

OPTIMAL EVASIVE STRATEGIES FOR GROUPS OF
INTERACTING AGENTS WITH MOTION
CONSTRAINTS

WILLIAM LEWIS SCOTT, III

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: PROFESSOR NAOMI EHRLICH LEONARD

JANUARY 2017

© Copyright by William Lewis Scott, III, 2016.

All rights reserved.

Abstract

In this thesis we examine systems of pursuit and evasion with multiple evaders from several perspectives. Through the analysis of mathematical models and the study of field experiments we seek to understand how constraints on individual motion and sensing abilities affect outcomes on the level of the group and the individual.

We present a nonlinear model of pursuit and evasion on the plane for a single pursuer and two evaders. Control laws are defined so that each evader trades off between evasion and herding. We analyze the system dynamics in terms of relative shape variables and derive conditions on the control parameters and initial conditions that determine whether capture occurs.

We consider a system where evaders have heterogeneous limits on speed, turning rate, and lateral acceleration, versus a single pursuer with limited speed but no turning constraints. Optimal strategies are derived for the one-on-one differential game, and these form the basis of strategies for the multiple-evader system. Explicit analytic expressions for open-loop and state-feedback forms of optimal controls are derived in a related minimum-time problem for a single agent with the evader's motion constraints. For the multiple evader system, we propose a pursuer strategy of optimal target selection which leads to capture in bounded time and we prove how any evader not initially targeted can avoid capture with a *reactive evasion* strategy. We consider optimal strategies for agents with radius-limited sensing capabilities. In the case that evader turning rate is unbounded, we prove conditions for evader capture avoidance through a local strategy of *risk reduction*, which we show leads to group aggregation.

We present some preliminary results from a field experiment conducted at the Ol Pejeta Conservancy in Laikipia, Kenya in July 2014 to study evasive behaviors in herds of plains zebra under pursuit by an artificial predator, the “robo-lion.” Through manual tracking of video from an overhead camera we extract quantitative trajectory data for each zebra in

the herd and for the robo-lion. The observed zebra behaviors of efficient evasion and group alignment serve to motivate our modeling efforts.

Acknowledgements

First of all, and most of all, I would like to thank my advisor, Naomi Leonard, for her thoughtful guidance, enthusiastic encouragement, and unending support during my time at Princeton. It has been a great pleasure to collaborate with someone who takes great pride and responsibility in her role as a world-class scientist, writer, teacher, and mentor. She gave me the freedom to explore a wide variety of research problems that I found interesting, and I always knew I could count on her to provide valuable insight and fruitful suggestions on any topic of discussion. I am fortunate to have had the opportunity to work with Naomi as a teaching assistant for two classes. Her clear and logical teaching style will serve as a model for me as I pursue a career in academia. I am immensely grateful to Naomi for helping me to grow as a scientific writer. At first I dreaded receiving my paper drafts back completely covered in red ink, but I soon came to greatly appreciate Naomi's keen attention to detail and uncompromising pursuit of clarity during the editing process. The past five years have gone by too quickly, and I am sad to leave when there is still so much to learn.

Naomi has done a wonderful job of attracting a diverse and talented group of students to the Dynamical Control Systems Lab. I want to acknowledge lab members that I looked up to when I was getting started, including Paul Reverdy, Darren Pais, George Young, Tian Shen, and Dan Swain. I would like to thank all the present lab members, Peter Landgren, Liz Davison, Bec Gray, Renato Pagliara Vasquez, Anthony Savas, and Desmond Zhong for taking an interest in my work and helping me through the difficult process of finishing my thesis. I also would like to thank the postdocs in our lab, Biswadip Dey, Kayhan Ozcimder, Karla Kvaternik, and Vaibhav Srivastava, who have served as mentors and role models and encouraged me to continue in academia as a postdoc myself.

I thank Simon Levin for serving on my Ph.D. committee and for reading my thesis, and I thank Philip Holmes for serving on my Ph.D. committee. I always enjoyed meeting with Simon and Phil. Their stories and insightful questions helped me to see my research problems from new perspectives, and to place my work into a broader context. I thank

Jeremy Kasdin for serving as a thesis reader and offering valuable comments. I also thank Clancy Rowley for serving as an examiner at my FPO, and for helping me to prepare for the subject component of my general exam.

I am deeply grateful to Daniel Rubenstein for our collaboration over the past few years. Our discussions about animal behavior had a major role in guiding the direction of my research at Princeton. I especially thank him for inviting me to travel to Kenya to take part in the “robo-lion” experiment. Staying at Mpala to learn about zebras was truly a once in a lifetime experience that I will cherish. I want to also thank the students who helped out with the robo-lion project. Thank you to Kaia Tombak for dutifully aiming the camera while I had the fun job of piloting the robo-lion, and thank you to Eric Principato for literally saving the whole project when he was able to fix the robo-lion’s malfunctioning wiring. Thank you to Blair Costelloe for building the robo-lion, and for serving as my travel guide as I made my way to Kenya for the first time. Thank you to Marina Latif, Pria Louka, Julie Pourtois, and Lisa Sheridan for helping with the painstaking work of manually tracking the zebra trajectories in the videos.

Success in graduate school would have been impossible without some friends to share in the ups and downs. I met so many amazing people at Princeton, and I am happy to count many friends among them. To Katie Fitch, A J Riggs, and Anthony DeGennaro, who shared my struggles studying for the general exam in dynamics and controls; to Scott Dawson, Steve Atkinson, Ryan Davis, Brendan Andrade, Matt Plasek, and the others in my cohort in MAE; to Sudhir Raskutti, Max Hirschberger, Catherine Offord, Ari Strandburg-Peshkin, and the others I met outside of class, I am grateful for our friendship and the adventures we shared during our time together at Princeton.

I would like to acknowledge Nuno Martins and Shinkyu Park from the University of Maryland, and Kyler Abernathy, Greg Marshall, and Konrad Aschenbach from the National Geographic Society for their collaboration on the NSF “Cittercam” project.

I extend a special thanks to Jill Ray, our MAE Graduate Program Administrator, who has been a great help to me in navigating through all the necessary forms and paperwork of graduate school.

When I was in college at the University of Delaware, two professors in particular fostered my interest in research and guided me towards pursuing graduate education. I am grateful to Erik Thostenson for hiring me to work in his composite materials lab. I owe a tremendous debt of gratitude to Ioannis Poulakakis for setting me on the path towards graduate school. As a first year professor, he took me under his wing and tutored me in dynamics and controls so I could work on an independent research project. He encouraged me to take difficult classes in my senior year so I could have a strong graduate school application, and pushed me to apply to all the best schools (including Princeton). He is the one who introduced me to Naomi, and for that I am grateful.

Finally, thank you to Mom and Dad and Megan and Brian (and Maisy) for always believing in me. You guys are the best!

This dissertation carries the number T-3334 in the records of the Department of Mechanical and Aerospace Engineering.

In memory of Mommom,
Ruth Ann Scott (1929-2011)

Contents

Abstract	iii
Acknowledgements	v
List of Tables	xiii
List of Figures	xiv
1 Introduction	1
1.1 Overview of topics	1
1.1.1 Pursuit and evasion	1
1.1.2 Minimum-time optimal control	3
1.1.3 Pursuit and evasion with multiple evaders	4
1.1.4 Field studies of pursuit in nature	6
1.2 Outline	7
2 Three agent model of pursuit, evasion, and herding	9
2.1 Introduction	9
2.2 Steered-particle model	11
2.3 Simplified first-order model	13
2.3.1 Length dynamics	17
2.3.2 Angle dynamics	17
2.3.3 Parallel motion equilibria	18
2.4 Mother caribou in pure evasion	19

2.4.1	Avoiding mother-calf separation	19
2.4.2	Disadvantage of pure evasion by the calf	21
2.4.3	When the bear ignores the calf	22
2.5	Simulations and estimated solution to a two-player game	26
3	Minimum-time trajectories for steered agent with constraints on speed, lateral acceleration, and turning rate	29
3.1	Introduction	29
3.2	Problem statement and system dynamics	31
3.3	Extremal trajectories from Pontryagin’s minimum principle	32
3.3.1	Switching functions and generic control inputs	35
3.3.2	Singular control inputs	36
3.3.3	Multiply-singular control	38
3.4	Families of optimal trajectories	39
3.4.1	Trajectories ending in forward motion	40
3.4.2	Trajectories ending in fast turn	42
3.5	The Optimal Trajectory	44
3.5.1	Trajectory parameterized by switching times	44
3.5.2	Trajectory-type partition	45
3.5.3	Optimal switching times for each compound trajectory type	46
3.6	State-feedback formulation of optimal control law	55
3.7	Special cases for large and small values of μ	56
3.7.1	Relaxed acceleration constraint	56
3.7.2	Highly constrained lateral acceleration	59
3.8	Final Remarks	60
4	Optimal evasive strategies for multiple interacting agents with motion constraints	61

4.1	Introduction	61
4.2	Problem statement and equations of motion	64
4.3	Pursuit and evasion with two agents	65
4.3.1	Terminal conditions	67
4.3.2	Optimal trajectories for pursuit and evasion	69
4.3.3	Evader control switching times	71
4.3.4	Evader state-feedback control law	72
4.3.5	Condition on capture radius	75
4.3.6	The case of the slower pursuer	76
4.4	Reactive evasion with multiple evaders	78
4.4.1	Evader domain of danger and target avoidance	78
4.5	Risk minimization under limited sensing	84
4.5.1	Relaxing constraints on evader turning rate	84
4.5.2	Local target selection	88
4.5.3	Local evasion strategy	89
4.5.4	Risk reduction phase	89
4.5.5	Local reactive evasion phase	91
4.6	Discussion	94
5	Zebra experiment	95
5.1	Background	97
5.1.1	Plains zebra	97
5.1.2	Previous work	98
5.2	Equipment and methods	100
5.3	Data processing	101
5.4	Results and analysis	102
5.4.1	Speed	107
5.4.2	Initiation of movement	108

5.4.3	Distance to the robo-lion	109
5.4.4	Dynamics of heading alignment	112
5.4.5	Locations within the herd	114
5.4.6	Dynamics of mother-foal pairs	115
5.5	Conclusion	119
6	Final remarks	120
6.1	Conclusions	120
6.2	Future directions	123
6.2.1	Improving models of pursuit and evasion	123
6.2.2	Use of robotics in biological research	125
A	Data processing for the robo-lion zebra experiment	126
A.1	Video stabilization	126
A.2	Track extraction	129
A.3	Coordinate transformation	129
A.4	Trajectory smoothing	133
	Bibliography	136

List of Tables

3.1	Range of possible segment durations for each minimum-time trajectory type	46
3.2	Parametric description of minimum-time trajectory-type partition boundaries	48
5.1	Robo-lion experiment trial summary	103
5.2	Average and maximum speeds for each trial	107
A.1	Image to world transformation parameters for each trial	133

List of Figures

2.1	System states for three-agent mother-calf-bear steered-particle model	11
2.2	Simulated trajectories in the simplified first-order model for various parameter values	15
2.3	Diagram of shape variables in the three-agent model	16
2.4	Range of parameters w_c and w_t such that the mother and calf caribou cannot be separated by the bear when starting with the mother closer to the calf . .	22
2.5	Coordinate diagram for trajectories with $w_m = w_t = 0$, $w_c = 1$	24
2.6	Simulation of capture probability under varying control parameters	27
3.1	Admissible control inputs for steered agent with limits on speed, turning rate and lateral acceleration as described in Sec. 3.2.	31
3.2	Diagram of control switching points for families of minimum-time trajectories ending in forward motion and fast turn	41
3.3	Trajectory-type partition in minimum-time problem	47
3.4	Minimum-time trajectories with varying lateral acceleration constraint . . .	49
3.5	Minimum time-to-reach surfaces with varying lateral acceleration constraint	50
3.6	Control switching regions for state-feedback control in the minimum-time problem	55
3.7	Trajectory-type partition and optimal trajectories for minimum-time problem for extreme values of the lateral acceleration constraint parameter	57

4.1	Optimal trajectories for one-on-one pursuit and evasion in reduced coordinates	73
4.2	Optimal pursuer heading θ_p^* in reduced coordinates	74
4.3	Time-to-capture T_{cap} surface in reduced coordinates	75
4.4	Optimal trajectories in reduced coordinates in the case of a slower pursuer $\bar{v}_p < \bar{v}_e$	77
4.5	Simulation of reactive evasion for ten evaders with motion constraints	83
4.6	Comparison of the “slowing” and “spiral” reactive evasion strategies	87
4.7	Weighted-Voronoi domain of danger partition for evaders with relaxed turning constraints	93
5.1	The robo-lion device in side and front view	100
5.2	Smoothed trajectories for trials 13-1 through 14-1	104
5.3	Smoothed trajectories for trials 14-3 through 15-3	105
5.4	Smoothed trajectories for trials 15-5 through 16-5	106
5.5	Zebra activity states before and during robo-lion chase	110
5.6	Time spent with head up before initiating movement across zebra classes	111
5.7	Scaled distance to robo-lion at start of movement	111
5.8	Minimum scaled distance to the robo-lion	112
5.9	Herd alignment and evasiveness order parameters over time for trial 16-1	114
5.10	Average forward position in herd for zebras arranged by trial and by zebra class	116
5.11	Trajectories of mother-foal pairs with the robo-lion	118
A.1	Illustration of horizon detection and homography transformation	130

Chapter 1

Introduction

In this thesis we examine systems of pursuit and evasion with multiple evaders from several perspectives. Through the analysis of mathematical models and the study of field experiments we seek to understand how constraints on individual motion and sensing abilities affect outcomes on the level of the group and the individual. This chapter presents background on the different approaches we consider in Section 1.1, and includes a detailed outline of the thesis in Section 1.2.

1.1 Overview of topics

1.1.1 Pursuit and evasion

The mathematical study of pursuit and evasion has a rich history stretching back to the 18th century with Bouguer's famous pirate ship pursuit problem [56]. It still remains an active area of study today due to the multitude of possible applications and variations, both in understanding biological systems and in the design of control algorithms for engineered robotic systems.

In the most general form, a mathematical model of pursuit and evasion features two agents, a pursuer and an evader, and is defined by equations of motion for each agent along

with control laws specifying their control inputs, and some terminal condition on the state that defines when capture has occurred. The goal of the researcher is to characterize how different system parameters and initial conditions affect the outcome of the chase: whether it will end in capture or escape.

Given the equations of motion and a capture condition, the study of mathematical models of pursuit and evasion can be categorized into two broad approaches. The first is to propose the control laws for the pursuer and evader *a priori*, and then to analyze the dynamics in the sense of a nonlinear system. The control laws used in this approach are often inspired by observations of a biological system. The second approach is to define an objective function to be minimized by the pursuer and maximized by the evader, and then to apply techniques from optimal control theory to generate optimal control laws for the pursuer and evader. This approach is known as a “differential game” and was first formalized by Isaacs in 1965 [44]. Hybrid approaches that look for an optimal strategy against an opponent with a known control law have also been considered.

In this thesis we use both approaches to study different pursuit and evasion systems: In Chapter 2 we define control laws in a two-evader system that trade off between classical evasion and herding strategies, and in Chapter 4 we consider a differential game for an evader with motion constraints inspired by terrestrial animal locomotion.

Many of the classical pursuit and evasion problems of the first type that can be solved in closed form are collected in Nahin’s book [56]. These involve the well known pursuer control laws of “classical pursuit” and “classical evasion” where the pursuer aligns its velocity in the direction of the evader at all times and the evader aligns its velocity away from the pursuer, respectively. The other traditional strategy is called “constant bearing,” where an agent aligns itself such that its opponent is kept at a constant relative angle. The constant bearing strategy is the optimal strategy for a pursuer to intercept an evader moving in a straight line at constant speed.

Pursuit and evasion control laws have been studied extensively for steered-particle systems, in which agents are modeled with unicycle dynamics on the plane with control inputs of speed and turning rate. Recent studies have analyzed “motion camouflage” strategies employed by dragonflies and bats that take advantage of the limited visual capabilities of their insect prey [45]. The selection of pursuit and evasion control laws for steered-particle agents has been studied as an evolutionary game [78, 59].

A biologically inspired analysis of pursuit and evasion with acceleration constraints in [41] suggests that a more agile but slower evader can escape from a fast pursuer with limited lateral acceleration by veering to the side at the last moment. These dynamics are adapted to apply to schooling fish in [79], where they draw the distinction between “avoidance” strategies that aim to avoid detection by the pursuer, and “evasion” strategies to avoid capture when detected.

The differential game approach has traditionally been applied to systems with engineering applications, such as those with car-like or airplane-like dynamics. Perhaps the most famous differential game is the so-called “homicidal chauffeur” system featuring a fast car-like pursuer with constant speed and limited turning radius against a slower but agile evader with no limits on turning, where the evader can escape capture indefinitely by moving to the side when the pursuer nears [44]. This was extended to the “game of two cars” where both agents have limited turning radii [55, 34].

Several recent papers examine differential games featuring steered agents with turning constraints, such as a differential-drive pursuer with omnidirectional evader [65], the same but with the roles switched [64], and an omnidirectional pursuer with a car-like evader [27].

1.1.2 Minimum-time optimal control¹

Minimum-time problems have long been a source of fascination for the mathematics community, reaching as far back as the *brachistochrone* problem that arguably led to the creation of

¹The discussion in this section is adapted from [68].

optimal control theory and the calculus of variations [72]. The ongoing progress in the availability and capability of mobile robotic systems has seen an increase in interest in robotic minimum-time motion planning problems.

Much work has been done to characterize time-optimal trajectories for car-like systems with limited turning radius, notably the “Dubins vehicle,” which has a constant forward speed, and the Reeds-Shepp vehicle, which also allows for reverse motion. See [69] for a review, and [71] for detailed derivation of optimal paths with fixed terminal headings. Recently geometric methods were applied to minimum-time problems for differential drive vehicles with limited wheel speeds in [7], and to extremal trajectories for more general constraints in [6]. These methods have also been applied to minimum-time trajectories for omni-directional wheeled robots [5], and minimum-wheel-rotation trajectories for the differential drive [14].

The solutions for a minimum-time problem in a given system are closely related to solutions for differential games featuring agents with those dynamics. In the case that the evader remains motionless, the optimal pursuer strategy is exactly equivalent to the problem of reaching the location of the evader in minimum time. We utilize this relationship in our analysis of a minimum-time problem and differential game both featuring an agent with constraints on speed, turning rate, and lateral acceleration in Chapters 3 and 4.

1.1.3 Pursuit and evasion with multiple evaders²

Many of the approaches used to study one-on-one pursuit and evasion systems can be adapted to include multiple evaders.

The original definition of a differential game assumes two players with opposite aims: the pursuer minimizes the objective function and the evader maximizes it. One approach that has been used to accommodate a system with multiple evaders is to consider the multiple evaders as a single team working together with all of their controls dictated by a single leader.

²The discussion in this section is adapted from [67].

Systems with multiple evaders have been studied for various objective functions, including the time to capture all evaders [12, 51]. In the differential game of “cooperative defense” of [32, 31], an objective function is defined such that the evader team aims to surround the pursuer before capture.

The system of “cyclic pursuit” is one in which multiple agents each act as both pursuer and evader, each individual targets the next one, forming a loop. With the appropriate choice of control law, such as constant bearing [33], limit cycles and relative shape equilibria can be achieved.

The study of pursuit and evasion for multi-agent systems is often motivated by problems in wildlife management, where sensitivity of animal group behavior to environmental change needs to be addressed, and by problems in distributed control of mobile robotic networks, where coordination can be advantageous in maneuvers that require approaching or avoiding a directed signal.

Predator avoidance has long been considered a key factor in animal group formation, first studied mathematically for identical evaders on the plane in the “selfish herd” model of Hamilton [39]. In the selfish herd model, a predator randomly chooses a point on the plane and targets the individual nearest that point. Hamilton shows that individuals can reduce their risk of capture by moving closer to their neighbors. The selfish herd model has been extended to include evolutionary dynamics, which lead to formation of large groups [62, 83]. Several numerical studies have examined properties of group motion in multiple evader systems where biologically inspired strategies are chosen a priori: on the plane [43, 48, 86], in discrete space [75], in three dimensions [76], with multiple pursuers [1], and with strategies based on observations of crabs and shorebirds [77]. Non-spatially explicit game theoretic models of multiple evader systems have been posed for both homogeneous evaders [19], and heterogeneous evaders [25, 26]. Recent studies of evasive behavior in different animal species are reviewed in [23].

1.1.4 Field studies of pursuit in nature

There has been much interest in recent years in understanding the hunting strategies used by large mammalian predators. The development of low cost and high resolution Global Positioning System (GPS) technology allows for unprecedented accuracy in the measurement of animal trajectories; however, limited battery capacity and the difficulty of deploying collars on multiple animals necessitate trade offs in the design of experiments. One main approach has been to place a single collar on a predator and record with high temporal resolution to understand individual pursuit instances. A second main approach has been to place multiple collars with low measurement frequency (but longer battery life) to study trends in habitat use.

Cheetah hunting behavior has been studied with combination GPS and accelerometer sensors [81, 82, 40, 35]. Puma energetics during attacks were characterized by accelerometer measurements [80]. A GPS study on lions and plains zebras showed that zebras use lower quality grazing areas to avoid lions [9].

A hybrid approach can provide rich data as well, as in [52] where aerial LiDAR (Light Detection And Ranging) scans of foliage were combined with GPS collars on lions to study the effect of prey visibility on hunting strategies in different habitats.

Collecting direct video footage of an event allows the researcher to track all the members of an animal group at once. However, processing the raw video data to extract quantitative trajectories can be significantly more complex than with GPS sensors. In certain controlled environments, automated tracking software can be used to automatically detect positions and poses of animals within each video frame. For an extensive review of automated tracking in biology, see [21]. Outdoor video tracking studies have focused on aerial species, in part because of the high contrast between individuals and the sky facilitates automatic tracking. However the use of multi-view video systems necessary to reconstruct three dimensional trajectories poses additional issues. Examples include studies of collective flight trajectories in large groups of starlings [8] and bats [46]. The evasive actions of schools of zebrafish in

response to an artificial predator have also been studied, for a predator moving on a straight line path [70], and a predator controlled through a real-time feedback control system based on visual tracking of the fish [73].

1.2 Outline

In Chapter 2 we present a nonlinear three agent model for a single pursuer and two evaders inspired by predation by bears of calves in caribou herds. We identify the pursuer as a bear and the two evaders as a caribou mother and her calf. The three agents are given equal speeds, and we consider the pursuer to be successful if it is able to reach a point on the line connecting the two evaders, in effect separating the calf from its mother. Strategies are defined such that a single parameter for each agent represents a trade off between two goals. The pursuer trades off between targeting each of the two evaders for classical pursuit. The evaders each trade off between a herding strategy of approaching each other, and an evasion strategy of moving away from the pursuer. We analyze the system dynamics in terms of relative shape variables and derive conditions for stability of relative equilibria. In the case of the mother using a pure evasion strategy we prove conditions such that the bear cannot come between the mother and the calf.

In Chapter 3 we present a more complicated motion model of a steered agent with constraints on speed, angular turning rate, and lateral acceleration. As a first step towards understanding pursuit and evasion under these motion constraints, we analyze the problem of reaching a point on the plane in minimum time with free terminal heading angle. We use Pontryagin's minimum principle to derive extremal controls and characterize families of optimal trajectories. We solve for open-loop control switching times for any destination point on the plane. In addition, we present a state-feedback formulation of the optimal control based on the relative position of the destination in a body-fixed frame. Special cases

for relaxed and extreme values of the lateral acceleration constraint are also considered, with minimum-time trajectories derived in both open-loop and state-feedback form in each case.

Chapter 4 presents the turning-constrained motion model of Chapter 3 in the context of pursuit and evasion. We use the theory of differential games to derive optimal strategies in a one-on-one game for an omnidirectional pursuer with limited speed against an evader with the motion model of Chapter 3. The strategies of the one-on-one game are used as building blocks to derive strategies in a system with a single pursuer and a group of evaders with heterogeneous motion constraints. In the multiple-evader system we propose for the pursuer a strategy of *optimal target selection* such that the evader that could be captured in minimum time in a one-on-one setting is chosen as the target. For evaders, we prove that in the case of all-to-all sensing any evader that is not the current target can always choose from a set of *reactive evasion* control inputs in order to avoid capture. The currently targeted evader must use the optimal evasive strategy from the one-on-one game to delay its capture for as long as possible. We also consider the case in which the pursuer and evaders have radius-limited sensing. We relax the constraint on evader turning rate, and we propose a local strategy of *risk reduction*. We prove that any evader that is not the target can avoid capture using the risk reduction strategy.

Chapter 5 discusses a field experiment conducted in 2014 to study the evasive behaviors of plains zebra in response to an artificial predator.

We conclude with a discussion of future directions in Chapter 6.

Chapter 2

Three agent model of pursuit, evasion, and herding*

2.1 Introduction

The mathematical study of pursuit and evasion has a rich history stretching back to the 18th century with Bouguer’s famous pirate ship pursuit problem [56]. Although the basic concept of pursuit and evasion is not new, it still remains an active area of study due to the multitude of possible applications and variations. The study of pursuit and evasion for multi-agent systems is motivated by problems in wildlife management, where sensitivity of animal group behavior to environmental change needs to be addressed, and by problems in distributed control of mobile robotic networks, where coordination can be advantageous in maneuvers that require approaching or avoiding a directed signal.

Feedback laws can be used to great effect in describing the interactions among animals, both in herding and pursuit behaviors. One-on-one pursuit and evasion has been studied extensively for steered-particle systems with feedback control laws: recent studies look at “motion camouflage” strategies employed by dragonflies and bats [45], cyclic pursuit [33], and

*The content of this chapter appears verbatim in [66].

pursuit and evasion strategy selection as an evolutionary game [78, 59]. Voronoi diagram approaches have been useful for systems with multiple pursuers [4], and in strategies for trapping an evader in a limited environment [42]. Systems with two evaders have been studied under the framework of differential game theory, in the “successive pursuit” of [12], and more recently the “cooperative defense” of [32, 31].

As a step towards our goal of developing a framework to examine the role of herding among evaders in pursuit-evasion dynamics, we focus in the present chapter on studying how a heterogeneous system with three agents behaves under a combination of feedback control laws inspired by the cohesion/repulsion (herding) feedback rules of [60], and pursuit and evasion feedback rules of [78] and [59].

An example from nature of a herding pair of evaders and a single pursuer is a mother caribou and her calf fleeing from a predator. For the woodland caribou (*Rangifer tarandus caribou*) of Northeastern Canada, population growth is heavily influenced by calf mortality due to predation by bear, wolf, coyote, and lynx. In the wild, female caribou usually produce one calf per year, in the early summer. Calves typically stay close to the mother and are most vulnerable to predation in the first month of life as they struggle to keep pace with the adults [11, 54].

The aim of the present work is to define a mathematical model for predation in which (1) a predator (e.g. a bear) pursues one of the two evaders (e.g., a mother or a calf caribou) or a point along the line that connects the two, and (2) each evader chooses a strategy that is a convex combination of evasion and herding. In Section 2.2 of this chapter we discuss a steered-particle model for the bear and caribou that features pursuit, evasion, and herding control laws. In Section 2.3 we introduce a simplified first-order model that captures much of the same behavior. Reduced-order shape dynamics are presented, and classes of equilibria are defined. In Section 2.4 we study the case of the mother using a pure evasion strategy and prove conditions such that the bear cannot come between the mother and the calf. For special parameter values, we provide a stability analysis and analytic solutions. In Section

2.5 we present numerical results that suggest existence of optimal strategies and a way in which the system could be viewed as a differential game.

2.2 Steered-particle model

Our model is motivated by the interactions among a bear, a mother caribou, and her calf. The bear is a pursuer in the classical sense, choosing its target along the line between the two caribou. Each caribou is an evader and a herder; its strategy is a convex combination of classical evasion of the bear and herding with the other caribou.

The equations of motion are based on a steered-particle model: the input u_j controls the angular velocity of agent j . In the case of constant speeds considered here, this is equivalent to choosing the instantaneous curvature of the trajectory. The mother and calf are taken to have unit speed, and the bear has a speed $v \geq 1$. The agents are taken to evolve on the complex plane, so system states are the position vectors $\mathbf{r}_j \in \mathbb{C}$ and heading angles θ_j measured counterclockwise from the real axis, as shown in Fig. 2.1. The equations of motion

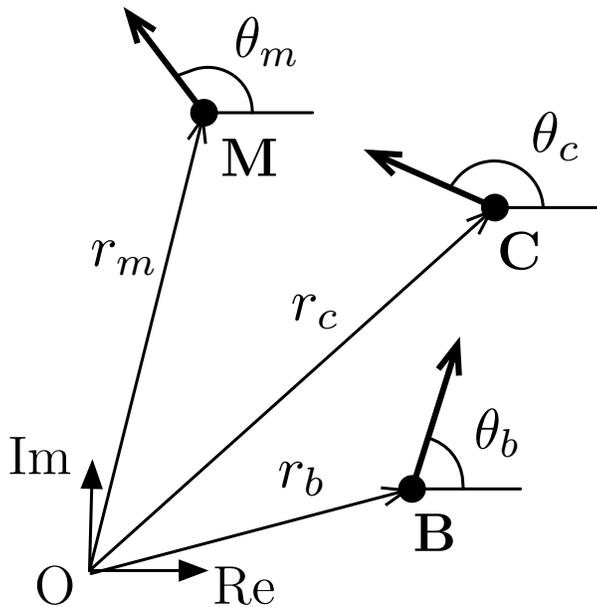


Figure 2.1: System states for steered-particle model of bear (B), mother (M), and calf (C).

are given by

$$\begin{aligned}
\dot{\mathbf{r}}_b &= v e^{i\theta_b}, & \dot{\theta}_b &= u_b, \\
\dot{\mathbf{r}}_m &= e^{i\theta_m}, & \dot{\theta}_m &= u_m, \\
\dot{\mathbf{r}}_c &= e^{i\theta_c}, & \dot{\theta}_c &= u_c.
\end{aligned} \tag{2.1}$$

The feedback control law for the bear is taken from the ‘‘classical pursuit’’ law of [78]:

$$\begin{aligned}
u_b &= -\eta \left\langle \frac{\mathbf{r}_p}{|\mathbf{r}_p|}, i e^{i\theta_b} \right\rangle, \\
\text{with } \mathbf{r}_p &= \mathbf{r}_b - [w_t \mathbf{r}_c + (1 - w_t) \mathbf{r}_m],
\end{aligned} \tag{2.2}$$

where \mathbf{r}_p is the vector from the target point to the bear’s position. The target point is a point along the line connecting the calf and the mother parameterized by a target parameter, $w_t \in [0, 1]$; when $w_t = 1$, the bear targets the calf, when $w_t = 0$ it targets the mother, and for intermediate values of w_t the bear targets a point between the two. $\eta > 0$ is a gain which, when high enough, guarantees convergence to a ‘‘pursuit manifold’’ in finite time as discussed in [78].

Feedback laws for the mother and calf are taken to be a convex combination of evasion from the bear and herding with the other caribou. The weights $w_m, w_c \in [0, 1]$ represent the mother’s and calf’s reliance on herding, respectively:

$$\begin{aligned}
u_m &= w_m u_{m,herd} + (1 - w_m) u_{m,evade} \\
u_c &= w_c u_{c,herd} + (1 - w_c) u_{c,evade}.
\end{aligned} \tag{2.3}$$

The evading rule $u_{j,evade}$ is the ‘‘classical evasion’’ law from [59]:

$$u_{j,evade} = -\eta \left\langle \frac{\mathbf{r}_b - \mathbf{r}_j}{|\mathbf{r}_b - \mathbf{r}_j|}, i e^{i\theta_j} \right\rangle. \tag{2.4}$$

The herding rule comes from the repulsion-orientation-attraction (ROA) laws for group motion presented in [60], with concentric non-overlapping zones (see also [36]). When the distance between calf and mother is less than repulsion radius r_r , a repulsion rule is used. With only two agents herding, the repulsion rule takes the form of the evasion rule (2.4). For a distance greater than r_r but less than orientation radius r_o , an orientation rule is used which steers the two agents towards alignment. For a distance greater than r_o but less than attraction radius r_a , an attracting rule is used, which, for only two agents herding, takes the form of the pursuit rule (2.2). For a distance greater than r_a , no interaction occurs and $u_{j,herd} = 0$. For the calf, then, the herding feedback rule is

$$u_{c,herd} = \begin{cases} -\eta \left\langle \frac{\mathbf{r}_m - \mathbf{r}_c}{|\mathbf{r}_m - \mathbf{r}_c|}, ie^{i\theta_c} \right\rangle, & \text{if } |\mathbf{r}_c - \mathbf{r}_m| < r_r, \\ \left\langle ie^{i\theta_c}, e^{i\theta_m} \right\rangle, & \text{if } r_r \leq |\mathbf{r}_c - \mathbf{r}_m| < r_o, \\ \eta \left\langle \frac{\mathbf{r}_m - \mathbf{r}_c}{|\mathbf{r}_m - \mathbf{r}_c|}, ie^{i\theta_c} \right\rangle, & \text{if } r_o \leq |\mathbf{r}_c - \mathbf{r}_m| < r_a, \\ 0, & \text{if } |\mathbf{r}_c - \mathbf{r}_m| \geq r_a. \end{cases} \quad (2.5)$$

The herding rule for the mother $u_{m,herd}$ is the same with indices c and m switched.

For the remainder of this chapter, the bear's velocity is taken to be the same as the caribou ($v = 1$), in order to study steady-state behavior of the system.

2.3 Simplified first-order model

In [78] it is shown that under the classical pursuit steering law (2.2) with high enough gain, the ‘‘pursuit manifold’’ of the system is reachable within finite time, such that the pursuer will be traveling in a direction directly towards the target. Similarly the classical evasion control law will bring the states towards a corresponding evasion manifold where the evader travels directly away from the pursuer.

If we limit the caribou herding interaction to just the attraction mode, the caribou control law becomes a linear interpolation between classical evasion of the bear and classical pursuit of the other caribou. With high gains on the control inputs, the heading dynamics will quickly settle on the desired directions. We may simplify the model by eliminating the heading dynamics, instead taking the instantaneous desired direction of each agent as the control input.

In this way, the parameters w_t , w_c , and w_m serve to interpolate between two different direction vectors. For convenience, the relative vector between agent positions is written as $\mathbf{r}_{jk} = \mathbf{r}_j - \mathbf{r}_k$. The bear's parameter w_t defines its direction of travel as a convex combination of the direction towards the calf (\mathbf{r}_{cb}) and the direction towards the mother (\mathbf{r}_{mb}). The calf's parameter w_c defines its direction of travel as a convex combination of the direction towards the mother (\mathbf{r}_{mc}) and the direction away from the bear (\mathbf{r}_{cb}). Similarly, the mother's parameter w_m defines its direction of travel as a convex combination of the direction towards the calf (\mathbf{r}_{cm}) and the direction away from the bear (\mathbf{r}_{mb}). The equations of motion for the three agents thus become

$$\begin{aligned}\dot{\mathbf{r}}_m &= \hat{\mathbf{u}}_m, \\ \dot{\mathbf{r}}_b &= \hat{\mathbf{u}}_b, \\ \dot{\mathbf{r}}_c &= \hat{\mathbf{u}}_c,\end{aligned}\tag{2.6}$$

with unit-length direction inputs

$$\begin{aligned}\hat{\mathbf{u}}_m &= e^{i[w_m \angle \mathbf{r}_{cm} + (1-w_m) \angle \mathbf{r}_{mb}]}, \\ \hat{\mathbf{u}}_b &= e^{i[w_t \angle \mathbf{r}_{cb} + (1-w_t) \angle \mathbf{r}_{mb}]}, \\ \hat{\mathbf{u}}_c &= e^{i[w_c \angle \mathbf{r}_{mc} + (1-w_c) \angle \mathbf{r}_{cb}]}.\end{aligned}\tag{2.7}$$

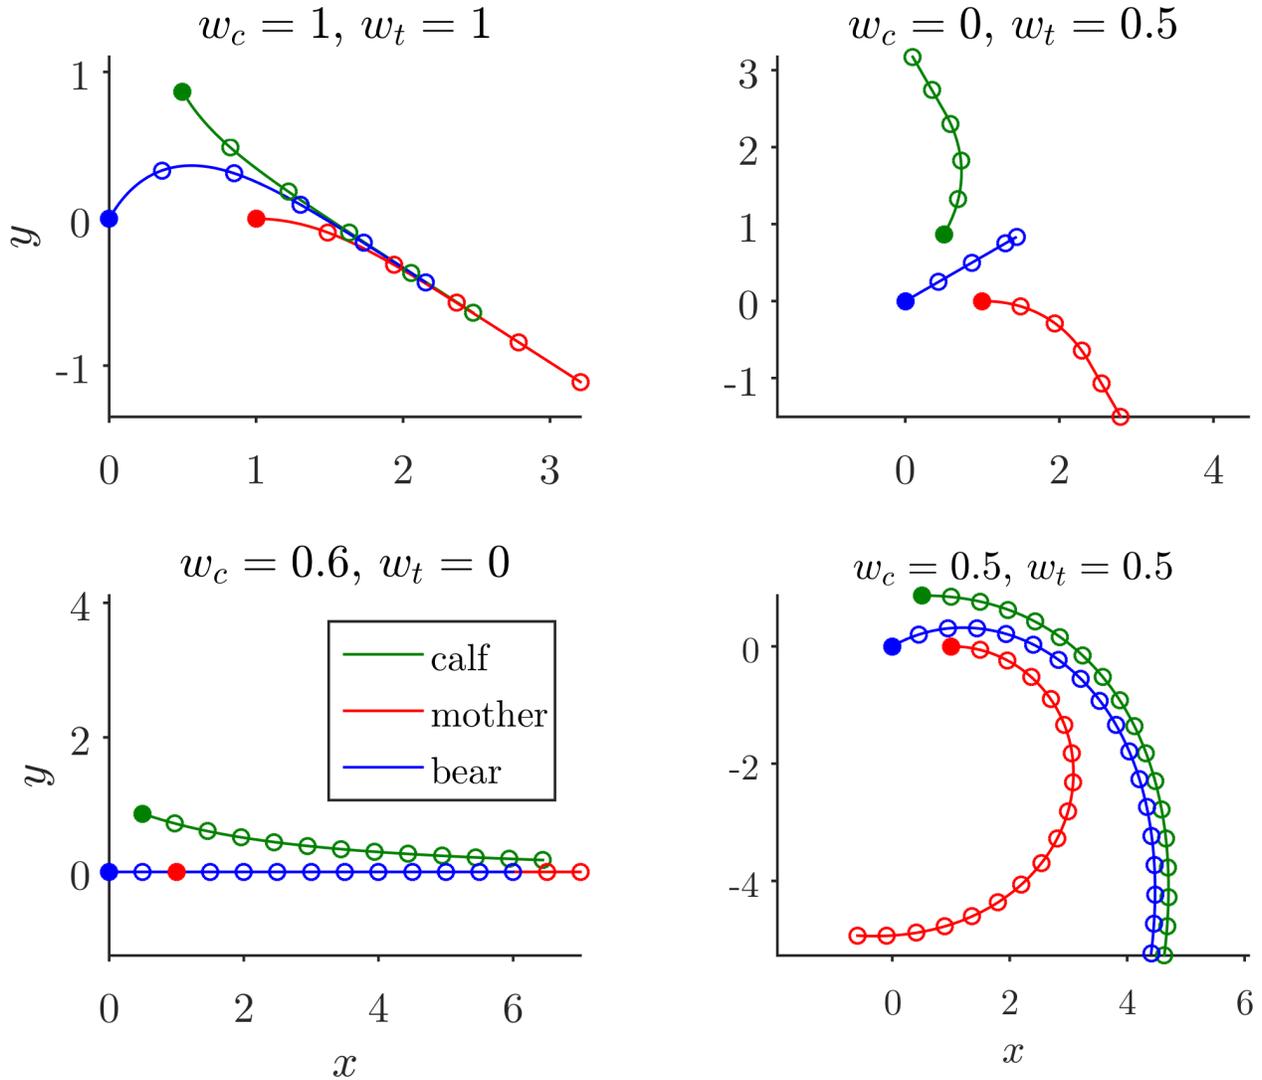


Figure 2.2: Simulated trajectories in the simplified first-order model for various parameter values illustrating different end conditions. Initial conditions (shown as filled circles) are identical for all cases, with the three agents equidistant from each other. All agents have unit speed, and open circles denote locations at 0.5 s intervals. Top left: reaches B–C–M “safe calf” configuration. Top right: M–B–C “splitting” configuration. Bottom left: “parallel motion” configuration. Bottom right: M–B–C “splitting” configuration following a looping path.

It is important to note that the control laws become undefined when any two agents are coincident, due to their dependence on the heading angles of the relative vectors \mathbf{r}_{jk} .

Figure 2.3 illustrates the effect of the control weighting parameters on the agent trajectories. Under identical initial conditions different weights can lead to positive and negative outcomes for the calf.

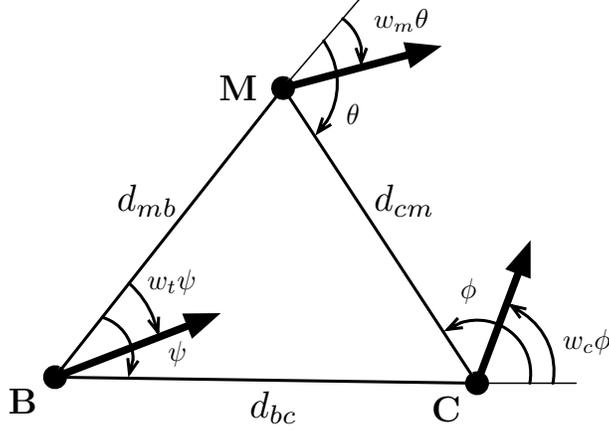


Figure 2.3: Shape variables $\psi, \phi, \theta, d_{bc}, d_{cm}, d_{mb}$ describing relative positions of the bear (B), the mother caribou (M) and the caribou calf (C), along with directions of motion as determined by parameters $w_t, w_c,$ and w_m .

Since the control laws rely solely on the relative vectors between the three agents, the relative motion of the agents is independent of the positions of the agents with respect to a global coordinate frame. Thus it is possible to reduce the order of the system further and directly study the dynamics of “shape variables,” which describe the triangle formed by the agents, as shown in Fig. 2.3.

By the geometry of the triangle, the three angles ψ, ϕ, θ are related by

$$\pi + \psi = \phi + \theta, \quad (2.8)$$

and the side lengths d_{bc}, d_{cm}, d_{mb} can be related to the angles by the law of sines:

$$\frac{d_{bc}}{\sin(\theta)} = \frac{d_{cm}}{\sin(\psi)} = \frac{d_{mb}}{\sin(\phi)}. \quad (2.9)$$

2.3.1 Length dynamics

By projecting the motion of the agents along each of the sides of the triangle, we can find how the side lengths change over time:

$$\begin{aligned}
 \dot{d}_{bc} &= \cos(w_c \phi) - \cos((1 - w_t)\psi), \\
 \dot{d}_{cm} &= -\cos((1 - w_m)\theta) - \cos((1 - w_c)\phi), \\
 \dot{d}_{mb} &= \cos(w_m \theta) - \cos(w_t \psi).
 \end{aligned} \tag{2.10}$$

2.3.2 Angle dynamics

By a similar method, we can project the motion of each agent along the axis normal to a side to find how that side of the triangle rotates in time. Then the dynamics of the individual angles are given by the difference between the rotation of the sides constituting each angle:

$$\begin{aligned}
 \dot{\psi} &= [\sin(w_t \psi) - \sin(w_m \theta)] / d_{mb} - [\sin(w_c \phi) - \sin((1 - w_t)\psi)] / d_{bc}, \\
 \dot{\phi} &= [\sin((1 - w_c)\phi) - \sin((1 - w_m)\theta)] / d_{cm} - [\sin(w_c \phi) - \sin((1 - w_t)\psi)] / d_{bc}, \\
 \dot{\theta} &= [\sin(w_t \psi) - \sin(w_m \theta)] / d_{mb} - [\sin((1 - w_c)\phi) - \sin((1 - w_m)\theta)] / d_{cm}.
 \end{aligned} \tag{2.11}$$

One must note, however, that a triangle can be defined by three side lengths, two angles and a side length, or two side lengths and an angle, but not by three angles alone. So to study the behavior of this system one should choose three appropriate variables (e.g. d_{bc} , ψ , and ϕ), and eliminate the others using the constraint equations (2.10)-(2.11) to be left with three first-order equations in three variables.

2.3.3 Parallel motion equilibria

Several classes of “shape equilibria” corresponding to parallel motion of all agents can be found at fixed points of the length dynamics ($\dot{d}_{bc} = \dot{d}_{cm} = \dot{d}_{mb} = 0$). These occur when the following three equations are satisfied:

$$\begin{aligned} w_c \phi &= (1 - w_t)\psi \\ w_m \theta &= w_t \psi \\ (1 - w_c)\phi + (1 - w_m)\theta &= \pi. \end{aligned} \tag{2.12}$$

Since the side lengths are not involved in these equations, any shape configuration that is a similar triangle to an equilibrium configuration will also be an equilibrium configuration.

Adding these three equations yields the angle constraint (2.8): using that constraint to solve for one angle we are left with two linear equations in two angles and three parameters. For any pair of angles, (2.12) gives a one-parameter family of equilibrium solutions for w_m , w_c , and w_t . Special cases where the three agents are collinear are discussed below:

B-C-M: When the calf is in between the bear and mother, $\psi = \phi = 0$ and $\theta = \pi$. This configuration is only an equilibrium when $w_m = 0$, such that the mother caribou’s strategy is pure evasion of the bear.

B-M-C: When the mother is in between the bear and calf, $\psi = \theta = 0$ and $\phi = \pi$. In symmetry with the previous case, this configuration is only an equilibrium when $w_c = 0$, such that the calf’s strategy is pure evasion.

M-B-C: When the bear is between the two caribou, $\psi = \phi = \theta = \pi$. This is a parallel motion equilibrium when $w_t = w_m = 1 - w_c$. In this case, the direction of travel is not necessarily along the line formed by the agents (as in the previous two cases), but at an angle of $w_t\pi$ from the line.

The next sections present analysis of the behavior of the dynamics in the case $w_m = 0$.

2.4 Mother caribou in pure evasion

It has been observed that in the heat of a predation event, a mother caribou may make her own safety her priority and focus on evasion, only changing her course to go back for her calf once the threat has passed [20]. The calf is expected to follow, but it does not always do so and may become separated from its mother, which makes the calf very vulnerable. This situation can be modeled with (2.7) by setting the mother's parameter to be $w_m = 0$ (pure evasion).

In this case, the length dynamics simplify to

$$\begin{aligned}\dot{d}_{bc} &= \cos(w_c \phi) - \cos((1 - w_t)\psi), \\ \dot{d}_{cm} &= \cos(\phi - \psi) - \cos((1 - w_c)\phi), \\ \dot{d}_{mb} &= 1 - \cos(w_t \psi),\end{aligned}\tag{2.13}$$

and we can note that $\dot{d}_{mb} > 0$ for all $\psi, w_t \neq 0$.

2.4.1 Avoiding mother-calf separation

We prove conditions on w_c and w_t for the model (2.13) where $w_m = 0$ such that the bear can never come directly between the calf and mother (M-B-C configuration) from initial conditions satisfying $d_{cm} < d_{mb}$. For the same set of initial conditions in the special case that the calf 'ignores' the bear completely and uses pure herding ($w_c = 1$), the bear can never come directly between the calf and mother.

Theorem 2.4.1. *Consider the system (2.13) corresponding to $w_m = 0$, and suppose that $0 < d_{cm} < d_{mb}$ at time $t = t_0$. If w_c, w_t are such that the inequality*

$$-\cos(2\psi) - \cos((1 - w_c)(\pi - \psi)) + \cos(w_t \psi) < 1\tag{2.14}$$

is satisfied for all $\psi \in [0, \pi/2)$, then the system will never reach an M-B-C configuration at any future time $t > t_0$. An M-B-C configuration is defined such that the distances satisfy $d_{cm} = d_{mb} + d_{bc}$ with $d_{mb} > 0$ and $d_{bc} > 0$. In case $d_{bc} = 0$, $d_{cm} = 0$, or $d_{mb} = 0$, the dynamics are assumed to terminate.

In the special case that $w_c = 1$ (the calf ignores the bear) for the same initial conditions, the system will never reach an M-B-C configuration at any time $t > t_0$ for any $w_t \in [0, 1]$.

Proof. By definition, at an M-B-C configuration, $d_{cm} > d_{mb} > 0$. Hence, if we can show that $d_{cm} \leq d_{mb}$ for all time, then we have shown that the system can never reach an M-B-C configuration.

Consider system configurations in which $d_{cm} = d_{mb}$; these form a surface δS that separates the space of length configurations into $S^+ = \{d_{cm} > d_{mb}\}$ and $S^- = \{d_{cm} < d_{mb}\}$. Thus, we are done if we can show that solutions stay in $S^- \cup \delta S$.

On δS , if $d_{bc} = 0$, $d_{cm} = 0$, or $d_{mb} = 0$, then the dynamics are terminated. Where $d_{bc} \neq 0$, $d_{cm} \neq 0$, and $d_{mb} \neq 0$ on this surface, the agents form an isosceles triangle with $\phi = \pi - \psi$ and $\psi \in (0, \pi/2)$ and a line (B-M-C configuration) when $\phi = \pi - \psi$ and $\psi = 0$. Substituting $\phi = \pi - \psi$ in (2.13) gives

$$\begin{aligned}\dot{d}_{cm} &= -\cos(2\psi) - \cos((1 - w_c)(\pi - \psi)), \\ \dot{d}_{mb} &= 1 - \cos(w_t \psi).\end{aligned}\tag{2.15}$$

The inequality condition (2.14) is simply the condition that $\dot{d}_{cm} < \dot{d}_{mb}$ whenever $d_{cm} = d_{mb} \neq 0$ and $d_{bc} \neq 0$.

If $d_{cm} < d_{mb}$ at time $t = t_0$, then the state is in S^- , and by continuity the system cannot reach S^+ without first passing through δS . If the system reaches δS and $d_{cm} = d_{mb} = 0$ or $d_{cb} = 0$, then the dynamics terminate. If the system reaches δS and $d_{cm} = d_{mb} \neq 0$ and $d_{bc} \neq 0$ and the inequality condition (2.14) holds, then $\dot{d}_{cm} < \dot{d}_{mb}$ and the dynamics must

remain in $S^- \cup \delta S$. Thus, the system can never reach an M-B-C configuration when starting with $d_{cm} < d_{mb}$.

In the special case with $w_c = 1$ where the calf uses a pure herding strategy and ignores the bear, the inequality condition (2.14) simplifies to

$$\cos(w_t \psi) - \cos(2\psi) < 2, \quad (2.16)$$

which always holds for $\psi \in [0, \pi/2)$. Thus for $w_c = 1$, the system can never reach an M-B-C configuration when starting with $d_{cm} < d_{mb}$, regardless of the value of the bear's parameter w_t . \square

Fig. 2.4 shows numerical calculations of the range of parameters w_c and w_t for which the inequality (2.14) holds. Note that when $w_c \leq 1/4$ there is no guarantee, for any $w_t \in [0, 1]$, that an M-B-C configuration will be avoided.

2.4.2 Disadvantage of pure evasion by the calf

In the case that both mother and calf use a pure evasion strategy, it can be shown that the calf will always become separated from the mother unless starting from a configuration with $\psi = 0$. When $w_m = w_c = 0$, the dynamics of the system simplify greatly, and can be described in terms of lengths d_{bc} , d_{mb} , and angle ψ by

$$\begin{aligned} \dot{d}_{bc} &= 1 - \cos((1 - w_t)\psi) \\ \dot{d}_{mb} &= 1 - \cos(w_t \psi) \\ \dot{\psi} &= \frac{\sin(w_t \psi)}{d_{mb}} + \frac{\sin((1 - w_t)\psi)}{d_{bc}}. \end{aligned} \quad (2.17)$$

When starting with $\psi = 0$, the system can be in either B-C-M, or B-M-C collinear equilibria. Otherwise for all initial conditions with d_{bc} , d_{mb} , $\psi \neq 0$, the three variables will increase

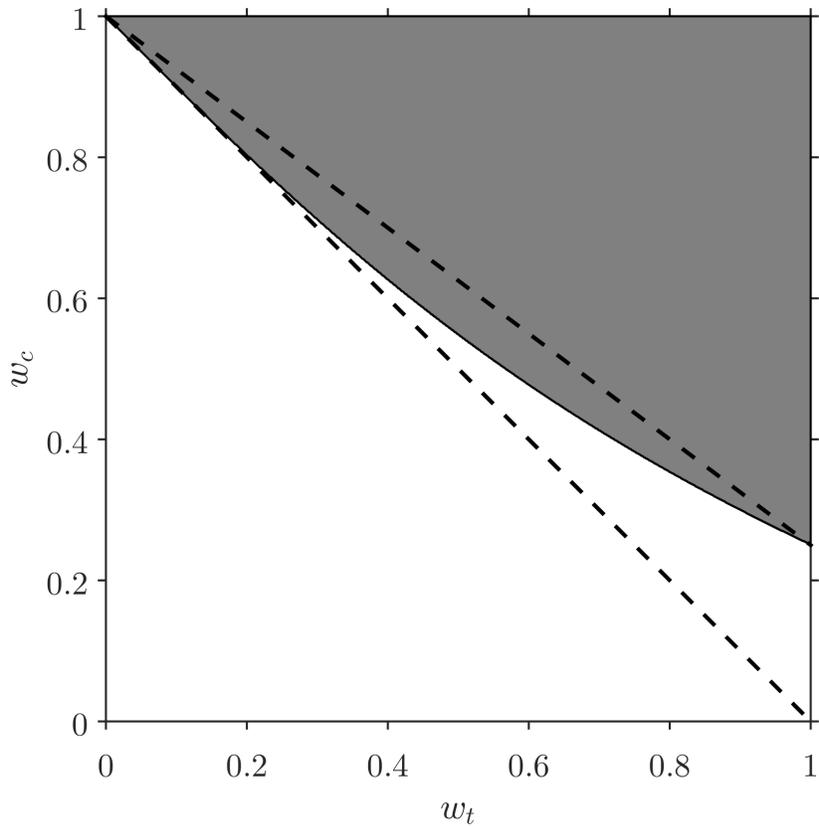


Figure 2.4: Range of parameters w_c and w_t such that the mother and calf caribou cannot be separated by the bear when starting with $d_{cm} < d_{mb}$. The shaded area is the safe range computed from (2.14). Dashed lines $w_c = 1 - w_t$ and $w_c = 1 - \frac{3}{4} w_t$ are shown for reference

monotonically, with ψ eventually approaching π , taking the system towards an M–B–C configuration with the caribou separated from each other by the bear.

2.4.3 When the bear ignores the calf

Fig. 2.4 suggests an advantage for the bear to target the mother rather than the calf in order to separate the calf from its mother. Here, we study the dynamics in the case that $w_m = w_t = 0$, i.e., the bear purely pursues the mother and the mother purely evades the bear. The length dynamics for d_{mb} simplify to $\dot{d}_{mb} = 0$, so the distance between the mother and the bear remains constant throughout the trajectory. This reduces the shape dynamics

to two dimensions, which can be described in terms of ϕ , and ψ with

$$\begin{aligned}\dot{\phi} &= \frac{\sin(\phi)}{d_{mb}} \left[\frac{\sin(\psi) - \sin(w_c \phi)}{\sin(\phi - \psi)} + \frac{\sin((1 - w_c)\phi) - \sin(\phi - \psi)}{\sin(\psi)} \right] \\ \dot{\psi} &= \frac{\sin(\phi)}{d_{mb}} \left[\frac{\sin(\psi) - \sin(w_c \phi)}{\sin(\phi - \psi)} \right].\end{aligned}\tag{2.18}$$

A line of equilibria exists where $\psi = w_c \phi$, which corresponds to parallel motion of all three agents, with the bear directly following the mother and the calf off to the side, or between them.

We compute the Jacobian for points along the equilibria line $\psi = w_c \phi$ by substituting for ψ . Because this is a line of equilibria, one eigenvalue must be zero with its corresponding eigenvector pointing along the line of equilibria. Since the Jacobian is a 2×2 matrix, the other eigenvalue must be given by the trace of the Jacobian, which simplifies to

$$\lambda = \frac{\sin(\phi) [(1 - w_c) \sin(2w_c \phi) - w_c \sin(2(1 - w_c)\phi)]}{2d_{mb} \sin(w_c \phi) \sin((1 - w_c)\phi)}.\tag{2.19}$$

The stability of the equilibria line is thus determined by the sign of this eigenvalue. Since the leading term and each term in the denominator of (2.19) are positive for $w_c \in (0, 1)$ and $\phi \in (0, \pi)$, we need only consider the expression in brackets:

$$f(w_c, \phi) = (1 - w_c) \sin(2w_c \phi) - w_c \sin(2(1 - w_c)\phi).\tag{2.20}$$

We show for all $\phi \in (0, \pi)$ that $\lambda > 0$ for $w_c \in (0, 1/2)$ and $\lambda < 0$ for $w_c \in (1/2, 1)$. First note that $f(w_c, 0) = 0$. The partial derivative of $f(w_c, \phi)$ with respect to ϕ is given by

$$\frac{\partial f}{\partial \phi} = 4w_c(1 - w_c) \sin(\phi) \sin((1 - 2w_c)\phi).\tag{2.21}$$

This derivative is positive for $w_c \in (0, 1/2)$ and negative for $w_c \in (1/2, 1)$ for all $\phi \in (0, \pi)$. $f(w_c, \phi)$ is monotonically increasing in ϕ for $w_c \in (0, 1/2)$, so $f(w_c, \phi) > f(w_c, 0) = 0$, and

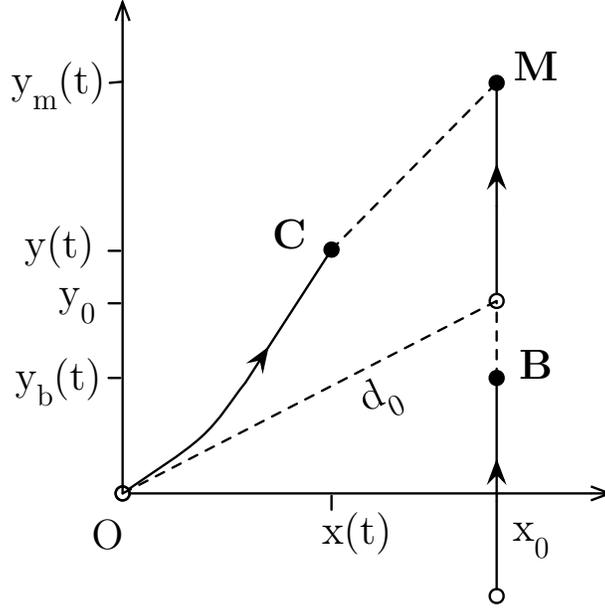


Figure 2.5: Coordinate diagram for trajectories with $w_m = w_t = 0$, $w_c = 1$. Open circles show initial positions, and closed circles show positions after time t .

we can conclude that the eigenvalue λ must be positive in that range. Similarly $f(w_c, \phi)$ is monotonically decreasing in ϕ for $w_c \in (1/2, 1)$, so λ must be negative in that range. Thus the line of parallel-motion equilibria is stable for $w_c \in (1/2, 1)$ and unstable for $w_c \in (0, 1/2)$ for all $\phi \in (0, \pi)$. The change in stability of the parallel equilibria at $w_c = 1/2$ suggests the presence of a local bifurcation.

In the special case of $w_m = w_t = 0$ and $w_c = 1$, the bear and the mother caribou travel in a straight line with constant d_{mb} , and the calf directly pursues the mother. This configuration is equivalent to the classic Bouguer problem of a pirate ship in classical pursuit of a merchant ship that is traveling in a straight line. Analytic solutions for certain initial conditions are presented in [56, 10], and the following steady-state analysis for the caribou system is based closely on the method of [56].

Consider an inertial frame with origin fixed at the initial position of the calf, with the positive y -axis in the direction of the vector from the bear's initial position to the mother's (see Fig. 2.5). Let (x_0, y_0) be the initial coordinates of the mother on this frame, and let $d_0 = d_{cm,0}$ be the initial distance from mother to calf at time $t = t_0$.

Mother and bear both travel along the line $x_m = x_0$, with the mother's position at time t given by

$$y_m(t) = y_0 + t, \quad (2.22)$$

and the bear follows below at a constant distance.

The calf's trajectory traces out a curve $y = f(x)$ such that at each point its tangent will pass through the current position of the mother, since the calf is engaging in classical pursuit. So at each point in time,

$$y' = \frac{dy}{dx} = \frac{y_m - y}{x_0 - x} = \frac{y_0 - y + t}{x_0 - x}. \quad (2.23)$$

The agents move at unit speed, so the arclength of the calf's curve is simply the elapsed time. By solving (2.23) for t and setting it equal to the formula for arclength of the calf's curve we arrive at an integro-differential equation, which, when solved at initial condition $y'|_{x=0} = y_0/x_0$, yields

$$y' = \frac{1}{2} \left(\frac{y_0 + d_0}{x_0 - x} + \frac{x_0 - x}{y_0 + d_0} \right). \quad (2.24)$$

At any given time, the distance from the calf to the mother is

$$d_{cm}^2 = (x_0 - x)^2 + (y_m - y)^2 = (x_0 - x)^2 \left[1 + \left(\frac{y_m - y}{x_0 - x} \right)^2 \right]. \quad (2.25)$$

Recalling (2.23), the final term in the brackets in (2.25) is simply $(y')^2$, and thus we arrive at an equation for d_{cm} as a function of x :

$$d_{cm}^2 = x_0^2 \left[\frac{1}{2} \left(1 - \frac{x}{x_0} \right)^2 + \frac{1}{4} \left(\frac{y_0 + d_0}{x_0} \right)^2 + \frac{1}{4} \left(\frac{x_0}{y_0 + d_0} \right)^2 \left(1 - \frac{x}{x_0} \right)^4 \right]. \quad (2.26)$$

By inspection of Fig. 2.5, it is clear that as t grows large, x approaches x_0 , so the steady-state distance is simply

$$d_{cm,ss} = \lim_{t \rightarrow \infty} d_{cm}^2 = \lim_{x \rightarrow x_0} d_{cm}^2 = \left(\frac{y_0 + d_0}{2} \right)^2, \quad (2.27)$$

or in terms of the initial distances,

$$d_{cm,ss} = \frac{(d_{cm,0} + d_{mb,0})^2 - d_{bc,0}^2}{4 d_{mb,0}}. \quad (2.28)$$

Then for $d_{cm,ss} < d_{mb,0}$, the system will end up in a B–C–M configuration, and for $d_{cm,ss} > d_{mb,0}$, the system will end up in an M–B–C configuration. Setting $d_{cm,ss} = d_{mb,0}$, we can find the locus of initial conditions where the calf ends up directly on the bear.

2.5 Simulations and estimated solution to a two-player game

Trajectories were computed in Matlab using the forward-Euler method of integration on the first-order model (2.6)-(2.7). Leaving the initial distance from bear to calf constant at $d_{bc} = 10$, the initial position of the mother was varied over an equally-spaced grid on the circular area centered at the calf with radius 10, (i.e. only initial conditions where $d_{cm} \leq d_{bc}$).

Trials were run with $\Delta t = 0.1$ s for 2500 s, or until the agents reached a collinear configuration. “Capture” of the calf was taken to occur when $d_{bc} \leq 1$ or $\psi \geq \pi/2$ at the final time, with the assumption that in the wild a calf separated from its mother would eventually tire and be captured.

Leaving $w_m = 0$ constant (mother using pure evasion), the parameters w_c and w_t were varied over the range $[0,1]$ in increments of 0.05. The percentage of trials ending in capture for each particular w_c, w_t combination gives a measure of the “fitness” of those strategies against one another, which can be seen in Fig. 2.6.

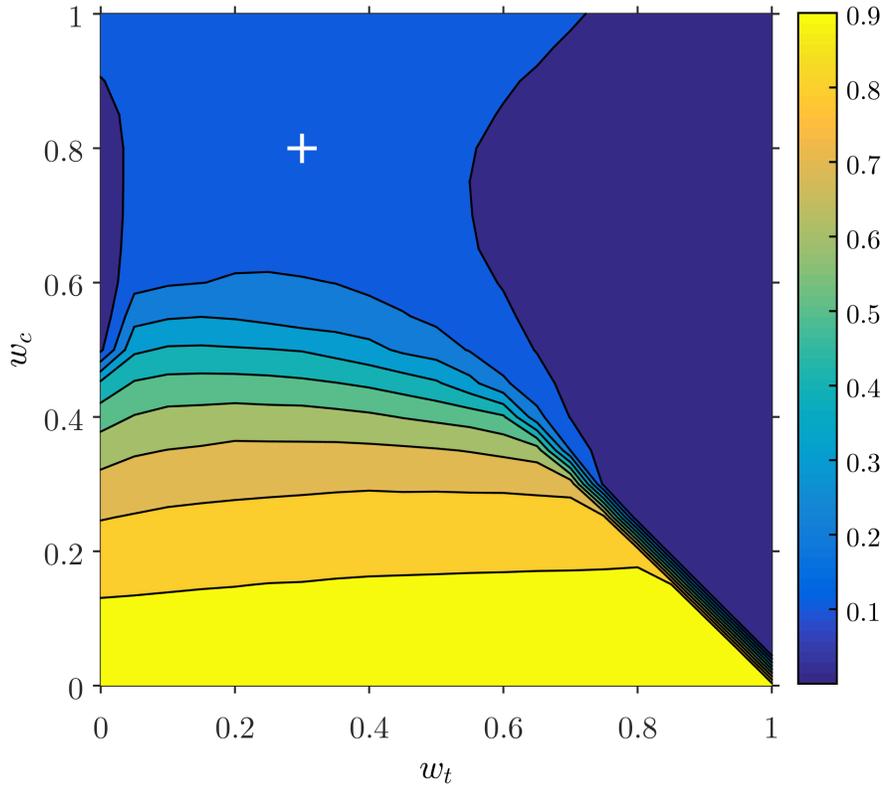


Figure 2.6: Fraction of capture events for different combinations of w_t and w_c , with $w_m = 0$ for initial conditions satisfying $d_{cm} < d_{bc}$. The estimated location of the saddle point is denoted by the cross at $w_t = 0.3$, $w_c = 0.8$, which has a capture fraction of 0.137.

Consider a zero-sum game in which the bear and calf each choose a value for their control parameter before the system is run with $w_m = 0$ for an unknown random initial condition from the set described above, with $d_{bc} = 10$ and the location of the mother chosen randomly from the circle about the calf of radius 10. Here the bear aims to choose a value of w_t that maximizes the chance of capture, and the calf aims to choose a value of w_c that minimizes the chance of capture. A saddle point in this w_t - w_c strategy space shown in Fig. 2.6 represents a solution to the game: equilibrium strategies for the two players such that each individual would decrease its fitness by choosing a different strategy. From the simulations, we estimate a saddle point at $w_t = 0.3$, $w_c = 0.8$, with a capture fraction of 0.137.

In the present chapter we presented and analyzed a system of pursuit and evasion where agent strategies were chosen a priori based on inspiration from biological systems, with

explicit parameters representing the tradeoff between evasion and herding behavior. In Chapter 4, we take a different approach, considering a system of pursuit and evasion with multiple evaders from an optimal control perspective. Using the framework of the *differential game*, we derive optimal strategies in the form of state-feedback control laws dependent on agent relative positions and constraints on individuals' speed and turning rate. In the case of agents with limited sensing, we show that self-interested evaders can use a strategy of approaching a slower neighboring agent to reduce the risk of capture, which parallels our result here from Theorem 2.4.1 that showed how the calf agent can benefit from using a pure herding strategy.

Chapter 3

Minimum-time trajectories for steered agent with constraints on speed, lateral acceleration, and turning rate*

3.1 Introduction

Minimum-time problems have long been a source of fascination for the mathematics community, reaching as far back as the *brachistochrone* problem that arguably led to the creation of optimal control theory and the calculus of variations[72]. The ongoing progress in the availability and capability of mobile robotic systems has seen an increase in interest in robotic minimum-time motion planning problems.

Much work has been done to characterize time-optimal trajectories for car-like systems with limited turning radius, notably the “Dubins vehicle,” which has a constant forward speed, and the Reeds-Shepp vehicle, which also allows for reverse motion. See [69] for a review, and [71] for detailed derivation of optimal paths with fixed terminal headings.

*This chapter is adapted from a paper by W.L. Scott and N.H. Leonard that has been submitted to the ASME Journal of Dynamic Systems, Measurement and Control for review, with additional content in Section 3.7 adapted from [68].

In this chapter, we solve the problem of reaching a desired point on the plane in minimum time for a “steered agent” agent with constraints on forward speed, angular turning rate, and lateral acceleration. Our model differs from both the Dubins and Reeds-Shepp models in that it allows for the agent to rotate in place with zero forward speed.

Here the lateral acceleration constraint creates coupling between the two inputs of speed and turning rate, such that the agent must slow down to achieve a higher turning rate and vice versa. This constraint is chosen with regards to legged locomotion. A study of the kinematics of horses during polo games and track racing[74] indicates that grip strength and limb force limits constrain the maximum lateral acceleration during a turn, such that the horses must decrease their speed in anticipation of tight turns.

We are motivated by the desire to create a movement model with meaningful parameters that is analytically tractable for a single agent and applicable in the study of terrestrial animal motion as well as robotic motion. We aim to apply our results for individual optimal trajectories as building blocks for improving models of collective motion, such as the model of group evasion from a pursuer presented in [67], which does not include limits on turning rate.

Our analysis is based closely on the geometric methods of Balkcom and Mason, which were applied to differential drive vehicles with limited wheel speeds in [7], and to extremal trajectories for more general constraints in [6]. These methods have also been applied to minimum-time trajectories for omni-directional robots[5], and minimum wheel-rotation for the differential drive[14].

In Section 3.2 we present the formal problem statement and system equations of motion. Section 3.3 derives extremal control inputs according to Pontryagin’s minimum principle. In Section 3.4 we prove conditions on the possible families of optimal trajectories, and in Section 3.5 we solve for open-loop control switching times for all cases. Section 3.6 presents a state-feedback formulation of the optimal control based on the relative position of the

destination in a body-fixed frame. Finally we examine special limiting cases of the lateral acceleration parameter in Section 3.7 and conclude in Section 3.8.

3.2 Problem statement and system dynamics

We consider the problem of the minimum-time trajectory for a steered agent on the plane to reach a desired destination point $(x_1, y_1)^T \in \mathbb{R}^2$ with no fixed final heading. The state of the agent consists of its position $(x, y)^T \in \mathbb{R}^2$ and heading angle $\theta \in \mathbb{S}$, evolving under the equations of motion,

$$\dot{x} = v \cos \theta,$$

$$\dot{y} = v \sin \theta,$$

$$\dot{\theta} = \omega.$$

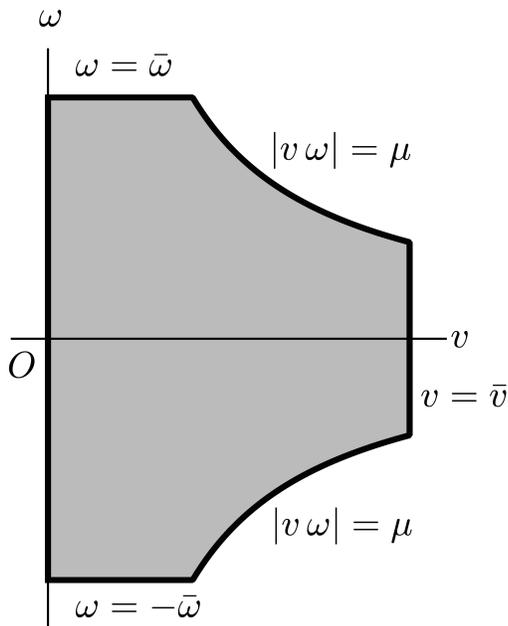


Figure 3.1: Admissible control inputs for steered agent with limits on speed, turning rate and lateral acceleration as described in Sec. 3.2.

The control inputs $\mathbf{u} = (v, \omega)^T$ consist of the forward speed $v \in \mathbb{R}$ and the turning rate $\omega \in \mathbb{R}$. We impose the following constraints on the control input:

1. *Limited speed:* Let $\bar{v} > 0$ be the maximum speed. The speed control must satisfy $v \leq \bar{v}$ for all time.
2. *No reverse motion:* Speed must satisfy $v \geq 0$ for all time, such that the agent never moves in reverse.
3. *Limited turning rate:* Let $\bar{\omega}$ be the maximum turning rate. Then the turning control must satisfy $|\omega| \leq \bar{\omega}$ for all time.
4. *Limited lateral acceleration:* Let μ represent the maximum lateral acceleration (turning traction limit). The inputs v and ω must satisfy $|v\omega| \leq \mu$ for all time. We assume that $\mu < \bar{v}\bar{\omega}$ so that the lateral acceleration constraint is active on part of the boundary of the control domain.

Let Ω be the set of all admissible inputs \mathbf{u} satisfying the above constraints, as shown in Fig. 3.1.

3.3 Extremal trajectories from Pontryagin’s minimum principle

To solve for the optimal trajectories we begin by using Pontryagin’s minimum principle to find families of *extremal* trajectories that satisfy necessary conditions on optimality. We follow the method used in [7], which solved for optimal trajectories for differential drive robots with the same equations of motion as the current system, but with different constraints on the inputs. This leads to different switching functions and extremal controls. In the subsequent section we use boundary constraints to characterize which of the extremal trajectories are optimal under different conditions.

The single agent minimum-time problem is described by

$$t_1 = \min_{\mathbf{u}(t) \in \Omega} \int_0^T 1 dt \quad \text{s.t.} \quad \dot{\mathbf{q}}(t) = \mathbf{f}(\mathbf{q}, \mathbf{u}), \quad \psi(t_1) = 0,$$

where $\mathbf{q} = (x, y, \theta)^T$, $\mathbf{f}(\mathbf{q}, \mathbf{u}) = (v \cos \theta, v \sin \theta, \omega)^T$, and $\psi(t) = (x(t) - x_1)^2 + (y(t) - y_1)^2$, with initial conditions $\mathbf{q}(0) = (x_0, y_0, \theta_0)^T$.

The control Hamiltonian for the system dynamics is

$$\begin{aligned} H(\boldsymbol{\lambda}, \mathbf{q}, \mathbf{u}) &= \boldsymbol{\lambda} \cdot \mathbf{f}(\mathbf{q}, \mathbf{u}) + 1 \\ &= \lambda_x v \cos \theta + \lambda_y v \sin \theta + \lambda_\theta \omega + 1, \end{aligned}$$

where the adjoint vector $\boldsymbol{\lambda} = (\lambda_x, \lambda_y, \lambda_\theta) \in \mathbb{R}^3$ represents the partial derivative of the value of the cost function (in this case the minimum time remaining to reach the destination) with respect to the system state.

Pontryagin's minimum principle states that extremal control $\mathbf{u}^*(t)$ satisfies the following necessary condition:

$$H(\boldsymbol{\lambda}(t), \mathbf{q}(t), \mathbf{u}^*(t)) = \min_{\mathbf{u}(t) \in \Omega} H(\boldsymbol{\lambda}(t), \mathbf{q}(t), \mathbf{u}(t)) = 0.$$

The dynamics of the adjoint vector are governed by $\dot{\boldsymbol{\lambda}} = -\frac{\partial H}{\partial \mathbf{q}}$, which can be written as

$$\begin{aligned} \dot{\lambda}_x &= 0, \\ \dot{\lambda}_y &= 0, \\ \dot{\lambda}_\theta &= \lambda_x v \sin \theta - \lambda_y v \cos \theta. \end{aligned}$$

From this we see that λ_x and λ_y are constant over time. Thus, as noted in [7], since $\dot{\lambda}_\theta = \lambda_x \dot{y} - \lambda_y \dot{x}$, λ_θ can be directly integrated as

$$\lambda_\theta = \lambda_x y - \lambda_y x - \rho,$$

for some constant of integration ρ .

Since the terminal heading θ_f is free, we must impose that $\lambda_\theta(t_1) = 0$. This allows us to solve for ρ : $\lambda_\theta(t_1) = \lambda_x y_1 - \lambda_y x_1 - \rho = 0$. Thus,

$$\lambda_\theta(t) = \lambda_x(y(t) - y_1) - \lambda_y(x(t) - x_1). \quad (3.1)$$

For a nontrivial optimal trajectory, the adjoint vector must be nonzero at all times, implying that λ_x and λ_y cannot both be zero. Without loss of generality, let $\lambda_x = \lambda_0 \sin \gamma$ and $\lambda_y = -\lambda_0 \cos \gamma$ for some unknown angle γ , with the constant $\lambda_0 \neq 0$ computed from the terminal condition

$$H(\lambda(t_1), \mathbf{q}(t_1), \mathbf{u}^*(t_1)) = 0.$$

We next derive the extremal controls. Following the derivation in [7], define the following two functions of the state:

$$\begin{aligned} \eta(x, y) &= (x - x_1) \cos \gamma + (y - y_1) \sin \gamma, \\ \text{and } \beta(\theta) &= \theta - (\gamma + \pi/2). \end{aligned}$$

The level set $\eta = 0$ describes a line in the x - y plane passing through the destination point $(x_1, y_1)^T$. β describes the agent's heading relative to the direction along the $\eta = 0$ line.

Now we can write the Hamiltonian in terms of the inputs and these functions of the state:

$$H = \lambda_0(-v \cos \beta + \omega \eta) + 1.$$

This makes it straightforward to apply Pontryagin’s minimum principle to find extremal controls as a function of the state. Since this is a minimum-time problem with the dynamics linear in the control inputs, the optimal control will be of a *bang-bang* type, always taking values along the control constraint surfaces shown in Fig. 3.1.

3.3.1 Switching functions and generic control inputs

Given our constraints on the control inputs, we need to determine which value of \mathbf{u} will minimize the Hamiltonian for each point in the state space. We follow the same procedure as in [68], except that the additional constraint on lateral acceleration prompts a third switching function.

We define three switching functions that can be used to determine the control input $\mathbf{u}^* = (v^*, \omega^*)$ that will minimize the Hamiltonian for a given state:

$$\phi_1(\mathbf{q}) = \cos \beta(\mathbf{q}), \tag{3.2}$$

$$\phi_2(\mathbf{q}) = -\eta(\mathbf{q}), \tag{3.3}$$

$$\phi_3(\mathbf{q}) = \bar{\omega}|\eta| - \bar{v} \cos \beta(\mathbf{q}). \tag{3.4}$$

Let $\text{sgn}(z)$ be defined as the standard sign function for $z \in \mathbb{R}$:

$$\text{sgn } z = \begin{cases} -1, & z < 0, \\ 0, & z = 0, \\ 1, & z > 0. \end{cases}$$

On time intervals for which the switching functions are nonzero, the corresponding extremal controls are called generic controls. These fall into three categories based on the signs of the switching functions. The generic control inputs along with their corresponding extremal trajectories are as follows, for initial state $\mathbf{q}(0) = (x_0, y_0, \theta_0)^T$. In each case $\theta(t) = \theta_0 + \omega^*t$.

1. *Rotation:* When $\phi_1 < 0$, the agent rotates in place at maximum turning rate: $v^* = 0$ and $\omega^* = \bar{\omega} \operatorname{sgn}(\phi_2)$. Here $\mathbf{q}(t) = (x_0, y_0, \theta_0 + \omega^*t)^T$.
2. *Slow turn:* When $\phi_1 > 0$ and $\phi_3 > 0$, the agent moves forward with low speed while turning at the maximum rate: $v^* = \mu/\bar{\omega}$ and $\omega^* = \bar{\omega} \operatorname{sgn}(\phi_2)$. The agent moves on a circular arc with radius $R_s = \mu/\bar{\omega}^2$, with

$$\begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} + B(\theta_0) \begin{pmatrix} R_s \sin(\bar{\omega}t) \\ \operatorname{sgn}(\phi_2)R_s(1 - \cos(\bar{\omega}t)) \end{pmatrix}, \quad (3.5)$$

where $B(\theta)$ is the standard rotation matrix

$$B(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}.$$

3. *Fast turn:* When $\phi_1 > 0$ and $\phi_3 < 0$, the agent moves forward at maximum speed while turning at a lower rate: $v^* = \bar{v}$ and $\omega^* = \operatorname{sgn}(\phi_2)\mu/\bar{v}$. The agent moves on a circular arc with radius $R_f = \bar{v}^2/\mu$, with

$$\begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} + B(\theta_0) \begin{pmatrix} R_f \sin(\mu t/\bar{v}) \\ \operatorname{sgn}(\phi_2)R_f(1 - \cos(\mu t/\bar{v})) \end{pmatrix}, \quad (3.6)$$

3.3.2 Singular control inputs

At times where some $\phi_i = 0$, there exist multiple inputs that minimize the Hamiltonian. When the state arrives at such a switching surface, the control may have an instantaneous switching if the state instantaneously traverses the switching surface, or an interval of singular control where the state remains on the switching surface for some time interval. We must examine each switching surface separately.

Forward motion

When $\phi_2 = 0$ with $\phi_1 > 0$, at some time t' the agent is on the switching surface between fast turn left and fast turn right. Geometrically, this places the agent on the $\eta = 0$ line with $\cos \beta > 0$, and the Hamiltonian is minimized by $v = \bar{v}$ with ω taking any value in $[-\mu/\bar{v}, \mu/\bar{v}]$. For $\cos \beta \neq 1$, the agent is not aligned with the $\eta = 0$ line. So the positive speed will bring it off of the line at the next moment, which would manifest as an instantaneous switching from fast turn left to fast turn right, or vice-versa, without an extended singular control. In the case that $\cos \beta = 1$, the agent is arriving on the $\eta = 0$ line and tangent to it. Forward motion with $\omega = 0$ would keep it on the line for further time $t > t'$. This represents an interval of singular control, which as we show is part of many time optimal trajectories.

Consider the case that the agent starts with its heading in the direction of the destination point. It is clear that the minimum-time trajectory to reach that point is forward motion at maximum speed. This illustrates that solutions to the minimum-time problem can contain singular intervals of forward motion in this system.

Slow turn and fast turn

When $\phi_3 = 0$ and $\phi_1 > 0$ at some time t' , the agent is on the switching surface between fast turn and slow turn with direction determined by the sign of ϕ_2 . On this switching surface, the Hamiltonian is minimized by two possible values of the control input: a fast turn $\mathbf{u} = (\bar{v}, \pm\mu/\bar{v})^T$ or a slow turn $\mathbf{u} = (\mu/\bar{\omega}, \pm\bar{\omega})^T$ (positive for left turn).

If we can show that the derivative of ϕ_3 is nonzero on the switching surface, that will imply that no singular trajectory segments can exist along that surface. Taking the derivative of ϕ_3 with respect to time, we find

$$\dot{\phi}_3 = \sin \beta (v\bar{\omega} \operatorname{sgn}(\eta) + \omega\bar{v}).$$

Noting that $\text{sgn}(\omega^*) = -\text{sgn}(\eta)$ in all cases, we can substitute to get a necessary condition for slow-fast singular control:

$$\dot{\phi}_3 = \text{sgn}(\eta) \sin \beta (v^* \bar{\omega} - |\omega^*| \bar{v}) = 0.$$

Control values for neither slow turn nor fast turn can make the term in parentheses zero, so the only possibility for a singular control interval is the case where $\sin \beta = 0$. Taking an additional time derivative, we find

$$\ddot{\phi}_3 = \omega \text{sgn}(\eta) \cos \beta (v \bar{\omega} - |\omega| \bar{v}),$$

which is nonzero for both fast and slow turn inputs. Thus $\dot{\phi}_3$ can only vanish for a single instant, and no extended singular interval is possible on the $\phi_3 = 0$ switching surface.

Rotation and slow turn

For states on the switching surface $\phi_1 = 0$, the Hamiltonian is minimized by $\omega^* = \bar{\omega} \text{sgn}(\phi_2)$ with v freely chosen as any value in the range $[0, \mu/\bar{\omega}]$. Taking the derivative of ϕ_1 with respect to time,

$$\begin{aligned} \dot{\phi}_1 &= -\omega \sin \beta \\ &= \pm \omega, \text{ for } \cos \beta = 0. \end{aligned}$$

Thus on the switching surface the derivative is always nonzero for the minimizing control $\mathbf{u} = \mathbf{u}^*$, so there can be no singular interval for the $\phi_1 = 0$ switching surface.

3.3.3 Multiply-singular control

We consider the situation that the state lies on multiple switching surfaces simultaneously. If both ϕ_1 and ϕ_2 (and subsequently ϕ_3) are zero for a given state, then $H = 1$. One of

the conditions for the application of the minimum principle in a minimum-time problem is that the Hamiltonian is constant with $H = 0$. So we can conclude that a minimum-time trajectory can never reach a state that lies on multiple switching surface in this system. The trivial case corresponds to the agent starting at the destination point.

3.4 Families of optimal trajectories

Now that we have enumerated the types of extremal trajectory segments, the task is to show which combination of extremal segments make up the minimum-time trajectory to a given destination point. We examine the possible terminal conditions and integrate backwards in time to find switching conditions compatible with the terminal constraints. In Section 3.5 we show how to reach any point in the plane by one of these “nominal trajectories.”

Theorem 3.4.1. *Minimum-time trajectories must end in either a fast turn or forward motion segment.*

Proof. The terminal condition states that $\phi_2 = 0$ at the final time, and fast turning and forward motion are the only two controls that can bring the trajectory to (or remain on) that surface. □

Theorem 3.4.2. *(Theorem 2 from [68]) If the heading of the agent is aligned with the baseline vector from the agent to its destination at time t' , the minimum-time trajectory for $t > t'$ consists of forward motion at maximum speed.*

Proof. The shortest distance between two points on the plane is a straight line. The agent can move at its maximum speed \bar{v} while traveling in a straight line in the direction of its heading. Thus the minimum-time trajectory to a point on the heading tangent line in front of the agent is forward motion at maximum speed. □

3.4.1 Trajectories ending in forward motion

For a trajectory to end with a forward motion segment, it must have $\eta(t_1) = 0$ and $\cos \beta(t_1) = 1$ at the terminal time t_1 . These conditions will hold for the duration of the forward motion segment, no matter how long it lasts. From Section 3.3.2, we know that a forward motion segment can only be preceded by a fast turn. Suppose that at some time the control switches from fast turn to forward motion. To compute the maximum duration of a fast turn leading to forward motion, we integrate backwards in time to find the switching times corresponding to $\phi_3 = 0$, (see Fig. 3.2 on the left).

Using (3.2) and (3.6) to solve for the value of the switching functions over a fast turn trajectory ending with $\eta = 0$ and $\cos \beta = 1$, we find that $\phi_3 = 0$ for a fast turn of duration $\bar{\tau}_f = \bar{\theta}_f \bar{v} / \mu$, where

$$\bar{\theta}_f = \cos^{-1} \left(\frac{R_f}{R_f + b} \right), \quad \text{with } b = \bar{v} / \bar{\omega}. \quad (3.7)$$

Continuing backwards in time with a slow turn segment using (3.2) and (3.5), we find that $\phi_1 = 0$ for a slow turn of duration $\bar{\tau}_s = \bar{\theta}_s / \bar{\omega}$ with

$$\bar{\theta}_s = \pi/2 - \bar{\theta}_f = \sin^{-1} \left(\frac{R_f}{R_f + b} \right).$$

Continuing further back in time we have a rotation segment. Fig. 3.2 on the left shows the state of the agent relative to the $\eta = 0$ line at the times of control switching for the backwards-in-time trajectory described above.

From Bellman's principle of optimality, we know that subsets of a trajectory at different starting points (but sharing an endpoint) will also be optimal trajectories themselves. So these switching intervals $\bar{\tau}_f$ and $\bar{\tau}_s$ allow us to define the family of all trajectories that end with a forward motion segment of nonzero length.

This family of trajectories consists of all trajectories of the following types (for both left and right turns)

- F : Forward motion only

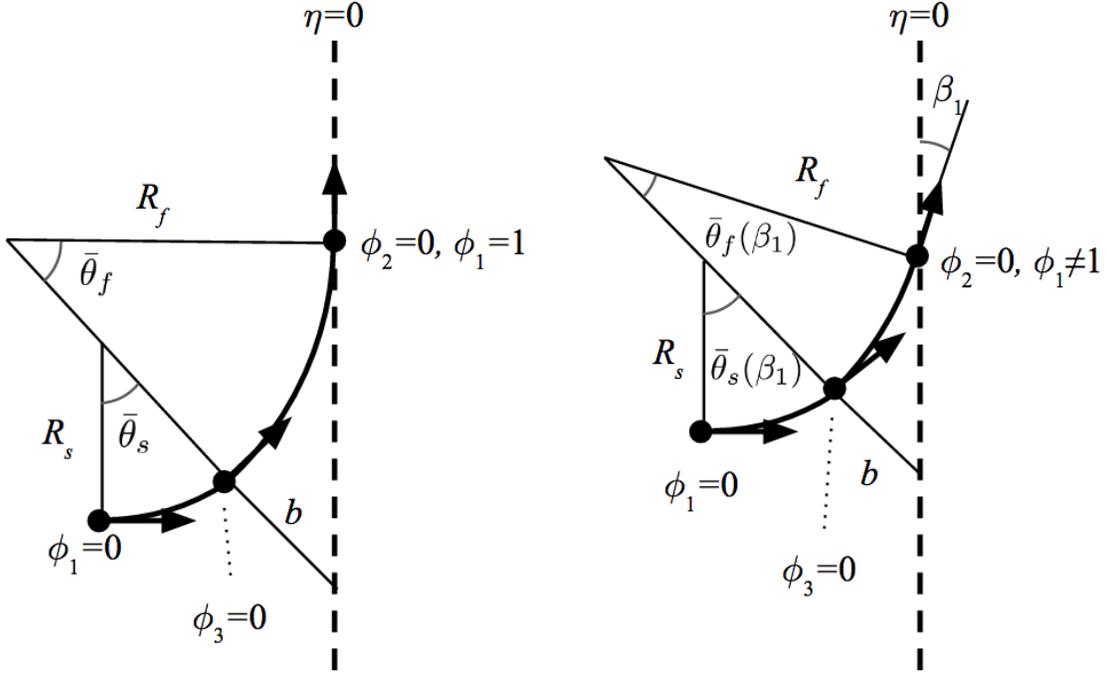


Figure 3.2: Illustration of system state relative to the $\eta = 0$ line (shown dashed) at points where a switching function ϕ_i reaches zero. Left: trajectory ending in forward segment ($\beta_1 = 0$). Right: trajectory ending in fast turn segment.

- $T_f F$: Fast turn of up to $\bar{\tau}_f$ duration followed by some forward motion
- $T_s T_f F$: Slow turn of up to $\bar{\tau}_s$ duration, followed by fast turn of $\bar{\tau}_f$ duration, followed by some forward motion
- $RT_s T_f F$: Rotation, followed by slow and fast turns of duration $\bar{\tau}_s$ and $\bar{\tau}_f$ respectively, followed by some forward motion

Theorem 3.4.3. *The maximum rotation duration for a minimum-time $RT_s T_f F$ trajectory with forward motion duration of τ_d is given by $\bar{\tau}_r = \bar{\theta}_r / \bar{\omega}$, where*

$$\bar{\theta}_r = \pi - \tan^{-1} \left(\frac{\bar{v}\tau_d + R_s + (R_f - R_s) \sqrt{1 - \left(\frac{b}{R_f + b}\right)^2}}{b} \right).$$

Proof. Consider a left-turning trajectory of type RT_sT_fF with a rotation duration of $\tau_r = 0$ and forward motion duration τ_d . Starting from the origin $\mathbf{q}(0) = (0, 0, 0)^T$, we use (3.5) and (3.6) to calculate the endpoint of the trajectory at time $t_1 = \bar{\tau}_s + \bar{\tau}_f + \tau_d$ as

$$\begin{pmatrix} x(t_1) \\ y(t_1) \end{pmatrix} = \begin{pmatrix} b \\ \bar{v}\tau_d + R_s + (R_f - R_s)\sqrt{1 - \left(\frac{b}{R_f+b}\right)^2} \end{pmatrix}.$$

Increasing the duration of the initial rotation segment will cause the endpoint to rotate about the origin, and a duration of $\bar{\tau}_r$ as defined above brings the destination to the negative x -axis. By symmetry, a right-turning trajectory with the same segment durations would bring the agent to that same point in the same amount of time. From this we can determine that a left-turning trajectory with rotation longer than $\bar{\tau}_r$ would put the destination at a point that can be reached in less time with a right-turning trajectory. Thus a minimum-time trajectory of type RT_sT_fF cannot include a rotation segment of duration greater than $\bar{\tau}_r$. Trajectories with rotation segments of duration exactly $\bar{\tau}_r$ correspond to destinations lying directly behind the initial position of the agent. \square

3.4.2 Trajectories ending in fast turn

Trajectories with $\eta(t_1) = 0$ and $\beta(t_1) \neq 0$ end in a fast turn. We can again integrate backwards in time to find families of optimal trajectories, but in this case the switching angles will be a function of the terminal value of $\beta(t_1)$. The state of the agent relative to the $\eta = 0$ line at the times of control switching for this backwards-in-time trajectory is illustrated in Fig. 3.2 on the right.

Let $\beta_1 = \beta(t_1)$. Using (3.2) and (3.6) to solve for the value of the switching functions over a fast turn trajectory ending with $\eta = 0$ and $\cos \beta_1 \neq 1$, we find that the maximum duration of a fast turn segment (corresponding to $\phi_3 = 0$ at the start) is $\bar{\tau}_f(\beta_1) = \bar{\theta}_f(\beta_1)\bar{v}/\mu$

with

$$\bar{\theta}_f(\beta_1) = -\beta_1 + \cos^{-1} \left(\frac{R_f}{R_f + b} \cos \beta_1 \right). \quad (3.8)$$

Preceding a fast turn segment of maximum length, we can have a slow turn segment lasting up to $\bar{\tau}_s(\beta_1) = \bar{\theta}_s(\beta_1)/\bar{\omega}$ with

$$\bar{\theta}_s(\beta_1) = \pi/2 - \bar{\theta}_f(\beta_1) - \beta_1 = \sin^{-1} \left(\frac{R_f}{R_f + b} \cos \beta_1 \right). \quad (3.9)$$

And again there can be a rotation segment prior to full-length slow turn and fast turn segments for a given value of β_1 .

This family of trajectories comprises all trajectories of the following types (for both left and right turns)

- T_f : Fast turn of up to $\bar{\tau}_f$ duration only
- $T_s T_f$: Slow turn of up to $\bar{\tau}_s(\beta_1)$ duration, followed by fast turn of $\bar{\tau}_f(\beta_1)$ duration, for some $\beta_1 \in [0, \pi/2]$.
- $RT_s T_f$: Rotation, followed by slow and fast turns of duration $\bar{\tau}_s(\beta_1)$ and $\bar{\tau}_f(\beta_1)$ respectively, for some $\beta_1 \in [0, \pi/2]$.

Theorem 3.4.4. *The maximum rotation duration for a time-optimal $RT_s T_f$ trajectory with $\beta(t_1) = \beta_1$ is given by $\bar{\tau}_r = \bar{\theta}_r/\bar{\omega}$, where*

$$\bar{\theta}_r = \pi - \tan^{-1} \left(\frac{R_s - R_f \cos(\bar{\theta}_s + \bar{\theta}_f) + (R_f - R_s) \cos \bar{\theta}_s}{R_f \sin(\bar{\theta}_s + \bar{\theta}_f) - (R_f - R_s) \sin \bar{\theta}_s} \right),$$

with $\bar{\theta}_f = \bar{\theta}_f(\beta_1)$ and $\bar{\theta}_s = \bar{\theta}_s(\beta_1)$ as defined in (3.8) and (3.9).

Proof. Follows similarly to the proof of Theorem 3.4.3. □

By Theorems 3.4.3 and 3.4.4 we have determined seven minimum-time trajectory types: F , $T_f F$, $T_s T_f F$, $RT_s T_f F$, T_f , $T_s T_f$, and $RT_s T_f$, which we illustrate in Table 3.1. In the next

section we show that these seven trajectory types cover the space of destinations. We show how to determine the optimal trajectory type and switching times given a destination point.

3.5 The Optimal Trajectory

Here we present the explicit form of the minimum-time trajectories. The seven types of trajectory (in each turning direction) from Theorems 3.4.3 and 3.4.4 correspond to the different possible combinations of rotation, slow turn, fast turn, and forward trajectory segments (Table 3.1). We first describe the partition of the plane into regions for the different trajectory types. We then present the explicit form of the optimal trajectory for each type individually.

3.5.1 Trajectory parameterized by switching times

Given the initial position $P_0 = (x_0, y_0)^T$ and heading θ_0 at time t_0 and a set of time durations spent in each control segment, we can write the trajectory explicitly as a sequence of translations and rotations about different points. Let τ_r, τ_s, τ_f , and τ_d represent the time spent in rotation, slow turn, fast turn, and forward segments, respectively. The switching times and the final time are denoted

$$t_{rs} = t_0 + \tau_r$$

$$t_{sf} = t_{rs} + \tau_s$$

$$t_{fd} = t_{sf} + \tau_f$$

$$t_1 = t_{fd} + \tau_d.$$

Note that the time to reach the destination is given by the sum of the four segment durations:

$$t_1 - t_0 = \tau_r + \tau_s + \tau_f + \tau_d.$$

The heading angle $\theta(t)$ at time t , for initial heading $\theta(0) = \theta_0$, is given by the following expression, with $+$ for left turn, $-$ for right turn:

$$\theta(t) = \begin{cases} \theta_0 \pm \bar{\omega}(t - t_0), & \text{for } t_0 \leq t < t_{sf} \\ \theta_0 \pm (\bar{\omega}(\tau_r + \tau_s) + (\mu/\bar{v})(t - t_{sf})), & \text{for } t_{sf} \leq t < t_{fd} \\ \theta_0 \pm (\bar{\omega}(\tau_r + \tau_s) + (\mu/\bar{v})\tau_f) & \text{for } t_{fd} \leq t \leq t_1. \end{cases}$$

For convenience, we define the headings at switching times as $\theta_{rs} = \theta(t_{rs})$ and so on.

We can write the agent's position at a given time as the sum of vectors for each segment. Define the "turning vectors" $T_s^\pm(t)$ and $T_f^\pm(t)$ as the translation due to fast turn and slow turn segments of duration t that start from the origin $\mathbf{q}(0) = (0, 0, 0)^T$, as defined in (3.5) and (3.6). Superscripts $+$ and $-$ denote left and right turns, respectively. Also define the "forward motion vector" $F(t)$ as $F(t) = (\bar{v}t, 0)^T$.

The position of the agent $P = (x, y)^T$ at time t is given by the following expression (with $+/-$ on θ and T omitted for clarity):

$$P(t) = \begin{cases} P(t_0), & t_0 \leq t < t_{rs}, \\ P(t_0) + B(\theta_{rs})T_s(t - t_{rs}), & t_{rs} \leq t < t_{sf}, \\ P(t_0) + B(\theta_{rs})T_s(\tau_s) + B(\theta_{sf})T_f(t - t_{sf}), & t_{sf} \leq t < t_{fd}, \\ P(t_0) + B(\theta_{rs})T_s(\tau_s) + B(\theta_{sf})T_f(\tau_f) + B(\theta_{fd})F(t - t_{fd}), & t_{fd} \leq t \leq t_1. \end{cases} \quad (3.10)$$

3.5.2 Trajectory-type partition

To calculate switching times for an open-loop optimal trajectory, we must first determine which trajectory type can be used to reach the destination. For a given set of initial conditions $\mathbf{q}(t_0) = (x_0, y_0, \theta_0)^T$, the plane can be partitioned into regions according to which trajectory

Table 3.1: Range of possible segment durations for each trajectory type. For the entries containing β_1 , the range is $\beta_1 \in [0, \pi/2]$.

Type	τ_r	τ_s	τ_f	τ_d
F	0	0	0	$[0, \infty]$
$T_f F$	0	0	$[0, \bar{\tau}_f]$	$[0, \infty]$
$T_s T_f F$	0	$[0, \bar{\tau}_s]$	$\bar{\tau}_f$	$[0, \infty]$
$RT_s T_f F$	$[0, \bar{\tau}_r(\tau_d)]$	$\bar{\tau}_s$	$\bar{\tau}_f$	$[0, \infty]$
T_f	0	0	$[0, \bar{\tau}_f]$	0
$T_s T_f$	0	$[0, \bar{\tau}_s(\beta_1)]$	$\bar{\tau}_f(\beta_1)$	0
$RT_s T_f$	$[0, \bar{\tau}_r(\beta_1)]$	$\bar{\tau}_s(\beta_1)$	$\bar{\tau}_f(\beta_1)$	0

type can be used to reach destination points in a given region. The set of trajectory types in Table 3.1 together cover the plane for all possible destinations.

The boundaries between different trajectory types are described in Table 3.2 parameterized by segment durations. These boundary curves are either formed by a circular arc, a straight line, or in the case of the $RT_s T_f - T_s T_f$ boundary a polynomial curve described by $y = g(x)$ for $x \in [0, b]$ with

$$g(x) = R_s + (R_f - R_s) \sqrt{1 - \left(\frac{x}{R_s + b}\right)^2} - R_f \sqrt{1 - \left(\frac{x}{b}\right)^2}. \quad (3.11)$$

An example of the trajectory type partition is shown in Fig. 3.3 for initial condition $\mathbf{q}(t_0) = (0, 0, 0)^T$. Here the x -axis separates left turning from right turning trajectories. The positive x -axis is itself a trajectory-type region corresponding to the forward-only trajectory type. The negative x -axis also separates left turning from right turning trajectories, but in this case destinations lying there can be reached in equal time from either left or right turning trajectories with the same segment durations.

3.5.3 Optimal switching times for each compound trajectory type

Here we derive the open-loop optimal control segment durations τ_r , τ_s , τ_f , and τ_d for trajectories in each of the regions defined above. For all, we assume the agent starts at the origin

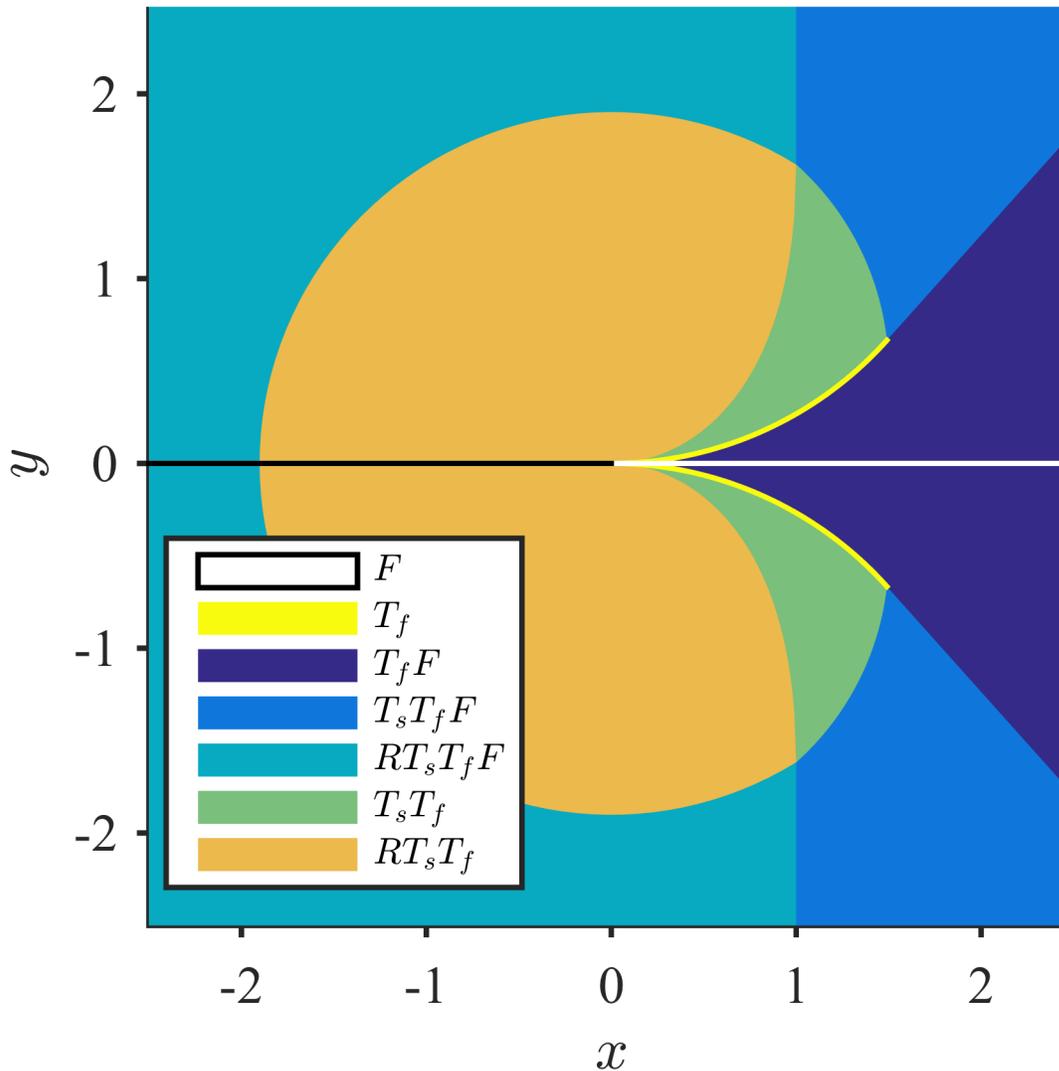


Figure 3.3: Partition of the plane shows the sequence of control segments for different destination regions for minimum-time trajectories of an agent starting at the origin with heading in the direction of the positive x -axis and $\bar{v} = 1$, $\bar{\omega} = 1$, and $\mu = 0.5$. Destinations with $y > 0$ use left turn controls, and those with $y < 0$ use right turn controls. For those destinations along the positive x axis, no turning is necessary. For those on the negative x axis, left and right turning trajectories take the same amount of time.

Table 3.2: Parametric description of trajectory type partition boundaries. For the entries containing β_1 , the range is $\beta_1 \in [0, \pi/2]$.

Boundary	τ_r	τ_s	τ_f	τ_d
$RT_sT_fF-RT_sT_f$	$[0, \bar{\tau}_r(0)]$	$\bar{\tau}_s$	$\bar{\tau}_f$	0
$T_sT_fF-T_sT_f$	0	$[0, \bar{\tau}_s]$	$\bar{\tau}_f$	0
$T_sT_fF-T_fF$	0	0	$[0, \bar{\tau}_f]$	0
$RT_sT_fF-T_sT_fF$	0	$\bar{\tau}_s$	$\bar{\tau}_f$	$[0, \infty]$
$T_sT_fF-T_fF$	0	0	$\bar{\tau}_f$	$[0, \infty]$
$RT_sT_f-T_sT_f$	0	$\bar{\tau}_s(\beta_f)$	$\bar{\tau}_f(\beta_f)$	0
left T_fF -right T_fF	0	0	0	$[0, \infty]$
left RT_sT_f -right RT_sT_f	$\bar{\tau}_r(\beta_f)$	$\bar{\tau}_s(\beta_f)$	$\bar{\tau}_f(\beta_f)$	0
left RT_sT_fF -right RT_sT_fF	$\bar{\tau}_r(\tau_d)$	$\bar{\tau}_s$	$\bar{\tau}_f$	$[0, \infty]$

$\mathbf{q}(t_0) = (0, 0, 0)^T$, and that the destination lies in the upper half-plane so that left turning controls ($\omega \geq 0$) are used.

“Compound” trajectories are those that feature more than one control segment. For each compound trajectory type, there are two unknown values to solve for, as indicated in Table 3.1. Compound trajectories ending in a forward motion segment have unknown durations for the initial and final segments. Compound trajectories ending in a fast turn segment have the initial segment duration and the parameter β_1 as unknowns.

The general strategy to solve for the control segment durations is to first write the equation for the destination in terms of the segment durations as in (3.10). The equation is rearranged such that the initial segment duration appears only in a rotation matrix pre-multiplying one of the sides. Taking the 2-norm thus removes the initial unknown angle, allowing us to solve for the duration of the final segment. We then use substitution to solve for the other unknown.

For convenience we parameterize segment durations according to the change in heading or distance travelled, letting $\theta_r = \bar{\omega}\tau_r$, $\theta_s = \bar{\omega}\tau_s$, $\theta_f = (\mu/\bar{v})\tau_f$, and $d = \bar{v}\tau_d$.

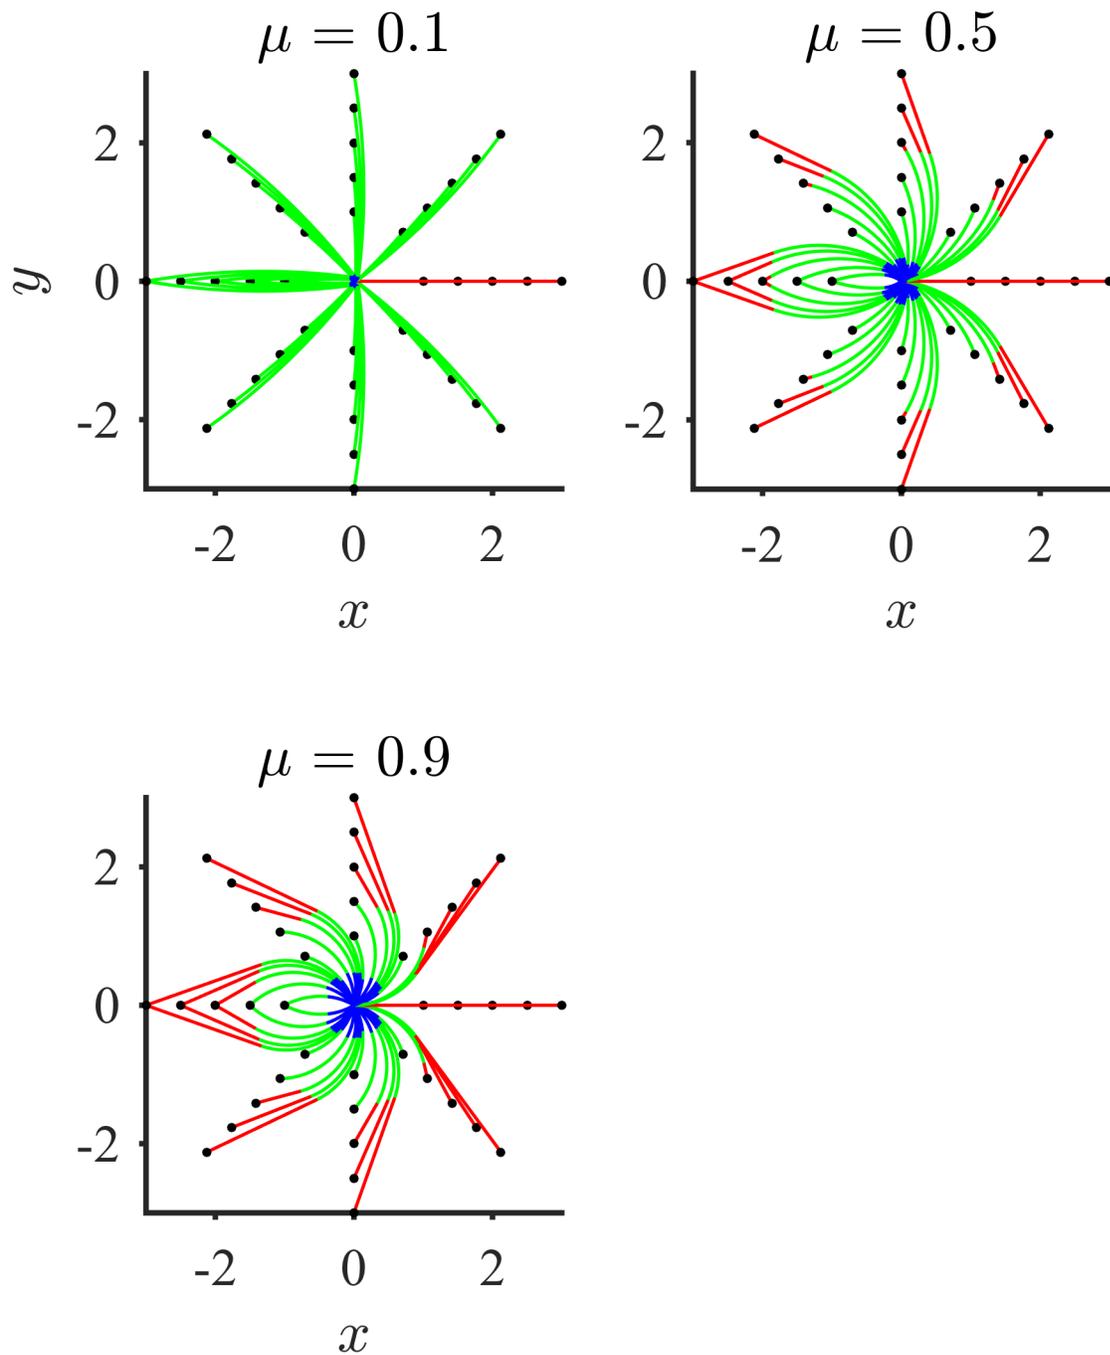


Figure 3.4: Minimum-time trajectories for an agent starting at the origin with heading in the direction of the positive x -axis with $\bar{v} = 1$ and $\bar{\omega} = 1$. Trajectories segments are colored by type: blue is slow turn, green is fast turn, and red is forward motion.

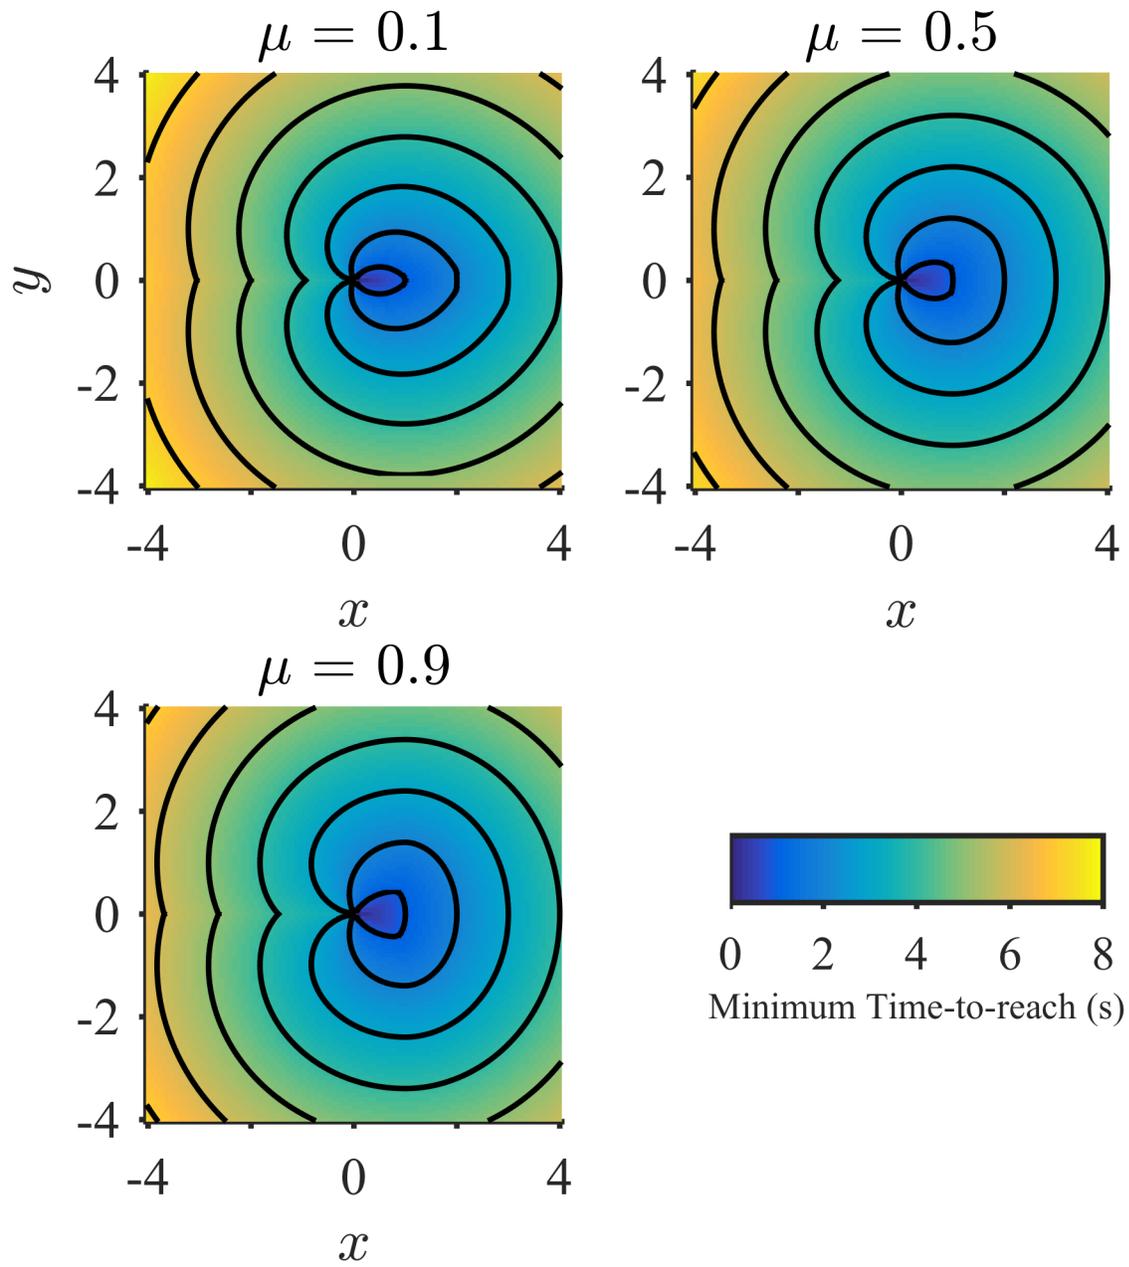


Figure 3.5: Minimum time-to-reach surfaces for an agent starting at the origin with heading in the direction of the positive x -axis with $\bar{v} = 1$ and $\bar{\omega} = 1$, for different values of the lateral acceleration constraint μ . Isochron (equal time-to-reach) lines are drawn at 1 second intervals.

$T_s T_f$ trajectory

At the destination $P(t_1) = (x_1, y_1)^T$ we have

$$\begin{aligned} P(t_1) &= T_s(\tau_s) + B(\theta_s)T_f(\tau_f) \\ &= \begin{pmatrix} R_f \sin(\theta_s + \theta_f) + (R_s - R_f) \sin \theta_s \\ R_s - R_f \cos(\theta_s + \theta_f) + (R_f - R_s) \cos \theta_s \end{pmatrix}. \end{aligned}$$

Note that the left turning vector T_s can be decomposed into

$$T_s(\tau_s) = \begin{pmatrix} 0 \\ R_s \end{pmatrix} + B(\theta_s) \begin{pmatrix} 0 \\ -R_s \end{pmatrix}$$

so that our destination equation can be rewritten with θ_s only appearing within a rotation matrix:

$$P(t_1) - \begin{pmatrix} 0 \\ R_s \end{pmatrix} = B(\theta_s) \left(T_f(\tau_f) - \begin{pmatrix} 0 \\ R_s \end{pmatrix} \right).$$

Taking the 2-norm of both sides removes the dependence on θ_s , and we can solve for

$$\theta_f = \cos^{-1} \left(\frac{R_f^2 + (R_f - R_s)^2 - x_1^2 - (y_1 - R_s)^2}{2R_f(R_f - R_s)} \right).$$

Then θ_s is the clockwise angle about point $(0, R_s)^T$ from $T_f(\theta_f)$ to our destination:

$$\theta_s = \text{atan2}(y_1 - R_s, x_1) - \text{atan2}(R_f(1 - \cos \theta_f) - R_s, R_f \sin \theta_f),$$

where $\text{atan2}(y, x)$ is the standard two-input inverse tangent function with range $[-\pi, \pi]$.

$T_f F$ trajectory

Following the same method, we have

$$P(t_1) = T_f(\theta_f) + B(\theta_f)F(d)$$

$$P(t_1) - \begin{pmatrix} 0 \\ R_f \end{pmatrix} = B(\theta_f) \left(F(d) - \begin{pmatrix} 0 \\ R_f \end{pmatrix} \right).$$

Again taking the norm, we solve for the forward distance

$$d = \sqrt{x_1^2 - R_f^2 + (y_1 - R_f)^2},$$

with

$$\theta_f = \text{atan2}(y_1 - R_f, x_1) - \text{atan2}(-R_f, d).$$

$T_s T_f F$ trajectory

For the $T_s T_f F$ trajectory type, we have

$$P(t_1) = T_s(\theta_s) + B(\theta_s)(T_f(\bar{\theta}_f) + B(\bar{\theta}_f)F(d)).$$

We can rewrite the above equation to isolate θ_s into a single rotation matrix:

$$\begin{pmatrix} x_1 \\ y_1 - R_s \end{pmatrix} = B(\theta_s) \left(\begin{pmatrix} 0 \\ R_f - R_s \end{pmatrix} + B(\bar{\theta}_f) \begin{pmatrix} d \\ -R_f \end{pmatrix} \right). \quad (3.12)$$

Taking the norm of both sides results in $0 = d^2 + c_1 d + c_2$, a quadratic equation in d with

$$c_1 = 2(R_f - R_s) \sin \bar{\theta}_f$$

$$c_2 = R_f^2 + (R_f - R_s)^2 - 2R_f(R_f - R_s) \cos \bar{\theta}_f - x_1^2 - (y_1 - R_s)^2.$$

We take the positive root as our solution for d . Then to calculate the angle θ_s we can simply subtract the angle of the vector on the right hand side of (3.12) from the angle of the vector on the left hand side of (3.12):

$$\theta_s = \text{atan2}(y_1 - R_s, x_1) - \text{atan2}(R_f - R_s + d \sin \bar{\theta}_f - R_f \sin \bar{\theta}_f, d \cos \bar{\theta}_f + R_f \sin \bar{\theta}_f).$$

RT_sT_fF trajectory

This trajectory type includes slow turn and fast turn segments of full possible durations of $\bar{\tau}_s$ and $\bar{\tau}_f$ respectively. Let $P_{sf} = (x_{sf}, y_{sf})^T$ represent the vector traversed during a T_sT_f trajectory with those durations. Then

$$P_{sf} = T_s(\bar{\tau}_s) + B(\bar{\theta}_s)T_f(\bar{\tau}_f) = \begin{pmatrix} b \\ g(b) \end{pmatrix},$$

with g from (3.11). Then

$$P(t_1) = \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} = B(\theta_r) \left(P_{sf} + \begin{pmatrix} 0 \\ d \end{pmatrix} \right).$$

By taking the norm of both sides we can solve for d :

$$d = \sqrt{x_1^2 + y_1^2 - x_{sf}^2 - y_{sf}^2}.$$

Then the rotation angle is given by

$$\theta_r = \text{atan2}(y_1, x_1) - \text{atan2}(y_{sf} + d, x_{sf}).$$

RT_sT_f trajectory

This is the single trajectory type which eludes an explicit analytical solution, although the equations involved are well behaved and simple to solve numerically.

We must find the value $x_c \in [0, b]$ which satisfies the following relation

$$x_1^2 + y_1^2 = x_c^2 + y_c^2 :$$

where $y_c = g(x_c)$ with g from (3.11). Geometrically $(x_c, y_c)^T$ is the point on the RT_sT_f - T_sT_f boundary surface which is the same distance from the origin as the destination point. The distance is smoothly and monotonically increasing with x_c , so it should not pose any difficulties to calculate numerically. Once found, the rotation angle is given by

$$\theta_r = \text{atan2}(y_1, x_1) - \text{atan2}(y_c, x_c),$$

and θ_s, θ_f are found according to the solution for a T_sT_f trajectory presented in Section 3.5.3, using $(x_c, y_c)^T$ as the destination.

Fig. 3.4 shows minimum-time trajectories to various destination points under different values of the lateral acceleration limit μ , including examples for different trajectory types.

By calculating the optimal switching times at every destination point on the plane, we can build a map of the time-to-reach under minimum-time control as a function of the destination. In Fig. 3.5 we illustrate how the minimum time-to-reach is affected by varying the μ parameter of the acceleration constraint. As the value of μ decreases, the time-to-reach increases for all destinations except those reachable by a forward motion-only trajectory.

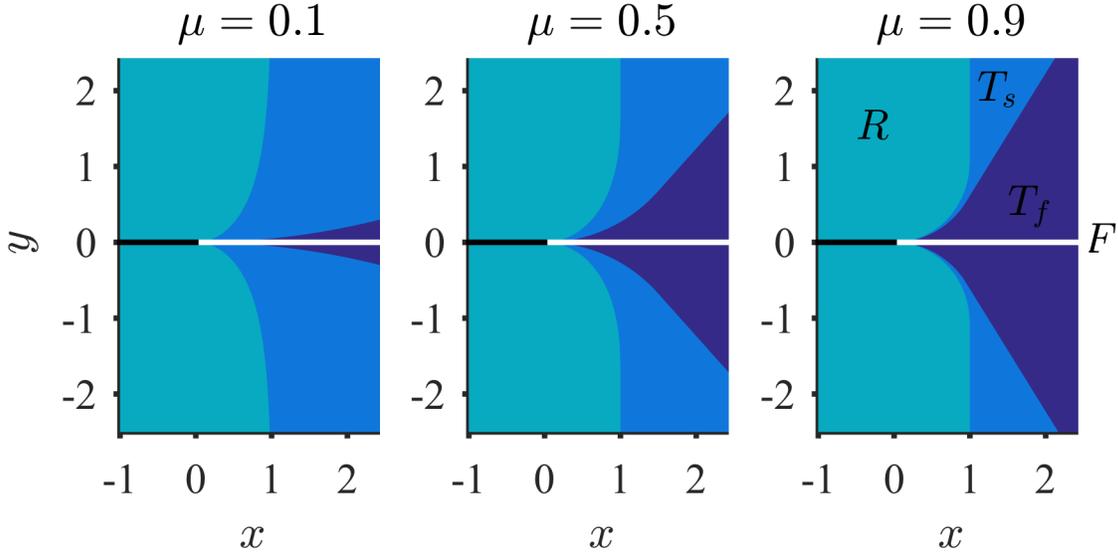


Figure 3.6: Control switching regions for destination relative to agent at the origin with heading in the direction of the positive x -axis with $\bar{v} = 1$ and $\bar{\omega} = 1$. The lateral acceleration parameter μ is varied, from left, $\mu = 0.1$, 0.5 , and 0.9 . Destinations in the upper half plane ($y > 0$) are optimally reached by left turn maneuvers, and in the lower half plane by right turn maneuvers.

3.6 State-feedback formulation of optimal control law

The optimal minimum-time trajectories can be described in terms of a state-feedback law where we consider the dynamics of the location of the destination point in a body-fixed frame with the agent at the origin, heading in the direction of the positive x -axis. The coordinate transformation from world to body-fixed is given by

$$\begin{pmatrix} x_{rel} \\ y_{rel} \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} x_1 - x \\ y_1 - y \end{pmatrix}.$$

The optimal control consists of the following rules:

1. If destination is on the positive x -axis, go forward.
2. Else if destination is in a trajectory-type region with fast turn as the initial segment, go in a fast turn in the appropriate direction.
3. Else if destination is in a trajectory-type region with slow turn as the initial segment, go in a slow turn in the appropriate direction.
4. Else, rotate in the appropriate direction.

Fig. 3.6 illustrates the state-feedback control-type regions under different values of the lateral acceleration constraint μ . As μ decreases, the edges separating the fast turn and slow turn control regions approach the positive x -axis.

3.7 Special cases for large and small values of μ

Here we examine two limiting cases in the minimum-time problem. We first consider relaxing the constraint on lateral acceleration. We then look at the limiting case for very low μ , which has parallels to the problem of a forward-only differential drive vehicle. Control-type regions and optimal trajectories for both cases are shown in Fig. 3.7.

3.7.1 Relaxed acceleration constraint

For $\mu \geq \bar{v}\bar{\omega}$, the lateral acceleration constraint does not affect the boundary of the permissible control space. Here the controls are limited to the rectangle $v \in [0, \bar{b}]$, $\omega \in [-\bar{\omega}, \bar{\omega}]$. In this case, the extremal control is determined through only two switching functions, namely ϕ_1 and ϕ_2 from the general system. In effect, the slow turn and fast turn extremals merge into a single turn trajectory with $v = \bar{v}$, $\omega = \pm\bar{\omega}$, and radius $R = \bar{v}/\bar{\omega}$.

Following the analysis of Section 3.4, we find that there are five trajectory types: forward (F), turn and forward (TF), rotate, turn, and forward (RTF), turn (T), and rotate and turn

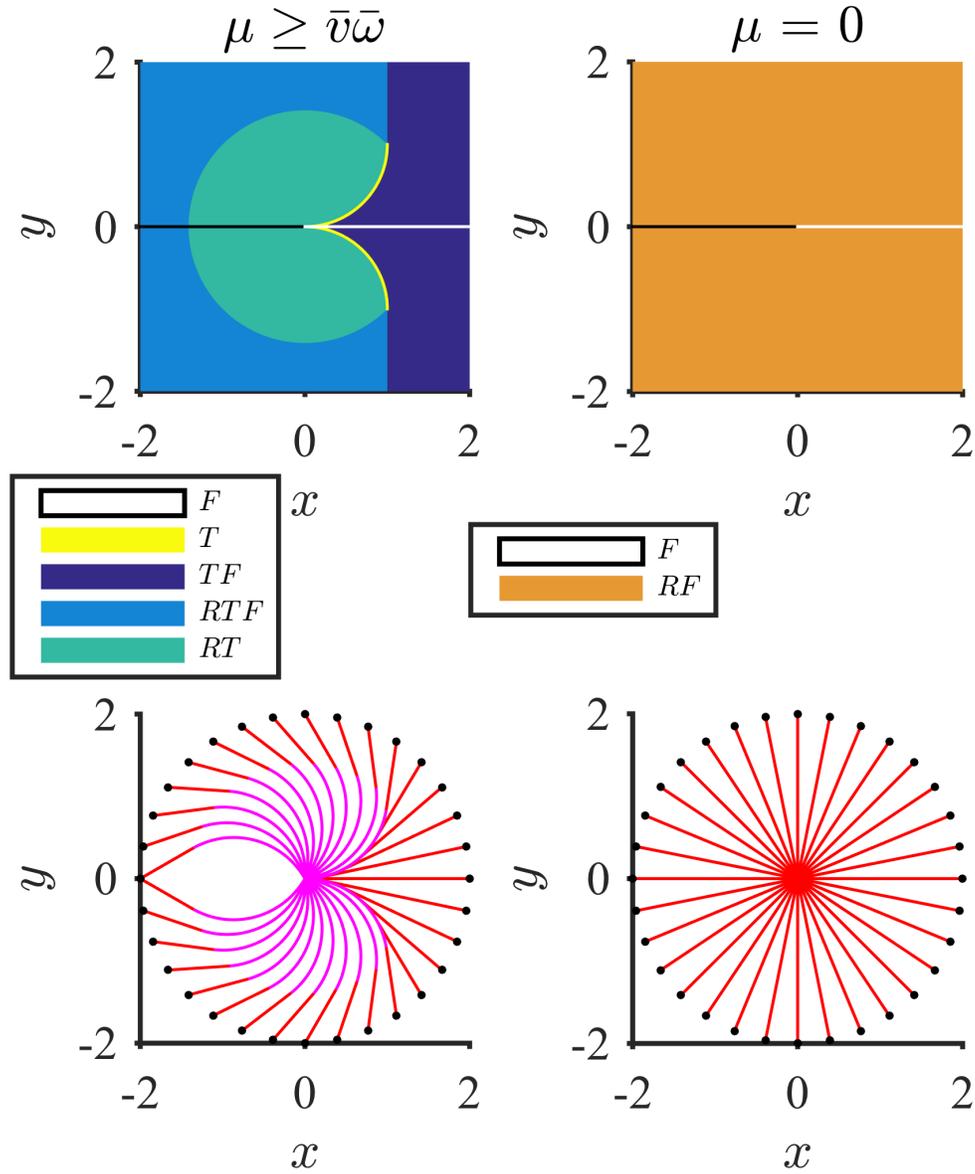


Figure 3.7: Trajectory-type partition (top) and optimal trajectories (bottom) for extreme values of μ for an agent starting at the origin with heading in the direction of the positive x -axis with $\bar{v} = 1$ and $\bar{\omega} = 1$. Left: $\mu \geq \bar{v}\bar{\omega}$, so fast and slow turns are equivalent, both having radius of b . Right: $\mu = 0$, with only rotation and forward motion possible. Trajectories segments are colored by type: magenta is turn, and red is forward.

(*RT*). Destinations on the positive x -axis can be reached by F trajectories. Destinations on the quarter circle

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} R \sin \theta \\ R(1 - \cos \theta) \end{pmatrix}, \quad \theta \in [0, \pi/2] \quad (3.13)$$

can be reached by a turn-only trajectory. The arc given by (3.13) serves as the boundary between TF and RT trajectory types. The boundary between TF and RTF trajectory types is a vertical line $x = R, y \geq R$. Lastly the boundary between RT and RTF trajectory types is the circular arc

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \sqrt{2}R \sin \theta \\ \sqrt{2}R(1 - \cos \theta) \end{pmatrix}, \quad \theta \in [\pi/4, \pi].$$

For a more detailed discussion of this system without the acceleration constraint see [68]. The optimal open-loop segment durations for an agent starting at the origin $\mathbf{q}(0) = (0, 0, 0)^T$ with destination $(x_1, y_1)^T$ are summarized below for each trajectory type. τ_{rot} , τ_{turn} , and τ_{for} are the durations of the rotate, turn, and forward segments, respectively.

- *RTF*: $\tau_{rot} = \theta_{rot}/\bar{\omega}$, $\tau_{turn} = \frac{\pi}{2\bar{\omega}}$, and $\tau_{for} = d_{for}/\bar{v}$, with

$$\begin{aligned} \theta_{rot} &= \text{atan2}(y_1, x_1) - \text{atan2}(R' + d_{for}, R'), \\ d_{for} &= -R + \sqrt{x_1^2 + y_1^2 - R^2}. \end{aligned}$$

- *TF*: $\tau_{rot} = 0$, $\tau_{turn} = \theta_{turn}/\bar{\omega}$, and $\tau_{for} = d_{for}/\bar{v}$, with

$$\begin{aligned} d_{for} &= -R + \sqrt{x_1^2 + y_1^2 - R^2}, \\ \theta_{turn} &= \text{atan2}(y_1 - R', x_1) - \text{atan2}(-R', d_{for}). \end{aligned}$$

- *F*: $\tau_{rot} = \tau_{turn} = 0$ and $\tau_{for} = x_1/\bar{v}$.

- *RT*: $\tau_{rot} = \theta_{rot}/\bar{\omega}$, $\tau_{turn} = \theta_{turn}/\bar{\omega}$, and $\tau_{for} = 0$, with

$$\theta_{turn} = 2 \sin^{-1} \left(\sqrt{x_1^2 + y_1^2} / (2R') \right),$$

$$\theta_{rot} = \text{atan2}(y_1, x_1) - \theta_{turn}/2.$$

- *T*: $\tau_{rot} = \tau_{for} = 0$ and $\tau_{turn} = \sin^{-1}(x_1/R')/\bar{\omega}$.

3.7.2 Highly constrained lateral acceleration

For $\mu = 0$, the acceleration constraint is equivalent to constraining either v or ω to be zero at any given time. We can interpret this as the slow turn merging with rotation, and the fast turn merging with forward motion. The admissible control space becomes a “T” shape, such that the agent can either rotate, or move forward, but not both at the same time. Under those constraints, the extremal controls are specified by the signs of two switching functions, ϕ_2 and ϕ_3 from the general system. We find that time optimal trajectories consist of rotating in place until facing the destination, then moving forward at full speed. The time to reach a destination point $(x_1, y_1)^T = (d_f \cos \alpha, d_f \sin \alpha)^T$ starting from the origin is simply

$$t_f = \tau_{rot} + \tau_{for} = |\alpha|/\bar{\omega} + d_f/\bar{v}.$$

Interestingly, this control scheme is also optimal for a differential-drive robot constrained to only go forward, with input constraints $v \geq 0$ and $v/\bar{v} + |\omega|/\bar{\omega} \leq 1$. For that system, the switching functions and extremal trajectories are the same. The only difference is that the extreme corners of the control space are connected by a straight line in the differential drive case, rather than a concave curve for the limited acceleration system.

3.8 Final Remarks

We have derived optimal control laws for an agent with constraints on speed, angular turning rate, and lateral acceleration in the problem of reaching a destination point in minimum time with free terminal heading. The optimal control laws were presented in both open-loop and feedback control formulations, with analytic expressions for the optimal trajectories.

These control laws and the related time-to-reach surfaces can be used as a building block for problems involving multiple agents. The minimum time problem with free terminal heading is also closely related to the two-player differential game of pursuit and evasion. The evader aims to avoid capture for as long as possible, which is achieved in some cases by using a minimum-time trajectory to reach a point in the space with a lower time-to-reach for the evader than the pursuer. We explore the problem of a single pursuer facing multiple evasive agents without turning constraints in [67].

In Chapter 4 we analyze a differential game of pursuit and evasion between an “omnidirectional” pursuer with limited speed and an evader with the motion constraints described in the current chapter. We show that the optimal evader trajectories in that system are equivalent to the minimum time trajectories presented here, for an appropriately chosen destination point where capture occurs.

Chapter 4

Optimal evasive strategies for multiple interacting agents with motion constraints*

4.1 Introduction

We consider a system with a single pursuer and multiple heterogeneous evader agents moving on the plane. The goal of the pursuer is to capture an evader in the minimum time possible. The goal of each individual evader is to first of all avoid capture, and if that is not achievable to delay capture for as long as possible. The pursuer has limited speed, and the evaders have limits on speed, angular turning rate, and lateral acceleration. To analyze this system we make use of the framework of *differential games* introduced by Isaacs [44]. We examine optimal strategies for the one-on-one pursuit-evasion differential game under these motion constraints, and use those as the building blocks for strategies in the system of multiple evaders.

*This chapter is adapted from a paper by W.L. Scott and N.H. Leonard that has been submitted to Automatica for review, with additional content in Section 4.5.1 from [67].

In the multiple-evader system we propose for the pursuer a strategy of *optimal target selection* such that the evader that could be captured in minimum time in a one-on-one setting is chosen as the target. For evaders, we prove that in the case of all-to-all sensing any evader that is not the current target can always choose from a set of *reactive evasion* control inputs in order to avoid capture. The currently targeted evader must use the optimal evasive strategy from the one-on-one game to delay its capture for as long as possible. We also consider the case in which the pursuer and evaders have radius-limited sensing. We relax the constraint on evader turning rate, and we propose a local strategy of *risk reduction*. We prove that any evader that is not the target can avoid capture using the risk reduction strategy. An early version of this relaxed case appeared in [67].

Predator avoidance has long been considered a key factor in animal group formation, first studied mathematically for identical evaders on the plane in the “selfish herd” model of Hamilton [39]. Our present investigation into the dynamics of a group of self-interested evaders draws on the spirit of a selfish herd. However, it is differentiated through the use of continuous-time dynamics for both pursuer and evaders, and the inclusion of heterogeneity among the evaders. Our investigation is also motivated by the problem of designing dynamics for group formation in engineered multi-agent systems. We prove a decentralized control law for a collective response to a moving threat that accounts for practical motion constraints of individual agents. Our results also provide a control mechanism for spontaneous aggregation.

Hamilton’s model has been extended to include evolutionary dynamics, which lead to formation of large groups [62, 83]. Cooperative evader strategies have been studied as differential games in systems where all evaders are captured [12, 51], and in systems where evaders have defensive capabilities [32, 31]. Several numerical studies have examined properties of group motion in multiple-evader systems where biologically inspired strategies are chosen a priori: on the plane [43, 48, 86], in discrete space [75], in three dimensions [76], with multiple pursuers [1], and with strategies based on observations of crabs and shorebirds [77]. Non-spatially explicit game theoretic models of multiple-evader systems have been posed for both

homogeneous evaders [19], and heterogeneous evaders [25, 26]. Generalized Voronoi diagrams have been used to analyze systems where evader strategies are known by the pursuers [4]. In the present work, we show that standard Voronoi diagrams fall short when turning rate is constrained because in this case optimal risk reduction depends on relative headings as well as relative positions.

We first analyze the system for an “omnidirectional” pursuer agent with limited speed against a group of evader agents with limits on speed, angular turning rate, and lateral acceleration. These evader constraints are chosen with regards to legged locomotion. A study of the kinematics of horses during polo games [74] indicates that grip strength and limb force limits constrain the maximum lateral acceleration during a turn.

Several recent papers examine differential games featuring steered agents with turning constraints, such as a differential-drive pursuer with omnidirectional evader [65], the same but with the roles switched [64], and an omnidirectional pursuer with a car-like evader [27]. These types of dynamics have also been studied in the context of minimum-time problems for a single agent, for the fixed-speed *Reeds-Shepp* vehicle [71], and for a differential-drive vehicle with limited wheel speed [7], and our own work on an agent with limited turning rate that can rotate in place [68]. A biologically inspired analysis of pursuit and evasion with acceleration constraints in [41] suggests that a more agile but slower evader can escape from a fast pursuer with limited lateral acceleration by veering to the side at the last moment. Recent studies of evasive behavior in different animal species are reviewed in [23].

We start with the problem statement and system equations of motion in Section 4.2. In Section 4.3 we derive optimal trajectories and an evader feedback-control law for the one-on-one differential game with evader turning constraints. In Section 4.4 we describe the optimal strategies for the multiple-evader system and illustrate with simulations. We introduce limits on sensing radius in Section 4.5 and examine the problem of evader *risk minimization* in the absence of the pursuer. We address the case where the constraints on evader turning rate and angular acceleration are relaxed. Finally we conclude with discussion in Section 4.6.

4.2 Problem statement and equations of motion

We consider a system with a single pursuer agent P and a heterogeneous group of n evader agents E_i . The pursuer agent P is modeled as an agent that can freely move in any direction with maximum speed \bar{v}_p , position $\mathbf{r}_p(t) \in \mathbb{R}^2$ at time t , and velocity control input $\mathbf{u}_p(t) = (v_{x_p}(t), v_{y_p}(t))^T \in \mathbb{R}^2$ with $\|\mathbf{u}_p(t)\|_2 \leq \bar{v}_p$ for all t . Evaders are modeled as steered agents with inputs of speed $v_i(t) \in \mathbb{R}$ and turning rate $\omega_i(t) \in \mathbb{R}$, written in combination as $\mathbf{u}_i(t) = (v_i(t), \omega_i(t))^T$. An evader's state at time t is composed of its position on the plane $\mathbf{r}_i(t) \in \mathbb{R}^2$ along with its heading angle $\theta_i(t) \in \mathbb{S}^1$.

For each evader agent E_i , we impose the following motion constraints:

- *Forward motion:* Speed must satisfy $v_i(t) \geq 0$ for all time t , such that the agent never moves in reverse.
- *Limited speed:* Let $\bar{v}_i > 0$ be the maximum speed. The speed control must satisfy $v_i(t) \leq \bar{v}_i$ for all time t .
- *Limited turning rate:* Let $\bar{\omega}_i > 0$ be the maximum turning rate. The turning control must satisfy $|\omega_i(t)| \leq \bar{\omega}_i$ for all time t .
- *Limited lateral acceleration:* Let μ_i represent the maximum lateral acceleration (turning traction limit). The inputs $v_i(t)$ and $\omega_i(t)$ must satisfy $|v_i(t)\omega_i(t)| \leq \mu_i$ for all time t . We further impose the condition that $\mu_i < \bar{v}_i \bar{\omega}_i$ so that the lateral acceleration constraint is active on part of the boundary of the control domain.

Let Ω_i be the set of all admissible vector-valued input functions \mathbf{u}_i satisfying the above constraints. Refer to Fig. 3.1 for an illustration of the constraints.

The system equations of motion are

$$\begin{aligned}
 \dot{\mathbf{r}}_p &= \mathbf{u}_p, \\
 \dot{\mathbf{r}}_i &= \begin{pmatrix} v_i \cos \theta_i \\ v_i \sin \theta_i \end{pmatrix}, \\
 \dot{\theta}_i &= \omega_i, \text{ for } i = 1, 2, \dots, n.
 \end{aligned} \tag{4.1}$$

In this system, we define the pursuer's goal to be the capture of a single evader in minimum time. We define the goal of each evader to avoid capture altogether, or if that is not achievable then to delay capture for as long as possible.

4.3 Pursuit and evasion with two agents

Consider the system above with a single evader denoted by the subscript e , where the pursuer has a higher maximum speed $\bar{v}_p > \bar{v}_e$. In this case the pursuer can always guarantee eventual capture. To determine optimal strategies for each agent, we formulate the problem as a *differential game* with the time to capture as the payoff. The two agents' goals are directly opposed: the pursuer aims to minimize the time-to-capture while the evader aims to maximize time-to-capture. We define capture as the condition that the distance between the agents is equal to a capture radius, $l \geq 0$. The standard form of the pursuit-evasion differential game is described by the value function with unity integral cost $L = 1$:

$$T_{cap} = \min_{\mathbf{u}_p \in \Omega_p} \max_{\mathbf{u}_e \in \Omega_e} \int_0^T 1 dt,$$

under dynamics $\dot{\mathbf{q}}(t) = \mathbf{f}(\mathbf{q}, \mathbf{u}_p, \mathbf{u}_e)$, where

$$\mathbf{q} = \begin{pmatrix} x_p \\ y_p \\ x_e \\ y_e \\ \theta_e \end{pmatrix}, \quad \mathbf{f}(\mathbf{q}, \mathbf{u}_p, \mathbf{u}_e) = \begin{pmatrix} v_{x_p} \\ v_{y_p} \\ v_e \cos \theta_e \\ v_e \sin \theta_e \\ \omega_e \end{pmatrix},$$

and terminal condition $\psi(T) = 0$ for

$$\psi(t) = (x_p(t) - x_e(t))^2 + (y_p(t) - y_e(t))^2 - l^2.$$

Define the adjoint vector as a row vector,

$$\boldsymbol{\lambda} = \frac{\partial}{\partial \mathbf{q}} T_{cap} = (\lambda_{x_p}, \lambda_{y_p}, \lambda_{x_e}, \lambda_{y_e}, \lambda_{\theta_e}). \quad (4.2)$$

The control Hamiltonian for the game has the form

$$\begin{aligned} H(\boldsymbol{\lambda}, \mathbf{q}, \mathbf{u}_p, \mathbf{u}_e) &= \boldsymbol{\lambda} \cdot \mathbf{f}(\mathbf{q}, \mathbf{u}_p, \mathbf{u}_e) + 1 \\ &= \lambda_{x_p} v_{x_p} + \lambda_{y_p} v_{y_p} + \lambda_{x_e} v_e \cos \theta_e + \lambda_{y_e} v_e \sin \theta_e + \lambda_{\theta_e} \omega_e + 1. \end{aligned} \quad (4.3)$$

Then optimal control inputs \mathbf{u}_p^* , \mathbf{u}_e^* are specified by the ‘‘Main equation’’ of Isaacs [44]:

$$H(\boldsymbol{\lambda}, \mathbf{q}, \mathbf{u}_p^*, \mathbf{u}_e^*) = \min_{\mathbf{u}_p \in \Omega_p} \max_{\mathbf{u}_e \in \Omega_e} H(\boldsymbol{\lambda}, \mathbf{q}, \mathbf{u}_p, \mathbf{u}_e) = 0.$$

The adjoint equations of motion are $\dot{\boldsymbol{\lambda}} = -\frac{\partial H}{\partial \mathbf{q}}$, and so $\dot{\lambda}_{x_p} = \dot{\lambda}_{y_p} = \dot{\lambda}_{x_e} = \dot{\lambda}_{y_e} = 0$ and $\dot{\lambda}_{\theta_e} = \lambda_{x_e} v_e \sin \theta_e - \lambda_{y_e} v_e \cos \theta_e$. We note (as in [7]) that since $\dot{\lambda}_{\theta_e} = \lambda_{x_e} \dot{y}_e - \lambda_{y_e} \dot{x}_e$, λ_{θ_e} can

be directly integrated:

$$\lambda_{\theta_e} = \lambda_{x_e} y_e - \lambda_{y_e} x_e - \rho, \quad (4.4)$$

for some constant of integration ρ .

4.3.1 Terminal conditions

We start by defining a parameterization of the capture surface:

$$\mathbf{q}(T) = \mathbf{h}(\mathbf{s}) = (s_1 + l \cos s_4, s_2 + l \sin s_4, s_1, s_2, s_3)^T.$$

The value of the game, T_{cap} , does not depend directly on the terminal state, so all its partial derivatives with respect to the terminal surface are zero:

$$0 = \frac{\partial T_{cap}}{\partial s_i} = \boldsymbol{\lambda}(T) \cdot \frac{\partial \mathbf{h}}{\partial s_i}, \quad \text{for } i = 1, 2, 3, 4.$$

This gives us four conditions on the adjoint at the terminal time $t = T$:

$$\begin{aligned} 0 &= \lambda_{x_p} + \lambda_{x_e}, \\ 0 &= \lambda_{y_p} + \lambda_{y_e}, \\ 0 &= \lambda_{\theta_e}, \\ 0 &= l(\lambda_{y_p} \cos s_4 - \lambda_{x_p} \sin s_4). \end{aligned} \quad (4.5)$$

Note that these constraints give us information about the relative magnitudes of the adjoint values, but not the absolute magnitudes. For convenience, we define the following normalized adjoint values:

$$\hat{\lambda}_{x_p} = \lambda_{x_p} / \lambda_0 = \cos(s_4),$$

$$\begin{aligned}\hat{\lambda}_{y_p} &= \lambda_{y_p}/\lambda_0 = \sin(s_4), \\ \hat{\lambda}_{x_e} &= \lambda_{x_e}/\lambda_0 = -\cos(s_4), \\ \hat{\lambda}_{y_e} &= \lambda_{y_e}/\lambda_0 = -\sin(s_4),\end{aligned}$$

where $\lambda_0 = \sqrt{\lambda_{x_p}^2(T) + \lambda_{y_p}^2(T) + \lambda_{x_e}^2(T) + \lambda_{y_e}^2(T)}$.

To determine the magnitude λ_0 of the terminal adjoint vector, we use the equation for the Hamiltonian (4.3) at the terminal time, under the optimal control inputs for each agent:

$$H(\boldsymbol{\lambda}(T), \mathbf{h}(\mathbf{s}), \mathbf{u}_p^*, \mathbf{u}_e^*) = 0 = \lambda_0(\hat{\lambda}_{x_p} v_x^* + \hat{\lambda}_{y_p} v_y^* + \hat{\lambda}_{x_e} v_e^* \cos s_3 + \hat{\lambda}_{y_e} v_e^* \sin s_3) + 1.$$

The pursuer's optimal control to minimize H is given by $v_x^*(T) = -\bar{v}_p \hat{\lambda}_{x_p}$ and $v_y^*(T) = -\bar{v}_p \hat{\lambda}_{y_p}$. The evader's optimal control depends on the location of the terminal state on the terminal surface:

$$v_e^*(T) = \begin{cases} \bar{v}_e, & \cos(s_4 - s_3) < 0, \\ 0, & \cos(s_4 - s_3) \geq 0. \end{cases}$$

Thus we can solve for λ_0 through substitution:

$$\lambda_0 = \begin{cases} (\bar{v}_e \cos(s_4 - s_3) + \bar{v}_p)^{-1}, & \cos(s_4 - s_3) < 0, \\ \bar{v}_p^{-1}, & \cos(s_4 - s_3) \geq 0. \end{cases} \quad (4.6)$$

The “useable part” of the capture surface is defined as the set of points \mathbf{s} where the pursuer can force the state to penetrate the surface:

$$\min_{\mathbf{u}_p} \max_{\mathbf{u}_e} \boldsymbol{\lambda}(\mathbf{s}) \cdot \mathbf{f}(\mathbf{h}(\mathbf{s}), \mathbf{u}_p, \mathbf{u}_e) < 0. \quad (4.7)$$

Since we assume a faster pursuer with $\bar{v}_p > \bar{v}_e$, the entire capture surface comprises the useable part. In Section 4.3.6 we briefly consider the case of a slower pursuer.

4.3.2 Optimal trajectories for pursuit and evasion

Given the state at the time of capture, the trajectories for each agent can be integrated backwards in time based on the optimal controls corresponding to the adjoint vectors as computed above. These trajectories can be used to determine the evader control switching surfaces as a function of the relative position of the pursuer by examining the relative configuration of the agents at the evader switching times.

It is clear from (4.3) that the terms in H corresponding to evader and pursuer control inputs are independent. Thus we can apply Pontryagin's minimum principle for each agent independently and derive optimal trajectories given the proper boundary conditions. Since $\dot{\lambda}_{x_p} = \dot{\lambda}_{y_p} = 0$, the adjoint entries for the pursuer remain constant throughout, and the pursuer will use a constant control input. To minimize H , the optimal path of the pursuer is to follow a straight line at full speed \bar{v}_p in the direction opposite its associated adjoint vector $(\lambda_{x_p}, \lambda_{y_p})^T$.

Suppose, without loss of generality, that the evader is captured while at the origin, with its heading along the positive x -axis. The state at capture is given by $\mathbf{q}(T) = (l \cos s, l \sin s, 0, 0, 0)^T$, for some capture angle $s \in [-\pi, \pi]$. The value of the adjoint vector at capture can be computed from (4.5) and (4.6) through substitution with $s = s_4 - s_3$.

For trajectories ending at this capture point, the pursuer's optimal control is a constant vector $\mathbf{u}_p^* = (-\bar{v}_p \cos s, -\bar{v}_p \sin s)^T$. Integrating backwards in time with $\tau = T - t$, the pursuer's trajectory is a straight line going away from the capture point at the origin:

$$\begin{pmatrix} x_p \\ y_p \end{pmatrix} = \begin{pmatrix} (l + \bar{v}_p \tau) \cos s \\ (l + \bar{v}_p \tau) \sin s \end{pmatrix}.$$

The evader control input $\mathbf{u}_e^* = (v_e^*, \omega_e^*)$ that maximizes H depends not only on the terminal conditions but also on the current state of the system. To determine the extremal

evader control, we define three state-dependent switching functions:

$$\begin{aligned}
\phi_1(\mathbf{q}) &= -\cos(\theta_e - s), \\
\phi_2(\mathbf{q}) &= x_e \sin s - y_e \cos s, \\
\phi_3(\mathbf{q}) &= \bar{\omega}_e |\phi_2| - \bar{v}_e \phi_1.
\end{aligned} \tag{4.8}$$

Let $\text{sgn}(z)$ be the standard sign function for $z \in \mathbb{R}$:

$$\text{sgn } z = \begin{cases} -1, & z < 0, \\ 0, & z = 0, \\ 1, & z > 0. \end{cases}$$

On time intervals for which the switching functions are nonzero, the corresponding extremal controls are called generic. These fall into three categories based on the signs of the switching functions:

- *Rotation:* When $\phi_1 < 0$, the agent rotates in place: $v_e^* = 0$ and $\omega_e^* = \bar{\omega}_e \text{sgn}(\phi_2)$.
- *Slow turn:* When $\phi_1 > 0$ and $\phi_3 > 0$, the agent moves forward with low speed while turning at the maximum rate: $v_e^* = \mu_e / \bar{\omega}_e$ and $\omega_e^* = \bar{\omega}_e \text{sgn}(\phi_2)$. The agent moves on a circular arc with radius $R_s = \mu_e / \bar{\omega}_e^2$.
- *Fast turn:* When $\phi_1 > 0$ and $\phi_3 < 0$, the agent moves forward at maximum speed while turning at a lower rate: $v_e^* = \bar{v}_e$ and $\omega_e^* = \text{sgn}(\phi_2) \mu_e / \bar{v}_e$. The agent moves on a circular arc with radius $R_f = \bar{v}_e^2 / \mu_e$.

In the case that $s = \pm\pi$, we have that $\phi_1 = 1$ and $\phi_2 = 0$ at capture. The evader control input that maximizes H is not unique: any input with $v_e = \bar{v}_e$ and $\omega_e \in [-\mu_e / \bar{v}_e, \mu_e / \bar{v}_e]$ is maximizing. Integrating backwards in time, any control with $\omega_e \neq 0$ will immediately cause the evader to leave the $\phi_2 = 0$ switching surface, bringing it into a generic fast turn segment.

However, should the evader use a control of $v_e = \bar{v}_e$ and $\omega_e = 0$ for an extended interval, it will remain on the switching surface. This *forward motion* evader control is optimal only when the pursuer is directly behind the evader, such that both agents are moving in the direction of the baseline vector from the pursuer to the evader. Once started, forward motion continues until capture.

4.3.3 Evader control switching times

To calculate switching times for the evader optimal control, we integrate the equations of motion backwards in time from capture at time $t = T$. Let $T_{cap} = T - t$ be the time remaining until capture along a specific retro-time trajectory with pursuer and evader agent each using its optimal control. The evader's optimal trajectory will have some combination of rotation, slow turn, fast turn, and forward segments based upon the value of s at capture. For $\sin s > 0$ at capture, the evader will use right turning controls ($\omega_e \leq 0$), and for $\sin s < 0$ at capture, left turning controls.

For $\cos s \geq 0$ at capture, the evader's trajectory consists only of rotation. The pursuer moves in a straight line directly towards the evader.

For $-1 < \cos s < 0$ at capture, the evader's trajectory ends in a fast turn. Proceeding backwards in time from capture by integrating the equations of motion (4.1) using fast turn input for the evader, the state crosses the $\phi_3 = 0$ switching surface at the time given by $T_{cap} = \tau_f(s) = \theta_f(s)\mu_e/\bar{v}_e$, where

$$\theta_f(s) = |s| - \cos^{-1} \left(\frac{\bar{v}_e \bar{\omega}_e}{\bar{v}_e \bar{\omega}_e + \mu_e} \cos s \right).$$

At that time, the evader control switches to a slow turn in the same direction, for duration of $\tau_s(s) = \theta_s(s)/\bar{\omega}_e$, where $\theta_s(s) = |s| - \pi/2 - \theta_f(s)$. It is at that point that the evader crosses the $\phi_1 = 0$ switching surface, and switches to rotation control in the same direction.

When $\cos s = -1$ at capture, the evader's trajectory can end in a forward segment, or a fast turn in either direction. This single end condition corresponds to a whole family of optimal trajectories with varying time spent in the forward segment, τ_d , and either right or left turns. Going backwards in time from capture, at time $T_{cap} = \tau_d$ the evader switches to a fast turn, up to the maximum duration $\bar{\tau}_f = \tau_f(\pi)$, at which point it switches to a slow turn of up to the maximum duration $\bar{\tau}_s = \tau_s(\pi)$, then it switches to rotation.

4.3.4 Evader state-feedback control law

We define a set of *reduced coordinates* $(x_{rel}, y_{rel})^T \in \mathbb{R}^2$ describing the position of the pursuer relative to a frame fixed on the evader, with the positive x -axis in the direction of the evader's heading:

$$\begin{pmatrix} x_{rel} \\ y_{rel} \end{pmatrix} = \begin{pmatrix} \cos \theta_e & \sin \theta_e \\ -\sin \theta_e & \cos \theta_e \end{pmatrix} \begin{pmatrix} x_p - x_e \\ y_p - y_e \end{pmatrix}. \quad (4.9)$$

Consider the set of all optimal trajectories $\mathbf{q}(t)$ described above in Sections 4.3.2 and 4.3.3, along with their associated optimal controls $\mathbf{u}_p^*(t)$ and $\mathbf{u}_e^*(t)$ for all $s \in [-\pi, \pi]$ and $\tau_d \geq 0$. Any ordered pair of capture angle and time-to-capture (s, T_{cap}) (along with forward duration τ_d for $s = \pm\pi$) corresponds to a point in the reduced space, (x_{rel}, y_{rel}) , along with an associated optimal evader input and optimal pursuer input (transformed to the reduced coordinate frame). By collecting the set of all of these optimal trajectories together in reduced space, we can derive the optimal controls in state-feedback form: this is the inverse mapping that takes us from a point in the reduced space, pursuer position relative to the evader, to the associated optimal control.

Since the evader has a discrete set of possible optimal controls, its state-feedback control law maps a region of the reduced space to each input. Fig. 4.1 shows optimal trajectories in the reduced coordinates, with color denoting the optimal evader control input at each point. The locus of points in trajectories where evader control switching occurs together form the switching surfaces for the evader, shown as black lines in the figure. The negative

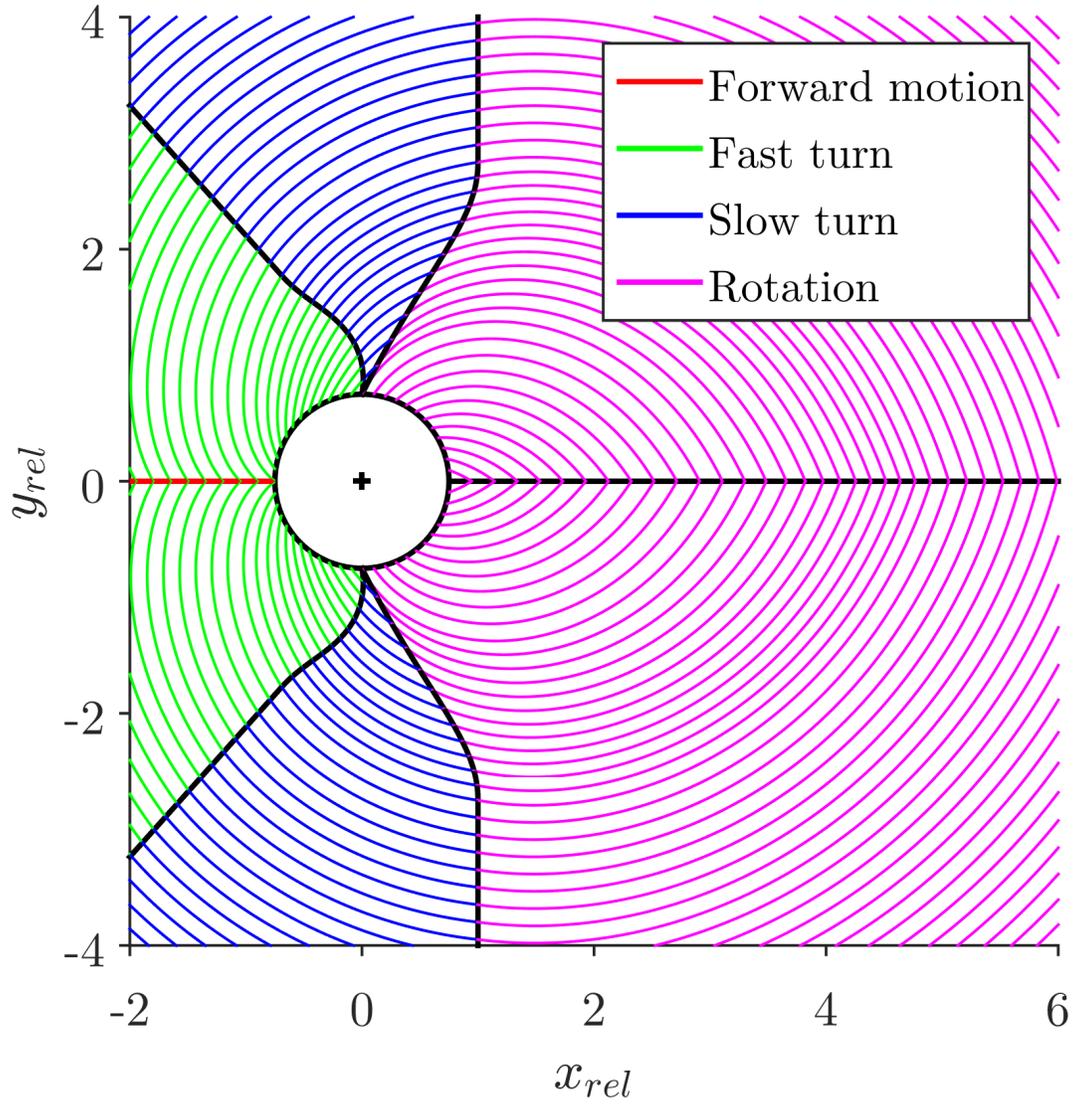


Figure 4.1: Optimal trajectories in reduced coordinates (4.9). Color denotes optimal evader control input, with left turns for $y_{rel} < 0$ and right turns for $y_{rel} > 0$. Here $\bar{v} = 1$, $\bar{\omega} = 1$, $\mu = 0.5$, and $v_p = 1.5$. The capture radius is set to the minimum value of $l = l_c$ from (4.11).

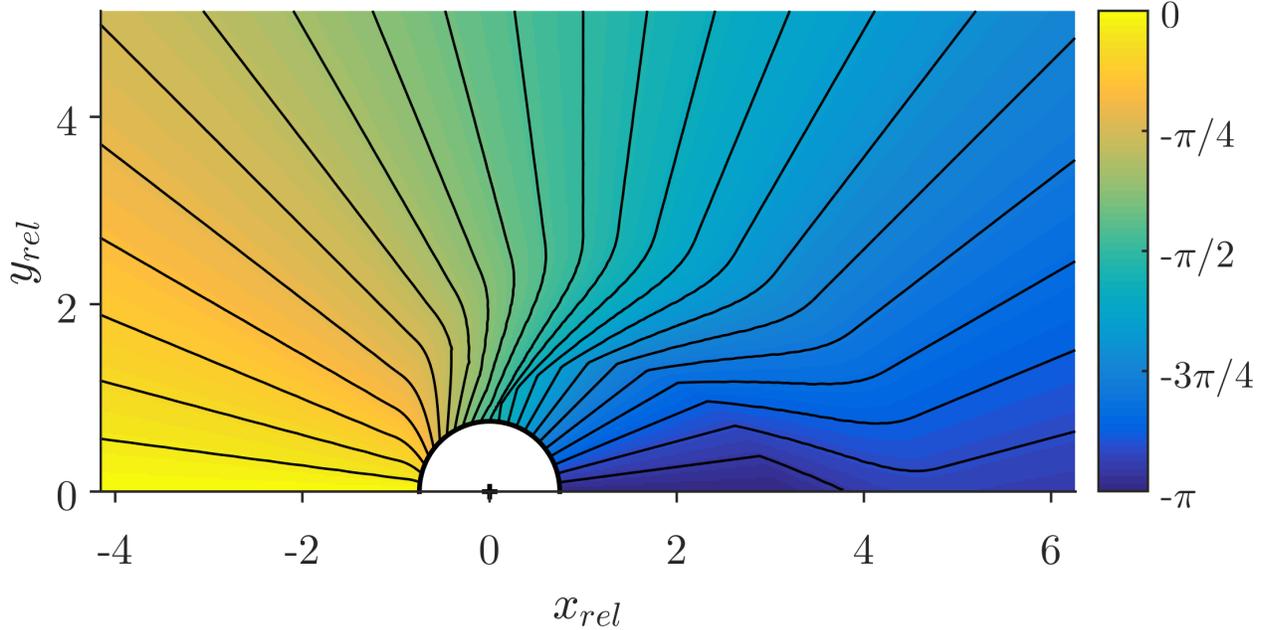


Figure 4.2: Optimal pursuer heading θ_p^* in reduced coordinates (4.9). Lines are drawn at $\pi/24$ radian intervals. Here $\bar{v} = 1$, $\bar{\omega} = 1$, $\mu = 0.5$, and $v_p = 1.5$ for capture radius $l = l_c$ from (4.11). Note that the pursuer headings take the opposite value on the lower half plane, such that $\theta_p^*(x_{rel}, -y_{rel}) = -\theta_p^*(x_{rel}, y_{rel})$.

x -axis in the reduced coordinates corresponds to evader forward motion with the pursuer directly behind it. This line constitutes what is known as a “universal surface” in differential game theory—it is an optimal trajectory that behaves such that other optimal trajectories run into it and flow along it like tributaries meeting a river. Conversely the positive x -axis is a “dispersal surface” in that games starting with the evader facing directly towards the pursuer can proceed either with left or right turning evader trajectories, leading to capture in an equal amount of time under optimal play by both agents.

Similarly the set of optimal trajectories in reduced coordinates also provides mappings from relative pursuer position to the optimal pursuer heading, shown in Fig. 4.2, and to the optimal time-to-capture T_{cap} , shown in Fig. 4.3. This mapping of the state to the value of the game plays a large role in the development of optimal strategies in the multiple-evader system presented in Section 4.4.

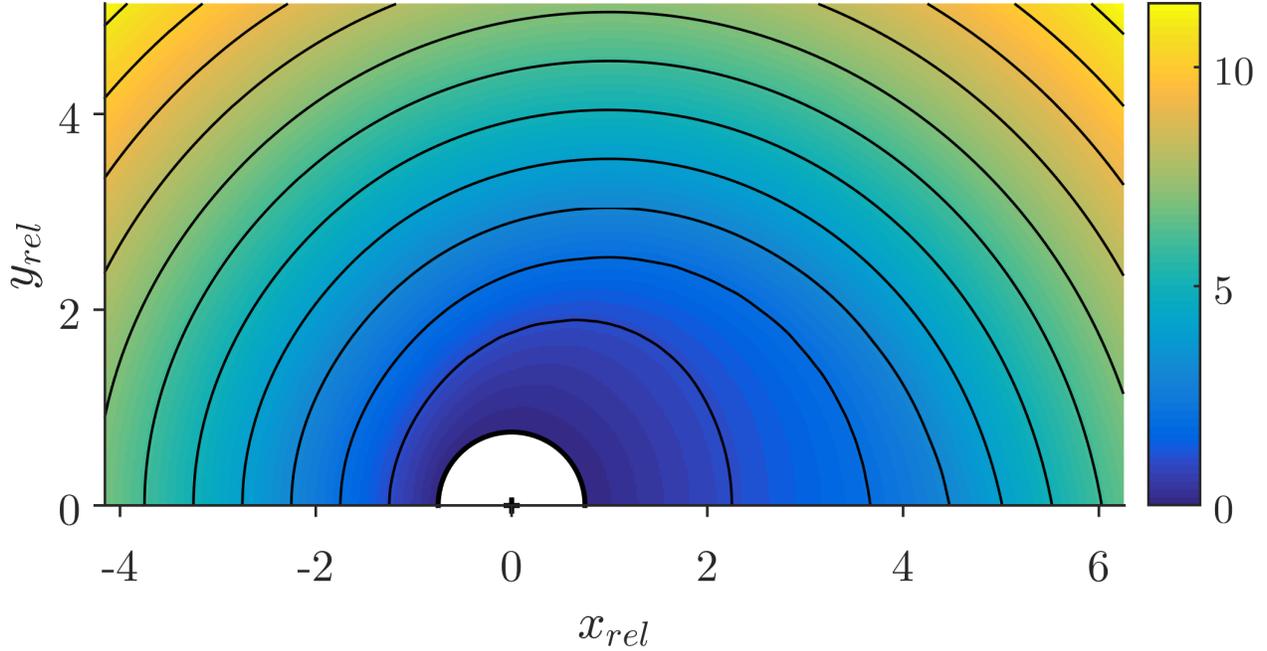


Figure 4.3: Time-to-capture T_{cap} surface in reduced coordinates (4.9) with $\bar{v} = 1$, $\bar{\omega} = 1$, $\mu = 0.5$, and $v_p = 1.5$ for capture radius $l = l_c$ from (4.11). The T_{cap} surface is mirrored about the x_{rel} -axis, such that $T_{cap}(x_{rel}, -y_{rel}) = T_{cap}(x_{rel}, y_{rel})$. Contours are drawn at one-second increments.

Remark 4.3.1. *The evader optimal trajectories are equivalent to the optimal trajectories for the problem of reaching a desired point on the plane in minimum time under these motion constraints. This holds similarly for the pursuer, since the fastest way to reach a point is to go directly towards it at full speed. The one-on-one game of pursuit and evasion is in a sense equivalent to the problem of choosing the point on the plane resulting in capture which maximizes the evader’s time-to-reach and using the optimal trajectory to reach that point. In [68] we derived minimum-time trajectories to reach a point on the plane for a steered agent with speed and turning rate constraints but with no limit on lateral acceleration.*

4.3.5 Condition on capture radius

The inverse mapping described in Section 4.3.4 is well defined if every state in the reduced coordinates maps to a unique time-to-capture and optimal control input for each agent. This will hold if the capture radius l is positive and sufficiently large. We find that this is satisfied

as long as the $s = \pi/2$ evader rotate-only trajectory does not intersect the slow turn-rotation switching surface, except where they meet at the capture surface.

We can solve for the condition on l by requiring that the slope of the slow turn-rotation switching surface is greater than that of the $s = \pi/2$ rotate-only trajectory in the upper half plane of reduced coordinates where they extend from the capture surface at $(x_{rel}, y_{rel})^T = (0, l)^T$:

$$\begin{aligned} \left. \frac{dy}{dx} \right|_{x_{rel}=0, y_{rel}=l}^{\text{switching surface}} &\geq \left. \frac{dy}{dx} \right|_{x_{rel}=0, y_{rel}=l}^{\text{rotate-only traj.}} \\ \Leftrightarrow \frac{2v_p\omega_e}{\mu + v_e\omega_e} &\geq \frac{vp}{l\omega_e}. \end{aligned} \quad (4.10)$$

The minimum capture radius l_c , such that (4.10) is satisfied for $l \geq l_c$, is

$$l_c = \frac{\mu + v_e\omega_e}{2\omega_e^2} = \frac{1}{2} \left(R_s + \frac{\bar{v}_e}{\bar{\omega}_e} \right). \quad (4.11)$$

4.3.6 The case of the slower pursuer

Consider a system where the top speed of the pursuer is less than that of the evader, $\bar{v}_p < \bar{v}_e$. In this case there exists a set of initial conditions that allow the evader to avoid capture under optimal play. We construct what is known as a “barrier” surface in the reduced state space that separates the x_{rel} - y_{rel} plane into distinct capture and non-capture regions. The barrier is formed by the optimal trajectory which reaches the capture surface at the boundary of the “useable part,” (UP) defined in (4.7).

For this system, the useable part corresponds to the points on the capture surface where $-\bar{v}_p + \bar{v}_e \cos s \leq 0$. In the case of a faster pursuer, the entire capture surface is useable. For a slower pursuer, however, the useable part is the range $|s| \leq \cos^{-1}(-\bar{v}_p/\bar{v}_e)$. To construct the barrier, we simply integrate the optimal trajectory backwards in time from the two edges of the useable part, until they meet at the dispersal surface on the positive x_{rel} -axis.

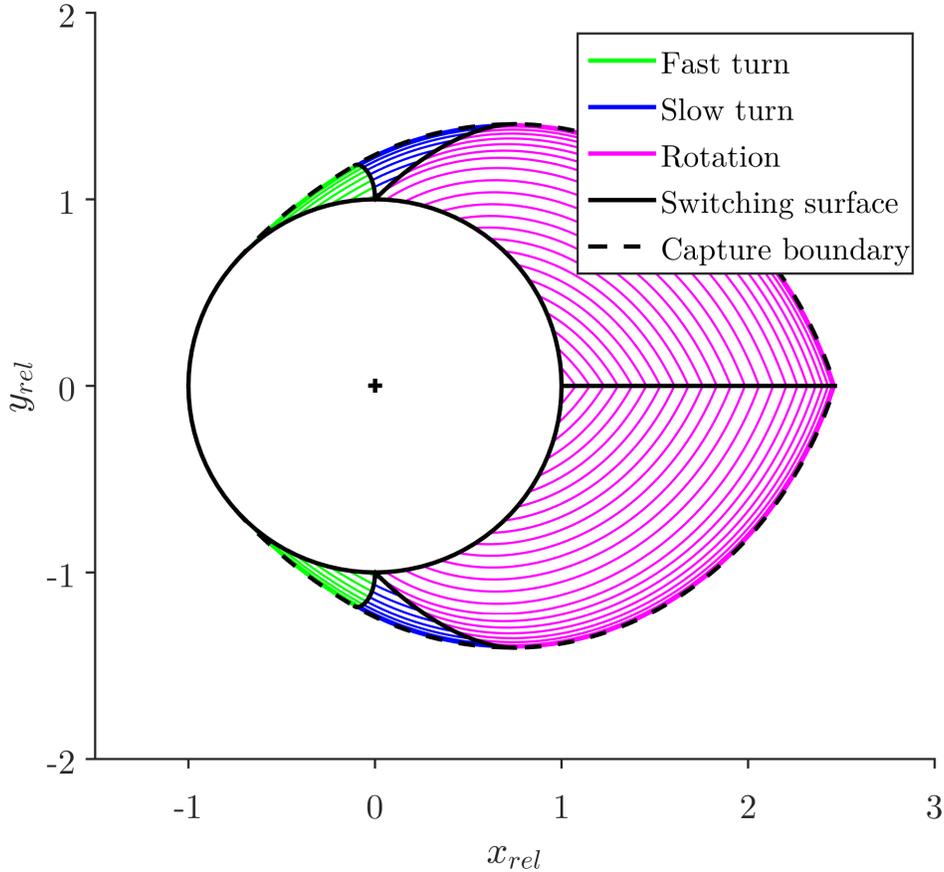


Figure 4.4: Optimal trajectories in reduced coordinates in the case of a slower pursuer $\bar{v}_p < \bar{v}_e$. Color denotes optimal evader control input, with left turns for $y_{rel} < 0$, right turns for $y_{rel} > 0$. Here $\bar{v}_e = 1$, $\bar{\omega}_e = 1$, $\mu = 0.5$, and $v_p = 0.75$, with capture radius $l = l_c$. Optimal strategies are not defined outside of the capture region.

Within the capture region, optimal trajectories are constructed by integrating backwards in time from points on the useable part of the capture surface, as derived in Sections 4.3.2 and 4.3.3. Optimal trajectories are not defined outside of the capture region, as there can be no value (time-to-capture) for a game that does not end in capture. Figure 4.4 illustrates the capture boundary surface along with optimal trajectories inside the capture region for a system with a slower pursuer. Capture is only achievable for the pursuer if it starts in the small capture region situated close to and in front of the evader.

4.4 Reactive evasion with multiple evaders

We now consider a system with a single fast pursuer and multiple evaders with heterogeneity in their control constraints. Here the pursuer’s goal is to capture any single evader in the minimum time possible. The goal of each evader is to avoid capture by the pursuer, or else delay capture for as long as possible. We show that the optimal strategy for the pursuer is to choose a target evader with the lowest bounded time-to-capture (T_{cap}) and using the optimal control strategy from the one-on-one differential game. The targeted evader must use its optimal strategy from the one-on-one game, but the others can use what we call a “reactive evasion” strategy that will guarantee that they do not become the target of the pursuer.

Theorem 4.4.1. *The optimal strategy for the pursuer in a multiple-evader system is to choose as its target the evader with the lowest time-to-capture and utilize the optimal strategy for the one-on-one game against that evader. This strategy will guarantee capture in time $T \leq \min_i T_{cap,i}$.*

Proof. Under optimal pursuer play in a one-on-one game against an evader E_i , the pursuer is guaranteed to capture the evader in time $T \leq T_{cap,i}$. Conversely any evader using optimal evasive control is guaranteed to avoid capture up to time $T_{cap,i}$. Thus, the best possible strategy for the pursuer to catch an evader in minimum time is to use the optimal strategy for the one-on-one system against the evader with the minimum value of $T_{cap,i}$. \square

4.4.1 Evader domain of danger and target avoidance

For any given pursuer location, the optimal target is the evader E_i that has the lowest value of $T_{cap,i}$. We can thus partition the plane into “domains of danger,” denoted \mathcal{D}_i , corresponding to the points on the plane where a given evader E_i has the minimum value of $T_{cap,i}$:

$$\mathcal{D}_i = \left\{ \mathbf{p} \in \mathbb{R}^2 \mid i = \underset{j}{\operatorname{argmin}} T_{cap,j} \right\}. \quad (4.12)$$

Fig. 4.5 shows an example of the domains of danger before and after a pursuit for a group of evaders of different maximum speeds and common turning rate and lateral acceleration constraints. From Theorem 4.4.1, at any given time during pursuit, the optimal target of the pursuer is the evader in whose domain of danger the pursuer currently resides. If an evader can keep the pursuer from entering its domain of danger, then it will not become a target and can avoid capture. We show that this is indeed possible for any evader not initially targeted, under the assumption that all agents can sense one another's relative position and orientation, and have knowledge of each other's motion constraints.

Remark 4.4.1. *For evader agents with no constraints on turning rate, the domain of danger partition takes on the form of a “multiplicatively-weighted Voronoi diagram,” which is described in Section 4.5. In the present more general system, the T_{cap} surface is not radially symmetric, so the domain of danger depends not only on distance to other evaders, but on their relative headings as well.*

Theorem 4.4.2. *Consider the multiple-evader system with dynamics (4.1) and a pursuer using optimal target selection and pursuit as in Theorem 4.4.1. For each evader not currently targeted by the pursuer there exists a set of control inputs that will guarantee that the evader will not become the target at any future time. If all non-targeted evaders employ such a strategy for the duration of a chase, the pursuer and targeted evader strategies are equivalent to those from the one-on-one game, with capture occurring at the initial value of T_{cap} for the targeted evader.*

Proof. Let E_g be the evader with the unique minimum value of T_{cap} at the start of the chase $t = t_0$, so $T_{cap,g}(t_0) = \min_i T_{cap,i}$ for $i \in \{1, 2, \dots, n\}$. Any other evader E_i has $T_{cap,i}(t_0) > T_{cap,g}(t_0)$, and it will remain untargeted as long as that inequality holds. If we can show that there exists a control input for E_i such that $\frac{d}{dt}T_{cap,i} \geq \frac{d}{dt}T_{cap,g}$ for all time $t \geq t_0$, then by continuity $T_{cap,i}(t) > T_{cap,g}(t)$ for all $t \geq t_0$ as well.

Recall that the adjoint vector represents the partial derivative of the game's value with respect to the state, $\boldsymbol{\lambda} = \frac{\partial}{\partial \mathbf{q}} T_{cap}$. So

$$\dot{T}_{cap} = \frac{d}{dt} T_{cap} = \frac{\partial}{\partial \mathbf{q}} T_{cap} \cdot \dot{\mathbf{q}} = \boldsymbol{\lambda} \cdot \mathbf{f}(\mathbf{q}, \mathbf{u}_p, \mathbf{u}_e). \quad (4.13)$$

In the one-on-one game under optimal control by both agents, $\dot{T}_{cap} = -1$ at all times, based on the definition of the optimal controls in (4.3). Thus $\dot{T}_{cap,g} = -1$.

Any deviation from the optimal input by the pursuer (for instance targeting some other evader) will necessarily lead to a slower decrease of T_{cap} , so a non-targeted evader E_i can simply use its one-on-one optimal control to guarantee $\dot{T}_{cap,i} \geq -1$, thus avoiding becoming a target. However, this is usually not strictly necessary: a range of different evader inputs can be used to keep $\dot{T}_{cap,i} \geq -1$. We define *reactive evasion* as any evader input that satisfies this inequality. In particular, non-targeted evaders can apply a reactive evasion strategy that makes efficient use of their input energy. We show next how evader E_i can compute a state-feedback control for reactive evasion.

Under optimal pursuit (of some target evader E_g) the pursuer's trajectory is a straight line at maximum speed \bar{v}_p . So we can suppose that an evader E_i has some estimate of the pursuer's current direction of travel based on its recent behavior. Let θ_p be the current direction of travel of the pursuer, and v_p its speed, such that $\mathbf{u}_p = (v_p \cos \theta_p, v_p \sin \theta_p)^T = (v_{x_p}, v_{y_p})^T$.

Given the current relative position of the pursuer, the evader E_i can compute its trajectory under optimal play, including the values of the adjoint variables that parameterize the switching functions. Let $\mathbf{u}_{p,i}^*$ be the pursuer's optimal control with respect to pursuit of E_i , with associated heading $\theta_{p,i}^*$. The evader must choose its input \mathbf{u}_i to satisfy

$$\dot{T}_{cap,i}(\mathbf{q}, \mathbf{u}_p, \mathbf{u}_i) \geq \dot{T}_{cap,i}(\mathbf{q}, \mathbf{u}_{p,i}^*, \mathbf{u}_i^*),$$

equivalently, to satisfy

$$\boldsymbol{\lambda} \cdot \mathbf{f}(\mathbf{q}, \mathbf{u}_p, \mathbf{u}_i) \geq \boldsymbol{\lambda} \cdot \mathbf{f}(\mathbf{q}, \mathbf{u}_{p,i}^*, \mathbf{u}_i^*). \quad (4.14)$$

In order to remove the explicit dependence on the adjoint $\boldsymbol{\lambda}$ from (4.14), we now solve for $\boldsymbol{\lambda}$ in terms of the state and optimal controls. From the derivation of the pursuer optimal control in Section 4.3.2, we have $\lambda_{x_p} = -\lambda_{x_i} = -\lambda_0 \cos \theta_{p,i}^*$ and $\lambda_{y_p} = -\lambda_{y_i} = -\lambda_0 \sin \theta_{p,i}^*$. Let $(x_i(T), y_i(T))^T$ be the location of the evader at capture at time $t = T$ under optimal one-on-one controls for each agent. The pursuer's position at time $t \leq T$ under optimal control $\mathbf{u}_{p,i}^*$ is

$$\begin{aligned} x_p(t) &= x_i(T) - (l + \bar{v}_p(T-t)) \cos \theta_{p,i}^*, \\ y_p(t) &= y_i(T) - (l + \bar{v}_p(T-t)) \sin \theta_{p,i}^*. \end{aligned} \quad (4.15)$$

From (4.4) and (4.5),

$$\lambda_{\theta_i}(t) = \lambda_0 [(y_i(t) - y_i(T)) \cos \theta_{p,i}^* - (x_i(t) - x_i(T)) \sin \theta_{p,i}^*]. \quad (4.16)$$

By adding and subtracting $(l + \bar{v}_p(T-t)) \cos \theta_{p,i}^* \sin \theta_{p,i}^*$ from (4.16) and substituting with (4.15), we have

$$\lambda_{\theta_i}(t) = \lambda_0 [(y_i(t) - y_p(t)) \cos \theta_{p,i}^* - (x_i(t) - x_p(t)) \sin \theta_{p,i}^*].$$

Finally substituting these expressions for the adjoint back into (4.14) and rearranging gives us

$$\begin{aligned} \bar{v}_p - v_p \cos(\theta_p - \theta_{p,i}^*) + (v_i - v_i^*) \cos(\theta_i - \theta_{p,i}^*) \\ + (\omega_i - \omega_i^*) [(y_i - y_p) \cos \theta_{p,i}^* - (x_i - x_p) \sin \theta_{p,i}^*] \geq 0. \end{aligned} \quad (4.17)$$

Thus in order to keep $\dot{T}_{cap,i} \geq -1$, the evader must choose its inputs v_i and ω_i such that they satisfy the linear inequality (4.17). \square

Remark 4.4.2. *An evader E_i is not in danger of becoming targeted until its $T_{cap,i}$ is close to that of the target evader, at which point it must begin to use a reactive evasion strategy. For instance evader E_i can wait to initiate reactive evasion until $T_{cap,i} \leq \min_j T_{cap,j} + \epsilon$ for some chosen buffer value $\epsilon > 0$. Until that point, the evader is free to use any control input, for instance a “herding” strategy of aligning heading and matching speed with local neighbors, such as one adapted from the attraction-orientation-repulsion zonal model of [18].*

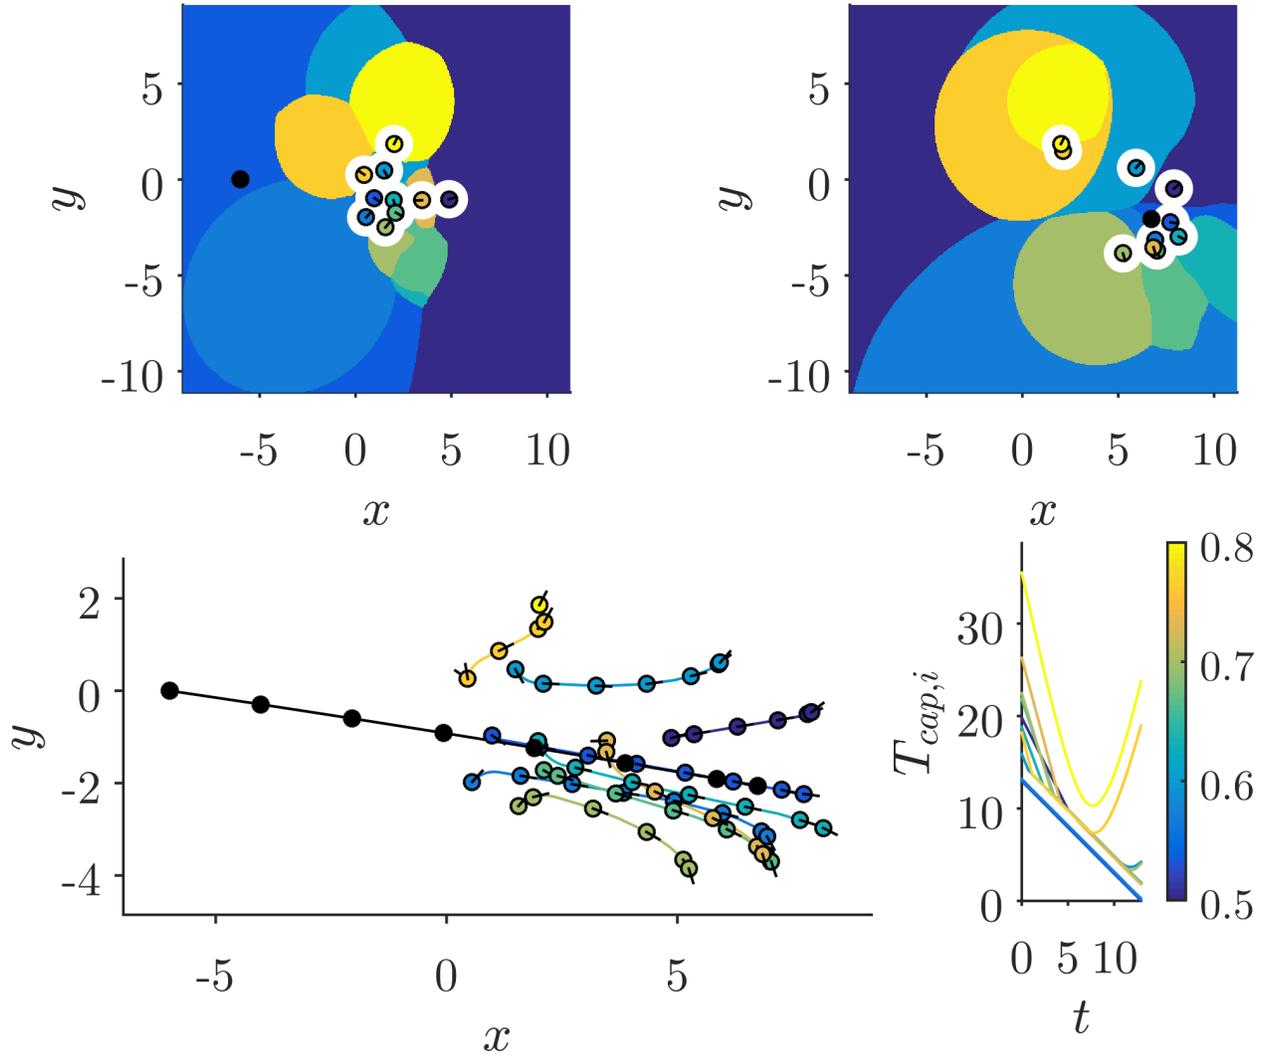


Figure 4.5: Simulation of reactive evasion with ten evaders. The evaders and their respective domains of danger are colored by evader speed, shown in colorbar at bottom right. Here $v_p = 1$, and for all evaders, $l = 1$, $\omega_i = 1$, and $\mu_i = 0.5$, with v_i evenly spaced in the range $[0.5, 0.8]$. Top left: evader domains of danger at initial time. Top right: evader domains of danger at capture. Pursuer is represented by filled black circle. Bottom left: agent trajectories, with snapshots showing evader headings every 2 seconds. Bottom right: evader $T_{cap,i}$ over the course of the chase. Note that the evader with the lowest T_{cap} has its value decreasing at a constant rate of -1 second per second until capture.

4.5 Risk minimization under limited sensing

In this section we consider the multiple-evader system in which the pursuer and evaders have a limited sensing range. Each agent can only make use of measurements of agents located inside its sensing range. The limited sensing range may be important for modeling large groups that are widely dispersed or that are tightly packed and suffer from occlusion. We define limits on sensing and adapt the pursuit and evasion strategies introduced in Section 4.4 to local strategies that address the uncertainty imposed by limited sensing.

In the local (sensing-limited) system, we define d_{sense} as the sensing radius for all agents. An agent's local sensing neighborhood consists of the set of agents within the sensing radius. Let d_{ij} be the distance between agents i and j . The local neighborhood of the pursuer is defined as

$$\mathcal{N}(P) = \{E_i | d_{ip} \leq d_{sense}\}. \quad (4.18)$$

The neighborhood of evader E_i is defined as

$$\mathcal{N}(E_i) = \{E_j | d_{ij} \leq d_{sense}\} \cup \{P | d_{ip} \leq d_{sense}\}. \quad (4.19)$$

For analytical tractability we relax the constraint on evader turning rate.

4.5.1 Relaxing constraints on evader turning rate

Consider the case in which the evader turning rates are not constrained. Without a limit on turning rate, an evader can reach any desired heading arbitrarily quickly, so it is equivalent to having the same dynamics as the “omnidirectional” pursuer. We will consider the terminal condition to be “point capture,” so that the game ends when the distance between pursuer and evader reaches zero.

Let $\mathbf{r}_p(t) = (x_p(t), y_p(t))^T \in \mathbb{R}^2$ be the pursuer position at time t , with velocity input $\mathbf{u}_p(t) \in \mathbb{R}^2$ and speed limit $\|\mathbf{u}_p(t)\|_2 \leq \bar{v}_p$ for all time t . Let $\mathbf{r}_i(t) = (x_i(t), y_i(t))^T \in \mathbb{R}^2$ be the

position of evader E_i at time t , with velocity input $\mathbf{u}_e(t) \in \mathbb{R}^2$ and speed limit $\|\mathbf{u}_i(t)\|_2 \leq \bar{v}_i$ for all time t . Assume a faster pursuer, i.e. $\bar{v}_p > \bar{v}_i$ for all $i = 1, 2, \dots, n$. The system dynamics are

$$\dot{\mathbf{r}}_p = \mathbf{u}_p, \quad \dot{\mathbf{r}}_i = \mathbf{u}_i, \quad \text{for } i = 1, 2, \dots, n. \quad (4.20)$$

Let $\mathbf{r}_{ip} = \mathbf{r}_i - \mathbf{r}_p$ be the *baseline vector* from the pursuer to evader E_i , with associated distance $d_{ip} = \|\mathbf{r}_{ip}\|_2$, and normalized unit vector $\hat{\mathbf{r}}_{ip} = \mathbf{r}_{ip}/d_{ip}$ defined for $d_{ip} > 0$. For the one-on-one pursuit-evasion game with these dynamics, as shown in [44], the adjoint vector and optimal controls for both pursuer and evader are constant, with $\boldsymbol{\lambda} = (\bar{v}_p - \bar{v}_i)^{-1}(-\hat{\mathbf{r}}_{ip}^T, \hat{\mathbf{r}}_{ip}^T)$, pursuer control $\mathbf{u}_p^* = \bar{v}_p \hat{\mathbf{r}}_{ip}$, and evader control $\mathbf{u}_i^* = \bar{v}_i \hat{\mathbf{r}}_{ip}$. The time to capture is simply $T_{cap,i} = d_{ip}/(\bar{v}_p - \bar{v}_i)$. These strategies are known as *classical pursuit* and *classical evasion*, respectively.

Theorem 4.4.1 and Theorem 4.4.2 apply to this system. Again the pursuer's optimal strategy is to target the evader with the minimum T_{cap} with the one-on-one pursuit strategy. The reactive evasion constraint for this system is derived as follows.

Let $\tilde{\theta}_{p,i}$ be the angle of the pursuer's motion measured counterclockwise relative to the baseline vector \mathbf{r}_{ip} for evader E_i , and let $\tilde{\theta}_i$ be the angle of the evader's motion relative to the same vector. The reactive evasion condition for this system is the following, from substitution of the adjoint $\boldsymbol{\lambda}$ and optimal controls \mathbf{u}_p^* and \mathbf{u}_i^* into (4.14):

$$\begin{aligned} \lambda_0 \hat{\mathbf{r}}_{ip} \cdot (\mathbf{u}_i - \mathbf{u}_p) &\geq \lambda_0 (\bar{v}_i - \bar{v}_p) \\ \Leftrightarrow v_i \cos \tilde{\theta}_i - v_p \cos \tilde{\theta}_{p,i} &\geq \bar{v}_i - \bar{v}_p. \end{aligned} \quad (4.21)$$

Setting the two sides of (4.21) equal specifies a family of possible evader reactive control inputs that trade off between the necessary speed and relative angle. We examine two possible control laws in more detail: “slowing evasion” where the evader heads directly away from the pursuer with $\tilde{\theta}_i = 0$, and “spiral evasion” where maximum speed $v_i = \bar{v}_i$ is used.

We note also that if $v_p \cos \tilde{\theta}_{p,i} \leq \bar{v}_p - \bar{v}_i$, then no evader control is necessary since $v_i = 0$ will satisfy the constraint.

Slowing evasion

For the slowing reactive evasion strategy, the evader chooses its heading according to classical evasion, such that it heads directly away from the pursuer with constant $\tilde{\theta}_i = 0$. The necessary speed of the evader is given by

$$v_i^{slowing} = \bar{v}_i + v_p \cos \tilde{\theta}_{p,i} - \bar{v}_p, \quad (4.22)$$

which is maximum at \bar{v}_i when the pursuer is heading directly towards the evader with $\tilde{\theta}_{p,i} = 0$, and smoothly decreases to zero as the pursuer relative angle increases.

The equations of motion can be solved explicitly for an evader using this strategy when the pursuer takes a straight line path. Consider a pursuer starting at the origin and moving at full speed in the direction of the positive x -axis, such that $\mathbf{r}_p(t) = (\bar{v}_p t, 0)^T$. We express the evader position in polar coordinates (d_i, ψ_i) relative to the pursuer, such that

$$\mathbf{r}_i(t) = \begin{pmatrix} x_i(t) \\ y_i(t) \end{pmatrix} = \mathbf{r}_p(t) + \begin{pmatrix} d_i(t) \cos \psi_i(t) \\ d_i(t) \sin \psi_i(t) \end{pmatrix}.$$

From initial conditions $d_i(0) = d_0$, $\psi_i(0) = \psi_0$, the evader's trajectory is described by

$$\begin{aligned} d_i(t) &= d_0 - (\bar{v}_p - \bar{v}_i)t \\ \psi_i(t) &= 2 \tan^{-1} \left(\tan \left(\frac{\psi_0}{2} \right) \left(1 - \frac{\bar{v}_p - \bar{v}_i}{d_0} t \right)^{\frac{-\bar{v}_p}{\bar{v}_p - \bar{v}_i}} \right), \end{aligned}$$

for $t = 0$ until

$$t = \frac{d_0}{\bar{v}_p - \bar{v}_i} \left(1 - \left(\cot \left(\frac{\psi_0}{2} \right) \sqrt{\frac{\bar{v}_i}{2\bar{v}_p - \bar{v}_i}} \right)^{\frac{\bar{v}_i - \bar{v}_p}{\bar{v}_p}} \right).$$

at which time $v_i^{slowing} = 0$ and the evader can stop.

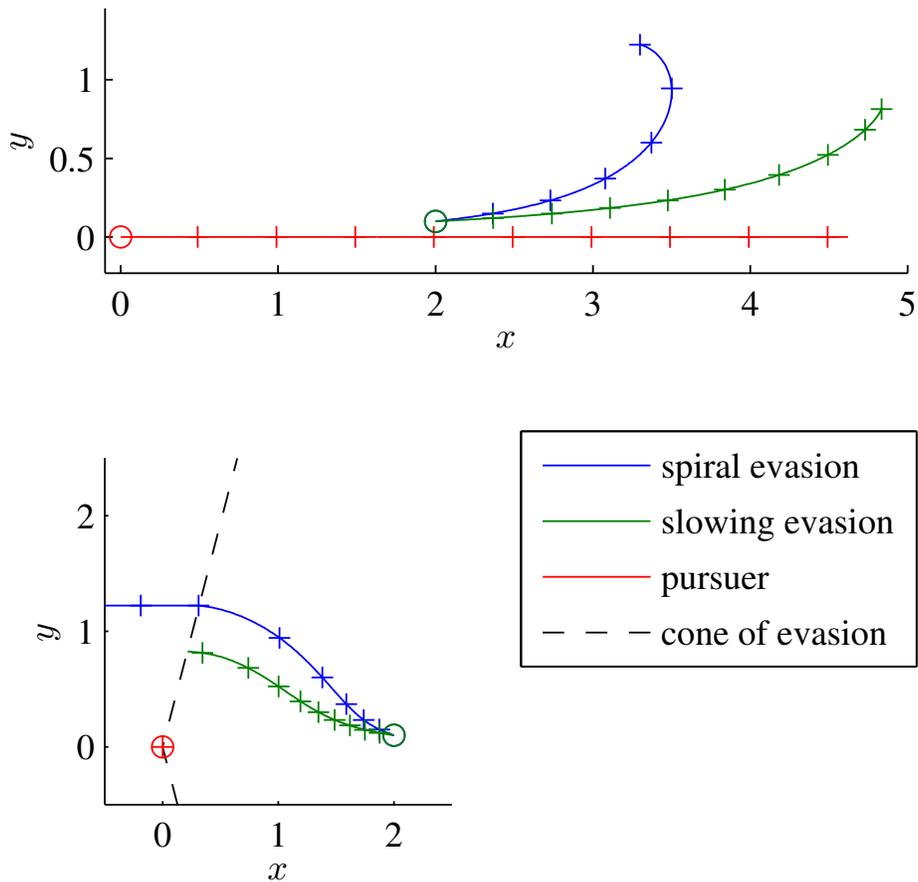


Figure 4.6: Comparison of the “slowing” and “spiral” reactive evasion strategies for pursuer with speed $\bar{v}_p = 1$ and evader with speed $\bar{v}_e = 0.75$. The pursuer starts at the origin $\mathbf{r}_p = (0,0)$ and travels along the x -axis, and the evader starts at $\mathbf{r}_e = (2,0.1)$. Circles indicate initial positions and crosses show positions at 0.5 second intervals, continuing until the slowing evader leaves the cone of evasion (reaches a point where its necessary speed is zero). Top: trajectories of the agents in the inertial frame. Bottom: the same trajectories in a frame relative to the pursuer. Spiral evasion leaves the cone of evasion at $t = 3.0$ s, and slowing evasion at $t = 4.6$ s. Note that the slowing evasion strategy leaves the cone of evasion at a smaller distance from the pursuer.

Spiral evasion

For the spiral reactive evasion strategy, the evader sets its speed to be constant at the maximum value $v_i = \bar{v}_i$, with the associated relative angle that satisfies equality for (4.21),

$$\tilde{\theta}_i^{spiral} = -\text{sgn}(\sin \tilde{\theta}_{p,i}) \cos^{-1} \left(\frac{v_p \cos \tilde{\theta}_{p,i} - \bar{v}_p + \bar{v}_i}{\bar{v}_i} \right).$$

The sign is chosen such that the evader is turning away from the direction of motion of the pursuer. For the duration of spiral reactive evasion the component of evader velocity normal to the evader-pursuer baseline vector is at the maximum value that satisfies (4.21). Consequently an evader using the strategy of spiral evasion leaves the zone where evasion is necessary in minimum time. Unfortunately it does not afford an analytic solution for trajectories as with slowing evasion. Though the spiral evasion strategy minimizes time spent evading, it may not be the most efficient in terms of energy use for an evader. Figure 4.6 shows a comparison of trajectories for an evader using the spiral and slowing strategies.

Slowing evasion and spiral evasion are just two examples out of a family of possible reactive evasion policies. Another possibility would be to set the angle to be $\tilde{\theta}_i$ to a nonzero constant, varying the speed to satisfy (4.21). This type of “constant bearing” strategy has been observed in fish [37], where evading individuals are thought to turn such that the pursuer remains at a constant position in the visual field.

4.5.2 Local target selection

Under local sensing, the pursuer must choose a control law based only on measurements of evaders in $\mathcal{N}(P)$. The optimal strategy is to target the evader in $\mathcal{N}(P)$ that can be captured in the minimum time:

$$\mathbf{u}_p = \bar{v}_p \hat{\mathbf{r}}_{gp}, \text{ where } g = \underset{E_j \in \mathcal{N}(P)}{\text{argmin}} T_{cap,j}. \quad (4.23)$$

During pursuit, if a new evader E_i moves into $\mathcal{N}(P)$ with $T_{cap,i} < T_{cap,g}$, the pursuer will switch to targeting E_i . This will decrease the remaining bound on time-to-capture. Thus the total time spent in pursuit will necessarily be less than or equal to the minimum local T_{cap} calculated at the start of pursuit. This local strategy will not lead to an underestimate of the total time spent pursuing, as the initial bound still holds throughout.

4.5.3 Local evasion strategy

Without evader turning constraints, the evader domain of danger partition as defined above in (4.12) is equivalent to a multiplicatively-weighted Voronoi diagram [2], where the weight on each evader's distance is given by $(\bar{v}_p - \bar{v}_i)$. This diagram generalizes the standard Voronoi domain of danger partition that represents Hamilton's selfish herd model of identical evaders [39].

However, unlike in the global sensing case, each evader E_i needs an evasive strategy when $P \notin \mathcal{N}(E_i)$. We propose a local evasion strategy for evader E_i that consists of 1) a *risk reduction phase* which is used when $P \notin \mathcal{N}(E_i)$, and 2) a *local reactive evasion phase* which is used when $P \in \mathcal{N}(E_i)$.

4.5.4 Risk reduction phase

We derive an evasive strategy for evader E_i when $P \notin \mathcal{N}(E_i)$ to reduce the size of its domain of danger. Let E_f and E_s be two evaders with $\bar{v}_f > \bar{v}_s$. We show that in this risk reduction phase E_f reduces the size of its domain of danger by moving close enough to E_s . The intuition is that were the pursuer to appear, E_s would be chosen over E_f as the target since E_f and E_s are approximately the same distance from the pursuer and E_s is slower than E_f .

Let $\mathbf{r}_{fs} = \mathbf{r}_f - \mathbf{r}_s$ be the baseline vector from E_s to E_f with associated distance $d_{fs} = \|\mathbf{r}_{fs}\|_2$. E_f 's domain of danger is the interior of the Apollonius circle [2] formed by the locus

of points where $T_{cap,f} = T_{cap,s}$. The circle has its center at

$$\mathbf{r}_{Apol,fs} = \mathbf{r}_f + \frac{(\bar{v}_p - \bar{v}_f)^2}{(\bar{v}_p - \bar{v}_s)^2 - (\bar{v}_p - \bar{v}_f)^2} \mathbf{r}_{fs} \quad (4.24)$$

and radius

$$R_{Apol,fs} = \frac{(\bar{v}_p - \bar{v}_s)(\bar{v}_p - \bar{v}_f)}{(\bar{v}_p - \bar{v}_s)^2 - (\bar{v}_p - \bar{v}_f)^2} d_{fs}. \quad (4.25)$$

$T_{cap,f}$ cannot be the minimum where $T_{cap,s} < T_{cap,f}$. The maximum distance from E_f to the boundary of the circle is

$$d_{App,fs} = \left(\frac{\bar{v}_p - \bar{v}_f}{\bar{v}_f - \bar{v}_s} \right) d_{fs}, \quad (4.26)$$

in the direction of \mathbf{r}_{fs} . Since $d_{App,fs}$ is proportional to d_{fs} , E_f may reduce this bound on the size of its domain of danger by approaching E_s and shrinking d_{fs} . Since $\bar{v}_f > \bar{v}_s$, E_f can always choose its velocity such that $\dot{d}_{fs} < 0$.

We define E_f as *risk minimized with respect to E_s* if its Apollonius circle with respect to E_s lies completely within its sensing range, satisfying the following condition:

$$d_{fs} + d_{App,fs} < d_{sense} \quad (4.27)$$

which is equivalent to

$$d_{fs} < \left(\frac{\bar{v}_f - \bar{v}_s}{\bar{v}_p - \bar{v}_s} \right) d_{sense}. \quad (4.28)$$

The risk reduction phase consists of each evader moving towards any chosen slower neighbor until either inequality (4.28) is satisfied or the pursuer enters the evader's sensing range and local reactive evasion is triggered. Fig. 4.7 illustrates how the domains of danger decrease in size during risk reduction.

4.5.5 Local reactive evasion phase

When pursuer P enters evader E_i 's sensing range, E_i only knows the T_{cap} of its neighbors in $\mathcal{N}(E_i)$, and the pursuer P chooses its target based only on the T_{cap} of its neighbors in $\mathcal{N}(P)$. In this context, an evader must use its best estimate of the pursuer's estimate of the minimum T_{cap} in order to decide when to begin its reactive evasion strategy. E_i should begin reactive evasion when $T_{cap,i} - T^*_{cap} \leq \epsilon$, for some $\epsilon > 0$ where T^*_{cap} is the minimum T_{cap} for the evaders in the set $\mathcal{N}(E_i)$. When $T_{cap,i} - T^*_{cap} > \epsilon$, E_i can remain in place.

The following theorem states that a risk minimized evader is guaranteed to avoid capture under local sensing if it also uses local reactive evasion.

Theorem 4.5.1. *For system (4.1) under local sensing with pursuer P using local pursuit strategy (4.23), let E_f and E_s be evaders with maximum speeds $\bar{v}_f > \bar{v}_s$. If E_f is risk minimized with respect to E_s at the time t_0 when P first enters $\mathcal{N}(E_f)$, then there exists a control input that guarantees that E_f will avoid capture.*

Proof. We must consider two cases separately. In the first case, P has targeted $E_g \neq E_f$ when it first enters $\mathcal{N}(E_f)$ E_f at time t_0 . In that case, $T_{cap,f} > T_{cap,g}$, so a local reactive evasion strategy may be used by E_f in order to avoid becoming the target.

In the second case, P targets E_f at time t_0 when it first enters $\mathcal{N}(E_f)$ and $d_{fp} = d_{sense}$. At time t_0 , P remains outside of the Apollonius circle defined by $T_{cap,f} = T_{cap,s}$ by the risk minimization condition (4.27)). Thus, P must switch to targeting E_s at the moment that E_s enters $\mathcal{N}(P)$ under its optimal target selection strategy. If we can guarantee that P will sense E_s before entering the circle, then P will switch its target before capturing E_f , and E_f will be able to use reactive evasion at that point to avoid capture. Under the local pursuit strategy with E_f as the target, $\dot{d}_{fp} < 0$, so P monotonically approaches E_f .

We will show that P cannot enter the circle without first sensing E_s , so long as E_f remains risk minimized. At the point when P senses E_s , it will switch to targeting E_s , and E_f will be able to use local reactive evasion to avoid becoming the target again.

Suppose that P does not sense E_s , i.e. $d_{sp} > d_{sense}$. Suppose also that P is on the boundary of the Apollonius circle: $d_{fp} = d_{App,fs}$. By the triangle inequality, $d_{sp} \leq d_{fp} + d_{fs}$. Substituting,

$$\begin{aligned} d_{sp} &\leq d_{App,fs} + d_{fs} \\ &\leq \left(\frac{\bar{v}_p - \bar{v}_f}{\bar{v}_f - \bar{v}_s} + 1 \right) d_{fs} = \left(\frac{\bar{v}_p - \bar{v}_s}{\bar{v}_f - \bar{v}_s} \right) d_{fs}. \end{aligned} \tag{4.29}$$

By the risk minimization condition (4.28), we have $d_{sp} < d_{sense}$. Thus by contradiction P cannot enter the Apollonius circle without first sensing E_s . \square

Consider a graph \mathcal{G} where evaders act as nodes, and an edge e_{ij} from evader E_i to evader E_j is present only if E_i is risk minimized with respect to E_j . This forms a directed graph with edges only going from a faster evader to a slower evader. Due to that hierarchy, any connected component must contain a spanning tree with the slowest evader in the component as the root.

Theorem 4.5.2. *Under the local evasion strategy, an evader can only be captured if it is the slowest evader in a connected component of \mathcal{G} .*

Proof. Let E_1 be the initial target of pursuer P under local sensing when P enters $\mathcal{N}(E_1)$ at time t_0 . If E_1 is risk minimized with respect to another evader E_2 at t_0 , then by the definition of \mathcal{G} it is not the slowest evader in its connected component, and by Theorem 4.5.1 the target of P will eventually switch to another evader. If E_1 is not risk minimized it must be the slowest evader within its connected component, and the other evaders will be able to use reactive evasion to avoid becoming the target, leading to the capture of E_1 \square

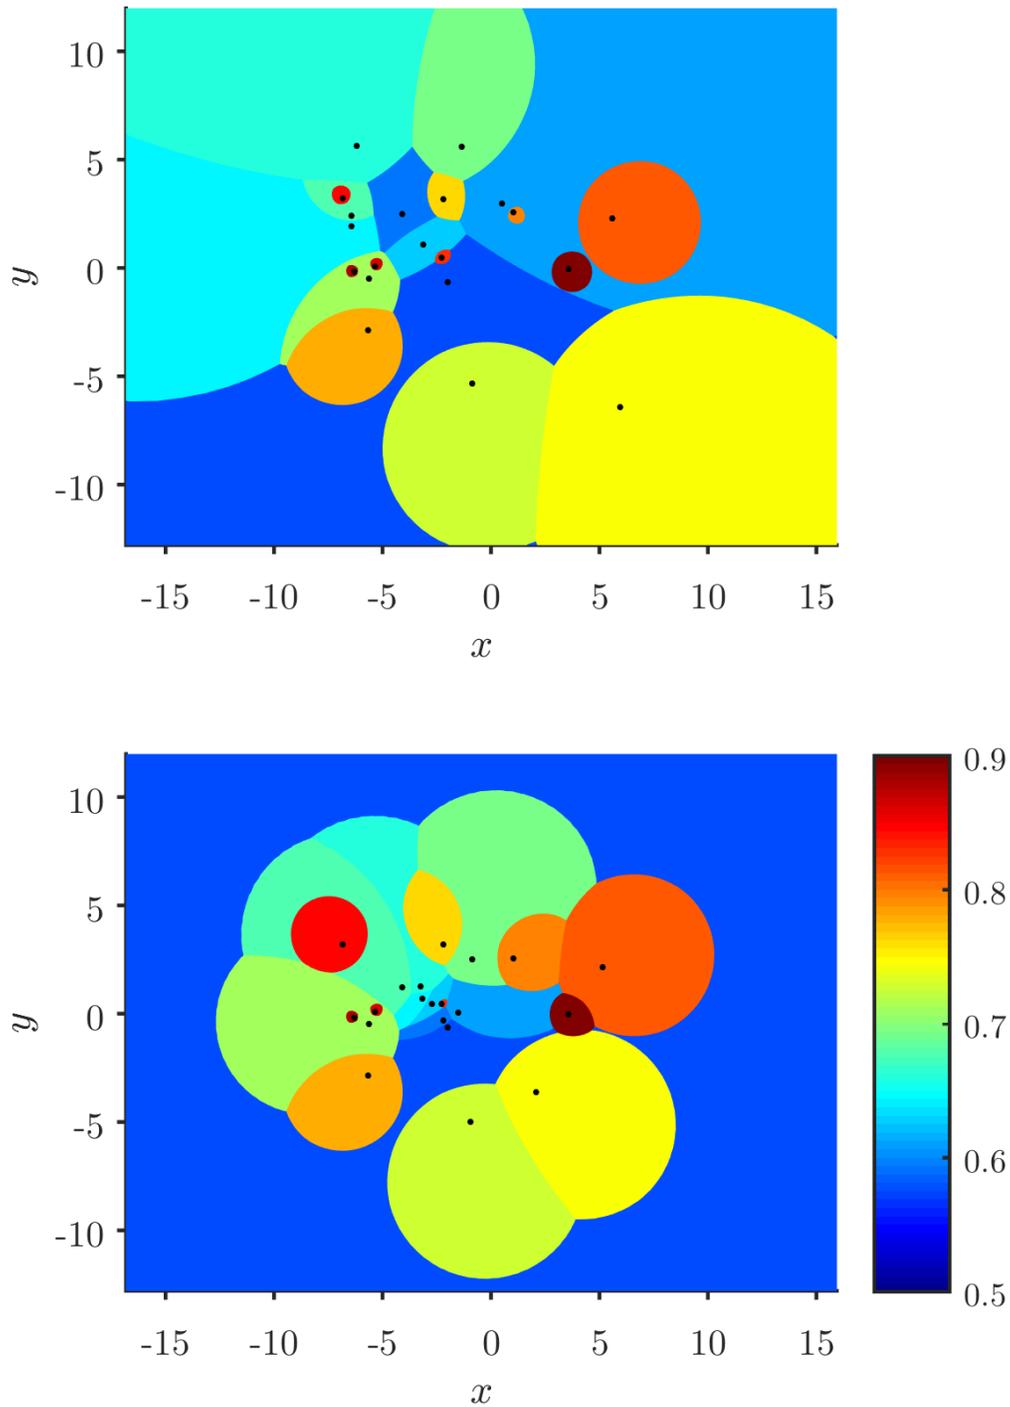


Figure 4.7: Weighted-Voronoi domain of danger partition, calculated for a pursuer with maximum speed $\bar{v}_p = 1$ and position not sensed by evaders. Each black dot denotes the position of an evader and the color of the surrounding cell (domain of danger) indicates the evader's maximum speed. Top: initially with random initial positions. Bottom: after running the risk reduction strategy for locally sensing evaders with sensing radius $d_{sense} = 10$.

4.6 Discussion

In this chapter, we have presented solutions for a pursuit-evasion differential game with practical motion constraints for the evader not considered before in this setting. The optimal strategies for the one-on-one game were used to analyze strategies in a system with multiple evaders. We showed that the optimal strategy for the pursuer is to focus on a single evader that can be captured in the minimum time, and that non-targeted evaders are always able to avoid capture by using a strategy of reactive evasion. We show how to compute decentralized state feedback reactive control laws. These reactive strategies allow an agent to remain still until necessary. The strategies also include herding behaviors. Because of the constraint on turning rate, the partition of the plane into domains of danger for the evaders depends not only on relative positions but also on relative headings.

In the case that agents have limited sensing range we have shown, for evaders without turning constraints, that a strategy for risk reduction provides a mechanism for group aggregation. This behavior could be leveraged in engineered multi-agent systems with limited sensing. For instance, if every agent (except for the slowest) stays close to at least one slower neighbor at all times, the sensing network will remain connected at all times based on the proof of Theorem 4.5.2.

A weakness of the current approach is the assumption that all agents have accurate knowledge of the motion constraints and states of other agents. To address this, the strategies will need to be adapted to uncertainties in the agent estimates of these system parameters. For example, the approach of [58], provides strategies for each agent in a three-player pursuit-evasion game, which are derived based on “worst case” values for uncertain parameters.

Chapter 5

Zebra experiment

This chapter describes a field experiment conducted at the Ol Pejeta Conservancy in Laikipia, Kenya in July 2014 to study evasive behaviors in herds of plains zebra under pursuit by an artificial predator, the “robo-lion.”

The robo-lion zebra experiment most closely relates to the the herd evasion model of Chapter 4, in that it features a single pursuer and a large group of heterogeneous evaders, and indeed the model and experiment were developed in tandem. Our result for risk minimization leading to group aggregation (presented in Section 4.5.1) sparked a discussion with Prof. Daniel Rubenstein in the Princeton University Department of Ecology and Evolutionary Biology, which led to our collaboration on the robo-lion zebra experiment. Observations from the experiment in turn led to the development of the motion model incorporating constraints on speed, turning rate, and lateral acceleration examined in Chapters 3 and 4.

Direct observation of predation has historically been difficult due to the low frequency of events, coupled with remote locations and low visibility at night or in wooded areas. By utilizing an artificial predator for our experiments, we can control the conditions for favorable data collection and ensure that the behavior of the pursuer is consistent across trials. To our knowledge, the experiment discussed in this chapter is the first example of the use of

trajectory tracking from direct video footage to study evasive behavior in a large terrestrial animal species in the field with an artificial predator.

Although the robo-lion is in some ways a poor substitute for a real predator (small size, low speed, mechanical sound), the zebras do react to it as if it were a true threat. Seventeen trials with the robo-lion were performed over the course of four days on different herds of plains zebra. Across these trials, several behavioral themes stand out:

- *Efficiency of movement:* The zebras do not start to move as soon as they notice the robo-lion, but wait until it comes closer. They match their speed to that of the robo-lion and do not exert more effort than necessary in evasion. Once the robo-lion stops, the zebras quickly stop and turn to face it.
- *Alignment with neighbors:* The zebras within a herd show a high degree of velocity alignment with one another, staying together as a group more than would be expected if each individual were to use a selfish classical evasion strategy.
- *Spatial sorting:* The different classes of zebra react differently to the threat. Mothers and foals spend less time vigilant before starting to move, and generally move to the front of the herd farther away from the pursuer. Males on the other hand stay to the rear in a seemingly defensive posture.

The results presented here are preliminary, and we do not aim to prove facts about the behavior of plains zebra, but rather to build insight into how features of our mathematical models and behaviors in biological systems are related. We will continue work on this project with a more thorough analysis to be published separately in the future, in which we will characterize sources of error in the data and assess the statistical significance of our results.

In Section 5.1, we provide background on the plains zebra and discuss previous work on pursuit and evasion in animals. Section 5.2 describes the equipment and procedures used to carry out the experiment. In Section 5.3 we discuss the process used to prepare the

data for analysis. We present our results and numerical analyses in Section 5.4, focusing on various aspects of the zebra behavior including initiation of movement, speed, distance to the robo-lion, heading alignment, relative locations within the herd, and the relationship between foals and their mothers during the chase. We finish with conclusions in Section 5.5.

5.1 Background

5.1.1 Plains zebra

The experiments were conducted in July 2014 at Ol Pejeta Conservancy in Laikipia, Kenya. Ol Pejeta Conservancy is a 360 km² fenced conservancy that features several grassy plains separated by *Acacia drepanolobium* and *Euclea divinorum* woodlands[17].

We focus our attention solely on plains zebras (*Equus burchellii*). Plains zebras are found in abundance at Ol Pejeta, allowing us to replicate our experiment on multiple different herds. Plains zebras form large herds made up of different classes of individuals, including young foals, their mothers, other adult females, and stallions, each of which may have different motion capabilities and strategies for evasion. The zebras are under a constant threat of predation from multiple predator species in the park, primarily lions (*Panthera leo*). Not unimportantly, the zebras at Ol Pejeta are accustomed to the presence of motor vehicles, allowing researchers to come close enough to film them.

Plains zebras are large-bodied ungulates that spend the majority of their time grazing. Social structure in plains zebra populations is organized in two layers. The fundamental social unit for plains zebra is a “harem” consisting of a single adult male (the “stallion”) and several unrelated adult females, along with their young offspring. The secondary structure of a “herd” is formed where multiple harems come together. Herds can also include groups of “bachelor males” made up of adult males that have not yet formed a harem of their own. While membership in a harem is stable over months or years, herds fuse and fission on a timescale of hours or days [29].

Plains zebra foals are of the “follower” type, (as opposed to “hiding”), standing and walking within an hour of birth and remaining with their mothers for between one and two years [57]. As they reach sexual maturity, females are recruited to join another harem, and males leave to join groups of bachelor males before starting their own harems.

Plains zebras at Ol Pejeta are subject to predation by lions (*Panthera leo*) and spotted hyenas (*Crocuta crocuta*). Leopards (*Panthera pardus*) and cheetahs (*Acinonyx jubatus*) typically favor smaller prey but have been observed to prey on zebra foals. Harems with more young foals tend to come together with each other to form a herd, though predator resistance for a large herd is not much more than in a harem with eight individuals [63].

Previous studies of plains zebra at Ol Pejeta suggest that lion presence can influence the zebras’ choice of habitat. The zebras tend to favor more wooded areas at night, away from the open plains where lions are more active [84]. Zebras staying on the plains at night modify their behavior to move more often, and with more turns as a possible anti-predator tactic [30].

5.1.2 Previous work

There has been much interest in recent years in understanding the hunting strategies used by large mammalian predators. The development of low cost and high resolution Global Positioning System (GPS) technology allows for unprecedented accuracy in the measurement of animal trajectories, however limited battery capacity and the difficulty of deploying collars on multiple animals necessitate tradeoffs in the design of experiments. Two main approaches have been to place a single collar on a predator and record with high temporal resolution to understand individual pursuit instances, or to place multiple collars with low measurement frequency (but longer battery life) to study trends in habitat use.

Cheetah hunting behavior has been studied with a combination of GPS and accelerometer sensors [81, 82, 40, 35]. Puma (*Puma concolor*) energetics during attacks were characterized by accelerometer measurements [80]. Low frequency GPS collars were used to study space

use by the competing predators: cheetah, lion, and spotted hyena [13]. A GPS study on lions and plains zebras showed that zebras use lower quality grazing areas to avoid lions [9]. The “ZebraNet” project focused on understanding social dynamics in plains zebra herds, and allowed communication between GPS collars [85, 84] over long timescale.

A hybrid approach can provide rich data as well, as in [52] where aerial LiDAR (Light Detection And Ranging) scans of foliage were combined with GPS collars on lions to study the effect of prey visibility on hunting strategies in different habitats.

Collecting direct video footage of an event allows the researcher to track all the members of an animal group at once. However, processing the raw video data to extract quantitative trajectories can be significantly more complex than with GPS sensors. In certain controlled environments, automated tracking software can be used to automatically detect positions and poses of animals within each video frame. For an extensive review of automated tracking in biology, see [21]. Outdoor video tracking studies have focused on aerial species, in part because of the high contrast between individuals and the sky facilitates automatic tracking. However the use of multi-view video systems necessary to reconstruct three dimensional trajectories poses additional issues. Examples include studies of collective flight trajectories in large groups of starlings (*Sturnus vulgaris*) [8] and bats (*Myotis velifer*) [46].

Use of robotic devices in animal behavior experiments is reviewed in [47]. The use of a robot disguised as an animal to influence group behavior in a leadership role has been investigated successfully in several species, including stickleback fish (*Gasterosteus aculeatus* L.) [28], zebrafish (*Danio rerio*) [61], and cockroaches (*Periplaneta americana*) [38].

The evasive actions of schools of zebrafish in response to an artificial predator have also been studied, for a predator moving on a straight line path [70], and a predator controlled through a real-time feedback control system based on visual tracking of the fish [73].

The experiment described in this chapter serves as a proof of concept that robotic artificial predator experiments can be performed in the field with large mammals, using video footage for quantitative tracking.



Figure 5.1: The “robo-lion” remote-controlled artificial predator device, measuring approximately 40 cm wide and 60 cm long. Left: side view. Right: front view.

5.2 Equipment and methods

The “robo-lion” device is a remote controlled car (Traxxas E-Maxx model 3903) with a puppet in the form of a lion mounted on top, shown in Fig. 5.1. The tan color, shaggy mane, and upraised tail serve as visual cues to suggest the form of a predatory lion (*Panthera leo*). The robo-lion is much smaller than any lion, and has a lower top speed. The top speed reached by the robo-lion in our experiments is estimated to be 8.8 m/s, during trial 14-4 (though this was an outlier, as the mean top speed across all trials is approximately 4.2 m/s). The top speed of a lion during a prey chase is reported as 13.9 m/s in [24]. Nonetheless, the zebras did respond to the approaching robo-lion with evasive action in all trials.

Two cameras were used to capture the trials: a Canon VIXIA HF R500 digital camcorder with zoom, and a wide-angle GoPro HERO3: White Edition with fixed focal length. Both cameras record 1080p resolution video at 30 frames per second. The camcorder was mounted on a tripod at eye level in the vehicle. It was used primarily to collect closeup views of each zebra before the chase began that were viewed afterward to determine the sex and approximate age of each individual. The camcorder footage during the chase was also consulted during the tracking process to help disambiguate between zebras that crossed paths with

each other. The GoPro was mounted at the top of a sturdy 6.81 m tall fishing pole. The footage from the GoPro was used for trajectory tracking.

For each trial with the robo-lion we used the following procedure:

1. Drive to find a herd of plains zebras that includes at least one young foal less than a year old.
2. Park the vehicle as close as possible without causing the herd to flee.
3. Set up the cameras and start filming, keeping as quiet as possible.
4. Wait ten minutes until the zebras have gone back to grazing.
5. Drive the robo-lion in a straight line path towards the herd, aiming for the foal when possible.

The chase would continue for approximately 30 seconds, or until the robo-lion was out of range of the controller, or out of view of the operator. In some cases we attempted a second chase after a short pause, but in practice these occurred too far away from the camera to be able to track them reliably. The analysis presented in this chapter truncates the tracks at the time that the robo-lion starts its second chase.

5.3 Data processing

The process used to extract meaningful quantitative data from the raw GoPro video consists of several steps. First, the videos from the GoPro are run through a stabilization algorithm to remove the effects of camera movement due to wind. Second, the positions of the zebras and the robo-lion are annotated manually to create the “track” data for each video. Next the parameters for the transformation from the image coordinate frame to a “ground” coordinate frame are estimated based on the location of the horizon within the video frame. Lastly, the trajectory data is processed with a smoothing algorithm to remove noise from the manual track data. The steps are described in detail in Appendix A.

Zebras were categorized into four classes based on close analysis of the camcorder footage and photographs taken prior to the robo-lion chase:

1. *Foals*: Zebras less than one year old.
2. *Mothers*: Adult females observed to nurse a foal, or otherwise associate closely with a specific foal.
3. *Males*: Adult males, either the stallion of a harem or a member of a bachelor group.
4. *Others*: Any zebras not included in the other classes. These are primarily adult non-lactating females, but may include adult males that could not be identified visually.

Zebra behavior at each video frame was categorized into one of four states:

1. *Head down*: Zebra is either lying on the ground or standing with its head at ground level, either grazing or standing still.
2. *Head up*: Zebra is standing stationary with its head up, vigilant.
3. *Moving*: Zebra is walking or running with its head up.
4. *Out of view*: Zebra is not visible in the video frame.

The timings of state changes for an individual zebra are accurate to within one tracked video frame in either direction (± 0.2 s).

5.4 Results and analysis

Tracking data was recorded and analyzed for a total of seventeen trials with the robo-lion taking place on July 13th through 16th, 2014. Table 5.1 summarizes the number of zebras of each class tracked for each trial. The total number of tracks differs from the maximum number of zebras in a single frame in some trials due to occlusion: If a zebra disappears

Table 5.1: Summary of tracking data from trials of the robo-lion zebra experiment.

date	trial	duration (s)	foals	mothers	males	others	total tracks	max zebras in a frame
13	1	40.0	1	1	1	7	10	10
13	2	59.8	0	2	1	8	11	7
13	3	30.0	0	0	1	33	34	24
13	6	59.8	1	1	1	3	6	6
13	8	48.2	0	1	1	5	7	7
14	1	40.0	1	2	0	11	14	8
14	3	41.8	0	1	2	4	7	7
14	4	30.0	3	3	2	16	24	21
14	5	42.6	0	0	2	7	9	9
15	1	19.8	1	1	2	7	11	11
15	2	60.0	2	0	1	15	18	18
15	3	45.0	0	0	1	30	31	31
15	5	54.8	1	2	2	11	16	16
15	6	30.0	0	0	0	10	10	10
16	1	90.0	2	2	2	19	25	24
16	3	40.0	1	1	0	1	3	3
16	5	24.0	1	1	3	4	9	9

behind an obstacle for some frames, it is tracked as a new zebra when it reappears if the identity is not clear to avoid the possibility of mislabeling.

The smoothed tracked trajectories are illustrated in Figs. 5.2, 5.3, and 5.4. In order to compare between trials, the coordinates in a given trial are scaled by the average speed of the robo-lion for that trial. Multiple trials are shown on the same set of axes for comparison purposes only: in actuality trials took place kilometers apart from one another in separate areas of the Ol Pejeta conservancy. We see in general that the zebra herds tend to stay together, with individuals following paths roughly parallel to the approaching robo-lion's heading.

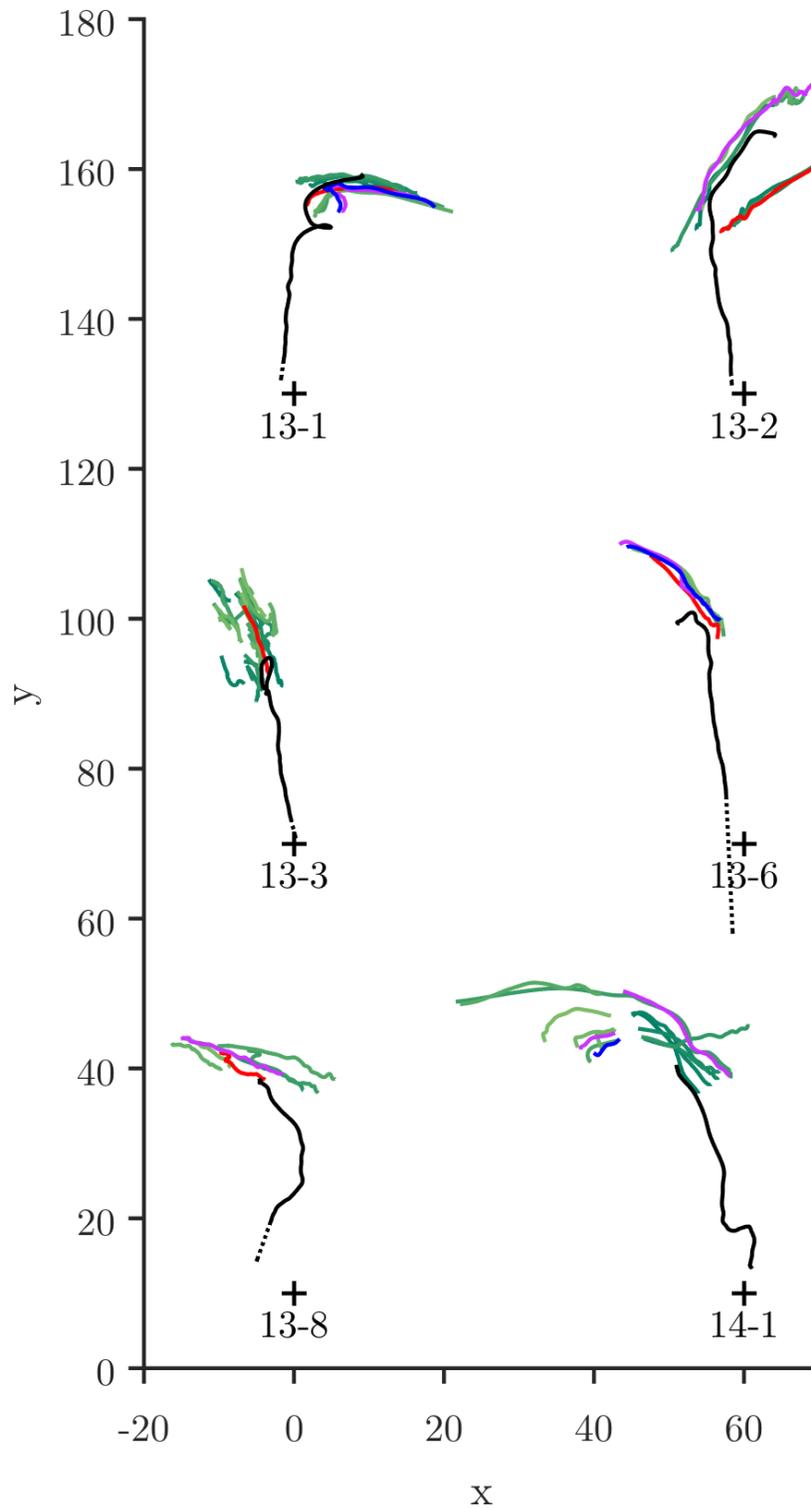


Figure 5.2: Smoothed trajectories for trials 13-1 through 14-1 in coordinates scaled by the average speed of the robo-lion. Color denotes zebra class; see Fig. 5.4 for legend. Dotted black lines indicate extrapolated robo-lion positions before it enters the camera view. Crosses represent the origin of the world coordinate frame (camera location) for each trial.

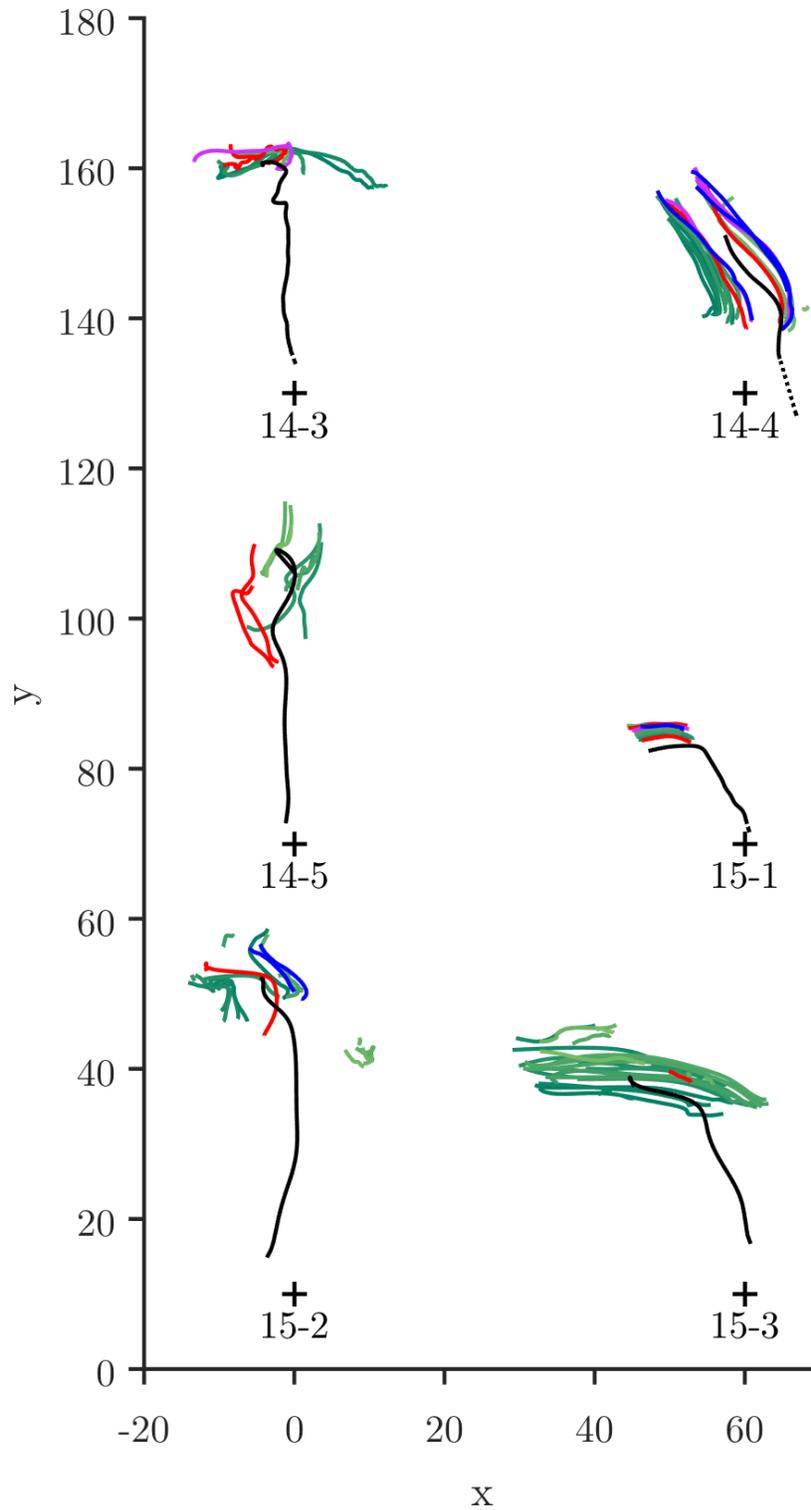


Figure 5.3: Smoothed trajectories for trials 14-3 through 15-3 in coordinates scaled by the average speed of the robo-lion. Color denotes zebra class; see Fig. 5.4 for legend. Dotted black lines indicate extrapolated robo-lion positions before it enters the camera view. Crosses represent the origin of the world coordinate frame (camera location) for each trial.

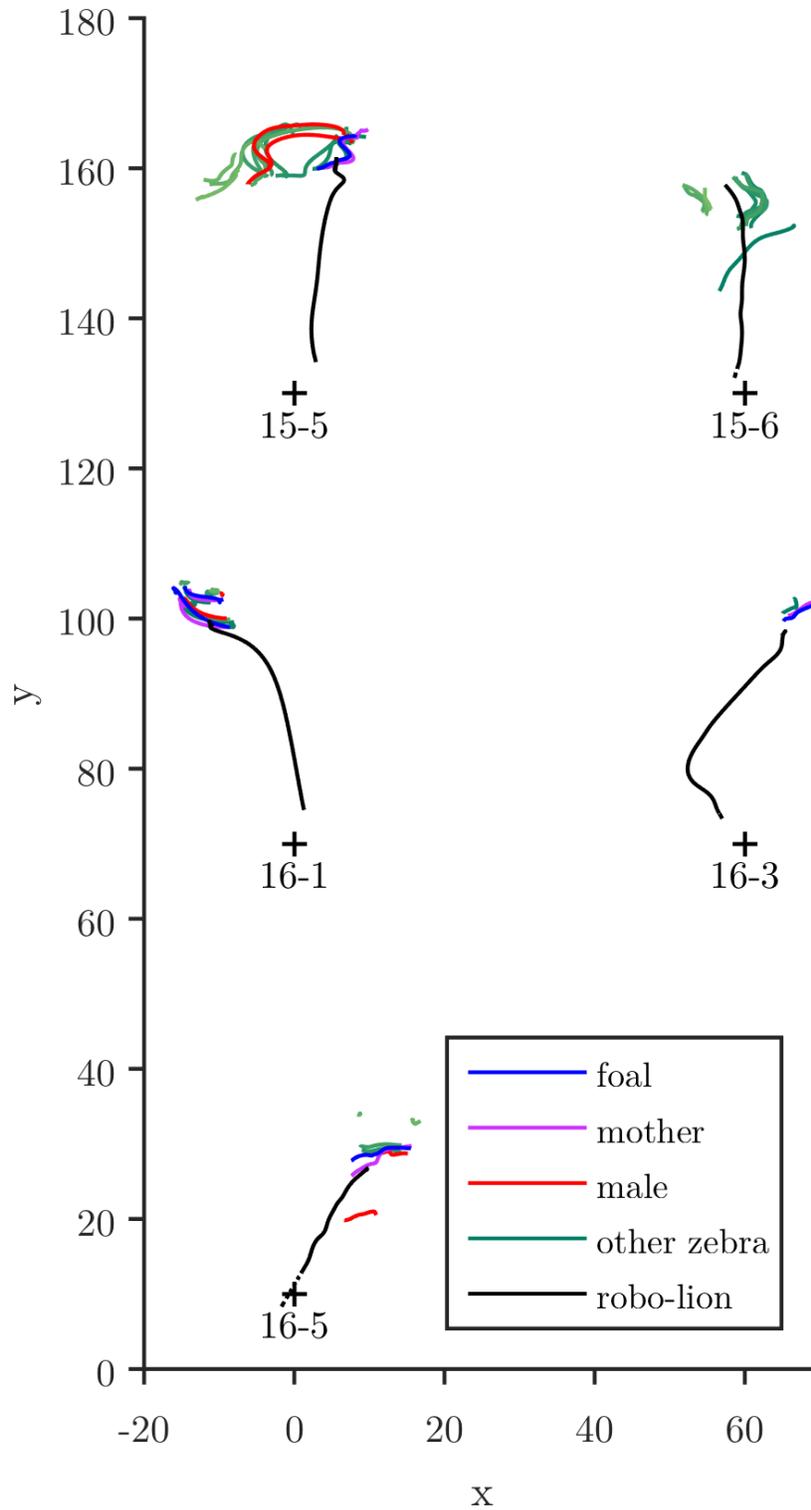


Figure 5.4: Smoothed trajectories for trials 15-5 through 16-5 in coordinates scaled by the average speed of the robo-lion. Color denotes zebra class. Dotted black lines indicate extrapolated robo-lion positions before it enters the camera view. Crosses represent the origin of the world coordinate frame (camera location) for each trial.

Table 5.2: Average, standard deviation σ , and maximum estimated speeds in m/s for each trial, including only video frames where an individual’s activity state is “moving.” Dashes indicate that no foals were present in the trial.

day	trial	robo-lion			foals			all zebras		
		mean	σ	max	mean	σ	max	mean	σ	max
13	1	2.5	0.6	4.0	2.6	1.2	4.8	2.1	1.4	5.4
13	2	4.0	1.0	5.7	-	-	-	2.7	1.5	8.2
13	3	3.3	1.5	6.7	-	-	-	1.9	1.6	15.1
13	6	2.2	1.1	3.8	0.8	0.7	2.5	0.8	0.6	2.8
13	8	1.1	0.5	2.0	-	-	-	0.9	0.5	3.3
14	1	3.1	1.1	5.0	3.2	2.5	7.1	4.3	2.3	10.7
14	3	1.9	0.8	3.6	-	-	-	1.5	0.8	3.9
14	4	7.8	0.7	8.8	8.1	3.6	15.7	7.2	4.1	16.2
14	5	3.8	1.3	5.6	-	-	-	4.2	2.4	21.1
15	1	3.1	0.5	4.2	1.1	1.5	4.3	2.4	1.4	4.7
15	2	1.9	0.8	3.3	1.9	0.8	3.6	0.9	0.7	4.4
15	3	2.0	0.7	3.0	-	-	-	2.2	1.7	6.6
15	5	1.9	1.0	3.6	0.6	0.3	1.3	0.7	0.5	2.7
15	6	2.6	0.4	3.4	-	-	-	1.8	0.7	4.4
16	1	1.2	0.8	2.4	0.5	0.5	1.4	0.4	0.3	1.8
16	3	2.0	0.6	3.0	1.5	0.7	2.9	1.5	0.7	3.4
16	5	3.0	0.5	3.6	1.6	0.9	3.9	1.8	1.1	5.5

5.4.1 Speed

Though the manufacturer claims that the remote-controlled car used in the robo-lion can reach top speeds of greater than 30 miles per hour* (13.4 m/s), the top speed estimated from smoothed tracking data is 8.8 m/s, during trial 14-4 (though this was an outlier, as the mean top speed across all trials is approximately 4.2 m/s). The inconsistent and low speeds were likely caused by a combination of difficult grassy terrain and partial battery discharge. Nonetheless, in each trial the speed of the robo-lion during the chase is significantly less than the top speed of a plains zebra, reported as approximately 16.0 m/s in [24]. Table 5.2 summarizes the average and maximum speeds for the robo-lion, zebra foals, and all zebras for each trial.

*<https://traxxas.com/products/models/electric/emaxx-tsm>

In every trial with foals present, except for trial 15-5, the top speed of each foal is greater than the average speed of the robo-lion in that trial. This suggests that all of the zebras including young foals have the ability to outrun the robo-lion.

In order to compare results between trials, we normalize the smoothed trajectory data for a given trial by dividing the spatial coordinates of the individuals by the average robo-lion speed for that trial. One unit in these “scaled coordinates” represents the average distance traversed by the robo-lion in one second. For the remainder of this chapter, all analysis is based on scaled coordinates.

Using the scaled coordinates, we calculate the average speed of each zebra across the video frames when its state is “moving.” Combining the 17 trials, the average speeds with standard deviation for each class of zebra are: 0.64 ± 0.50 unit/s for foals, 0.60 ± 0.49 unit/s for mothers, 0.62 ± 0.48 unit/s for males, 0.69 ± 0.61 unit/s for others, and 0.67 ± 0.58 unit/s overall including all classes. Since these average speeds are all less than one in the scaled coordinates, it means that the zebras on average ran slower than the robo-lion over the course of the chase. This indicates efficiency of motion, in that the zebras did not run significantly faster than necessary to avoid capture by the robo-lion.

5.4.2 Initiation of movement

Zebra activity at each tracked frame was classified into one of four states: head down (grazing), head up (alert), moving, and out of view. When possible, we would set up the camera equipment and wait for the majority of the herd to go back to a head down state before deploying the robo-lion, so that we could gather data on the time spent alert before movement for each zebra. Fig. 5.5 shows the fraction of the visible zebras in each state over time for each trial.

For zebras in the head down state at the start of the trial, we calculated the time spent in the head up state before the initiation of movement. Individuals starting in the head up or moving state were excluded. Fig. 5.6 illustrates the differences between the four classes

of zebra with respect to time spent in the head up state. The mother zebras spent notably less time with their heads up before initiating movement, with all but one spending less than 8 seconds with their head up prior to movement. The scaled distance to the robo-lion at initiation of movement is shown for each zebra not initially moving in Fig. 5.7. There is no apparent difference between the different classes with respect to the distance at the start of movement.

5.4.3 Distance to the robo-lion

The success of a pursuit is measured by the minimum distance to the evader. A real lion must come within a few meters before it can leap forward and attempt to grab its prey. We calculate the minimum distance to each zebra attained by the robo-lion in the scaled coordinates, shown organized by trial in Fig. 5.8.

The minimum scaled distance to the robo-lion among all foals is 2.28 units. Among all mothers, the minimum scaled distance is 1.97 units. Among males it is 1.71 units, and among others it is 1.01 units. The minimum scaled distance to the robo-lion out of all trials occurs for zebra 5 (a non-mother adult female) in trial 13-1 at time $t = 33$ s, with a distance of 1.01 scaled units. At that time, zebra 5 is already in motion with a scaled speed of 1.78 units/s, so the robo-lion with an average scaled speed of 1 unit/s would never be able to catch up.

The variation in minimum distances between trials may be due largely to differences in herd vigilance. A predator that sneaks up on a herd without being detected can come closer than if it approaches while the herd is watching. For a predator with a lower top speed and less stamina than a zebra, such as a lion [24], capture will be impossible unless the predator can close the distance to its target before the target accelerates past the predator's speed. Once the zebra reaches the speed of its pursuer, the distance will start to increase, assuming they are both traveling in the same direction.

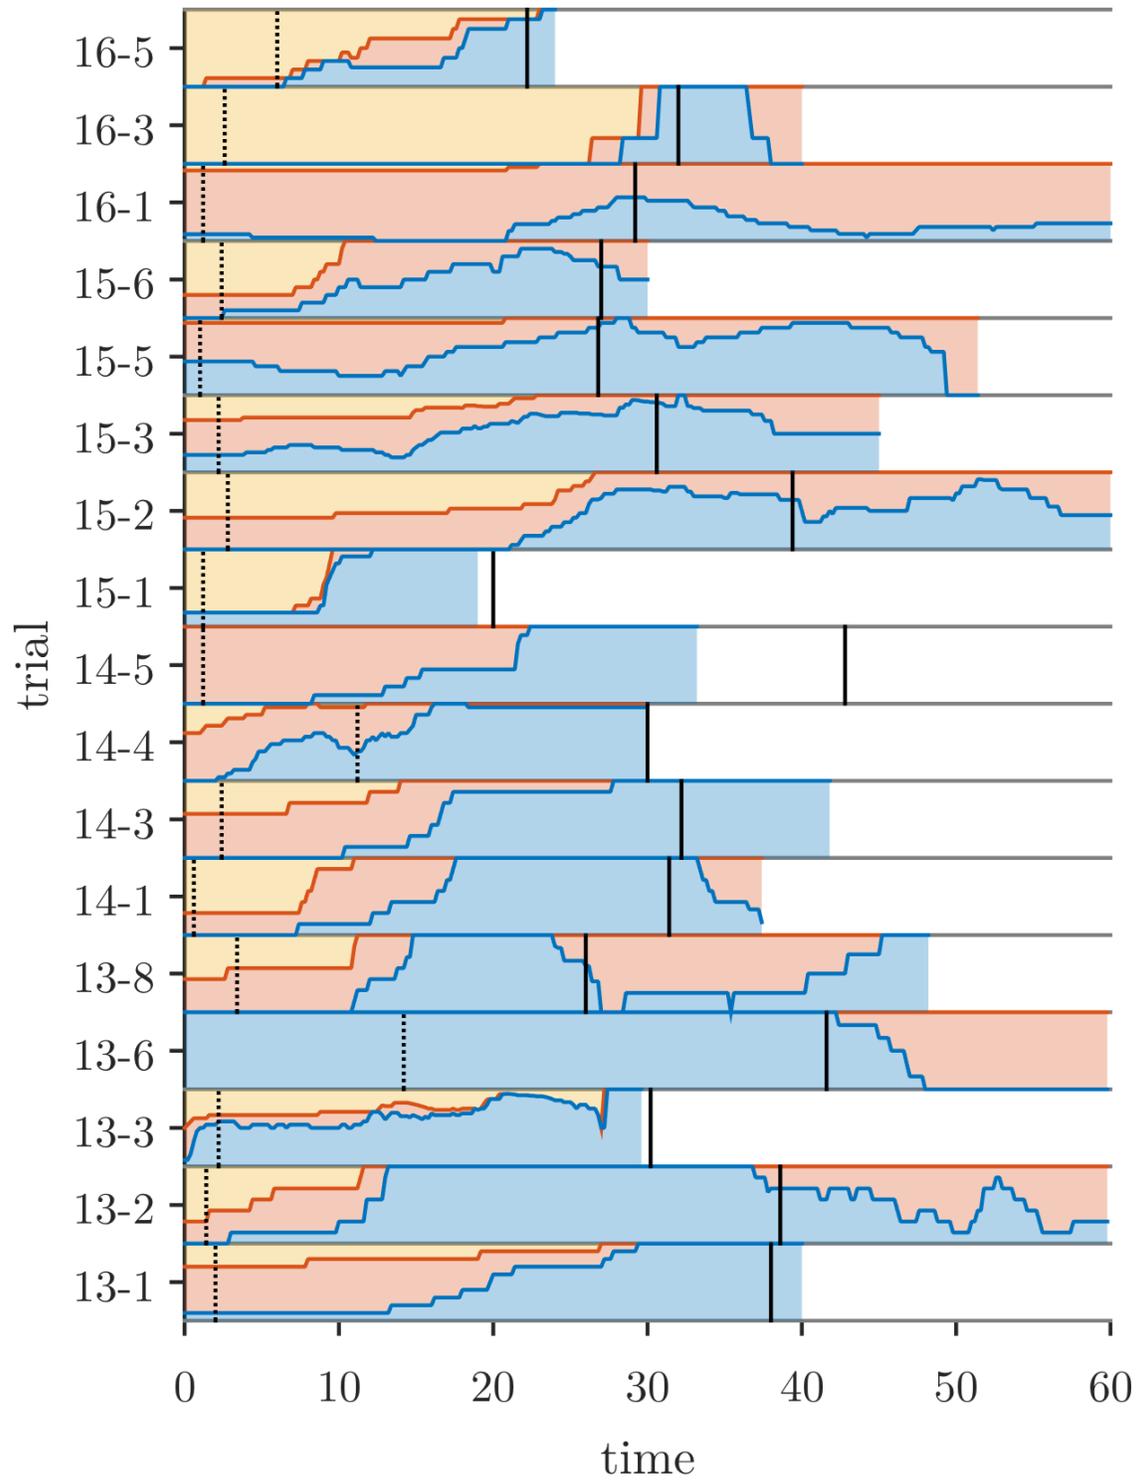


Figure 5.5: Fraction of visible zebras in each activity state over time. Colors denote activity state: Blue is moving, red is head up, and yellow is head down. The vertical dashed line shows the time when the robo-lion comes into view, and the solid vertical line shows the time when the robo-lion comes to a stop.

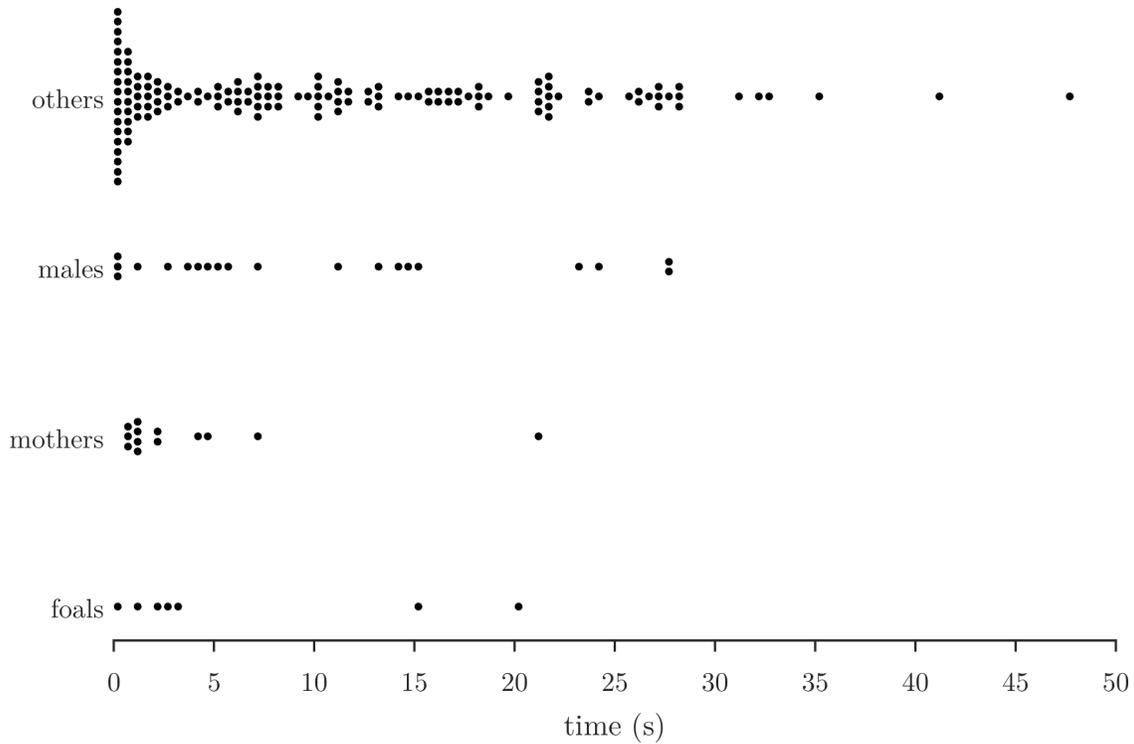


Figure 5.6: Time spent with head up before initiating movement for each individual in the four classes of zebra, combined from all tracked trials. Individuals with head up at the start are excluded. Each dot represents a single zebra. Data is binned into 0.5 s intervals. Head up durations are accurate to within two tracked video frames, or 0.4 s.

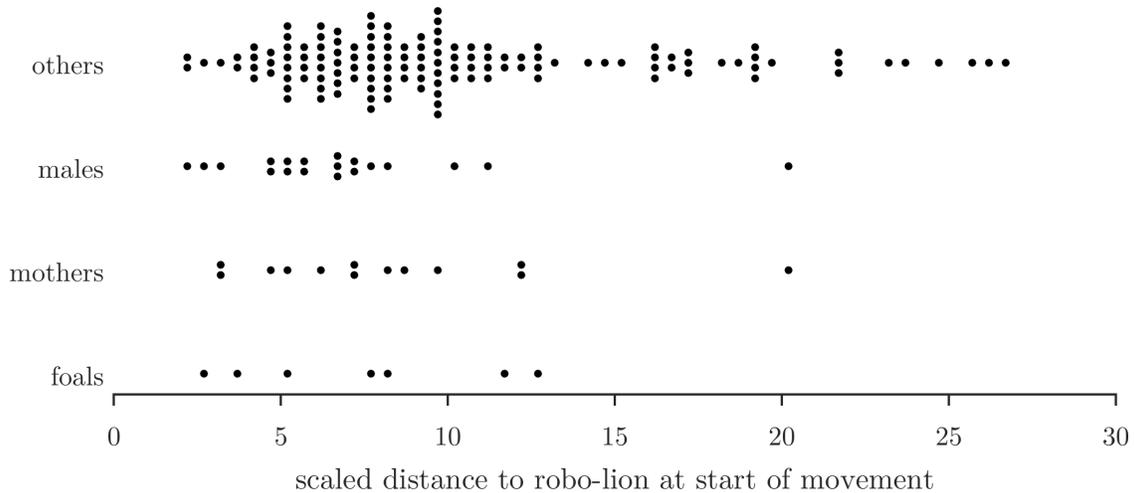


Figure 5.7: Scaled distance to robo-lion at start of movement for each individual in the four classes of zebra, combined from all tracked trials. Individuals moving at the start are excluded. Each dot represents a single zebra. Data is binned into 0.5 unit intervals.

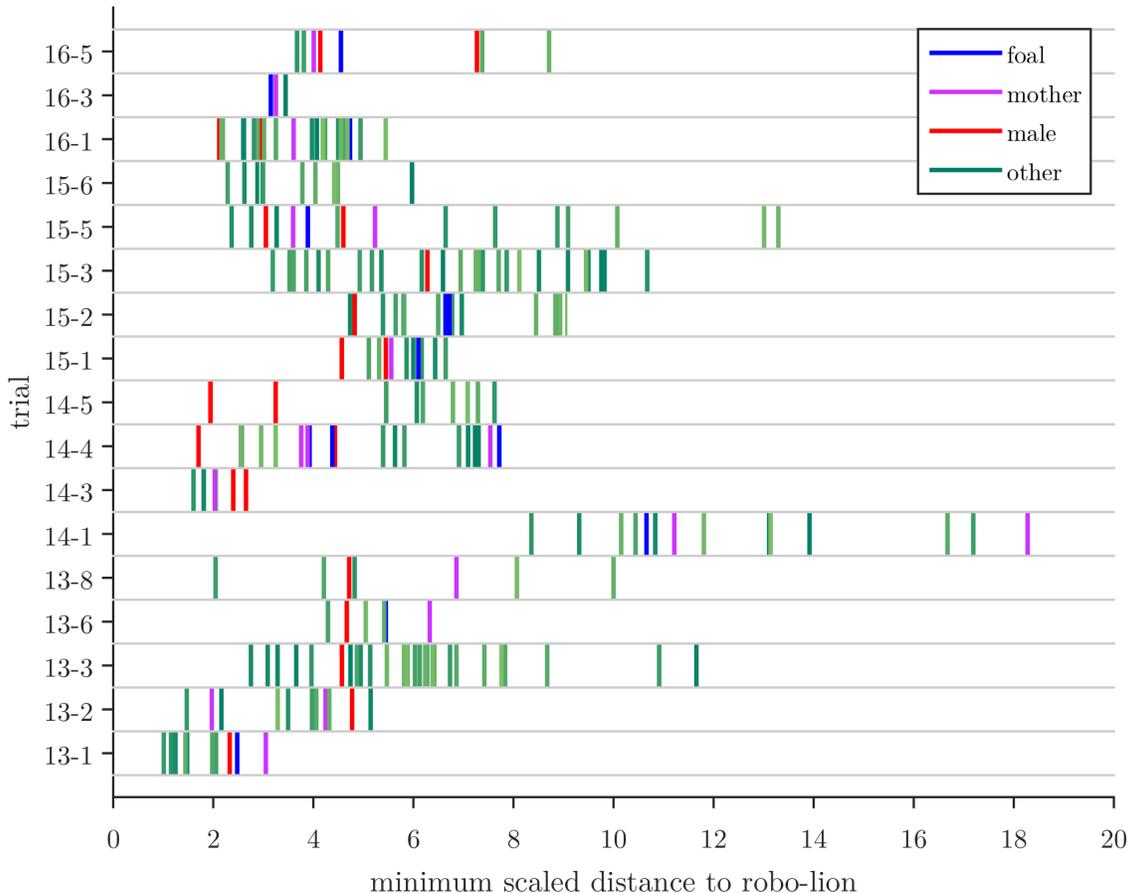


Figure 5.8: Minimum scaled distance to the robo-lion for individual zebras in each of the tracked trials.

5.4.4 Dynamics of heading alignment

In order to remain together while in motion, the zebras must align their heading directions. We can study the cohesiveness of the herd by computing order parameters that encode alignment, both among the zebras and in relation to the robo-lion. For zebra Z_i at time t with heading angle $\theta_i(t)$, let $\mathbf{t}_i = (\cos \theta_i, \sin \theta_i)^T$ be the unit vector in the direction of its heading, and let $\mathbf{r}_i(t)$ be its position in the scaled coordinates. Similarly let $\mathbf{t}_{lion} = (\cos \theta_{lion}, \sin \theta_{lion})^T$ for the robo-lion with heading $\theta_{lion}(t)$ at time t , and let $\mathbf{r}_{lion}(t)$ be its position in the scaled coordinates.

For a herd of n zebras Z_1 through Z_n , we define the “herd alignment” order parameter $\Phi_{herd}(t)$ at time t as

$$\Phi_{herd}(t) = \frac{1}{n} \left\| \sum_{i=1}^n \mathbf{t}_i(t) \right\|.$$

The herd alignment order parameter takes values in the range $[0, 1]$, with $\phi = 1$ if and only if all zebras in the herd are facing the same direction ($\mathbf{t}_i = \mathbf{t}_j$ for all zebras Z_i, Z_j within the herd).

To measure how closely the headings of the zebras match the heading of the robo-lion, we define the “herd–lion heading correlation” order parameter Φ_{lion} at time t as

$$\Phi_{lion}(t) = \left(\frac{1}{n} \sum_{i=1}^n \mathbf{t}_i(t) \right) \cdot \mathbf{t}_{lion}(t).$$

The herd–lion heading correlation takes values in the range $[-1, 1]$. We note that as it is defined, $\Phi_{lion} \leq \Phi_{herd}$ in all cases, with $\Phi_{lion} = 1$ if and only if all zebras in the herd are facing the same direction as the robo-lion ($\mathbf{t}_i = \mathbf{t}_{lion}$ for all zebras Z_i within the herd). Conversely, if and only if all lions are facing the opposite direction of the lion, then $\Phi_{lion} = -1$.

Lastly we define the “herd evasiveness” order parameter $\Phi_{evade}(t)$ at time t as

$$\Phi_{evade}(t) = \frac{1}{n} \sum_{i=1}^n \left(\mathbf{t}_i(t) \cdot \frac{\mathbf{r}_i(t) - \mathbf{r}_{lion}(t)}{\|\mathbf{r}_i(t) - \mathbf{r}_{lion}(t)\|} \right).$$

This order parameter measures on average how close the heading of each individual zebra in the herd matches the direction of classical evasion, the direction directly away from the robo-lion. $\Phi_{evade}(t) = 1$ only if each individual zebra is using a classical evasion strategy, and $\Phi_{evade}(t) = -1$ only if each zebra is facing directly towards the robo-lion.

There is an implicit tradeoff between herd alignment and herd evasiveness. The only situation where both Φ_{herd} and Φ_{evade} both equal one is if the positions of all of the zebras lie on the line of sight of the robo-lion, and all of the zebras and the robo-lion are facing in the same direction.

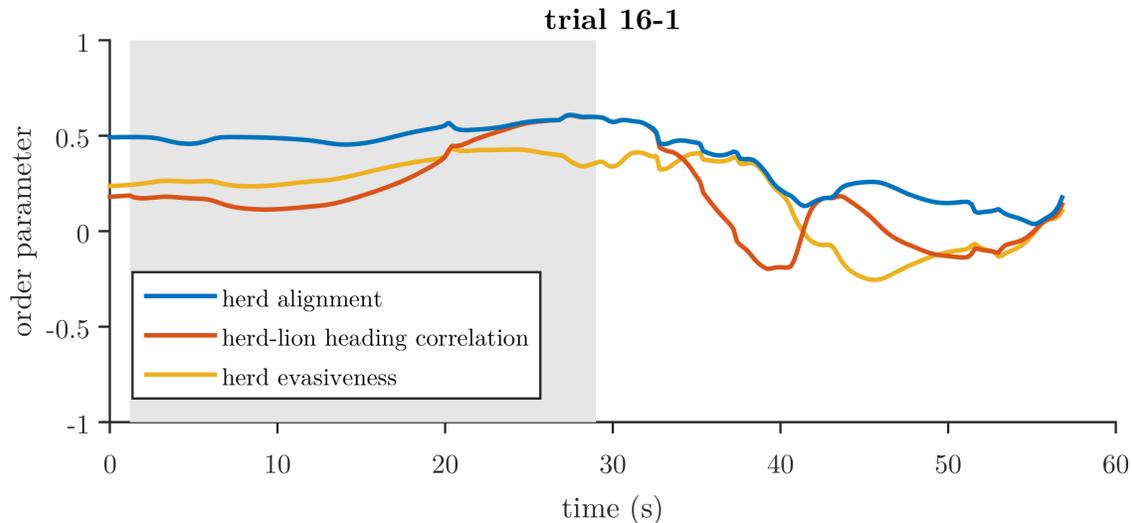


Figure 5.9: Herd alignment and evasiveness order parameters over time for trial 16-1, with a maximum $n = 23$ zebras in frame at one time. The shaded region represents video frames where the robo-lion is in view and moving. Robo-lion heading is extrapolated as constant outside the shaded region.

We find that for all tracked trials, the herd alignment order parameter is higher than the others when averaged over time. This suggests that the zebras are more concerned with herd cohesion than their own evasion. However, the fact that the robo-lion is no match for the zebras with respect to speed means that they can avoid capture without utilizing high-effort strategies like those discussed in Chapter 4.

Fig. 5.9 shows the timeseries of the three order parameters for trial 16-1. Herd alignment is highest throughout, with the other two varying during the chase. Once the robo-lion stops, the three order parameters approach zero as some members of the herd turn around to face the robo-lion in a defensive posture.

5.4.5 Locations within the herd

From watching the video footage, it appears that young foals and their mothers stay towards the front of the herd during the chase, and the male stallions seem to steer the group from the rear. To measure the extent of this spatial sorting, we define the “forward herd position”

parameter a which represents how far an individual is from the center of mass of its herd, projected onto the average direction of motion of the herd. For zebra Z_i in a herd of n zebras, its forward herd position at time t is calculated as

$$a_i(t) = (\mathbf{r}_i(t) - \mathbf{r}_{herd}(t)) \cdot \frac{\mathbf{v}_{herd}(t)}{\|\mathbf{v}_{herd}(t)\|}$$

where \mathbf{r}_{herd} is the herd center of mass with each zebra weighted equally, and $\mathbf{v}_{herd}(t)$ is the velocity of the center of mass, equivalent to the average velocity of zebras in the herd.

Fig. 5.10 illustrates the average forward position for each zebra, arranged by trial and by class. Taking the average within each class combined for all trials, the foals ($a = 1.37 \pm 2.60$) and mothers ($a = 1.78 \pm 3.27$) are more often found towards the front of the herd ($a > 0$), and males ($a = -0.17 \pm 2.60$) and others ($a = -0.31 \pm 4.45$) are found towards the rear ($a < 0$). Although the differences in means are small, with considerable standard deviations, it supports the theories that mothers and foals seek to keep the bulk of the herd between them and the pursuer, and that males lag towards the rear in order to defend their harem from the attacker.

There is a large apparent variation in herd size between different trials. In some cases this is due to some zebras not participating in the chase. For instance, in trial 14-4, one zebra does not regard the robo-lion as a threat, and remains stationary as the robo-lion passes it.

Additionally in some trials, the different harem groups making up the herd do not come together, seeming to evade from the robo-lion as separate independent units. It may be that if we study zebra positioning within an individual harem, rather than in the whole herd, then the spatial sorting of mothers, foals, and stallions will become more clear.

5.4.6 Dynamics of mother-foal pairs

We were especially interested in studying the behavior of mother and foal pairs to see how they react to a pursuer. From careful analysis of the camcorder footage taken before the

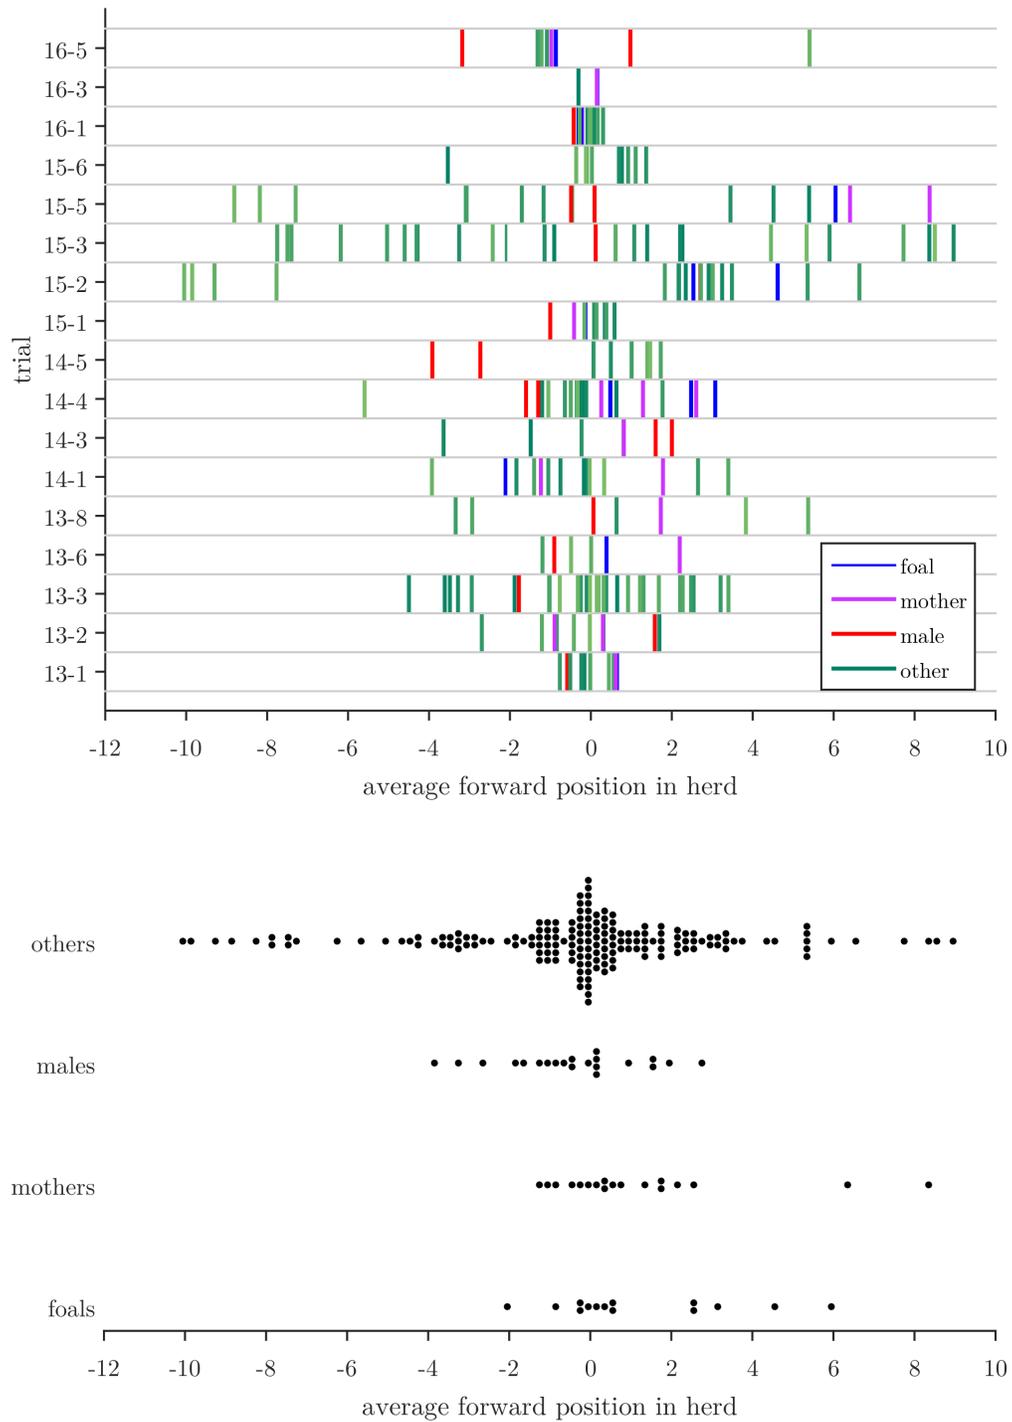


Figure 5.10: Average forward position in herd in scaled coordinates for each tracked zebra. Standard deviations for individual zebras are typically on the order of 1 or 2 units. Top: arranged by trial with color denoting zebra class. Bottom: arranged by zebra class, binned in 0.2 unit increments. Each dot represents a single zebra.

robo-lion chases, eleven mother-calf pairs were identified in the tracked zebra data across several trials: 13-1, 13-6, 14-4 (3 pairs), 15-1, 15-5, 16-1 (2 pairs), 16-3, and 16-5. Trajectories of the mother-foal pairs and the robo-lion are illustrated in Fig. 5.11.

In all cases, the mother-foal distance remained below the foal-lion and mother-lion distances at all times, and the ratio of mother-foal distance to foal-lion distance was never observed to be greater than 0.4. In the four cases that the mother initiates movement first, the foal follows by moving within 2.8 s. In four cases, the foal is moving at the start of the trial. In two cases the mother and foal initiate movement at the same time. The last case has the foal leading its mother in movement initiation by 2.4 s.

We calculate the relative angle between lion-zebra baseline vectors for the pairs, as in the angle ψ in the shape coordinates of the three-agent pursuit-evasion model in Chapter 2. This angle measures how close the robo-lion comes to a state with the robo-lion directly between the mother-foal pair, which would correspond to $\psi = 180^\circ$. Across all mother-foal pairs, we found the maximum relative heading difference to be only $\psi = 16.1^\circ$.

The data suggests that the zebra foals pay close attention to their mothers in the initiation of movement, and follow them during the chase.

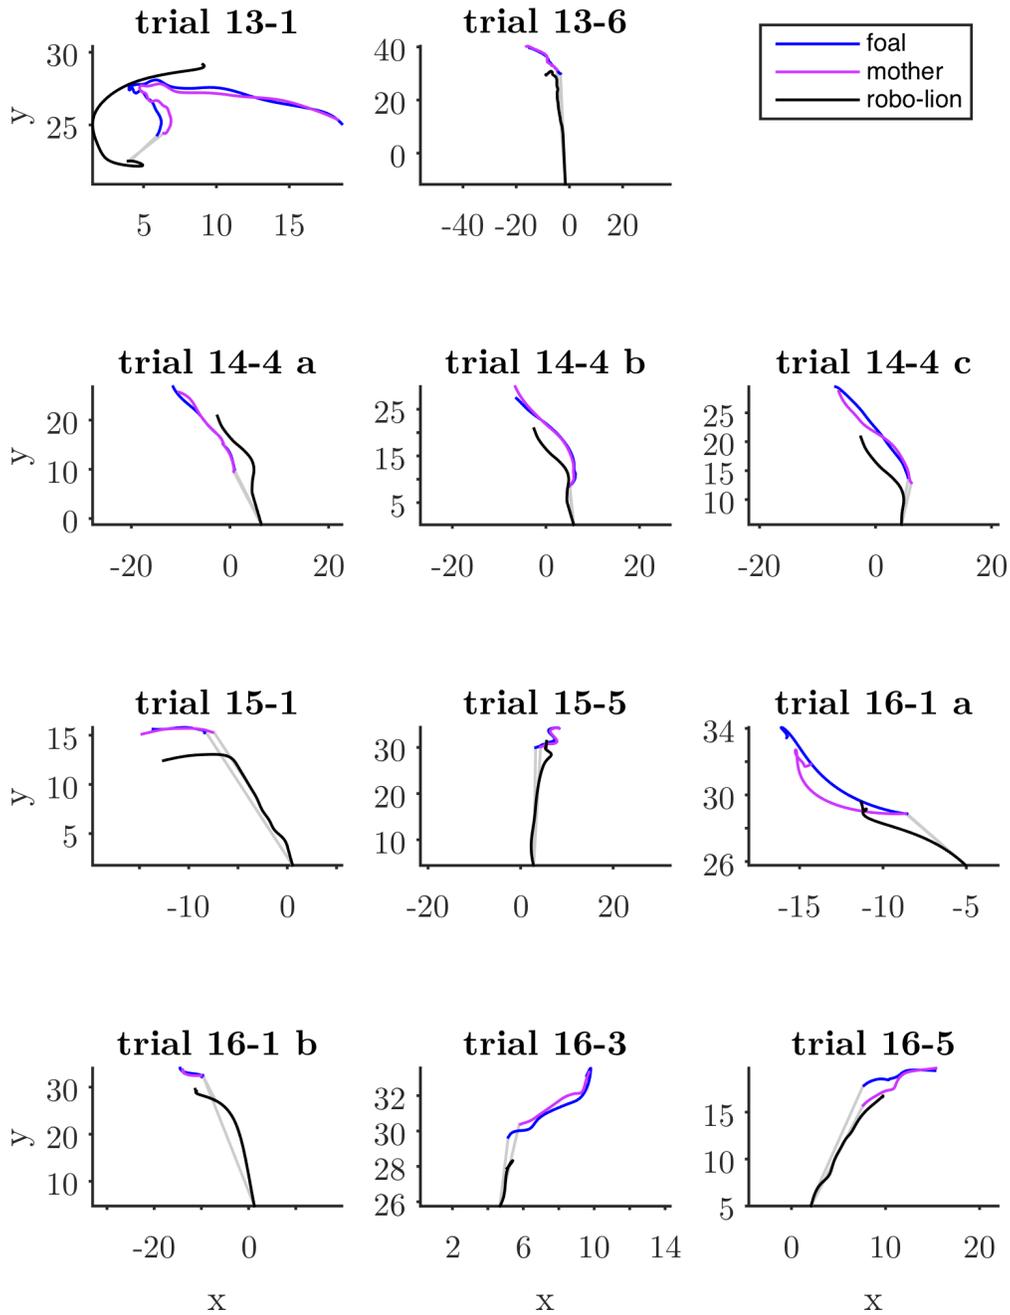


Figure 5.11: Trajectories of mother-foal pairs with the robo-lion, in the robo-lion speed-scaled coordinate frame. Trajectories begin at the time of first movement from either zebra, and end when the tracks end, or prior to robo-lion starting its second chase, whichever occurs first. Grey lines indicate the baseline vector from lion to zebra at the time when the zebra begins moving.

5.5 Conclusion

In this chapter we presented results and analysis for a field experiment to study the evasive behavior of plains zebra in response to an artificial predator, the robo-lion. We showed that it is possible to use video footage from a single camera to extract quantitative trajectory data for all members of the herd, which allows us to study many aspects of the herd behavior.

In many ways the zebras behaved in a manner consistent with our mathematical models, and the differences in behavior from our predictions serve as inspiration to develop new models. We found that the behavior of zebra foals and their mothers differed from the others in the herd with respect to movement initiation and relative position in the herd, in that they spent less time with their heads up before starting to move, and tended to be found in the front half of the herd over the course of the robo-lion chase. The mother and foal pairs stayed close together while evading, avoiding separation as in the model of Chapter 2.

The zebras did not apparently use their full effort to evade, only moving when necessary and as fast as necessary to keep a safe distance from the robo-lion. This is consistent with our model of reactive evasion from Section 4.4, however we did not account for the situation of group evasion from a pursuer with a lower top speed.

Though the results are still preliminary, this experiment has served as both a validation of our modeling approaches and inspiration for the development of new models of pursuit and evasion.

Chapter 6

Final remarks

In this thesis we have discussed several systems of pursuit and evasion with multiple evaders. We looked at feedback control laws trading off between herding and evasion in Chapter 2. We solved the problem of reaching a point on the plane in minimum time for an agent with constraints on speed, turning rate, and lateral acceleration in Chapter 3. In Chapter 4, we analyzed a differential game of pursuit featuring an evader with constraints on speed, turning rate, and lateral acceleration, against an omnidirectional, but speed-limited, pursuer, and extended that model to include multiple heterogeneous evaders, and derived strategies that allow individuals to avoid capture under different conditions. Finally in Chapter 5 we discussed preliminary results from a field experiment that looked at the evasive actions of plains zebra in response to pursuit from an artificial predator.

In Section 6.1 we summarize our results from each chapter, and in Section 6.2 we offer some ideas for future research in the mathematical modeling of pursuit and evasion, and in the use of robotic systems to aid in biological research in the field.

6.1 Conclusions

In Chapter 2 we presented a nonlinear model for a pursuer and two evaders inspired by predation in caribou. We defined strategies a priori in the form of feedback control laws

such that for each agent a single parameter encodes a tradeoff between two goals by linear interpolation. The pursuer trades off between targeting each of the two evaders for classical pursuit. The evaders each trade off between a herding strategy of approaching each other, and a classical evasion strategy of moving directly away from the pursuer.

We took inspiration from the example of caribou predation in nature to define the capture condition. A young caribou is most vulnerable when it is separated from its mother, so we label one evader as the calf and one as its mother, and define capture as the condition that the pursuer reaches a point between the two evaders.

By studying a model with constant speeds and no turning constraints on the agents, we were able to reduce the full system of three agents moving on the plane to a three-dimensional reduced system of shape-variables describing the relative positions of the agents. We showed that if the two evaders both use a pure evasion strategy, the pursuer is always able to separate them from each other. In the case of the mother using a pure evasion strategy we prove conditions such that the pursuer cannot come between the mother and the calf. We showed that as long as the mother is initially closer to the calf than to the pursuer, a calf strategy of pure herding will always avoid capture.

We next turned to the study of optimal strategies for an agent with constraints on speed, angular turning rate, and lateral acceleration. The constraints were chosen as a step towards a biologically relevant model of terrestrial animal locomotion while still keeping the steered-particle equations of motion. In Chapter 3 we solved the problem of reaching a point on the plane in minimum time for an agent with these motion constraints, and in Chapter 4 we considered a pursuit and evasion system with this type of agent as the evader. We demonstrated the connection between minimum-time problems and pursuit-evasion differential games by showing how the extremal trajectories generated through Pontryagin's minimum principle are identical in both systems.

For the minimum-time problem in Chapter 3, we derived analytic expressions for the optimal open-loop control switching times for any destination point on the plane. In addition,

we presented a state-feedback formulation of the optimal control based on the relative position of the destination in a body-fixed frame. Special cases for relaxed and extreme values of the lateral acceleration constraint were also considered, with minimum-time trajectories derived in both open-loop and state-feedback form in each case. For highly constrained lateral acceleration, we showed that the minimum-time trajectories resemble those found for differential-drive robots, consisting of rotation in place and forward straight-line motion.

In Chapter 4 we studied agents with constraints on speed, turning rate, and lateral acceleration as evaders against an omnidirectional pursuer with limited speed, both in a one-on-one setting as a differential game and in a system with multiple heterogeneous evaders interacting with each other. For the one-on-one differential game, we derived optimal trajectories for each agent for any initial conditions, and showed how feedback control laws could be computed based on the pursuer’s position relative to the evader in the reduced coordinates. We found that the optimal pursuer strategy is to follow a straight line path, and the optimal evader strategy is to follow a trajectory that is equivalent to the minimum-time trajectory to reach its location at capture.

We used the optimal strategies from the one-on-one differential game as building blocks to derive strategies in a system with a single pursuer and a group of evaders with heterogeneous motion constraints. We showed that the optimal strategy for the pursuer is to focus on a single evader that can be captured in the minimum time, and that non-targeted evaders are always able to avoid capture by using a strategy of *reactive evasion*. Under *reactive evasion*, we show that evaders are only forced to move evasively when they are close to becoming the target, which is equivalent to the pursuer approaching the edge of their domain of danger. This efficiency of motion gives the evaders the freedom to enact other behaviors at times when they are not in danger, such as a control law that brings evaders to a desired formation or maintains group cohesion.

We also considered the case in which the pursuer and evaders have radius-limited sensing. In the system with limited sensing, we showed that evaders without constraints on turning

rate can avoid capture through a strategy of risk reduction whereby each individual evader approaches a slower neighbor. The strategy provides a mechanism for group aggregation, as the selfish actions of each individual bring the group closer together as a whole. We suggested how the risk-reduction strategy could be adapted in engineered multi-agent systems with limited sensing to ensure that the sensing network remains connected at all times.

In Chapter 5 we presented preliminary results from a field experiment that used an artificial remote-controlled predator to study the evasive behaviors of plains zebra. We showed that it is possible to use video footage from a single camera to extract quantitative trajectory data for all members of the herd, which allows us to study many aspects of the herd behavior. In many ways the zebras behaved in a manner consistent with our mathematical models, and the differences in behavior from our predictions serve as inspiration to develop new models.

From analysis of the robo-lion chases, we identified a few consistent motifs in the evasive behavior of the zebras. The zebras were efficient in their movement, only starting when necessary and matching the speed of the robo-lion to keep a safe distance. The zebras showed a preference for herd cohesion over individual evasion by aligning their velocities with each other throughout the chase. Lastly, the different classes of zebras appeared to differ in their reactions to the robo-lion. The mother and foal pairs seemed to spend less time vigilant before deciding to move, and they were found in the front half of the herd more often than males, on average.

6.2 Future directions

6.2.1 Improving models of pursuit and evasion

There is still much work to be done in the study of mathematical models for pursuit and evasion. If we aim to design control laws for autonomous robots to implement pursuit and evasion, the development of more realistic motion models will be necessary. However, as

more complexity is added to a model, it becomes less analytically tractable, forcing the use of numerical methods to compute the optimal control inputs. While numerical methods can provide practical results, it is important to study simplified models as well in order to gain intuition into the structural features of the system. In this thesis we have focused our attention on simple models that afford analytic results while displaying some qualitative features of the biological systems that inspire us.

There are some clear paths forward to extend the work presented in this thesis. For the one-on-one differential game presented in Chapter 4, we assumed an omnidirectional pursuer with no constraints on its turning ability. Since the pursuer's heading was not included as a system state, we were able to reduce the system to two dimensions. If we adapt the model to include both pursuer and evader turning constraints, richer and more complex strategies would be possible, such as the veering maneuvers considered in [44, 41] that allow a slower but more agile evader to avoid capture. However, as more dimensions are added to the system it will be more difficult to find analytic solutions.

Similarly, adding constraints on forward acceleration would make the movement models of Chapters 3 and 4 more realistic while raising the complexity significantly, as the speeds of each agent would become additional states in the system.

A common feature of the mathematical models presented in this thesis is that each agent has exact knowledge of the system state, including the parameters of the motion constraints for each agent. In a realistic system, these quantities would be subject to uncertainties due to limitations in sensing abilities. To address this, the strategies will need to be adapted to accommodate uncertainties in the agent estimates of these system parameters. One approach is for evaders to build their strategy based on their worst-case estimate of the abilities of the pursuer, as in the model with uncertain speeds in [58]. In a sense, the evasive strategies developed in Chapter 4 represent a worst-case response to a pursuer with known speed but unknown limits on turning constraints, in that an evader strategy that avoids capture from an omnidirectional pursuer will also work on a pursuer with more constraints on its motion.

Control laws derived from pursuit and evasion can be relevant for the deployment of multi-agent robotic systems. In real world environments, robots might face pursuit, either from hostile agents or from natural predators mistaking the robots for their prey. Thus the ability to avoid capture while staying cohesive as a group will be important to successful operation. As suggested at the conclusion to Chapter 4, individual pursuit control laws can be adapted for problems of group aggregation and rendezvous.

6.2.2 Use of robotics in biological research

The use of robotic devices in biological field experiments presents a great opportunity to test theories and gather data that would otherwise be unavailable, both in the sense of bringing the behavior of interest closer to the researcher and in accessing remote locations to observe animals in a natural setting. Our robo-lion experiment presented in Chapter 5 serves as a proof of concept of the viability of the use of vision-based tracking along with artificial robotic predators to study evasive behaviors in terrestrial animals. The accuracy of the trajectory tracking was limited primarily by the low resolution of the camera, and the low viewing angle due to the distance from the herd. In future studies, higher spatial resolution can be achieved through the use of an unmanned aerial vehicle, such as a quadcopter, mounted with a high resolution camera to capture an aerial overhead view of the movement of the herd.

An appropriately designed robot can go to places that would be too dangerous for a person, opening up new opportunities to collect data about animal behavior in the field. Unobtrusive terrestrial vehicles could be utilized to stealthily follow a herd of animals throughout the day by posing as a member of a non-threatening species, for instance using a “robo-warthog” to monitor herds of plains zebra. Such a robot could safely approach the animals of interest closer than a person could, and with modern battery and solar technology a robot could endure long deployments and wait tirelessly to gather data. As automated vision-based tracking improves, it may even become possible to program a robot to follow a specific animal of interest without the need to install an obtrusive GPS collar.

Appendix A

Data processing for the robo-lion zebra experiment

A.1 Video stabilization

In order to achieve consistent tracking between video frames and remove distortion due to camera movement, it was necessary to process the overhead video with a stabilization algorithm. We used a custom algorithm based on the “SIFT Flow” dense optical flow algorithm [50] with additional smoothing along each row of pixels in the frame.

The wide-angle lens of the GoPro Hero3 introduces radial distortion that warps pixel positions near the edge of the frame, causing straight lines in the scene to become curves in the image. Using the functions of the Matlab Computer Vision System Toolbox, we estimated the camera’s distortion coefficients using a checkerboard grid pattern and corrected for the distortion in the stabilized video frames. For these stabilized and “rectified” frames, straight lines in the scene are displayed as straight lines in the image.

During the experiments, high wind combined with the flexibility of the camera-mounting pole to cause unsteady oscillations in the camera’s position and orientation. The GoPro Hero3 camera features a CMOS light sensor with a “rolling shutter.” This means that

during filming, each frame is generated by reading one row of pixels at a time from the sensor to build up the full image. When the camera is moving while filming, each row of pixels in a given frame will have a slightly different view of the scene, and must be stabilized independently.

Stabilization is useful for many video applications, including hand-held video, automotive dash-cam footage, and video taken from unmanned air vehicles such as remote controlled quadcopters. Previous work on video stabilization with rolling shutter focus on removing jitter for aesthetic reasons, and are unconcerned with the perfect registration of background between frames that is necessary for out tracking approach. Several algorithms have been proposed, including [3] that views stabilization as a temporal super-resolution problem, allowing them to separate moving objects from background jitter. In many cases the time-varying pose of the camera is estimated explicitly under a simplified motion model for camera movement within a frame, such as rectilinear motion [16], affine motion [15], or a piecewise quadratic model [49].

Simplified motion models allow for faster computation, but give less accurate results. For our purposes we seek the highest accuracy without regard for the time it takes to compute. Thus we use a dense optical flow algorithm to consider the movement of each pixel in the image individually, rather than a sparse feature detection scheme that considers a subset of points of interest to estimate global motion.

The procedure we follow for stabilization is as follows. Let F_0 and F_1 be two consecutive frames of video. Let $p_k(x, y)$ be the “pixel” in frame F_k at location (x, y) , where x is the column and y is the row. We apply the “SIFT Flow” dense correspondence algorithm from [50] to analyze pixel movement (or “optical flow”) between the frames. For each pixel $p_0(x, y)$ in F_0 , SIFT Flow returns an offset vector $\Delta(x, y) = (\Delta_x, \Delta_y)^T$ such that the pixel $p_1(x + \Delta_x, y + \Delta_y)$ is the pixel in frame in F_1 that most closely matches $p_0(x, y)$ according to the “Scale-invariant feature transform” (SIFT) descriptor. The SIFT descriptor, first

proposed in [53], is a mathematical description of the color gradients in the neighborhood of a given pixel in an image.

For each row y in F_0 , we use least-squares minimization to calculate coefficients $\mathbf{c}_y = (c_1, c_2, c_3, c_4, c_5)$ that minimize the error

$$E = \sum_x \|\Delta(x, y) - \mathbf{g}(x, \mathbf{c}_y)\| \quad (\text{A.1})$$

where $\mathbf{g}(x, \mathbf{c})$ is a smooth polynomial mapping from the column x and coefficients \mathbf{c} to the value of the offset vector in the form

$$\mathbf{g}(x, \mathbf{c}) = \begin{pmatrix} c_1 + c_2x \\ c_3 + c_4x + c_5x^2 \end{pmatrix}. \quad (\text{A.2})$$

Other forms of the mapping with higher order polynomials were considered, but the form in (A.2) that is linear for the column and quadratic for the row was found to produce results that were visibly indistinguishable with less computation time.

Once the coefficients have been computed for each row, the set of all smoothed pixel offsets are used to generate an inverse image transform, which is applied to frame F_1 to bring it into alignment with F_0 . The whole process is repeated for frames F_1 and F_2 , and so on, for all frames in the video.

The SIFT Flow algorithm works well for pixels in highly textured regions of an image with rich color gradients. Image areas covered by a single flat color, e.g. a clear blue sky, can yield low quality results. Since any one pixel in the sky is similar to any other, the offset vector returned by the algorithm may not correspond to the actual camera movement. To address that problem in videos with a clear sky, we segmented each frame into sky and ground regions using Matlab's *rgb2ind* color quantization function and discarded pixel offsets from the sky region. The overhead video from trials 13-8, 15-5, and 16-5 were processed in this way.

A.2 Track extraction

Due to the low contrast and changing size and shape of the zebras within the video frames, none of the off-the-shelf tracking programs that we tried were able to reliably track zebra positions automatically. Thus, the positions of each zebra and the robo-lion were recorded manually for each trial. The framerate was subsampled to five frames per second (down from thirty frames per second in the stabilized video) in order to reduce the workload. The track position data was entered using “Tracker” (physlets.org/tracker/), a free program from the NSF-funded Open Source Physics project. Princeton University undergraduates Julie Pourtois, Marina Latif, Pria Louka, and Lisa Sheridan assisted with the manual tracking.

In all, tracking data was recorded for seventeen trials with the robo-lion taking place on July 13 through 16, 2004. Table 5.1 summarizes the number of zebras of each class tracked for each trial.

A.3 Coordinate transformation

In order to analyze the tracked trajectories, we need to find a transformation from the tracked positions on the image to the actual positions on the plane of the ground. If we assume that the ground is flat and level with the horizon, then it is possible to construct a “homography” transformation (a mapping between two planes in a 3D space) from the image plane of the camera to the ground plane in the world by estimating the orientation (roll, pitch, and yaw angles) and height of the camera relative to the ground.

We use the position of the horizon in the first frame of the stabilized video to estimate the orientation of the camera in world coordinates. The horizon line in the image is detected automatically: First a Sobel filter is applied to highlight vertical gradients in the image, then a Hough transform is performed which detects the location of the dominant horizontal line.

The angle of the horizon line in the image provides a direct estimate of the roll angle θ_{roll} of the camera. The pitch angle θ_{pitch} is calculated from the vertical distance from the horizon

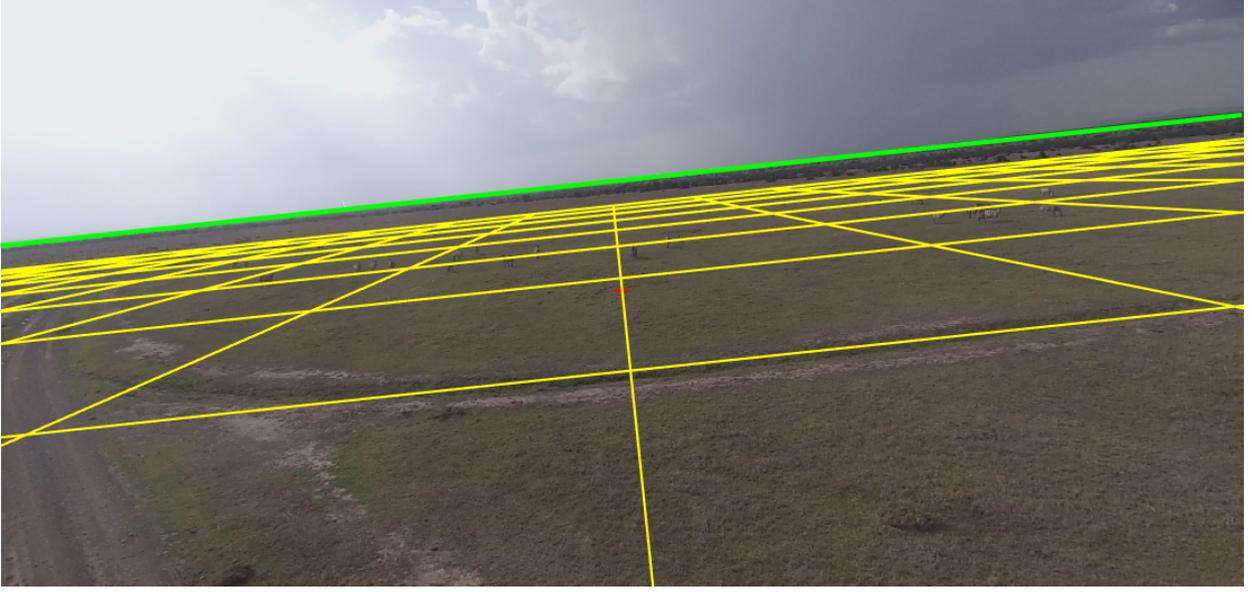


Figure A.1: This image illustrates the homography transformation from image to world coordinates for trial 14-4. The green line is the estimated position of the horizon. The red + indicates the center of the image, which is the origin of the image coordinate frame. Yellow lines show points on a square grid of lines spaced 25 m apart in world coordinates that have been transformed to the image coordinates and overlaid on the image. It is clear that spatial resolution in the depth direction (y -axis in world coordinates) is severely reduced for faraway points due to the low view angle.

line from the center of the picture. The yaw angle is a free parameter and we set it to zero, with the consequence that the y -axis in our world coordinates corresponds with the projection of the camera direction on the ground plane (i.e. the world (ground plane) coordinate frame is not necessarily aligned with magnetic North or any other absolute heading). The rotation matrix describing the camera's pose is then given by

$$R_{cam} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_{pitch} & -\sin \theta_{pitch} \\ 0 & \sin \theta_{pitch} & \cos \theta_{pitch} \end{bmatrix} \begin{bmatrix} \cos \theta_{roll} & -\sin \theta_{roll} & 0 \\ \sin \theta_{roll} & \cos \theta_{roll} & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Based on the estimate of the camera rotation matrix R_{cam} and a measurement of the height of the camera h_{cam} , the homography transform matrix H mapping points from image

to world coordinates is written as

$$H = -h_{cam} R_{cam} \begin{bmatrix} 1/f_x & 0 & 0 \\ 0 & 1/f_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & -w/2 \\ 0 & 1 & -h/2 \\ 0 & 0 & 1 \end{bmatrix},$$

where f_x and f_y are the image focal lengths, and w and h are the width and height of the image in pixels, respectively. The transformation from a point in image coordinates $\mathbf{r}_{im} = (x_{im}, y_{im})$ to a point in world coordinates $\mathbf{r}_{wor} = (x_{wor}, y_{wor})$ is given by

$$\begin{pmatrix} x'_{wor} \\ y'_{wor} \\ w'_{wor} \end{pmatrix} = H \begin{pmatrix} x_{im} \\ y_{im} \\ 1 \end{pmatrix} \quad (\text{A.3})$$

with

$$\mathbf{r}_{wor} = \begin{pmatrix} x_{wor} \\ y_{wor} \end{pmatrix} = \begin{pmatrix} x'_{wor}/w'_{wor} \\ y'_{wor}/w'_{wor} \end{pmatrix}. \quad (\text{A.4})$$

Let \mathbf{h} represent the inverse of this nonlinear transform function, such that $\mathbf{r}_{im} = \mathbf{h}(\mathbf{r}_{wor})$. It is computed the same as in (A.3) and (A.4), using the inverse matrix H^{-1} .

The low viewing angle of the GoPro camera (relative to an ideal overhead camera looking directly down at the ground) means that the track data has much higher resolution in the horizontal direction than the vertical direction. Thus the homography transformation is sensitive to small errors to the pitch angle. If the estimated pitch angle is too low, faraway points (points closer to the horizon in the image) appear more distant than they actually are, becoming increasingly skewed as the distance increases. Conversely if the estimated pitch angle is too high, distant points appear closer together than in reality.

In the absense of “ground truth” measurements of true distances in the scene, we look to the trajectory of the robo-lion as a source of data to perform camera pose calibration.

During chases, the controller of the robo-lion was set to maximum speed, with the operator attempting to follow a nominally straight path towards the zebra herd. If the camera pose is accurate, the speed of the robo-lion should be relatively constant throughout the chase, with some disturbance caused by the terrain.

We propose an algorithm for pitch angle correction that searches for the pitch angle which gives the smallest variation in robo-lion speed as a function of its distance to the camera in the world coordinates. We first smooth the robo-lion’s trajectory using the method described in Section A.4. Least squares regression is applied to fit a linear model to the relationship between distance to the camera d_l and the speed of the robo-lion v_l at each frame. If the slope coefficient is positive, the estimate of the pitch angle is lowered, and vice versa. This continues in an iterative process until an angle is found which cancels out the effect of distance on the speed.

Table A.1 presents the initial estimates and refined values for the camera parameters in each tracked trial.

Table A.1: Image to world transformation parameters for each trial. Camera height is the sum of the mounting pole height (6.81 m) and the height of the vehicle used for that trial: Toyota bus (0.96 m), Land Rover floor (0.96 m), and Land Rover on the spare tire (1.17 m).

day	trial	height (m)	roll	pitch	corrected pitch
13	1	7.77	7.00°	-99.06°	-102.38°
13	2	7.77	1.00°	-99.96°	-101.12°
13	3	7.77	-2.00°	-97.49°	-99.83°
13	6	7.77	-8.00°	-90.54°	-92.01°
13	8	7.77	-7.00°	-90.65°	-99.25°
14	1	7.77	1.00°	-98.65°	-98.99°
14	3	7.77	-0.00°	-95.72°	-100.68°
14	4	7.77	6.00°	-99.88°	-99.38°
14	5	7.77	5.00°	-98.16°	-99.01°
15	1	7.98	-4.00°	-106.49°	-110.53°
15	2	7.98	-1.00°	-98.73°	-99.63°
15	3	7.77	-5.00°	-107.47°	-110.64°
15	5	7.77	-6.00°	-104.51°	-110.11°
15	6	7.77	-1.00°	-101.26°	-104.48°
16	1	7.77	-2.00°	-105.76°	-112.23°
16	3	7.77	-7.00°	-101.15°	-104.83°
16	5	7.77	5.00°	-101.45°	-104.19°

A.4 Trajectory smoothing

To remove noise from the track data, we define a motion model and objective function for each agent and use numerical optimization to find a trajectory which minimizes the objective function. We adapt the jerk-minimization objective function of [22] to take into account the nonlinear effects of the homography transform used to map track data in the image coordinates to trajectories in the world coordinate frame.

We use a triple integrator system as our motion model. The system input is the rate of change of acceleration, known as “jerk.” For a zebra Z_i , let $\mathbf{r}_i(t) \in \mathbb{R}^2$ be its position in the world coordinate frame at time t , $\mathbf{v}_i(t) \in \mathbb{R}^2$ its velocity vector, $\mathbf{a}_i(t) \in \mathbb{R}^2$ its acceleration vector, and $\mathbf{u}_i(t) \in \mathbb{R}^2$ its jerk input vector. The motion model for zebra Z_i is

$$\dot{\mathbf{r}}_i = \mathbf{v}_i,$$

$$\begin{aligned}\dot{\mathbf{v}}_i &= \mathbf{a}_i, \\ \dot{\mathbf{a}}_i &= \mathbf{u}_i.\end{aligned}\tag{A.5}$$

Let \mathbf{z}_i^k be the tracked position of zebra Z_i in image coordinates for video frame k . The time between frames is constant at $\Delta t = 0.2$ s. Let $t = 0$ at frame 0. Then the time at frame k is defined as $t^k = k\Delta t$.

We define the objective function as a combination of position error and jerk input summed over the trajectory. We consider position error in image coordinates and jerk in world coordinates. This differs from [22] where position error is computed in world coordinates. Our modification is necessary for the current system because the noise is introduced to the system in the manual tracking data in the image frame, and the nonlinear homography transform would stretch the effect of the tracking errors so that they are no longer isotropic, accentuating errors in one direction over another.

The objective function for zebra Z_i in a track with $N + 1$ measurements \mathbf{z}_i^0 through \mathbf{z}_i^N is

$$J = \sum_{k=0}^N \|\mathbf{h}(\mathbf{r}_i(t^k)) - \mathbf{z}_i^k\|_2^2 + \lambda \int_0^T \|\mathbf{u}_i(t)\|_2^2 dt.$$

The parameter $\lambda > 0$ controls the relative weighting of position error and jerk. High values of λ will yield a smoother trajectory, but at the cost of a higher error. We used $\lambda = 0.001$ for all results presented in Chapter 5. It was chosen based on qualitative comparisons, since the high computational time of the smoothing algorithm prevented the implementation of a cross validation scheme to quantitatively assess goodness of fit.

To simplify the optimization procedure, we assume that the jerk input is constant between frames, such that $\mathbf{u}_i(t) = \mathbf{u}_i^k = \text{constant}$ for $t^k \leq t < t^{k+1}$. The objective function simplifies to

$$J = \sum_{k=0}^N \|\mathbf{h}(\mathbf{r}_i(t^k)) - \mathbf{z}_i^k\|_2^2 + \lambda \sum_{k=0}^{N-1} \|\mathbf{u}_i^k\|_2^2.$$

We use Matlab's built in unconstrained optimization solver *fminunc* to find a solution to the minimization problem. The state to optimize for a zebra track with $N + 1$ measurements \mathbf{z}_i^0 through \mathbf{z}_i^N is the set of jerk inputs \mathbf{u}_i^k for $k = 0, 1, 2, \dots, N$, along with initial position $\mathbf{r}_i(0)$, initial velocity $\mathbf{v}_i(0)$, and initial acceleration $\mathbf{a}_i(0)$. From the output of the solver, we can reconstruct the full trajectory. The same process is used for each zebra and the robo-lion in each of the tracked trials. For zebras that are stationary at the start of the trial, we smooth over a subset of tracked positions starting at the frame where the zebra initiates movement.

Bibliography

- [1] L. Angelani. Collective predation and escape strategies. *Physical Review Letters*, 109(11):118104, 2012.
- [2] F. Aurenhammer and H. Edelsbrunner. An optimal algorithm for constructing the weighted Voronoi diagram in the plane. *Pattern Recognition*, 17(2):251–257, 1984.
- [3] S. Baker, E. Bennett, S. B. Kang, and R. Szeliski. Removing rolling shutter wobble. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 2392–2399, 2010.
- [4] E. Bakolas and P. Tsiotras. Optimal pursuit of moving targets using dynamic Voronoi diagrams. In *Proc. IEEE Conf. Decision and Control*, pages 7431–7436, 2010.
- [5] D. J. Balkcom, P. A. Kavathekar, and M. T. Mason. Time-optimal trajectories for an omni-directional vehicle. *Int. J. Robotics Research*, 25(10):985–999, 2006.
- [6] D. J. Balkcom and M. T. Mason. Extremal trajectories for bounded velocity mobile robots. In *Proc. IEEE Int. Conf. Robotics and Automation*, volume 2, pages 1747–1752, 2002.
- [7] D. J. Balkcom and M. T. Mason. Time optimal trajectories for bounded velocity differential drive vehicles. *Int. J. Robotics Research*, 21(3):199–217, 2002.
- [8] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Empirical investigation of starling flocks: a benchmark study in collective animal behaviour. *Animal Behaviour*, 76(1):201–215, 2008.
- [9] F. Barnier, M. Valeix, P. Duncan, S. Chamaillé-Jammes, P. Barre, A. J. Loveridge, D. W. Macdonald, and H. Fritz. Diet quality in a wild grazer declines under the threat of an ambush predator. *Proc. Royal Society B: Biological Sciences*, 281(1785):20140446, 2014.
- [10] J. C. Barton and C. J. Eliezer. On pursuit curves. *Journal of the Australian Mathematical Society-Series B*, 41(3):358–371, 2000.
- [11] A. T. Bergerud. The population dynamics of Newfoundland caribou. *Wildlife Monographs*, (25):pp. 3–55, 1971.

- [12] J. V. Breakwell and P. Hagedorn. Point capture of two evaders in succession. *J. Optimization Theory and Applications*, 27(1):89–97, 1979.
- [13] F. Broekhuis, G. Cozzi, M. Valeix, J. W. McNutt, and D. W. Macdonald. Risk avoidance in sympatric large carnivores: reactive or predictive? *J. Animal Ecology*, 82(5):1098–1105, 2013.
- [14] H. Chitsaz, S. M. LaValle, D. J. Balkcom, and M. T. Mason. Minimum wheel-rotation paths for differential-drive mobile robots. *Int. J. Robotics Research*, 28(1):66–80, 2009.
- [15] W.-H. Cho and K.-S. Hong. Affine motion based CMOS distortion analysis and CMOS digital image stabilization. *IEEE Trans. Consumer Electronics*, 53(3):833–841, 2007.
- [16] J.-B. Chun, H. Jung, and C.-M. Kyung. Suppressing rolling-shutter distortion of CMOS image sensors by motion vector detection. *IEEE Trans. Consumer Electronics*, 54(4):1479–1487, 2008.
- [17] B. R. Costelloe and D. I. Rubenstein. Coping with transition: offspring risk and maternal behavioural changes at the end of the hiding phase. *Animal Behaviour*, 109:217–225, 2015.
- [18] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *J. Theoretical Biology*, 218(1):1 – 11, 2002.
- [19] R. Cressman and J. Garay. The effects of opportunistic and intentional predators on the herding behavior of prey. *Ecology*, 92(2):432–440, 2011.
- [20] L. Crisler. Observations of wolves hunting caribou. *Journal of Mammalogy*, 37(3):pp. 337–346, 1956.
- [21] A. I. Dell, J. A. Bender, K. Branson, I. D. Couzin, G. G. de Polavieja, L. P. Noldus, A. Pérez-Escudero, P. Perona, A. D. Straw, M. Wikelski, and U. Brose. Automated image-based tracking and its application in ecology. *Trends in Ecology & Evolution*, 29(7):417–428, 2014.
- [22] B. Dey and P. S. Krishnaprasad. Trajectory smoothing as a linear optimal control problem. In *Allerton Conf. Communication, Control, and Computing*, pages 1490–1497. IEEE, 2012.
- [23] P. Domenici and G. D. Ruxton. Prey behaviors during fleeing: escape trajectories, signaling and sensory defenses. In W. E. Cooper and D. T. Blumstein, editors, *Escaping from predators: and integrative view of escape decisions by prey*, pages 199–224. Cambridge University Press, 2015.
- [24] J. P. Elliott, I. M. Cowan, and C. S. Holling. Prey capture by the African lion. *Canadian Journal of Zoology*, 55(11):1811–1828, 1977.
- [25] I. Eshel. On a prey-predator nonzero-sum game and the evolution of gregarious behavior of evasive prey. *American Naturalist*, pages 787–795, 1978.

- [26] I. Eshel, E. Sansone, and A. Shaked. Gregarious behaviour of evasive prey. *J. Mathematical Biology*, 52(5):595–612, 2006.
- [27] I. Exarchos, P. Tsiotras, and M. Pachter. On the suicidal pedestrian differential game. *Dynamic Games and Applications*, 5(3):297–317, 2015.
- [28] J. J. Faria, J. R. G. Dyer, R. O. Clément, I. D. Couzin, N. Holt, A. J. W. Ward, D. Waters, and J. Krause. A novel method for investigating the collective behaviour of fish: introducing ‘Robofish’. *Behavioral Ecology and Sociobiology*, 64(8):1211–1218, 2010.
- [29] I. R. Fischhoff, S. R. Sundaresan, J. Cordingley, H. M. Larkin, M.-J. Sellier, and D. I. Rubenstein. Social relationships and reproductive state influence leadership roles in movements of plains zebra, *Equus burchellii*. *Animal Behaviour*, 73(5):825–831, 2007.
- [30] I. R. Fischhoff, S. R. Sundaresan, J. Cordingley, and D. I. Rubenstein. Habitat use and movements of plains zebra (*Equus burchelli*) in response to predation danger from lions. *Behavioral Ecology*, 18(4):725–729, 2007.
- [31] Z. E. Fuchs and P. P. Khargonekar. Encouraging attacker retreat through defender cooperation. In *Proc. IEEE Conf. Decision and Control*, pages 235–242, 2011.
- [32] Z. E. Fuchs, P. P. Khargonekar, and J. Evers. Cooperative defense within a single-pursuer, two-evader pursuit evasion differential game. In *Proc. IEEE Conf. Decision and Control*, pages 3091–3097, 2010.
- [33] K. S. Galloway, E. W. Justh, and P. S. Krishnaprasad. Geometry of cyclic pursuit. In *Proc. IEEE Conf. Decision and Control*, pages 7485–7490, 2009.
- [34] W. M. Getz and M. Pachter. Capturability in a two-target ‘game of two cars’. *J. Guidance, Control, and Dynamics*, 4(1):15–21, 1981.
- [35] S. Grünewälder, F. Broekhuis, D. W. Macdonald, A. M. Wilson, J. W. McNutt, J. Shawe-Taylor, and S. Hailes. Movement activity based classification of animal behaviour with an application to data from cheetah (*Acinonyx jubatus*). *PLoS ONE*, 7(11):e49120, 11 2012.
- [36] S. Gueron, S. A. Levin, and D. I. Rubenstein. The dynamics of herds: From individuals to aggregations. *J. Theoretical Biology*, 182:85–98, 1996.
- [37] S. J. Hall, C. S. Wardle, and D. N. MacLennan. Predator evasion in a fish school: test of a model for the fountain effect. *Marine Biology*, 91(1):143–148, 1986.
- [38] J. Halloy, G. Sempo, G. Caprari, C. Rivault, M. Asadpour, F. Tâche, I. Said, V. Durier, S. Canonge, J. M. Amé, et al. Social integration of robots into groups of cockroaches to control self-organized choices. *Science*, 318(5853):1155–1158, 2007.
- [39] W. D. Hamilton. Geometry for the selfish herd. *J. Theoretical Biology*, 31(2):295–311, 1971.

- [40] R. S. Hetem, D. Mitchell, B. A. de Witt, L. G. Fick, L. C. R. Meyer, S. K. Maloney, and A. Fuller. Cheetah do not abandon hunts because they overheat. *Biology Letters*, 9(5), 2013.
- [41] H. C. Howland. Optimal strategies for predator avoidance: the relative importance of speed and manoeuvrability. *J. Theoretical Biology*, 47(2):333–350, 1974.
- [42] H. Huang, W. Zhang, J. Ding, D. Stipanovic, and C. Tomlin. Guaranteed decentralized pursuit-evasion in the plane with multiple pursuers. In *Proc. IEEE Conf. Decision and Control*, pages 4835–4840, 2011.
- [43] Y. Inada and K. Kawachi. Order and flexibility in the motion of fish schools. *J. Theoretical Biology*, 214(3):371–387, 2002.
- [44] R. Isaacs. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. Dover, 1999. Originally published Wiley, 1965.
- [45] E. W. Justh and P. S. Krishnaprasad. Steering laws for motion camouflage. *Proc. Royal Society A: Mathematical, Physical and Engineering Science*, 462(2076):3629–3643, 2006.
- [46] Z. Kong, N. Fuller, S. Wang, K. Özcimder, E. Gillam, D. Theriault, M. Betke, and J. Baillieul. Perceptual modalities guiding bat flight in a native habitat. *Scientific Reports*, 6:27252, 2016.
- [47] J. Krause, A. F. T. Winfield, and J.-L. Deneubourg. Interactive robots in experimental biology. *Trends in Ecology & Evolution*, 26(7):369–375, 2011.
- [48] S. H. Lee, H. K. Pak, and T. S. Chon. Dynamics of prey-flock escaping behavior in response to predator’s attack. *J. Theoretical Biology*, 240(2):250–259, 2006.
- [49] Y. G. Lee and G. Kai. Fast-rolling shutter compensation based on piecewise quadratic approximation of a camera trajectory. *Optical Engineering*, 53(9):093101, 2014.
- [50] C. Liu, J. Yuen, and A. Torralba. SIFT flow: Dense correspondence across scenes and its applications. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 33(5):978–994, 2011.
- [51] S. Y. Liu, Z. Zhou, C. Tomlin, and K. Hedrick. Evasion as a team against a faster pursuer. In *Proc. American Control Conference*, pages 5368–5373, 2013.
- [52] S. R. Loarie, C. J. Tambling, and G. P. Asner. Lion hunting behaviour and vegetation structure in an African savanna. *Animal Behaviour*, 85(5):899–906, 2013.
- [53] D. G. Lowe. Object recognition from local scale-invariant features. In *Proc. IEEE Int. Conf. Computer Vision*, volume 2, pages 1150–1157. Ieee, 1999.
- [54] S. P. Mahoney and J. A. Schaefer. Long-term changes in demography and migration of Newfoundland caribou. *Journal of Mammalogy*, 83(4):957–963, 2002.

- [55] A. W. Merz. The game of two identical cars. *J. Optimization Theory and Applications*, 9(5):324–343, 1972.
- [56] P. J. Nahin. *Chases and Escapes: The Mathematics of Pursuit and Evasion*. Princeton University Press, 2012.
- [57] C. M. V. Nunez, C. S. Asa, and D. I. Rubenstein. Zebra reproduction: Plains zebra (*Equus burchelli*), mountain zebra (*Equus zebra*), and Grevy’s zebra (*Equus grevyi*). In A. O. McKinnon, E. L. Squires, W. E. Vaala, and D. D. Varner, editors, *Equine Reproduction*, volume 2, pages 2851–2865. Wiley-Blackwell, 2nd edition, 2011.
- [58] D. W. Oyler, P. T. Kabamba, and A. R. Girard. Dominance in pursuit-evasion games with uncertainty. In *Proc. IEEE Conf. Decision and Control*, pages 5859–5864. IEEE, 2015.
- [59] D. Pais and N. H. Leonard. Pursuit and evasion: Evolutionary dynamics and collective motion. In *Proc. AIAA Guidance, Navigation, and Control Conference*, 2010.
- [60] D. A. Paley, N. E. Leonard, R. Sepulchre, and I. D. Couzin. Spatial models of bistability in biological collectives. In *Proc. IEEE Conf. Decision and Control*, pages 4851–4856, 2007.
- [61] G. Polverino, N. Abaid, V. Kopman, S. Macrì, and M. Porfiri. Zebrafish response to robotic fish: preference experiments on isolated individuals and small shoals. *Bioinspiration & Biomimetics*, 7(3):036019, 2012.
- [62] T. C. Reluga and S. Viscido. Simulated evolution of selfish herd behavior. *J. Theoretical Biology*, 234(2):213–225, 2005.
- [63] D. I. Rubenstein. Ecology, social behavior, and conservation in zebras. In R. Macedo, editor, *Advances in the Study Behavior: Behavioral Ecology of Tropical Animals*, volume 42, pages 231–258. Elsevier Press, 2010.
- [64] U. Ruiz and R. Murrieta-Cid. A differential pursuit/evasion game of capture between an omnidirectional agent and a differential drive robot, and their winning roles. *Int. J. Control*, pages 1–16, 2016.
- [65] U. Ruiz, R. Murrieta-Cid, and J. L. Marroquin. Time-optimal motion strategies for capturing an omnidirectional evader using a differential drive robot. *IEEE Trans. Robotics*, 29(5):1180–1196, Oct 2013.
- [66] W. L. Scott and N. E. Leonard. Pursuit, herding and evasion: A three-agent model of caribou predation. In *Proc. American Control Conference*, pages 2978–2983, 2013.
- [67] W. L. Scott and N. E. Leonard. Dynamics of pursuit and evasion in a heterogeneous herd. In *Proc. IEEE Conf. Decision and Control*, pages 2920–2925, 2014.
- [68] W. L. Scott and N. E. Leonard. Time-optimal trajectories for steered agent with constraints on speed and turning rate. In *Proc. ASME Dynamic Systems and Control Conference*, 2016.

- [69] P. Souères and J.-D. Boissonnat. Optimal trajectories for nonholonomic mobile robots. In J.-P. Laumond, editor, *Robot Motion Planning and Control*, pages 93–170. Springer, 1998.
- [70] W. J. Stewart, A. Nair, H. Jiang, and M. J. McHenry. Prey fish escape by sensing the bow wave of a predator. *J. Experimental Biology*, 217(24):4328–4336, 2014.
- [71] H. J. Sussmann and G. Tang. Shortest paths for the Reeds-Shepp car: a worked out example of the use of geometric techniques in nonlinear optimal control. *Rutgers Center for Systems and Control Technical Report*, 10:1–71, 1991.
- [72] H. J. Sussmann and J. C. Willems. 300 years of optimal control: from the brachystochrone to the maximum principle. *IEEE Control Systems*, 17(3):32–44, 1997.
- [73] D. T. Swain, I. D. Couzin, and N. E. Leonard. Real-time feedback-controlled robotic fish for behavioral experiments with fish schools. *Proc. of the IEEE*, 100(1):150–163, 2012.
- [74] H. Tan and A. M. Wilson. Grip and limb force limits to turning performance in competition horses. *Proc. Royal Society of London B: Biological Sciences*, 278(1715):2105–2111, 2011.
- [75] R. Vabø and L. Nøttestad. An individual based model of fish school reactions: predicting antipredator behaviour as observed in nature. *Fisheries Oceanography*, 6(3):155–171, 1997.
- [76] R. Vabø and G. Skaret. Emerging school structures and collective dynamics in spawning herring: A simulation study. *Ecological Modelling*, 214(2):125–140, 2008.
- [77] S. V. Viscido, M. Miller, and D. S. Wethey. The response of a selfish herd to an attack from outside the group perimeter. *J. Theoretical Biology*, 208(3):315–328, 2001.
- [78] E. Wei, E. W. Justh, and P. S. Krishnaprasad. Pursuit and an evolutionary game. *Proc. Royal Society A: Mathematical, Physical and Engineering Science*, 465(2105):1539–1559, 2009.
- [79] D. Weihs and P. W. Webb. Optimal avoidance and evasion tactics in predator-prey interactions. *J. Theoretical Biology*, 106(2):189–206, 1984.
- [80] T. M. Williams, L. Wolfe, T. Davis, T. Kendall, B. Richter, Y. Wang, C. Bryce, G. H. Elkaim, and C. C. Wilmers. Instantaneous energetics of puma kills reveal advantage of felid sneak attacks. *Science*, 346(6205):81–85, 2014.
- [81] A. M. Wilson, J. C. Lowe, K. Roskilly, P. E. Hudson, K. A. Golabek, and J. W. McNutt. Locomotion dynamics of hunting in wild cheetahs. *Nature*, 498(7453):185–189, 2013.
- [82] J. W. Wilson, M. G. L. Mills, R. P. Wilson, G. Peters, M. E. J. Mills, J. R. Speakman, S. M. Durant, N. C. Bennett, N. J. Marks, and M. Scantlebury. Cheetahs, *Acinonyx jubatus*, balance turn capacity with pace when chasing prey. *Biology Letters*, 9(5), 2013.

- [83] A. J. Wood and G. J. Ackland. Evolving the selfish herd: emergence of distinct aggregating strategies in an individual-based model. *Proc. Royal Society B: Biological Sciences*, 274(1618):1637–1642, 2007.
- [84] P. Zhang, C. M. Sadler, T. Liu, I. Fischhoff, M. Martonosi, S. A. Lyon, and D. I. Rubenstein. Habitat monitoring with ZebraNet: design and experiences. In N. Bulusu and S. Jha, editors, *Wireless Sensor Networks: A Systems Perspective*, pages 235–257. Artech House, Norwood, MA, 2005.
- [85] P. Zhang, C. M. Sadler, S. A. Lyon, and M. Martonosi. Hardware design experiences in ZebraNet. In *Proc. Int. Conf. Embedded Networked Sensor Systems*, pages 227–238. ACM, 2004.
- [86] M. Zheng, Y. Kashimori, O. Hoshino, K. Fujita, and T. Kambara. Behavior pattern (innate action) of individuals in fish schools generating efficient collective evasion from predation. *J. Theoretical Biology*, 235(2):153–167, 2005.