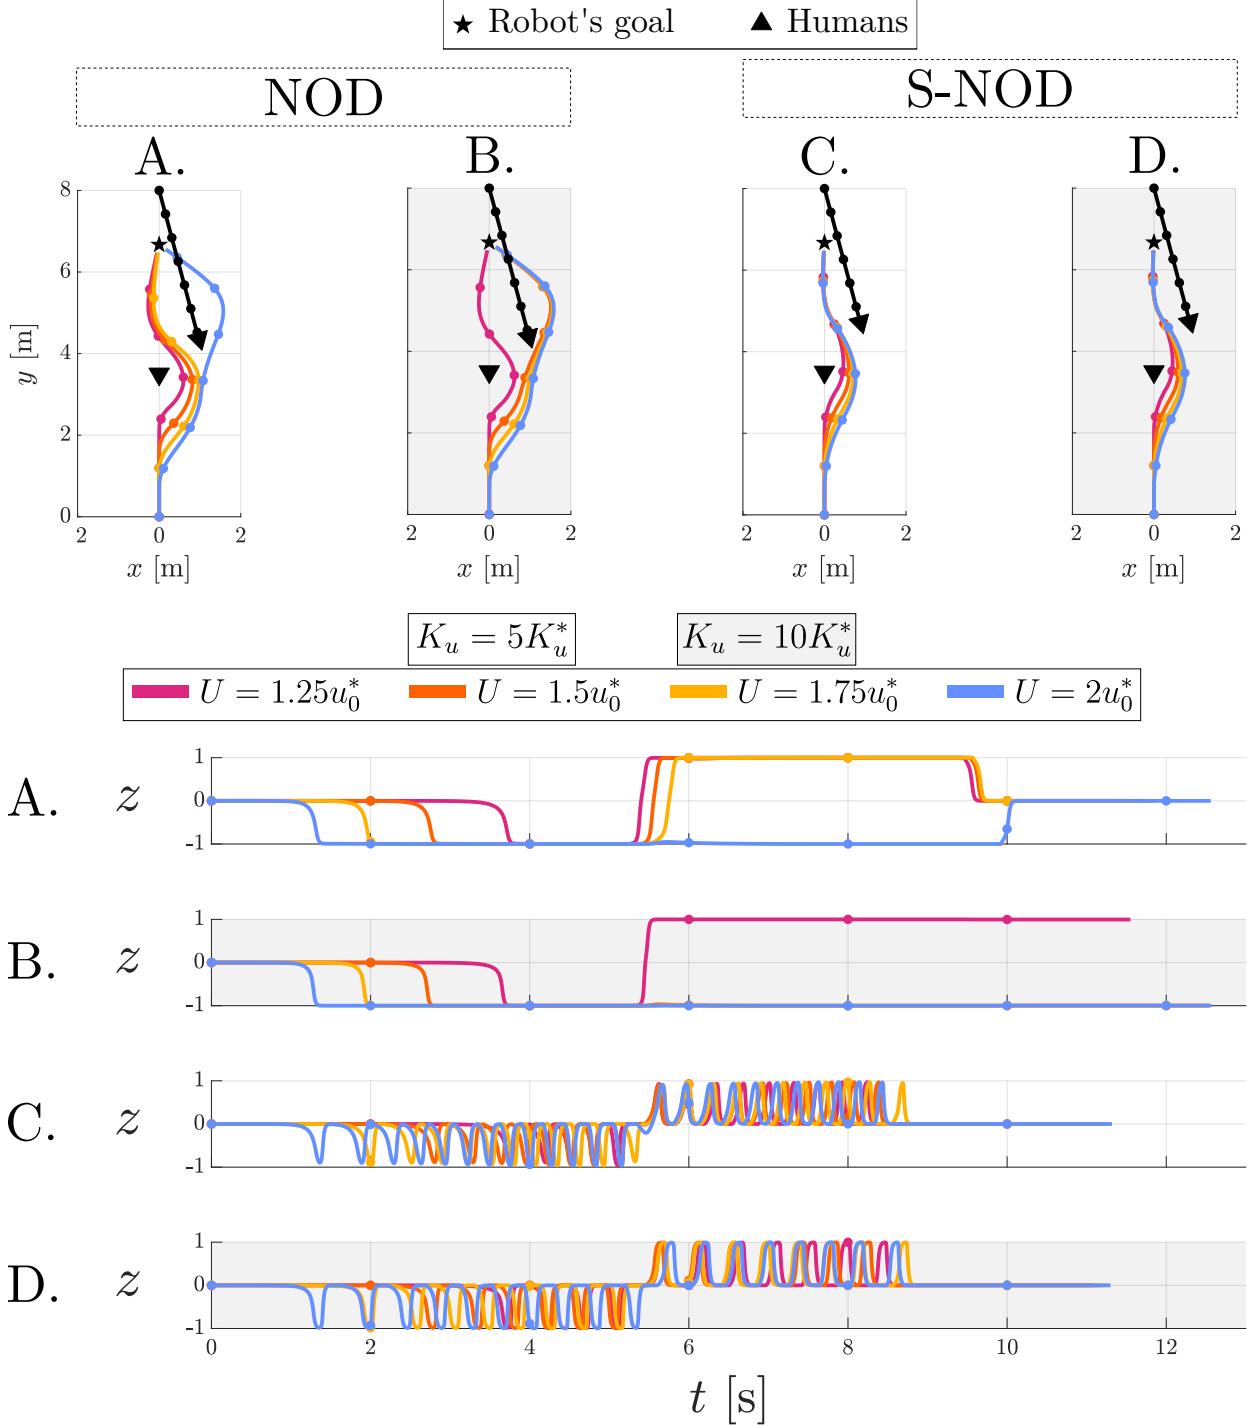


Figure 5.20: Trajectory and opinion comparisons between NOD and S-NOD social robots with various  $K_u$  and  $U$  values around two humans. (Top) The trajectories of a social robot. (Bottom) The opinion of the robot  $z$  over time  $t$  within each trial. Identifiers A.-D. link each trajectory to its opinion. Dots mark even temporal steps along each trajectory. White backgrounds correspond to  $K_u = 5K_u^*$  and gray shaded backgrounds correspond to  $K_u = 10K_u^*$ .

$U$	Model	$K_u$	Navigation Time [s]	$\min(D_1)$	$\min(D_2)$	Invasion Distance of Human 1 [m]	Invasion Distance of Human 2 [m]	$\ z\ $
$1.25u_0^*$	NOD	$5K_u^*$	11.54	0.55	1.00	1.37	0.42	33.79
		$10K_u^*$	11.55	0.58	0.99	1.38	0.44	45.28
	S-NOD	$5K_u^*$	11.07	0.44	0.70	1.28	0.69	13.22
		$10K_u^*$	11.06	0.43	0.70	1.28	0.69	15.05
$1.5u_0^*$	NOD	$5K_u^*$	11.74	0.68	0.94	1.22	0.52	36.67
		$10K_u^*$	12.52	0.88	0.71	1.22	0.75	47.41
	S-NOD	$5K_u^*$	11.19	0.60	0.71	1.29	0.71	15.07
		$10K_u^*$	11.17	0.59	0.70	1.29	0.71	16.53
$1.75u_0^*$	NOD	$5K_u^*$	11.86	0.80	0.88	0.89	0.62	38.79
		$10K_u^*$	12.55	1.01	0.74	0.89	0.71	49.10
	S-NOD	$5K_u^*$	11.27	0.70	0.71	1.24	0.72	16.75
		$10K_u^*$	11.25	0.68	0.70	1.24	0.71	18.41
$2u_0^*$	NOD	$5K_u^*$	12.56	1.09	0.76	0.66	0.69	41.49
		$10K_u^*$	12.57	1.09	0.76	0.66	0.69	50.39
	S-NOD	$5K_u^*$	11.32	0.75	0.72	1.16	0.72	17.78
		$10K_u^*$	11.30	0.75	0.73	1.17	0.73	19.65

Table 5.16: Performance metrics of the simulated navigation paths illustrated in Fig. 5.20 taken by NOD and S-NOD controlled robots across various values of  $U$  and  $K_u$  as the robot interacts with two humans.

We can observe a drastic difference in performance between NOD social robots with varying  $K_u$  and  $U$  values, whereas the S-NOD robot exhibits consistent behavior across these parameter changes.

We have noticed that the trajectories across both navigation models and all parameter options are no longer similar. After each trajectory chooses to dodge the first and stationary human towards the robot's right, the NOD robot may choose to continue to dodge the approaching second human further to the right or move towards the left, the latter being the behavior of all S-NOD paths. We no longer see a consistent pattern of behavior based on the  $U$  value when interacting with the second human. All S-NOD trajectories past the second human appear to collapse to the same path towards the robot's goal, whereas the NOD robot's path varies with both  $K_u$  and  $U$ . The slight difference between the S-NOD cases tracks from the theory presented in [26] and Chapter 4, i.e.,  $K_u$  does not require fine-tuning to drastically change the system's output behavior, so long as it is placed in the subcritical range with  $K_u > K_u^*$  the system will act with increased robustness and agility. However, for NOD with  $K_u = 5K_u^*$ , the NOD robot dodges the second human with a left turn for all  $U < 2u_0^*$  trajectories. Only when the NOD robot's opinion reaches  $U = 2u_0^*$  does it remain negative, causing the NOD robot to turn right to pass the second human before the opinion returns to zero. However, when  $K_u = 10K_u^*$ , the NOD robot dodges the second human with a left turn only for the  $U = 1.25u_0^*$  trajectory. Now with  $U > 1.25u_0^*$  the NOD robot's opinion remains negative, turning the robot right to pass the second human. In both of these instances, the opinion saturates and never returns to a nonzero indecision state ( $z \neq 0$ ).

It remains true that trajectories of the NOD social robot appear to provide more passing space, however

inefficiently, than the comparable trajectories of the S-NOD social robot. This is confirmed by the consistently lower minimum distances seen in the S-NOD cases within each group of experiments with the same  $U$  value. We reiterate that these results follow because the NOD robot is less agile than the S-NOD robot, and that neither algorithm explicitly constrains the robot to maintain distance from an approaching human. An additional mechanism for collision avoidance could be implemented to promote this behavior, thereby eliminating the only apparent advantage that a NOD robot has over an S-NOD robot.

As noted, the variation of  $K_u$  has a significant effect on the trajectories of the NOD robot. Its influence on the evolution of opinion for each robot can also be described. As the larger  $K_u$  value in the NOD social robot saturates to  $z = -1$  and remains after the robot has passed the human, only when  $U$  is sufficiently low can this opinion be reversed to induce a left turn. Even still, the opinion after the robot passes the second human ( $u_{0\ell} = 0$ ) remains saturated in an opinionated state ( $z \neq 0$ ). This leaves it susceptible once again to further inefficient turning behavior if new humans were introduced to the scene ( $u_{0\ell} \neq 0$ , reintroducing the opinionated turning by  $\beta$  in the heading update). The larger  $K_u$  value in the S-NOD social robot still causes broader and deeper spikes to the  $z = -1$  value, increasing in frequency as the human approaches. Still, due to its return to the sensitive  $z = 0$  state after each spike, the robot can quickly switch to a spiking  $z = 1$  value to turn the robot left around the second human.

We note the even more pronounced and consistently shorter navigation times of the S-NOD robot compared to its comparable NOD robot. Still, this efficiency comes at the expense of lower minimum distances maintained between passing humans and higher invasion distances for both humans in all cases where  $U > 1.5u_0^*$ .

Finally, we discuss the amount of control used within each experiment. There are still noticeable differences in the scale of control shown between the NOD and S-NOD robots, with the latter proving to require much less control signal to the robot. We see that a doubling of the  $K_u$  value from  $5K_u^*$  to  $10K_u^*$  results in an increase of total control within the range of 21% to 34% for the NOD robot, whereas this range is only 10.52% to 13.84% for the S-NOD robot.

We use the results of this section to provide insight into the consistent and agile performance of S-NOD and its benefits for social navigation compared to the more parameter-sensitive NOD model. This analysis provides an understanding of the parameter sensitivity and design trade-offs of each model when embodied on a social robot. While NOD does not perform poorly for social navigation, its weaknesses begin to show against S-NOD when the navigation scenario is more complex or cluttered. We further develop these results in the next section with a proof-of-concept demonstration of NOD and S-NOD robots navigating through five interactions.

## 5.4 Simulated multi-human navigation: NOD vs. S-NOD performance

We conclude the presentation of the advantages of the S-NOD algorithm over the NOD algorithm for social robot navigation with a representative simulation that models a complex navigation scenario involving a single robot and multiple humans. Based on the previous results in this chapter, we anticipate that a NOD robot will navigate inefficiently in complex scenarios, as it can become “stuck” in a fixed opinion and fail to remain responsive to the sequential decision-making demands along its path. A S-NOD robot, however, is repeatedly driven to an ultra-sensitive opinion state and therefore highly adaptable to new input from the environment. The only flaw observed in an S-NOD robot thus far has been its spatial intrusiveness, which is a consequence of its increased efficiency. We conclude with a short discussion on how this issue can be mitigated by integrating an additional collision avoidance algorithm to ensure safe navigation.

We present a simulated scenario with five moving humans placed along a robot’s path toward its goal. The robot is initially placed at  $(0, 0)$ m directly facing its goal at  $(0, 20)$ m. Note that this creates a significantly larger navigation area than previously seen in simulations and experiments in this thesis. The robot is programmed to move at a constant speed of  $V = 0.6$ m/s. For all subsequent simulations, the parameters within NOD equation (3.3) and S-NOD equation (5.2) held constant are as follows:  $\tau_z = 0.02$ ,  $d = 1$ ,  $a_{11} = a = 2$ ,  $\gamma = 3$ ,  $n = 2$ ,  $D_{crit} = 4$ m,  $\tau_{u_s} = 0.2$ ,  $K_{u_s} = 10$ ,  $k_\theta = 2$  and  $\beta = \pi/4$ . The robot was kept unbiased ( $b = 0$ ) for maximum flexibility in passing behavior.

There are five humans placed sporadically throughout the navigation area, each exhibiting different motions. Humans are tagged as Human  $\ell$  with  $\ell \in N_h = 1, \dots, 5$  according to their initial proximity to the robot. We detail these motions below:

- Human 1 begins at  $(-2, 2)$ m with heading  $\theta_{h_1} = 45^\circ$  and moves in that direction at a constant speed of  $V_{h_1} = 0.6$ m/s.
- Human 2 begins at  $(-0.5, 8)$ m with heading  $\theta_{h_2} = -55^\circ$  and moves in that direction at a constant speed of  $V_{h_2} = 0.7$ m/s.
- Human 3 begins at  $(0, 9)$ m with heading  $\theta_{h_3} = -120^\circ$  and moves in that direction with a sinusoidal shape at a constant speed of  $V_{h_3} = 0.7$ m/s.
- Human 4 begins at  $(-2, 16)$ m with heading  $\theta_{h_4} = -25^\circ$  and moves with a constant speed of  $V_{h_4} = 0.15$ m/s in that direction before making an abrupt  $50^\circ$  turn.

Model	NT [s]	PL [m]	CD [ $m^2$ ]	$\min(D_3)$	$\min(D_4)$	$ID_3$ [m]	$ID_4$ [m]	$\ z\ $
NOD	46.17	27.7	32,005.0	0.89	4.54	0.58	0	99.75
S-NOD	33.76	20.26	4,181.5	0.72	0.66	0.52	0.48	20.92

Table 5.17: Performance metrics for the NOD and S-NOD robots along the trajectories illustrated in Fig. 5.21. The minimum distances  $\min(D_{N_h})$  and their related invasion distances are only shared if the invasion distance is nonzero. Shorthand: NT = Navigation Time, PL = Path Length, CD = Centerline Deviation,  $ID_\ell$  = Invasion Distance of Human  $\ell$ .

- Human 5 begins at (3, 18)m with heading  $\theta_{h_5} = -140^\circ$  and moves with a constant speed of  $V_{h_5} = 0.4\text{m/s}$  in that direction before making an abrupt  $45^\circ$  turn.

We present this case where neither robot fails in their navigation, but moves with drastically different behavior along the path towards the goal. Recall that if the robot were undisturbed in the environment, it would move in a straight path to the goal along the  $x = 0\text{m}$  axis. We observe that within this cluttered environment with various moving behaviors, neither robot follows a straight path, but the S-NOD robot remains nearby at least. This efficiency is reflected in its results, summarized in Table 5.17.

We first qualitatively compare the trajectories taken by the robots.

The NOD robot begins its navigation with a robust positive opinion to turn towards the left, away from Human 1 as it crosses the robot's path. Upon interacting with Humans 2 and 3, the NOD robot changes to a negative opinion and turns towards the right. From this point on, the robot's sensitivity and its opinion do not fall to a nonzero value again. The robot maintains a large, arching right turn as it is captured by its sensitivity responding to Human 1, who is now moving in front of the robot's path. Only once the human moves sufficiently out of the robot's path (after nearly 25 seconds) does the robot begin to move towards its goal. This long, arching path to the goal results in a high navigation time, path length, and significant centerline deviation from the path to the goal. It does, however, reduce the opportunity to interact with and navigate near the other humans in the scene. This results in high minimum distance values between all humans (except for Human 1) and only one instance of personal space invasion, which is observed in Human 1. However, this behavior could only be helpful if there was no time sensitivity or control usage issue for the robot, as this behavior comes at the cost of a large amount of total control  $\|z\| = 99.75$ . The fact that the attention and opinion *never* relax to a zero value ( $z \neq 0$ ) leaves the NOD robot vulnerable to further inefficient, odd navigation behavior if another human comes into its view (making  $u_{0\ell} \neq 0$ , reintroducing the opinionated turning by  $\beta$  in the heading update of (3.3c)). The HRI literature, such as in [14], also suggests that this behavior would be uncomfortable for the interacting humans.

The S-NOD robot begins its navigation with a positive spiking opinion to turn towards the left, away from Human 1 as it crosses the robot's path. Upon interacting with Humans 2 and 3, the S-NOD robot

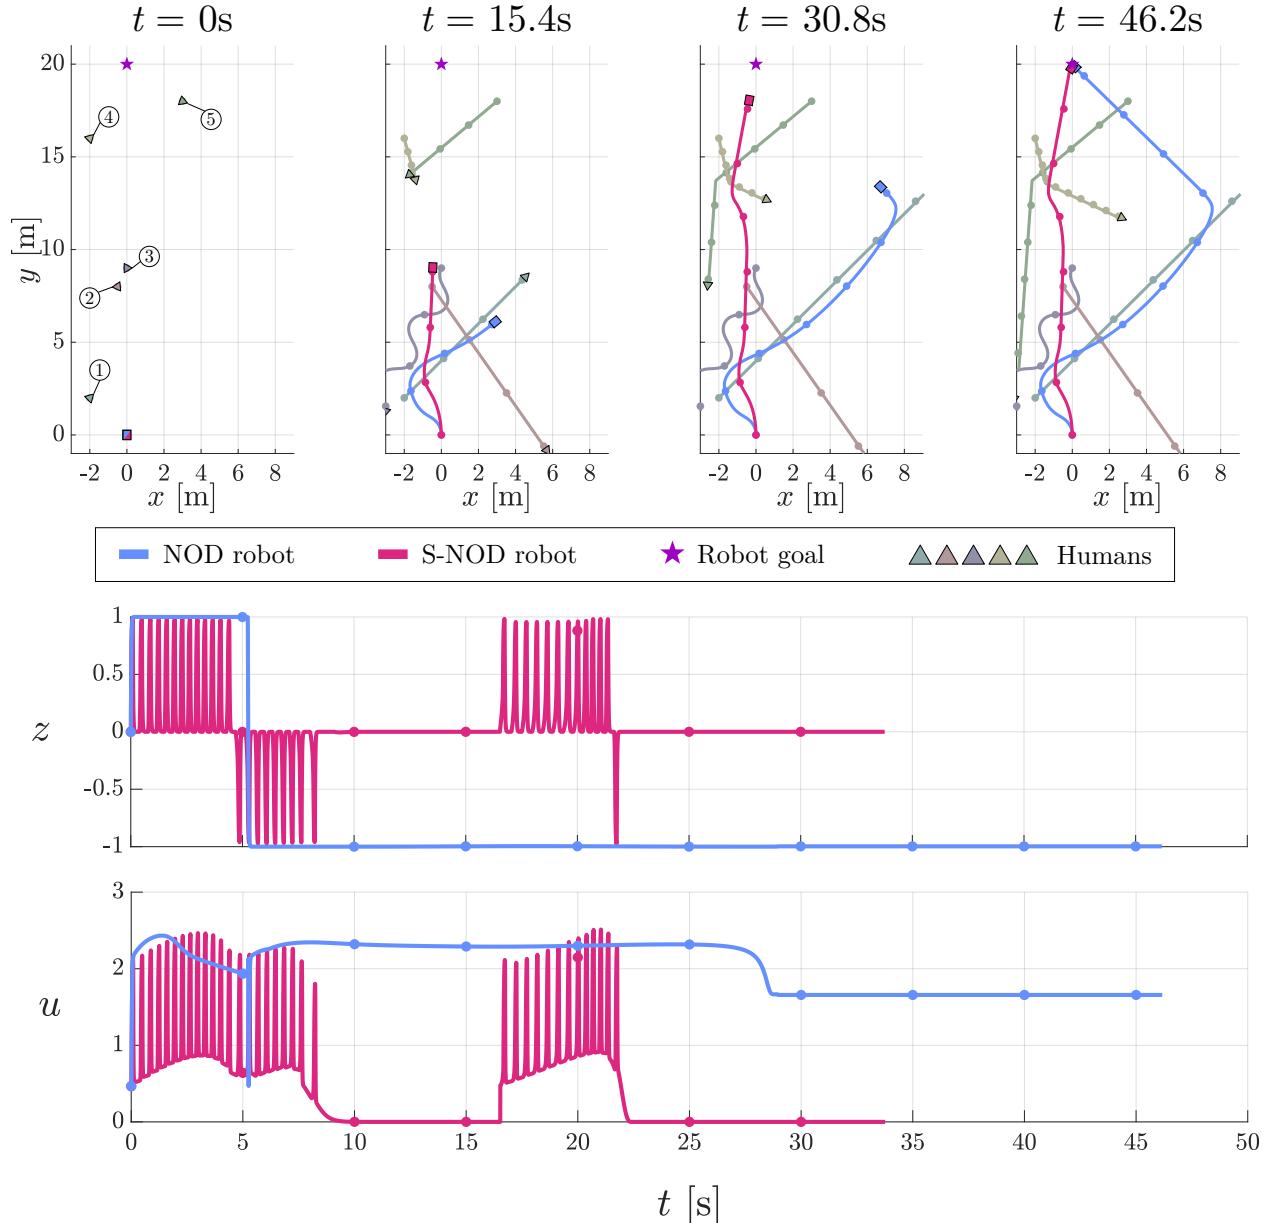


Figure 5.21: Superimposed simulations of NOD and S-NOD social robots navigating around five humans. (Top) Trajectory snapshots along the duration of each robot's navigation. On the left, the humans are numbered according to their initial proximity to the robot's initial location. (Middle) Opinion  $z$  over time  $t$  of both robots. (Bottom) Sensitivity  $u$  over time  $t$  of both robots. Dots along the trajectories and  $z$  and  $u$  traces mark even temporal steps.

changes to a negative opinion (as seen in the evolution, this is done more swiftly than that by the NOD robot) to turn towards the right. These turns are much smaller than those seen by the NOD robot, so the S-NOD robot now remains on a straight path toward its goal near the initially direct path on the  $x = 0$ m axis. The robot then begins to interact with Humans 4 and 5, who appear on the robot's left. Human 4's heading indicates that it is moving toward the right, while Human 5's heading indicates that it is moving toward the left. The robot forms a spiking positive opinion and turns left, away from the nearby Human 4. As it completes this turn, it grows in proximity to both Humans 4 and 5, which is reflected in the frequency increase of the opinion spikes. It experiences a single negative spike as Human 5 becomes the only interacting partner in its path, but once through the interaction, the robot travels directly toward the goal. This more direct path to the goal results in lower navigation time, path length, and centerline deviation from the path to the goal compared to that of the NOD robot. This does, however, allow for more nearby interactions with humans, as reflected in the lower minimum distance values between all humans, with two instances of personal space invasion observed in Humans 1 and 4. The control signal input is significantly lower than previously observed, at  $\|z\| = 20.95$ . The attention and opinion remain relaxed at zero value when not interacting with a robot, making the S-NOD robot more flexible and capable of responsive navigation behavior if another human comes into its view.

These behaviors were expected, given what we understand about S-NOD compared to NOD in theory, and we can interpret this in its embodied behavior. While NOD may not fail, a NOD robot will move with increasing inefficiency as sequential decision-making becomes necessary while navigating. Sequential decision-making is a crucial capability of a navigation algorithm, particularly as more humans are introduced into navigation scenarios. Further, the sloping, wide turns made by NOD to avoid humans may maintain safety, but sacrifice natural movement. As discussed in [14], humans prefer when robots exhibit cooperative behavior, so this kind of behavior by the NOD robot would be viewed as an uncomfortable navigation partner. In contrast, the S-NOD robot exhibits more understandable, goal-driven behavior. This does, however, increase the possibility of causing a collision when not making sloping turns. We believe that S-NOD is a low-level, computationally efficient method for dictating cooperative, natural, and deadlock-free navigation around humans. Its weaknesses can be addressed by incorporating a supplemental collision avoidance algorithm into the robot to maintain and ensure safety during navigation.

# Chapter 6

## Conclusions

### 6.1 Final remarks

In this dissertation, we presented the development, design, analysis, and validation of two forms of nonlinear opinion dynamics models (NOD and S-NOD) that control the navigation behavior of robots interacting with humans. Each model was defined with an emphasis on understanding the system's output behavior when tuning parameters. Both leverage the pitchfork bifurcation exhibited by the system dynamics to ensure the instability of the decision when the system is sufficiently sensitive.

The application of a Nonlinear Opinion Dynamics (NOD) model to social robot navigation was motivated by the biological inspiration behind the decision-making in the general Bizyaeva et al. model [1], an attribute we sought to translate to inherently collaborative, natural navigation behavior by the robot. In collaboration with María Santos, Shinkyu Park, and Naomi Ehrich Leonard, we adapted the NOD model to the social navigation setting and formally investigated the behavior of a physical robot navigating real spaces with humans. While the application was successful, we highlighted key elements of the control that made it so robust that it became resilient to the flexible opinion changes so crucial to its use.

To better serve the sequential nature of decision-making in a social robot navigating in real-time around multiple human movers, and to address the concern that opinions can become “stuck,” we adapted the NOD model to be spiking and event-based. In collaboration with Alessio Franci, Naomi Ehrich Leonard, and Ian Xul Belaustegui, we modified the NOD model by incorporating an additional slow-negative feedback mechanism inspired by the return to sensitivity exhibited in neuronal excitability. This led to the development of the Spiking Nonlinear Opinion Dynamics (S-NOD) model. The S-NOD model retains the fast and flexible decision-making of NOD and adds adaptability to the system by cycling solutions through the sensitive

bifurcation point, creating a more agile model. The contributions within each chapter of the dissertation are discussed in further detail below.

In Chapter 3, we presented a new and proactive approach to social robot navigation, leveraging an adaptation of the general Nonlinear Opinion Dynamics (NOD) model by Bizyaeva et al. [1]. This setting exploits the pitchfork bifurcation within the NOD system to guarantee deadlock-breaking passing behavior when the robot faces a navigation scenario with no obvious best passing direction. The robot can rapidly and reliably pass approaching human movers based on only their geometry, notably without requiring a complex model of human behavior. We analytically guaranteed deadlock-free navigation for the robot, leveraging linearization and bifurcation theory to find the system’s stable solutions. We found that the basal sensitivity  $u_0$  serves as a bifurcation parameter of the system, with its critical value  $u_0^* = d/a$  determined by the design of other parameters within the model. Once the robot’s basal sensitivity surpasses this critical value, indecision (in the form of a zero opinion state,  $z = 0$ ) becomes unstable. Therefore, we find that the robot’s parameters can be set to be sufficiently sensitive to its environment, allowing it to produce proactive navigation. This deadlock-breaking behavior is illustrated in simulations and further verified with physical experiments.

We defined the working environment for our specialized NOD model for social robot navigation and carefully described the adaptations made to the general form of Bizyaeva et al. [1]. We highlighted the relevant changes and improvements made from our initial publication of the NOD model [25] and then presented entirely new experimental results in Chapter 3.4. We systematically tested the NOD algorithm in experiments to serve many purposes. A single-robot, single-human experiment verifies the deadlock-breaking behavior of the NOD algorithm, navigating an unbiased robot out of the way of a human approaching it directly. These experiments also demonstrate that when interacting with a passing human, a robot with a biased passing preference can reliably overcome its predisposition and collaboratively pass the human mover. These experiments and their results further illustrate how the design parameter  $\beta$  in the robot navigation algorithm can tune the robot’s behavior and reveal the trade-off between navigation efficiency and spatial intrusiveness in the passing problem. Experiments with a single robot and two passing humans verify the flexibility of the NOD model, as the robot reliably modifies its trajectory according to both humans and successfully navigates through complex passing scenarios. The robot’s movements were collaborative with the humans it interacted with and appeared to have an inherent amount of safety in its navigation provided by the NOD control.

Based on the previous findings, we identified the need for the social robot’s opinion to return to the vicinity of its ultrasensitive bifurcation point, thereby adding agility without compromising the robustness of NOD. In Chapter 4, we design and present the Spiking Nonlinear Opinion Dynamics (S-NOD) model to satisfy this

need. We incorporated excitability into the NOD model by introducing an additional slow-negative variable that returns the system to an ultrasensitive state. We then specialized the S-NOD model for the social robot navigation setting. In this chapter, we demonstrated that S-NOD offers a first-of-its-kind *two-dimensional* framework for excitable (spiking) dynamics, enabling agile decision-making over two options. We explained how NOD can become too robust; however, the self-regulation of S-NOD recovers system flexibility and enables adaptive, event-based decision-making. We showed and analyzed the existence of limit cycles for certain parameter regimes in S-NOD. We presented a multi-agent S-NOD model and highlighted the potential agent (anti)synchronization, which is dependent on communication network parameters. We concluded this presentation of the formulation and analysis of S-NOD with an illustration of its agility in a simulated social robot navigation application.

In Chapter 5, we expanded on the application of S-NOD for social robot navigation to physical experiments with human movers, and simulations with a larger scale of both navigation space and number of humans. We presented the specialized S-NOD model for social robot navigation. The physical experiments of a robot controlled with S-NOD were constructed to mirror those within the NOD section. This allowed us to 1) systematically validate the deadlock-breaking behavior of S-NOD just as we had done with NOD, 2) validate the flexibility of S-NOD just as we had done for NOD, 3) provide an opportunity to directly compare the behaviors exhibited by the algorithms on a social robot, and 4) uncover the approximate tuning required to reproduce NOD robot behavior with an S-NOD robot. Experimental results show that S-NOD maintained the deadlock-free guarantee in navigation provided by NOD, and navigation performance was improved with increased efficiency and flexibility through event-based decision-making. Through a comparison of experiments, we explored the benefits and weaknesses of the S-NOD robot’s navigation behavior in relation to that of its NOD counterpart. We noted that while efficiency improves with S-NOD, it could sacrifice a comfortable level of passing distance, whereas a NOD robot would not. We noted, however, that the comfortable passing distance could likely be easily controlled with the addition of a well-studied mechanism from the literature for local obstacle avoidance. Through the example of a larger-scale five-human simulation, we demonstrated that while NOD provides a comfortable passing distance to humans, it does so with unnatural and puzzling behavior compared to the efficient and legible path of an S-NOD robot.

## 6.2 Future work

We believe there are numerous exciting opportunities for the extension and further development of the NOD and S-NOD theories, with various suggested improvements for social robot navigation.

There are promising developments in further theoretical analysis and understanding of S-NOD models

already. In a recent work [62] by my collaborators Ian Xul Belaustegui, Alessio Franci, and Naomi Ehrich Leonard, the S-NOD model is used as a controller and expands on the analysis of its behavior dependent on parameters. They identify the necessary conditions that cause the onset of periodic opinion spiking, analyzing the tunability of the spiking threshold of the opinion  $z$  as a function of the basal sensitivity  $u_0$  and the tunability of the spiking frequency with the magnitude of the input. As demonstrated in this dissertation, these results can be extended to an S-NOD controller for multiple agents. Additionally, the S-NOD controller can be further extended to accommodate more than two options. This presents the opportunity for spiking decision-making that occurs in as many directions as there are options in a system. There is much exciting future work to be done in studying S-NOD for multiple agents and multiple options.

There are various avenues for improvement to the NOD and S-NOD models applied to a social robot. We believe these models are powerful and demonstrably effective tools for decision-making in navigation settings, but may be best utilized as a low-computation controller to complement other navigation methods. These algorithms can reduce or eliminate the need for predictive or reinforcement learning models to maintain a long-term history of previous navigation interactions and/or train on complex models of human behavior. We hold that NOD and S-NOD best serve as decision-making algorithms with the added benefit of creating collaborative, natural movement when their control is used directly. However, these models are not themselves algorithms with collision avoidance. Robot motion can therefore be enhanced by allowing a different control to maintain a safe distance from a human while NOD or S-NOD maintains deadlock-breaking in further navigation.

We also believe that applying reinforcement learning to the NOD and S-NOD parameters may be a valuable avenue to better tailor the robot's behavior to its navigation environment. As observed in the two sets of single-robot, single-human experiments of S-NOD in Chapter 5.2.2.3, the parameters within the model can significantly impact the robot's turning behavior. Depending on the environment (e.g., a hallway, an open lobby, or a bustling warehouse), the robot may navigate more effectively with different navigation behaviors (e.g., maintaining a greater or lesser distance from nearby humans). Learning the parameter regimes that best fit these navigation environments can enhance the performance of social robots, thereby better equipping them for cooperative motion in each environment. Learning to adapt their parameters in response to a changing environment can keep these robots agile and adaptable. Already, Hu et al. [28] demonstrated the successful use of NOD alongside other methodologies to achieve safety-critical performance in autonomous racing. This offers a strong motivation to pursue similar directions of combining NOD to complement other methods in future work.

We believe that these nonlinear opinion models can be applied to various other tasks within a social robot, beyond navigation. The robot's behavior can be improved by utilizing the NOD or S-NOD mechanism to

control its speed, modulating it to move either faster or slower during a human interaction, thereby providing space, overtaking a slowed human, or enhancing safe navigation in a cluttered environment.

Finally, we note that the NOD and S-NOD models, specialized to social robots within this dissertation, concern scenes in which the number of potentially passing humans is known. In a non-laboratory environment, the robot will need to have significantly improved perceptual capabilities to better identify and observe humans within its space. There have been strides made in this setting by Amorim et al. in [63], which presents a spatially invariant form of perceptual decision-making. Its application to a robot within a crowd shows success in sensing and identifying through-paths between humans. In the reverse sense for our social robot, this can be used to determine the existence of humans, allowing their location and orientation to be assessed and used to form a collaborative opinion with NOD or S-NOD.

Opportunities abound for the formulations and applications of NOD and S-NOD. It has been an honor to aid in their conceptualization, development, and implementation thus far. I look forward to witnessing all future evolutions, refinements, improvements, and extensions to these powerful tools.

## **Part II**

# **Published work**

# Chapter 7

## Proactive opinion-driven robot navigation around human movers

CHARLOTTE CATHCART, MARÍA SANTOS, SHINKYU PARK, AND NAOMI EHRICH LEONARD

### Statement of contributions

The material presented in this chapter also appears in [25] and was presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) in Detroit, MI in 2023. I, Charlotte Cathcart, am the lead author of the manuscript. María Santos and Shinkyu Park are co-authors of the manuscript and acted in supervisory roles for the project. Naomi Ehrich Leonard is the principal investigator of this project. Naomi Ehrich Leonard and Shinkyu Park initially conceived of the NOD model for application in a robotic navigation system. All four authors participated in discussions on the work presented in this manuscript. I, Charlotte Cathcart, developed the final application and embodiment of NOD in a robotic navigation system, drawing from previous work by Shinkyu Park and Andrew Witmer. I, Charlotte Cathcart, led the writing of the manuscript, with significant writing and editing contributions by Naomi Ehrich Leonard and Shinkyu Park in Section 7.1. All four authors contributed to the drafting of the initial manuscript and assisted in the editing of the final manuscript. I, Charlotte Cathcart, ran all simulations, conducted all experiments, and created all figures within the manuscript. María Santos provided guidance on setting up hardware and experiments. All four authors, with significant guidance from María Santos and Shinkyu Park, assisted in formalizing the experimental approach to showcase NOD in a robotic navigation system. Naomi Leonard supervised the acquisition of funding for this project through ONR grant N00014-19-1-2556 and the

generosity of Lydia and William Addy '82 at Princeton University. I, Charlotte Cathcart, led the acquisition of approval of human-robot experiments from the Institutional Review Board (IRB) at Princeton University with guidance and editing from Naomi Ehrich Leonard.

## Abstract

We propose, analyze, and experimentally verify a new proactive approach for robot social navigation driven by the robot’s “opinion” for which way and by how much to pass human movers crossing its path. The robot forms an opinion over time according to nonlinear dynamics that depend on the robot’s observations of human movers and its level of attention to these social cues. For these dynamics, it is guaranteed that when the robot’s attention is greater than a critical value, deadlock in decision making is broken, and the robot rapidly forms a strong opinion, passing each human mover even if the robot has no bias nor evidence for which way to pass. We enable proactive rapid and reliable social navigation by having the robot grow its attention across the critical value when a human mover approaches. With human-robot experiments we demonstrate the flexibility of our approach and validate our analytical results on deadlock-breaking. We also show that a single design parameter can tune the trade-off between efficiency and reliability in human-robot passing. The new approach has the additional advantage that it does not rely on a predictive model of human behavior.

### 7.1 Introduction

Autonomous mobile robots are increasingly being used for tasks in settings such as warehouses and open public spaces where they will encounter human movers. In order to accomplish their tasks in these settings, the robots need to reliably and gracefully navigate around human movers. In this paper, we propose, analyze, and experimentally verify a new approach for the social navigation of a mobile robot. Fig. 7.1 shows experimental results of a mobile robot navigating around two human movers using the new approach.

We build on the nonlinear opinion dynamics model presented in [1] and propose an approach that allows a robot to rapidly form an *opinion* that represents the strength of its preference for which direction—left or right—it will use to pass each human mover crossing its path. This opinion, in turn, drives the robot’s motion, modifying its nominal path to reliably pass the human. A key to the opinion dynamics is that when the robot’s *attention* to social cues grows above a critical value, the neutral opinion to stay the course is destabilized and the robot rapidly forms a strong and stable opinion for moving in one of the two passing directions. Our approach is therefore to design dynamics for the robot’s attention that drive it above this

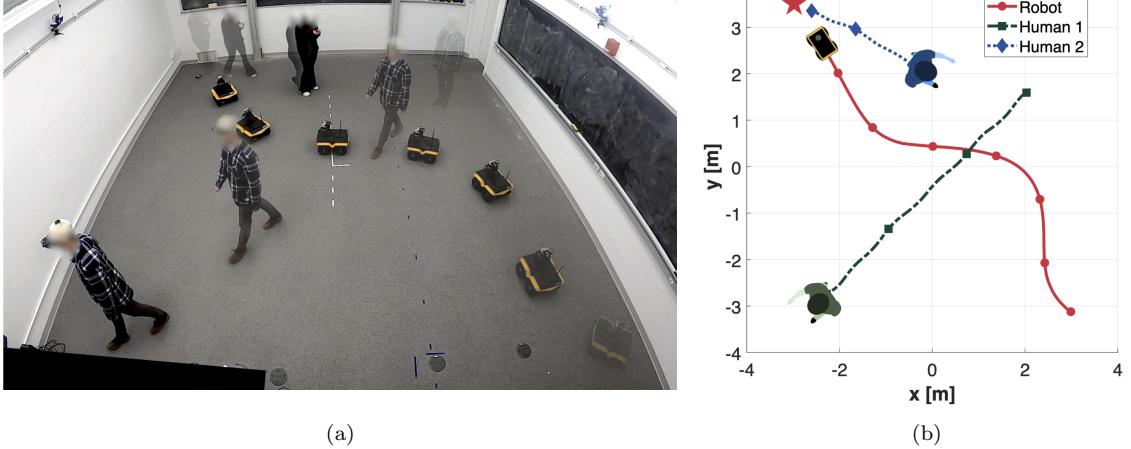


Figure 7.1: A robot using opinion-driven navigation to pass two humans. (a) Time-lapse of the experimental trial. (b) The full trajectories of the robot (red line) and two humans (blue and green lines) with temporal markers.

critical value when the robot senses a human mover approaching its path. The active control of attention yields a rapid and reliable passing motion in response to an approaching human mover; this renders our approach “proactive” rather than merely “reactive.”

Once the robot passes a human, its opinion with respect to that human is no longer relevant; the opinion quickly returns to its neutral value, allowing the robot to continue towards its destination. Likewise, the robot’s attention also goes to zero, making the robot ready for new potential conflicts with other movers. Figs. 7.1 and 7.2 provide experimental results of the robot navigating different encounters when traveling to a goal destination that is diagonally across an open space with two humans moving and pausing in a variety of scenarios.

Opinion dynamics are used to enable decision making in multi-agent systems in a range of tasks [17, 23, 64]. In the nonlinear opinion dynamics of [1], an agent’s opinion is influenced by the opinions of others when its attention exceeds a critical level. At this point the agents are guaranteed to form strong opinions (e.g., to agree on or coordinate among options), hence avoiding indecision, i.e., deadlock in their decision making. In the robot social navigation problem, we leverage the deadlock breaking guarantees of the coupled attention-opinion dynamics to ensure that, when necessary to avoid an approaching human mover, the robot will rapidly select and move in one of the two passing directions even if there is no indication from the human or the environment that one direction is better than the other, or if the robot’s bias for one direction or the other, if it has one, conflicts with the human’s chosen passing direction.

Of relevance to our work is the literature on robot social navigation (see recent survey articles [2–5, 65] and references therein), where a common theme is in investigating the design of navigation algorithms for

autonomous robots to safely and comfortably interact with the humans they encounter. Earlier work [6] in modeling human navigation behavior proposes a model based on the observation that the motion of pedestrians is subject to *social forces*. More recent works [7, 66] incorporate *social cues* into the social force model and the improved models are used to design robot navigation algorithms. The work of [67] proposes a constrained optimization approach to design a navigation algorithm that penalizes the robot when its behavior violates conventions observed in the human’s navigation. In [5], a reactive control policy is used to follow and maintain the passing sides observed by passing humans through social momentum. References such as [10, 68, 69] explain learning-based approaches that leverage the recent advancement in deep reinforcement learning to train mobile robots through multiple trial-and-error processes to safely navigate in human-populated areas.

Another important line of research in the social navigation literature is data-driven learning approaches that infer human navigation models from their demonstration data, and use the models to predict human motions and to design robot motion planners. The work of [70] leverages Bayesian learning to construct a motion model and personality characteristics of pedestrians, and use predicted pedestrian trajectories from the model for socially-aware robot navigation. Inverse Reinforcement Learning (IRL)-based approaches, for instance [9, 11, 69], take human demonstration data to estimate a utility function used in human navigation tasks, and use it to generate robot trajectories that imitate the demonstrated human motions. In particular, a recent relevant work [71] studies the effect of human-robot communication in social navigation and proposes an IRL-based robot planning framework to generate communication actions that maximize the robot’s transparency and efficiency.

Our work is distinct in that 1) it is proactive rather than reactive, 2) it does not require constructing a predictive model of human navigation as in IRL-based approaches, rather it only needs the robot to observe the position and moving direction of the human, and 3) our robot navigation model is analytically tractable so that we can establish a guarantee on deadlock-free decision making in the robot-human navigation. This contrasts with the reinforcement learning approaches, which are in general difficult to analyze, and existing reactive approaches, such as social force models, which do not provide the same deadlock-free guarantee.

In Section 7.2, we introduce the nonlinear opinion dynamics and propose a new model for robot navigation in a human-robot navigation setting. In Section 7.3, using tools from nonlinear dynamical systems theory, we discuss how the model ensures rapid deadlock-free robot navigation. To demonstrate and test the flexibility of our approach, we carry out experiments with two human participants and a mobile robot in a range of scenarios, which we report on in Section 7.4.1. We examine and validate the effectiveness of rapid deadlock-free navigation with further experiments in Section 7.4.2. We conclude with a discussion in Section 7.5.

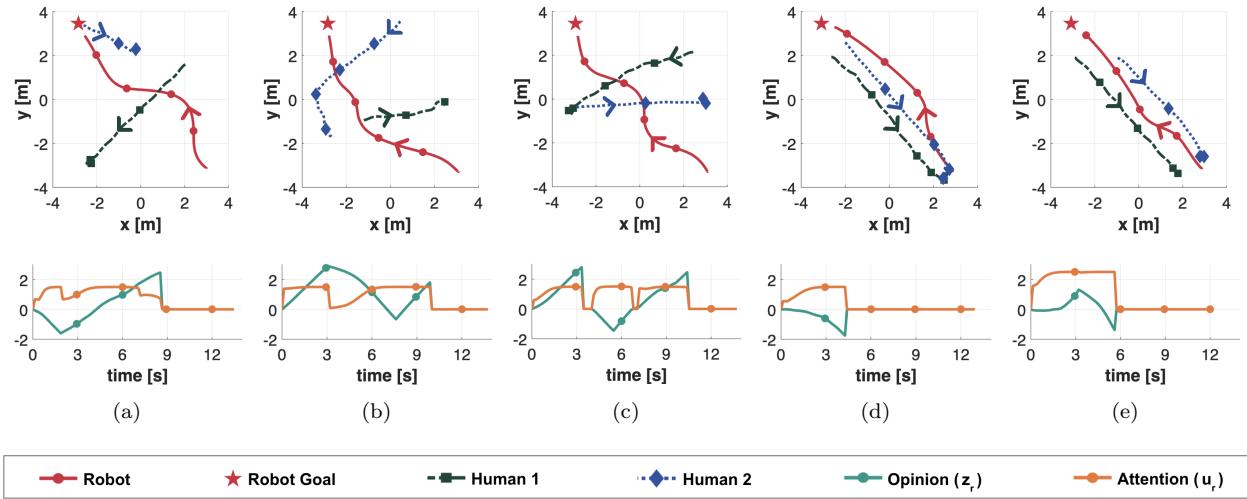


Figure 7.2: Multiple experimental trials with two humans and a robot using the new approach. The top row shows the complete trajectories of the robot (red line) and humans (green and blue lines) over the course of a trial as the robot moves toward its goal (red star). Each trajectory is marked with an arrow indicating the mover's direction. The bottom row shows the robot's opinion  $z_r$  (teal line) and attention  $u_r$  (orange line) over the course of the trial above it. Temporal markers (dots) are shown along spatial trajectories, opinion, and attention. See Section 7.4.1 for parameters used.

## 7.2 Nonlinear Opinion Dynamics in social navigation

We study a robot navigation problem where a robot approaches and passes human movers while traveling to its destination (see examples in Figs. 7.1 and 7.2). In this context, we want to enable the robot to repeatedly overcome human movers in a rapid and reliable fashion. We are also interested in tackling challenging scenarios such as the human-corridor passing problem [72–74] that may result in deadlock if, for example, both the robot and the human have conflicting passing biases. In these situations, a key objective is to ensure that the robot moves reliably around the human regardless of the human's awareness of the robot. It is also desirable that the robot moves efficiently around the human. However, reliability and efficiency are in tension: giving the human a lot of space may create reliably successful but inefficient passing whereas giving the human only a little space is efficient but creates less reliably successful passing.

To address these competing objectives, we propose a new dynamic model for robot navigation based on the nonlinear opinion dynamics of [1]. We review these dynamics in Section 7.2.1. We specialize the dynamics to proactive opinion-driven robotic navigation in Section 7.2.2 and show how a single design parameter can be used to control the reliability-efficiency trade-off. In Section 7.3, we provide analysis that shows how deadlock breaking is guaranteed.

### 7.2.1 Nonlinear Opinion Dynamics model

Consider a system of  $N_a$  agents forming opinions about two options. Let  $z_i \in \mathbb{R}$  be the opinion of agent  $i$ , which represents the strength of its preference for option 1 if  $z_i > 0$  and for option 2 if  $z_i < 0$ . It is indifferent, i.e., neutral, if  $z_i = 0$ . Strength of preference is  $|z_i|$ . The nonlinear opinion dynamics model, described below, explains how each agent  $i$  updates its opinion  $z_i$  continuously over time in response to its own opinion, the opinions of others  $z_k$ , and any internal bias or external stimulus  $b_i$ . Letting  $\dot{z}_i = dz_i/dt$ ,

$$\dot{z}_i = -d_i z_i + u_i \tanh \left( \alpha_i z_i + \gamma_i \sum_{k=1}^{N_a} a_{ik} z_k + b_i \right). \quad (7.1)$$

The opinion  $z_i$  can be interpreted as the discounted accumulation of *social influence* weighted by the parameter  $u_i \geq 0$ . The social influence is defined as the hyperbolic tangent function of the weighted sum of the opinion  $z_k$  of every agent  $k$  observed by agent  $i$  and a bias/stimulus  $b_i$ . The *resistance parameter*  $d_i > 0$  defines the rate of exponential discount in the accumulation of the social influence. The *attention*  $u_i \geq 0$  is a tuning variable, which can be adjusted to reflect the agent's (changing) effort to pay attention to the social influence. The parameter  $a_{ik} = 1$  if agent  $i$  can observe agent  $k$ ; otherwise,  $a_{ik} = 0$ . The parameters  $\alpha_i > 0$  and  $\gamma_i \in \mathbb{R}$  are weights defining how much influence  $z_i$  and  $z_k$ , respectively, have on agent  $i$ 's opinion update. If  $b_i > 0$  (resp.  $b_i < 0$ ), the bias is for option 1 (resp. for option 2). In case of no bias, we set  $b_i = 0$ .

### 7.2.2 Dynamic model for opinion-driven robot navigation

Building on (7.1), we propose a robot navigation model that forms an opinion to drive the robot's motor control in an uncrowded and uncluttered environment with human movers. We assume that the robot moves at a constant speed  $V_r$ , but can regulate its angular velocity. We represent the robot's position and heading angle as  $\mathbf{x}_r = (x_r, y_r)$  and  $\theta_r$ , respectively. For each human  $j$  that the robot can detect, we denote their speed  $V_{h_j}$ , position  $\mathbf{x}_{h_j} = (x_{h_j}, y_{h_j})$ , and heading angle  $\theta_{h_j}$ . Let  $\eta_{r_j}$  be the heading of the robot relative to the line between the robot and the human  $j$ . Let  $\eta_{h_j}$  be the heading of human  $j$  relative to the line between the robot and the human. See Fig. 7.3 for illustration of notation.

The robot focuses on the human mover  $j$  that minimizes  $\chi_j/\kappa_j$  where  $\chi_j = \|\mathbf{x}_r - \mathbf{x}_{h_j}\|$ ,  $\kappa_j = \cos \eta_{h_j}$ , and  $\eta_{h_j} \in (-\frac{\pi}{2}, \frac{\pi}{2})$ . This is the human who is most rapidly approaching the robot. We use  $\mathbf{x}_h(t)$ ,  $\eta_h(t)$ ,  $\chi(t)$ , and  $\kappa(t)$ , i.e., without index  $j$ , to refer to whichever human is the one most rapidly approaching the robot at time  $t$ .

We define  $z_r > 0$  (resp.  $z_r < 0$ ) as the robot's strength of preference for moving left (resp. right). When  $z_r = 0$ , the robot's opinion is neutral, i.e., it is indifferent to these options. Our approach does not

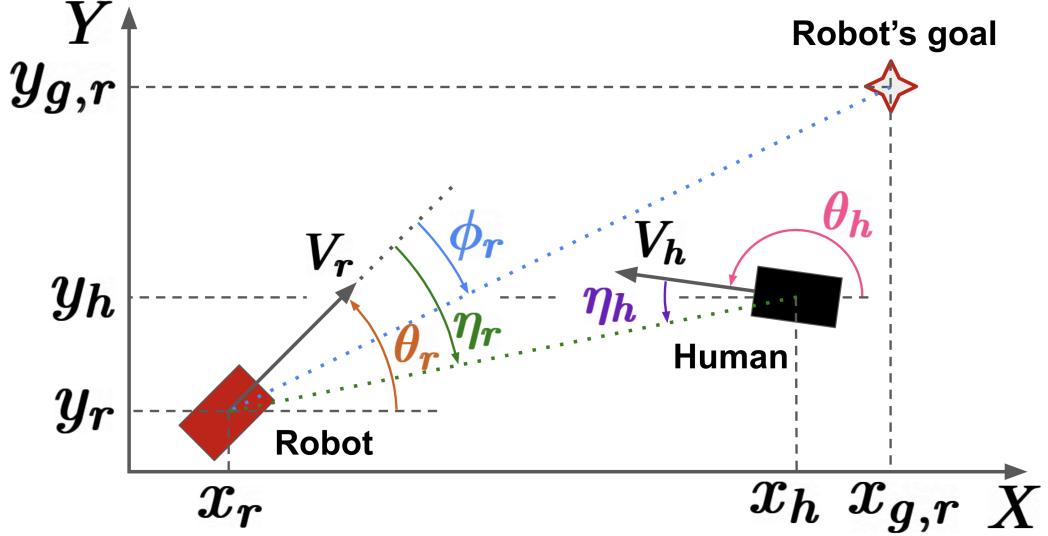


Figure 7.3: An illustration of notation for human-robot passing.

require any knowledge of a human model; however, we assume that the robot can measure  $\eta_h$  and use it as a proxy for the robot's perception of the human's opinion on direction as  $\hat{z}_h = \tan \eta_h$ <sup>1</sup>. This is unlike other approaches that require a longer-term prediction of human trajectories, such as [9, 11, 69, 71].

Our proactive opinion-driven robot navigation model specifies (a) how the robot's opinion  $z_r$  changes in response to its attention  $u_r$ , its current opinion, its estimate  $\hat{z}_h$  of the opinion of the focal human mover, and possibly a bias  $b_r$ ; (b) how the robot's attention  $u_r$  changes in response to  $\kappa$  and  $\chi$ ; and (c) how the robot's heading  $\theta_r$  changes as a function of its opinion  $z_r$  and the direction  $\phi_r$  to its goal:

$$\dot{z}_r = -d_r z_r + u_r \tanh(\alpha_r z_r + \gamma_r \hat{z}_h + b_r), \quad (7.2a)$$

$$\tau_u \dot{u}_r = -u_r + g(\kappa, \chi; R_r), \quad (7.2b)$$

$$\dot{\theta}_r = k_r \sin(\beta_r \tanh z_r + \phi_r), \quad (7.2c)$$

where  $d_r, \alpha_r, \gamma_r, \tau_u, R_r, k_r > 0$  and  $\beta_r \in (0, \frac{\pi}{2}]$  are design parameters. Note that (7.2a) is similar to (7.1) except the human's opinion  $z_h$  is replaced with the proxy  $\hat{z}_h = \tan \eta_h$ .

We design the attention dynamics (7.2b) so  $u_r$  grows quickly when a human mover gets close. Unless otherwise noted, we let  $\tau_u \rightarrow 0$  and define  $g$  using a Hill function to get

$$u_r = g(\kappa, \chi; R_r) = \underline{u} + (\bar{u} - \underline{u}) \left( \frac{(R_r \kappa)^n}{(R_r \kappa)^n + \chi^n} \right), \quad (7.3)$$

where  $0 \leq \underline{u} < \bar{u}$  and  $n > 0$ . The variable  $u_r$  increases from  $\underline{u}$  as the robot and human move closer towards

<sup>1</sup>We resort to [75–77] for the basis for estimating the human's navigation intent using their orientation.

collision, based on a critical distance parameter  $R_r > 0$ , and saturates at the value  $\bar{u}$ . This drives  $u_r$  above a critical value that destabilizes the neutral opinion  $z_r = 0$ , allowing the robot to rapidly form a strong opinion when a human mover approaches, and thus rapidly pass the human on one side or the other. In this sense our approach is proactive. See Section 7.3 for a rigorous analysis of the deadlock breaking.

To understand the role of design parameter  $\beta_r \in (0, \frac{\pi}{2}]$ , note that when  $z_r$  is sufficiently large so that  $\tanh z_r \approx 1$  (resp.  $-1$ ), (7.2c) steers the robot's heading angle an additional  $\beta_r$  radians in the counterclockwise (resp. clockwise) direction from the orientation to the goal location. Hence, we can tune  $\beta_r$  to prescribe how much the robot's heading angle should deviate from its direct path to its goal when it detects the human and forms a strong opinion on its passing direction. In this way the parameter  $\beta_r$  can be used to tune the reliability-efficiency trade-off as we show through the deadlock breaking human-robot experiments described in Section 7.4.2.

Our approach can be extended to incorporate path planning, e.g., to avoid driving the robot to a local minimum in the case of a cluttered environment. For example, this would be possible using a path planning approach such as the rapidly-exploring random tree (RRT) in place of (7.2c), with opinion  $z_r$  as an input. This would regulate not only the robot's angular velocity but also its moving speed.

### 7.3 Guarantee on deadlock-free navigation

A key contribution of our work is in guaranteeing deadlock-free navigation. We establish such a performance guarantee by analyzing the robot navigation model (7.2). In particular, we discuss how the robot can rapidly and reliably form a strong opinion to select one of the two options—move left ( $z_r > 0$ ) or right ( $z_r < 0$ )—and avoid colliding with a human, even when the human maintains a path straight for the robot and the robot has no bias ( $b_r = 0$ ) on which way to pass. To establish this, we use tools from nonlinear dynamical systems theory [43] to show that there is a deadlock-breaking *pitchfork bifurcation* in (7.2) when the robot's attention  $u_r$  reaches a critical level  $u_r^*$  (as it nears the human), corresponding to the destabilizing of the deadlock solution and the emergence of bi-stable solutions for moving left and for moving right.

We examine the challenging case in which the human does not react to the robot's movement. We validate our analysis through human-robot experiments in Section 7.4.2.

Suppose the robot is unbiased ( $b_r = 0$ ) and approaches a human who is walking straight towards it ( $\eta_h = 0$ ). In this setting, (7.2a) simplifies to

$$\dot{z}_r = -d_r z_r + u_r \tanh(\alpha_r z_r). \quad (7.4)$$

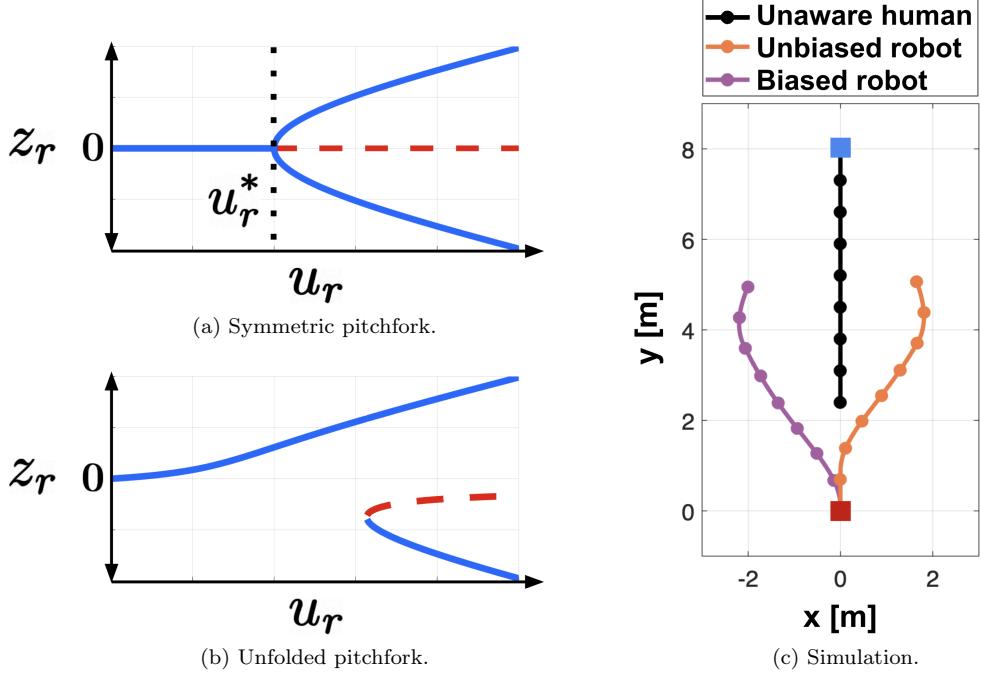


Figure 7.4: Analysis of deadlock breaking in the robot's opinion dynamics when the human moves straight towards the robot. (a) When the robot is unbiased ( $b_r = 0$ ), deadlock is broken as  $u_r$  increases above critical value  $u_r^*$ , where two stable (blue solid) symmetric opinionated solutions emerge and deadlock becomes unstable (red dashed). (b) When the robot is biased ( $b_r = 0.5$ ), the bifurcation “unfolds” where deadlock breaks but the likelihood of converging on one opinionated solution is greater than on the other. (c) Simulations of social navigation dynamics. Initial conditions for the robot and human indicated with red and blue boxes. Parameters of (7.2):  $d_r = \alpha_r = 0.1$ ,  $\gamma_r = 3$ ,  $\tau_u = 1$ ,  $g(\kappa, \chi; R_r) = \exp(\kappa(R_r - \chi))$  with  $R_r = 16$ ,  $k_r = 1$ , and  $\beta_r = \pi/4$ .

The neutral (deadlock) opinion  $z_r = 0$  is always an equilibrium solution of (7.4). However, we show that while for small values of attention  $u_r$  deadlock is a stable solution, for larger values of  $u_r$  it becomes unstable and two symmetric bi-stable solutions emerge corresponding to a strong opinion, one for going left and one for going right. This transition, illustrated in Fig. 7.4a as a plot of equilibrium values of  $z_r$  as a function of  $u_r$ , is called a *pitchfork bifurcation*.

To analyze the deadlock-breaking bifurcation, we linearize the nonlinear opinion equation (7.4) around the equilibrium  $z_r = 0$  and examine the eigenvalue  $\lambda = -d_r + \alpha_r u_r$  of the resulting linearization. The sign of  $\lambda$  governs the stability of the equilibrium  $z_r = 0$ . When  $\lambda < 0$  (resp.  $\lambda > 0$ ), then  $z_r = 0$ , and thus deadlock, is stable (resp. unstable).

The value of  $u_r$  corresponding to  $\lambda = 0$ , computed as  $u_r^* = d_r/\alpha_r$ , is thus the critical attention value. When the robot pays less attention ( $u_r < u_r^*$ ), then  $\lambda < 0$  and the robot remains in deadlock, attempting to go straight to its goal location. However, when the robot pays more attention ( $u_r > u_r^*$ ),  $\lambda > 0$  and deadlock becomes unstable. For  $u_r > u_r^*$  it can be shown that there are two additional symmetric equilibria  $z_r^{\text{eq}1} =$

$-z_r^{\text{eq}2} > 0$  that are both stable. These solutions correspond to a preference for going left ( $z_r = z_r^{\text{eq}1} > 0$ ), shown as the positive curve in blue in Fig. 7.4a, and a preference for going right ( $z_r = z_r^{\text{eq}2} < 0$ ), shown as the negative curve in blue in Fig. 7.4a. Note that the strength of preferences increases with increasing  $u_r > u_r^*$ . Because deadlock is unstable, the robot’s opinion will necessarily converge on one or the other opinionated solution. Which one it chooses will depend on initial conditions and noise.

When the robot is biased ( $b_r \neq 0$ ) or the human is approaching the robot obliquely ( $\eta_h \neq 0$ ), the pitchfork bifurcation *unfolds*, as illustrated in Fig. 7.4b. This implies that the robot prefers one side over the other when it passes the human mover. In particular, it can be shown that the robot prefers to move left if  $\gamma_r \tan \eta_h + b_r > 0$ , and right if  $\gamma_r \tan \eta_h + b_r < 0$ . Also, as we can observe from the diagram in Fig. 7.4b, where the robot has a bias  $b_r > 0$  for moving left, when  $u_r$  becomes sufficiently large, even though the robot favors left, if the robot is already moving right, it continues to move to this side. The analogous holds if  $b_r < 0$ .

We further illustrate the deadlock-breaking behavior with simulations in Fig. 7.4c. The human (trajectory in black) heads straight for the robot. In the unbiased case ( $b_r = 0$ ), the robot (trajectory in orange) moves straight just briefly before arbitrarily choosing to go right to pass around the human. This corresponds to behavior indicated by the negative blue curve in Fig. 7.4a. In the biased case ( $b_r > 0$ ), the robot (trajectory in purple) follows its bias and moves left, departing even sooner than it did in the unbiased case. This corresponds to the positive blue curve in Fig. 7.4b.

## 7.4 Experiments

We conducted two laboratory studies with human participants and one wheeled robot, a Clearpath Jackal UGV, moving in the 8m  $\times$  8m uncluttered space shown in Fig. 7.1a. We used a Vicon motion capture system to track the position and orientation of the robot and human movers who wore hats with a set of Vicon markers. The robot used the Vicon data to track the human movers. Our experimental goals are threefold: 1) to demonstrate the flexibility of the approach in that the robot can navigate a space while reliably interacting with multiple human movers in its path over a range of scenarios; 2) to validate the analysis of our algorithm, which shows that the robot is guaranteed to break deadlock, gracefully moving around an oncoming human mover even if the human is unaware of (or ignores) the robot and even if the robot has a bias that conflicts with the passing direction used by the human mover; and 3) to test our hypothesis that the trade-off between more efficient but less reliable passing and less efficient but more reliable passing can be controlled by the single parameter  $\beta_r$  in the robot’s algorithm (7.2).

### 7.4.1 Validation of flexibility of the approach

#### 7.4.1.1 Experimental setup

We ran a range of experimental trials each with a different scenario involving the robot and two human participants. In each trial, the robot and each of the humans were assigned a starting and goal location, which were selected to make the robot and human paths intersect. Human participants could walk along any path at any speed between their starting and goal locations.

In each trial, the robot was programmed to move at a constant speed of  $V_r = 0.75\text{m/s}$  towards its goal location while adjusting to human movement according to the navigation model (7.2) with attention dynamics specified by (7.3). At any given time, the robot considers only the closest nearby human (according to the measure  $\chi/\kappa$ ) seen within a distance of 20m and an angular range of  $(-\frac{\pi}{3}, \frac{\pi}{3})$  with respect to the robot's heading. If no humans are detected, the robot's attention and opinion are reset to their neutral value,  $u_r = z_r = 0$ . Results from five representative trials are shown in Fig. 7.2. The parameters for (7.2) were  $d_r = \alpha_r = 0.1$ ,  $\gamma_r = 4$ ,  $k_r = 1.5$ , and  $\beta_r = \pi/4$ . The parameters for (7.3) were  $\underline{u} = 0$  and  $R_r = n = 7$ . For trials in Fig. 7.2a-7.2d,  $\bar{u} = 1.5$  and for the trial in Fig. 7.2e,  $\bar{u} = 2.5$ .

#### 7.4.1.2 Results

Fig. 7.2 shows the resulting trajectories and the robot's opinion  $z_r$  and attention  $u_r$  over the full length of each trial. Temporal markers (dots) are included along the humans' trajectories and the robot's opinion and attention profiles. The top row shows how the robot navigates towards its goal while gracefully modifying its trajectory when encountering humans along its path. The bottom row shows how the robot's attention rises and falls in response to its proximity to a human. When the robot sees a human moving towards its left (resp. right), the opinion becomes negative (resp. positive) and the robot can be observed turning to its right (resp. left). When the robot sees no human to navigate around, its opinion is neutral and its go-to-goal behavior moves the robot towards its goal.

We observe in Fig. 7.2a-7.2c that the robot's opinion switches sign throughout each trial and that this is reflected in the robot's trajectory, which switches between turns to the left and turns to the right when it passes the human movers. The robot's attention rises and falls as the different participants are seen, maneuvered around, and passed by the robot. In Fig. 7.2d and 7.2e, the two human participants approach the robot side-by-side. However, the response of the robot is different in the two cases because the distance between the two participants is different. In Fig. 7.2d, the participants are close together and the robot passes to the right of both, whereas in Fig. 7.2e, the participants are further apart, and the robot navigates between them. This is a consequence of the proxy  $\tan \eta_h$  that has the same sign for each human mover in

the first case but different signs in the second case.

### 7.4.2 Validation of the deadlock breaking

#### 7.4.2.1 Experimental setup

Fixed pairs of starting and goal locations were assigned to the robot and a human participant. The human participant was asked to walk from (0m, 6.1m) to (0m, -1m), and the robot was programmed to navigate from (0m, 0m) to (0m, 6.1m). These locations were selected to make the robot and human move head-on toward one another.

The robot was programmed to move at a constant speed  $V_r = 0.7\text{m/s}$  toward its goal location, modifying its trajectory when encountering movers according to the navigation model (7.2) with parameters  $d_r = 0.5$ ,  $\alpha_r = 0.1$ ,  $\gamma_r = 3$ ,  $\tau_u = k_r = 1$ , and  $g(\kappa, \chi; R_r) = \exp(\kappa(R_r - \chi))$  with  $R_r = 11$ . We designed three cases corresponding to three different values of the robot's bias  $b_r$ : 1) unbiased ( $b_r = 0$ ), 2) biased to its left ( $b_r = 0.5$ ), and 3) biased to its right ( $b_r = -0.5$ ).

The participant was instructed to walk at their normal pace (their speed was recorded to be  $V_h = 1.09 \pm 0.03\text{m/s}$ ) towards their goal location according to one of three prompts: 1) go straight, 2) bear to the left, and 3) bear to the right.

We crossed the three cases for the robot and the three prompts for the human participant for a total of nine different trial configurations. We ran each of these nine different trial configurations five times for a total of 45 trials. Each of the 45 trials was run with  $\beta_r = \pi/4$  and  $\beta_r = \pi/6$  in (7.2) for a total of 90 trials.

#### 7.4.2.2 Results

Fig. 7.5 shows the resultant trajectories of the 90 trials organized by configuration on a  $3 \times 3$  grid. For a given configuration and value of  $\beta_r$  all five trials are plotted on the same graph. Trials where  $\beta_r = \pi/4$  are shaded in yellow and trials where  $\beta_r = \pi/6$  are unshaded. It can be observed that the robot navigated each trial configuration with similar path structure, regardless of the value of  $\beta_r$ .

In all the scenarios where the robot was unbiased (second row of Fig. 7.5), it successfully broke deadlock, verifying the guarantee of deadlock-free navigation provided by the model (7.2) and justified in the analysis of Section 7.3. In the trials when the human started directly facing the robot and continued walking straight ahead (UU), as in the simulation Fig. 7.4c, the robot quickly formed a strong opinion for one or the other direction. The robot chose to go left with about the same frequency that it chose to go right.

Having a bias allows the robot to rapidly form an initial opinion and break deadlock (turn left if  $b_r > 0$  or right if  $b_r < 0$ ). In the scenarios where the robot's bias was in conflict with the action taken by the human

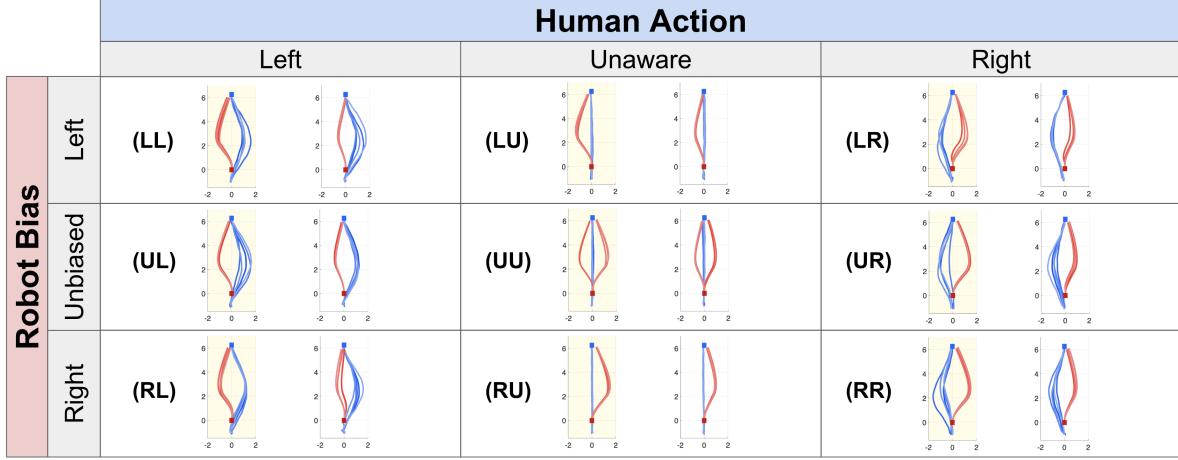


Figure 7.5: The trajectory data for five runs each of the nine trial configurations for the case  $\beta_r = \pi/4$  (shaded yellow) and for the case  $\beta_r = \pi/6$  (unshaded). Axes correspond to the  $xy$ -plane in meters. The robot paths are shown in red with a red box at the robot's starting position at about (0m, 0m). The human paths are shown in blue with a blue box at the human's starting position at about (0m, 6.1m). In trial configuration labels, L=left, U=unaware/unbiased, and R=right. Shorthand labels (eg. LL, LU) can be read as (robot bias, human action).

((LR) and (RL) in Fig. 7.5), the robot initially moved according to its bias but quickly adapted to the social cues given by the human and passed them in a cooperative fashion, i.e., matching the human movement and in opposition to its bias. This demonstration of flexibility provides evidence that the robot can reliably adjust its opinion to fit the social context in which it interacts with the human.

The results of Fig. 7.5 also provide evidence that a smaller  $\beta_r$  (unshaded plots) leads to more efficient (less time to goal) passing around the human as compared to a larger  $\beta_r$  (shaded plots). Fig. 7.6 provides further evidence of the role of  $\beta_r$  in tuning efficiency as the percent increase in length of the robot's path for the trials when  $\beta_r = \pi/4$  as compared to the case in which  $\beta_r = \pi/6$  was uniformly positive, at least 4% on average. Additionally, for each configuration, in trials with larger  $\beta_r$  the robot exhibited consistently higher maximum curvature along its path. Trials conducted with  $\beta_r = \pi/4$  showed an increase of approximately  $22.37\% \pm 6.71\%$  of the maximum curvature of the robot's trajectory as compared to the case  $\beta_r = \pi/6$ . This confirms that a robot with a larger  $\beta_r$  is less efficient.

Notably, Fig. 7.6 shows that the smallest percent increase in robot path length for the increase in  $\beta_r$  is in the UU case, when the robot was unbiased and the human unaware of the robot. This is consistent with the result that in this trial configuration, the robot took the most time to form a non-neutral opinion and turn to pass the human, which kept its paths in both  $\beta_r$  cases closer to the trial space's centerline than observed in other trial configurations.

Fig. 7.7 provides evidence that  $\beta_r$  tunes reliability and, together with the results of Fig. 7.6, that  $\beta_r$  tunes the efficiency-reliability trade-off, as hypothesized. The difference in the minimum distance recorded

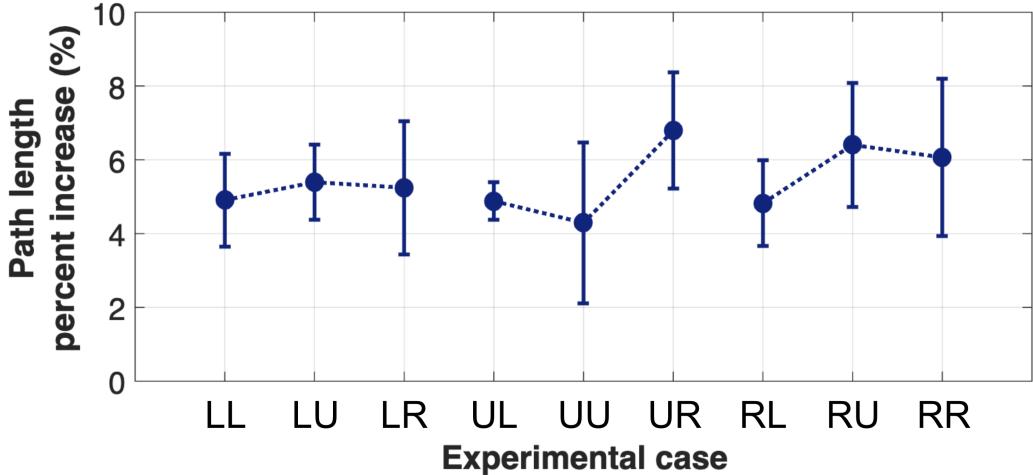


Figure 7.6: Percent increase of the robot’s path length for  $\beta_r = \pi/4$  compared to  $\beta_r = \pi/6$  for each of the nine configurations. Dotted lines link results associated with the same robot bias. L/U/R labels as in Fig. 7.5.

between the robot and human as they passed one another in each trial configuration for the different  $\beta_r$  values is shown in Fig. 7.7. The robot consistently came closer to the human along their paths for  $\beta_r = \pi/6$  as compared to  $\beta_r = \pi/4$ .

For each set of three configurations grouped by the robot’s bias, the robot came closest to the human whenever the human was unaware of the robot (i.e. LU, UU, RU). In the other configurations, the robot was able to cooperate with the human to form its opinion and pass the human like the human passed the robot. Without this cooperation, when the robot was the only participant in the passing, the passing distance was consistently smaller. The minimum distance in the case of the unbiased robot and unaware human was similar for the  $\beta_r = \pi/4$  and  $\beta_r = \pi/6$  trials. This suggests that this case is the most challenging for the robot independent of  $\beta_r$ . Still, the general decrease of the minimum distance between the robot and human that comes from a decrease in parameter  $\beta_r$  across all other configurations suggests that there is some design threshold where, once passed, the robot could not reliably navigate its way out of collision. Even if the robot’s algorithm is such that it can reliably form non-neutral opinions to break deadlock, the design parameters within the model must be sufficiently tuned for use in a real world dynamic context.

## 7.5 Discussion and final remarks

We present a new proactive approach to social robot navigation that leverages a nonlinear opinion dynamics model to enable a robot to rapidly and reliably pass approaching human movers, without requiring a model of human behavior. We show analytically and verify with human-robot experiments that this new navigation algorithm is guaranteed to break deadlock, even when the robot has no bias or evidence from the humans or

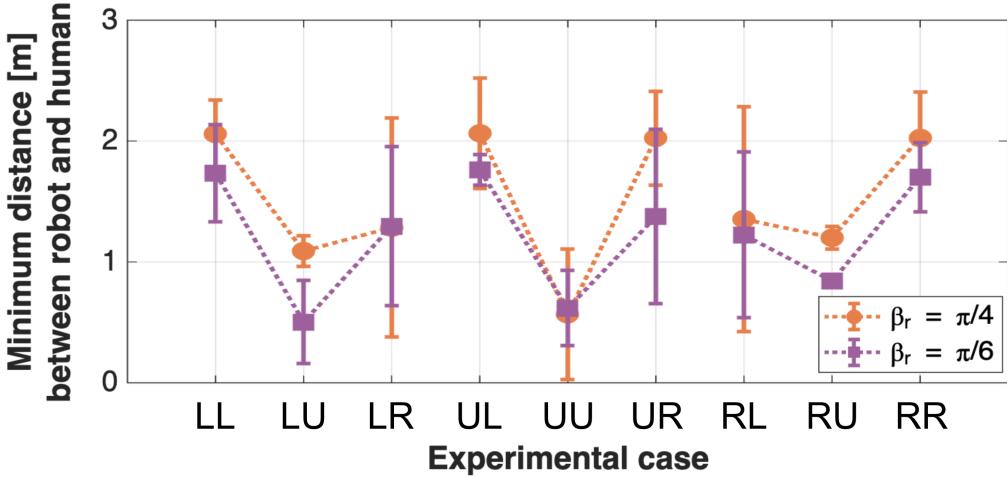


Figure 7.7: Average minimum distance between the robot and human for each of the nine configurations. Dotted lines link results associated with the same  $\beta_r$  value (orange line for  $\beta_r = \pi/4$ , purple line for  $\beta_r = \pi/6$ ) and robot bias. L/U/R labels as in Fig. 7.5.

the environment that one passing direction is better than the other. The experiments verify the flexibility of the approach with the robot reliably modifying its trajectory when encountering two human movers in its path. The experiments also verify that a robot with a bias for passing in one direction can still reliably pass the human mover even if the human chooses to pass in the direction that conflicts with the robot's bias. We show further how design parameters in the robot navigation algorithm can tune the robot's behavior, and verify in the experiments that parameter  $\beta_r$  tunes the efficiency-reliability trade-off in the passing problem. Future directions include extending the new approach to multi-robot social navigation in more complex scenarios, e.g., with more human movers and more cluttered environments. We also plan to investigate an extension that allows for increased attention to changes in context and tuning important trade-offs like efficiency versus reliability.

## Acknowledgment

The authors thank Anastasia Bizaeva and Alessio Franci for discussions on Section 7.3 and Giovanna Amorim and Justin Lidard for discussions on Section 7.4.

# Bibliography

- [1] A. Bizyaeva, A. Franci, and N. E. Leonard, “Nonlinear Opinion Dynamics With Tunable Sensitivity,” *IEEE Transactions on Automatic Control*, vol. 68, no. 3, pp. 1415–1430, Mar. 2023.
- [2] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, “Core Challenges of Social Robot Navigation: A Survey,” *J. Hum.-Robot Interact.*, vol. 12, no. 3, pp. 36:1–36:39, Apr. 2023.
- [3] R. Mirsky, X. Xiao, J. Hart, and P. Stone, “Conflict Avoidance in Social Navigation—a Survey,” *ACM Transactions on Human-Robot Interaction*, vol. 13, no. 1, pp. 1–36, Mar. 2024.
- [4] Y. Gao and C.-M. Huang, “Evaluation of Socially-Aware Robot Navigation,” *Frontiers in Robotics and AI*, vol. 8, Jan. 2022.
- [5] C. Mavrogiannis, P. Alves-Oliveira, W. Thomason, and R. A. Knepper, “Social Momentum: Design and Evaluation of a Framework for Socially Competent Robot Navigation,” *ACM Transactions on Human-Robot Interaction*, vol. 11, no. 2, pp. 1–37, Jun. 2022.
- [6] D. Helbing and P. Molnár, “Social Force Model for Pedestrian Dynamics,” *Physical Review E*, vol. 51, no. 5, pp. 4282–4286, May 1995.
- [7] A. K. Reddy, V. Malviya, and R. Kala, “Social Cues in the Autonomous Navigation of Indoor Mobile Robots,” *International Journal of Social Robotics*, vol. 13, no. 6, pp. 1335–1358, Sep. 2021.
- [8] M. Kamezaki, Y. Tsuburaya, T. Kanada, M. Hirayama, and S. Sugano, “Reactive, Proactive, and Inducible Proximal Crowd Robot Navigation Method Based on Inducible Social Force Model,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3922–3929, Apr. 2022.
- [9] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard, “Socially Compliant Mobile Robot Navigation via Inverse Reinforcement Learning,” *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, Sep. 2016.

[10] Y. F. Chen, M. Everett, M. Liu, and J. P. How, “Socially Aware Motion Planning With Deep Reinforcement Learning,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 1343–1350.

[11] B. Okal and K. O. Arras, “Learning Socially Normative Robot Navigation Behaviors With Bayesian Inverse Reinforcement Learning,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 2889–2895.

[12] P. Trautman and A. Krause, “Unfreezing the robot: Navigation in dense, interacting crowds,” in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2010, pp. 797–803.

[13] P. Trautman, J. Ma, R. M. Murray, and A. Krause, “Robot navigation in dense human crowds: The case for cooperation,” in *2013 IEEE International Conference on Robotics and Automation*. Karlsruhe, Germany: IEEE, May 2013, pp. 2153–2160.

[14] S. Saunderson and G. Nejat, “How Robots Influence Humans: A Survey of Nonverbal Communication in Social Human–Robot Interaction,” *International Journal of Social Robotics*, vol. 11, no. 4, pp. 575–608, Aug. 2019.

[15] A. F. Peralta, J. Kertész, and G. Iñiguez, “Opinion dynamics in social networks: From models to data.” In: T. Yasseri (Ed.), *Handbook of Computational Social Science*. Edward Elgar Publishing Ltd., 2023.

[16] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, and M. Starnini, “Modeling Echo Chambers and Polarization Dynamics in Social Networks,” *Physical Review Letters*, vol. 124, no. 4, p. 048301, Jan. 2020.

[17] M. A. Montes de Oca, E. Ferrante, N. Mathews, M. Birattari, and M. Dorigo, “Opinion Dynamics for Decentralized Decision-Making in a Robot Swarm,” in *Swarm Intelligence*, M. Dorigo, M. Birattari, G. A. Di Caro, R. Doursat, A. P. Engelbrecht, D. Floreano, L. M. Gambardella, R. Groß, E. Şahin, H. Sayama, and T. Stützle, Eds. Berlin, Heidelberg: Springer, 2010, pp. 251–262.

[18] F. A. Salem, R. d. S. Tchilian, S. R. D. Carvalho, and U. F. Moreno, “Opinion Dynamics over a Finite Set in Cooperative Multi-robot Systems: An Asynchronous Gossip-Based Consensus Approach,” *Congresso Brasileiro de Automática - CBA*, vol. 2, no. 1, Dec. 2020.

[19] M. Li, W. Bian, L. Chen, and M. Liu, “HiDeS: A higher-order-derivative-supervised neural ordinary differential equation for multi-robot systems and opinion dynamics,” *Frontiers in Neurorobotics*, vol. 18, Mar. 2024.

[20] R. Gray, A. Franci, V. Srivastava, and N. E. Leonard, “Multiagent Decision-Making Dynamics Inspired by Honeybees,” *IEEE Transactions on Control of Network Systems*, vol. 5, no. 2, pp. 793–806, Jun. 2018.

[21] V. Srivastava and N. E. Leonard, “Bio-Inspired Decision-Making and Control: From Honeybees and Neurons to Network Design,” in *2017 American Control Conference (ACC)*, May 2017, pp. 2026–2039.

[22] A. Franci, A. Bizyaeva, S. Park, and N. E. Leonard, “Analysis and Control of Agreement and Disagreement Opinion Cascades,” *Swarm Intelligence*, vol. 15, no. 1-2, pp. 47–82, Jun. 2021.

[23] A. Bizyaeva, G. Amorim, M. Santos, A. Franci, and N. E. Leonard, “Switching Transformations for Decentralized Control of Opinion Patterns in Signed Networks: Application to Dynamic Task Allocation,” *IEEE Control Systems Letters*, vol. 6, pp. 3463–3468, 2022.

[24] A. Bizyaeva, A. Franci, and N. E. Leonard, “Multi-Topic Belief Formation Through Bifurcations Over Signed Social Networks,” Jul. 2024.

[25] C. Cathcart, M. Santos, S. Park, and N. E. Leonard, “Proactive Opinion-Driven Robot Navigation Around Human Movers,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2023, pp. 4052–4058.

[26] C. Cathcart, I. X. Belaustegui, A. Franci, and N. E. Leonard, “Spiking Nonlinear Opinion Dynamics (S-NOD) for Agile Decision-Making,” *IEEE Control Systems Letters*, vol. 8, pp. 3267–3272, 2024, conference Name: IEEE Control Systems Letters. [Online]. Available: <https://ieeexplore.ieee.org/document/10817538>

[27] R. Kumar, A. Bhatti, and N. Yao, “Consensus Building in Human-robot Co-learning via Bias Controlled Nonlinear Opinion Dynamics and Non-verbal Communication through Robotic Eyes,” May 2025.

[28] H. Hu, J. F. Fisac, N. E. Leonard, D. Gopinath, J. DeCastro, and G. Rosman, “Think Deep and Fast: Learning Neural Nonlinear Opinion Dynamics from Inverse Dynamic Games for Split-Second Interactions,” Apr. 2025.

[29] T. M. Paine and M. R. Benjamin, “A Model for Multi-Agent Autonomy That Uses Opinion Dynamics and Multi-Objective Behavior Optimization,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 8305–8311.

[30] P. Dayan and L. F. Abbott, *Theoretical Neuroscience Computational and Mathematical Modeling of Neural Systems*, ser. Computational Neuroscience. Cambridge, Mass.: MIT Press, 2001.

[31] G. Balázs, A. van Oudenaarden, and J. J. Collins, “Cellular Decision Making and Biological Noise: From Microbes to Mammals,” *Cell*, vol. 144, no. 6, pp. 910–925, Mar. 2011.

[32] G. M. Süel, J. Garcia-Ojalvo, L. M. Liberman, and M. B. Elowitz, “An excitable gene regulatory circuit induces transient cellular differentiation,” *Nature*, vol. 440, no. 7083, pp. 545–550, Mar. 2006.

[33] J. Fromm and S. Lautner, “Electrical Signals and Their Physiological Significance in Plants,” *Plant, Cell & Environment*, vol. 30, no. 3, pp. 249–257, 2007.

[34] R. FitzHugh, “Mathematical models of excitation and propagation in nerve,” in *Biological Engineering*, H. P. Schwan, Ed. New York: McGraw-Hill Book Co., 1969, ch. 1, pp. 1–85.

[35] A. L. Hodgkin and A. F. Huxley, “A quantitative description of membrane current and its application to conduction and excitation in nerve,” *The Journal of Physiology*, vol. 117, no. 4, pp. 500–544, Aug. 1952.

[36] R. Sepulchre, “Spiking Control Systems,” *Proceedings of the IEEE*, vol. 110, no. 5, pp. 577–589, May 2022.

[37] C. Bartolozzi, G. Indiveri, and E. Donati, “Embodied Neuromorphic Intelligence,” *Nature Communications*, vol. 13, no. 1, p. 1024, Feb. 2022.

[38] A. Franci, G. Drion, and R. Sepulchre, “The Sensitivity Function of Excitable Feedback Systems,” in *2019 IEEE 58th Conference on Decision and Control (CDC)*, Dec. 2019, pp. 4723–4728.

[39] R. Sepulchre, G. Drion, and A. Franci, “Excitable Behaviors,” in *Emerging Applications of Control and Systems Theory: A Festschrift in Honor of Mathukumalli Vidyasagar*, R. Tempo, S. Yurkovich, and P. Misra, Eds. Cham: Springer International Publishing, 2018, pp. 269–280.

[40] N. J. Priebe, “Mechanisms of Orientation Selectivity in the Primary Visual Cortex,” *Annual Review of Vision Science*, vol. 2, pp. 85–107, Oct. 2016.

[41] J. P. Gallivan, C. S. Chapman, D. M. Wolpert, and J. R. Flanagan, “Decision-Making in Sensorimotor Control,” *Nature Reviews Neuroscience*, vol. 19, no. 9, pp. 519–534, Sep. 2018.

[42] N. E. Leonard, A. Bizyaeva, and A. Franci, “Fast and Flexible Multiagent Decision-Making,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, no. 1, pp. 19–45, 2024.

[43] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, ser. Applied Mathematical Sciences. New York, NY: Springer, 1983, vol. 42.

[44] M. Golubitsky and D. G. Schaeffer, *Singularities and Groups in Bifurcation Theory: Volume I*, ser. Applied Mathematical Sciences. New York, NY: Springer-Verlag, 1985, vol. 51.

[45] J. K. Hale and H. Koçak, *Dynamics and Bifurcations*, ser. Texts in Applied Mathematics, J. E. Marsden, L. Sirovich, M. Golubitsky, W. Jäger, and F. John, Eds. New York, NY: Springer, 1991, vol. 3.

[46] S. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*, 2nd ed., ser. A Chapman & Hall Book. Boca Raton London New York: CRC Press, 2015.

[47] “Jackal UGV - Small Weatherproof Robot - Clearpath.” [Online]. Available: <https://clearpathrobotics.com/jackal-small-unmanned-ground-vehicle/>

[48] “Jetson TX2 Module.” [Online]. Available: <https://developer.nvidia.com/embedded/jetson-tx2>

[49] “Puck by Velodyne.” [Online]. Available: <https://store.clearpathrobotics.com/products/puck>

[50] “vicon\_bridge - ROS Wiki.” [Online]. Available: [http://wiki.ros.org/vicon\\_bridge](http://wiki.ros.org/vicon_bridge)

[51] “Vicon Vero | Advanced Super Wide Motion Capture Camera.” [Online]. Available: <https://www.vicon.com/hardware/cameras/vero/>

[52] A. Witmer, “Development of Robot Navigation Systems,” Undergraduate Independent Work Thesis, Princeton University, Princeton, New Jersey, Apr. 2020.

[53] S. Kang, S. Yang, D. Kwak, Y. Jargalbaatar, and D. Kim, “Social Type-Aware Navigation Framework for Mobile Robots in Human-Shared Environments,” *Sensors*, vol. 24, no. 15, p. 4862, Jan. 2024, number: 15 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/1424-8220/24/15/4862>

[54] M. Walters, K. Dautenhahn, K. Koay, C. Kaouri, R. Boekhorst, C. Nehaniv, I. Werry, and D. Lee, “Close encounters: Spatial distances between people and a robot of mechanistic appearance,” in *5th IEEE-RAS International Conference on Humanoid Robots, 2005.*, Dec. 2005, pp. 450–455.

[55] E. Hall, E. Hall, and C. P. C. L. of Congress), *The Hidden Dimension*, ser. Anchor Books. Doubleday, 1969.

[56] N. Tsoi, J. Romero, and M. Vázquez, “How Do Robot Experts Measure the Success of Social Robot Navigation?” in *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI ’24. New York, NY, USA: Association for Computing Machinery, Mar. 2024, pp. 1063–1066.

[57] R. Sepulchre, G. Drion, and A. Franci, “Control Across Scales by Positive and Negative Feedback,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, no. 1, pp. 89–113, May 2019.

[58] G. Amorim, M. Santos, S. Park, A. Franci, and N. E. Leonard, “Threshold Decision-Making Dynamics Adaptive to Physical Constraints and Changing Environment,” in *2024 European Control Conference (ECC)*, Jun. 2024, pp. 1908–1913.

[59] K. Özcimder, B. Dey, R. J. Lazier, D. Trueman, and N. E. Leonard, “Investigating Group Behavior in Dance: An Evolutionary Dynamics Approach,” in *2016 American Control Conference (ACC)*, Jul. 2016, pp. 6465–6470.

[60] M. Krupa and P. Szmolyan, “Relaxation Oscillation and Canard Explosion,” *Journal of Differential Equations*, vol. 174, no. 2, pp. 312–368, Aug. 2001.

[61] J. Arbelaitz, A. Franci, N. E. Leonard, R. Sepulchre, and B. Bamieh, “Excitable Crawling,” 26th International Symposium on Mathematical Theory of Networks and Systems, May 2024.

[62] I. X. Belaustegui, A. Franci, and N. E. Leonard, “Tunable Thresholds and Frequency Encoding in a Spiking NOD Controller,” Apr. 2025.

[63] G. Amorim, A. Bizyaeva, A. Franci, and N. E. Leonard, “Spatially-Invariant Opinion Dynamics on the Circle,” *IEEE Control Systems Letters*, vol. 8, pp. 3231–3236, 2024.

[64] H. Hamann, “Opinion Dynamics With Mobile Agents: Contrarian Effects by Spatial Correlations,” *Frontiers in Robotics and AI*, vol. 5, Jun. 2018.

[65] J. Hart, R. Mirsky, X. Xiao, S. Tejeda, B. Mahajan, J. Goo, K. Baldauf, S. Owen, and P. Stone, “Using Human-Inspired Signals to Disambiguate Navigational Intentions,” in *Social Robotics*, A. R. Wagner, D. Feil-Seifer, K. S. Haring, S. Rossi, T. Williams, H. He, and S. Sam Ge, Eds. Cham: Springer International Publishing, 2020, pp. 320–331.

[66] H. Kivrak, F. Cakmak, H. Kose, and S. Yavuz, “Social Navigation Framework for Assistive Robots in Human Inhabited Unknown Environments,” *Engineering Science and Technology, an International Journal*, vol. 24, no. 2, pp. 284–298, Apr. 2021.

[67] R. Kirby, R. Simmons, and J. 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*. Toyama, Japan: IEEE, Sep. 2009, pp. 607–612.

[68] S. S. Samsani and M. S. Muhammad, “Socially Compliant Robot Navigation in Crowded Environment by Human Behavior Resemblance Using Deep Reinforcement Learning,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5223–5230, Jul. 2021.

[69] M. Kollmitz, T. Koller, J. Boedecker, and W. Burgard, “Learning Human-Aware Robot Navigation from Physical Interaction via Inverse Reinforcement Learning,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2020, pp. 11 025–11 031.

[70] A. Bera, T. Randhavane, R. Prinja, and D. Manocha, “SocioSense: Robot Navigation Amongst Pedestrians With Social and Psychological Constraints,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 7018–7025.

[71] Y. Che, A. M. Okamura, and D. Sadigh, “Efficient and Trustworthy Social Navigation via Explicit and Implicit Robot–Human Communication,” *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 692–707, Jun. 2020.

[72] E. Pacchierotti, H. Christensen, and P. Jensfelt, “Human-Robot Embodied Interaction an Hallway Settings: A Pilot User Study,” in *ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005.*, Aug. 2005, pp. 164–171.

[73] J. Thomas and R. Vaughan, “After You: Doorway Negotiation for Human-Robot and Robot-Robot Interaction,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2018, pp. 3387–3394.

[74] D. V. Lu and W. D. Smart, “Towards More Efficient Navigation for Robots and Humans,” in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov. 2013, pp. 1707–1713.

[75] S. M. Fiore, T. J. Wiltshire, E. J. C. Lobato, F. G. Jentsch, W. H. Huang, and B. Axelrod, “Toward Understanding Social Cues and Signals in Human–Robot Interaction: Effects of Robot Gaze and Proxemic Behavior,” *Frontiers in Psychology*, vol. 4, p. 859, Nov. 2013.

[76] V. V. Unhelkar, C. Pérez-D'Arpino, L. Stirling, and J. A. Shah, “Human-Robot Co-navigation Using Anticipatory Indicators of Human Walking Motion,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 6183–6190.

[77] P. Ratsamee, Y. Mae, K. Kamiyama, M. Horade, M. Kojima, and T. Arai, “Social Interactive Robot Navigation Based on Human Intention Analysis from Face Orientation and Human Path Prediction,” *ROBOMECH Journal*, vol. 2, no. 1, p. 11, Aug. 2015.