AUTONOMOUS VEHICLE MOTION PLANNING WITH ETHICAL CONSIDERATIONS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MECHANICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Sarah Marie Thornton

August 2018

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

J Gerdes, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Mykel Kochenderfer

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Marco Pavone

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(J. Christian Gerdes) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Mykel Kochenderfer)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Marco Pavone)

Approved for the Stanford University Committee on Graduate Studies

Abstract

Human drivers navigate the roadways by balancing values such as safety, legality, and mobility. An autonomous vehicle driving on the same roadways as humans likely needs to navigate based on similar values. For engineers of autonomous vehicle technology, the challenge is then to connect these human values to the algorithm design.

To address this challenge, a mapping of philosophical frameworks to mathematical frameworks is used in order to motivate various design choices in a motion planning algorithm. Deontological ethics parallels rule-based mathematical concepts while consequentialism parallels cost-based mathematical concepts. The philosophical theory of virtue ethics is also used to help motivate the relative weightings between the design objectives of path tracking, obstacle avoidance, and adherence to traffic laws. Experimental results of an autonomous vehicle navigating an obstructed two-lane roadway with a double yellow line demonstrate the implications of the various design choices in a model predictive steering controller.

In order to determine the success of the human values captured in an algorithm, the iterative methodology of value sensitive design (VSD) is used to formalize the connection of human values to engineering specifications. A modified VSD methodology is used to develop an autonomous vehicle speed control algorithm to safely navigate a pedestrian crosswalk. Two VSD iterations are presented that model the problem as a partially observable Markov decision process and use dynamic programming to compute an optimal policy to control the longitudinal acceleration of the vehicle based on the belief of a pedestrian crossing. The speed control algorithms are also tested in real-time on an experimental vehicle on a closed-road course.

Acknowledgments

Thank you to everyone who has helped me along this journey. The research presented here would not exist without those before me and without the support and encouragement from my colleagues, friends, and family. Please allow me to go into further detail about my gratitude.

Firstly, I would like to thank my advisor, Chris Gerdes. Thank you for allowing me to join your team and be a part of the special experience that is the Dynamic Design Lab at Stanford. I have learned so much from you over these few years: from vehicle dynamics to improvisation. Your sincerity towards teaching and mentorship is also both inspiring and cherishing to me. You never miss an opportunity to have everyone stand (even a room full of executives) and pretend to be a piece of squeaking tire tread. You also gave me many opportunities to improve my public speaking abilities. Most of all, I thank you for leading by example. It is not always easy, especially when asked to help our country's Department of Transportation for a year, but you always put your family first. You are always there to give support to us students when we need it most, but you also ensure that we have support amongst each other. Thank you for showing me what a true leader does.

I would also like to thank Mykel Kochenderfer. You have been a second advisor to me and welcomed me as an honorary SISLer. You have been so generous with your time and knowledge. Thank you for showing me how life can be modeled as a partially observable Markov decision process (POMDP). Your joy of POMDPs, probabilities, the Julia programming language, and PGFPlots is infectious. I am so glad to have been able to incorporate these in my research and thesis; all thanks to you.

The other members of my defense committee also have my gratitude. To Marco

Pavone, thank you for teaching me more about optimization in your optimal control class and for always pushing me to further understand the ethical implications in my work. To Dorsa Sadigh, thank you for bringing another perspective to campus with safe robotics. It was a pleasure to have you on my committee. To Markus Maurer, thank you for taking time from your sabbatical to be a part of my committee. Thank you for challenging the claims in my thesis; it helped me to clarify and acknowledge the limitations of my work.

There are many others I would like to thank. In particular, I would like to thank my sponsors. My work has been funded by the National Science Foundation and Ford Motor Company. I had the great fortune to work with Ford during my Master's, and it has been great to continue that relationship. Additionally, I would like to thank Patrick Lin and Jason Millar for showing me that the inquisitive spirit of the philosopher can help with designing safer and better technology. To Bryant Walker Smith, Steve Wu, and Bryan Casey, thank you for providing your expert legal advice. It helped to shape many discussions and thoughts that went into this thesis. I would also like to give my upmost gratitude to the staff of CARS and DDL who helped make all this research possible: Erina Dubois, Adele Tanaka, Larry Cathey, Stephen Zoeph, Jo Yuan, Sven Beiker, Ludivine Frezza, and Elizabeth Pearson. Thank you to SISLers for putting together POMDPs.jl and for the discussions. A big thank you to all the DDLers for being awesome and for all the discussions and code scrums.

Special thanks to all the spotters for my experiments over the years: Vivian Zhang, Gene Lewis, Matt Brown, Nathan Spielberg, John Alsterda, Tushar Goel, Erina Dubois, Larry Cathey, John Talbot, Jack Pigott, Oliver Stern, Stephany Cabrera, and Riley Patterson.

Last, and most definitely not least, I would like to thank my loving husband, Joseph Thornton, for always encouraging me to do my best and to go after new, challenging opportunities. I love you. And thank you to Michael, our ginger cat, for keeping my lap warm and fuzzy while studying for quals and for watching machine learning lectures with me.

Dedicated to my family.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

Driving allows humans to get from one destination to another in a timely and comfortable manner. People have places to be, other people to see, and things to do. Humans who drive enjoy the convenience, efficiency, and flexibility that automotive transportation provides in order to go about their day. This desire to travel illustrates that humans who drive find mobility to be important in their everyday lives. Yet, drivers are not the only humans that want mobility. A driver's desire for mobility can conflict with another road user's desire for mobility. For example, a pedestrian may want to cross the street within a crosswalk at the same time a driver is traveling through the roadway. This demonstrates a conflict over the value of mobility. Borning and Muller define the word "value" as "what a person or group of people consider important in life" [1]. Hence, mobility can be considered to be a human value.

Another value encountered while driving is safety. Imagine someone driving down a residential road and suddenly a toy ball rolls onto the street from behind a large family van parked on the street. The human driver cannot see if or when a person may run out to chase the ball. Nevertheless, human drivers realize that they just need to stop or slow down in order to maintain the safety of the scenario because that person may be someone's child. Valuing safety allows drivers to protect other humans and property from a two-ton moving vehicle.

A third value is legality. Traffic laws provide guidance, allowing humans to share the roadways with other drivers and vulnerable road users. Traffic laws tend to be written as strict rules, but human drivers often push the boundaries of these laws or break them entirely. For example, speed limits are phrased as a rigid upper limit on the maximum speed a vehicle should travel (California Vehicle Code §22348). However, judge instructions to juries on the traffic code suggest it is not so strict in practice [2]. Other apparently strict driving situations involve the adherence to double yellow lines, which are illegal to cross on California roadways according to California Vehicle Code §21460. This can conflict with the requirement to maintain a 3 ft gap when passing bicyclists (California Vehicle Code §21760) or when a large vehicle double parks on a road with a double yellow line. Human drivers value legality and manage to operate a vehicle amidst conflicts within legality.

When human drivers navigate the roadways, they do so by balancing values such as mobility, safety, and legality as evidenced by how they handle the above scenarios. There are, of course, many other values implicated while driving, such as care and respect for others, fairness and reciprocity, respect for authority, and trust and transparency. Depending on the scenario, one value may take precedence over another or the values may even conflict. Fortunately, human drivers have a way of determining what values are a priority at any given time.

1.2 Autonomous Vehicle Motion Planning

Autonomous vehicles entering the roadways are likely to encounter similar scenarios to human drivers, such as driving alongside other human drivers and vulnerable road users, carrying human occupants, and encountering objects or people suddenly appearing from behind occlusions. An autonomous vehicle will also have to navigate the roadways while balancing values such as mobility, safety, and legality. For engineers of autonomous vehicle technology, the challenge is to design motion planning algorithms in the midst of these competing human values.

According to the literature, there are many definitions of motion planning. This thesis defines motion planning as an algorithm that plans the lateral and longitudinal

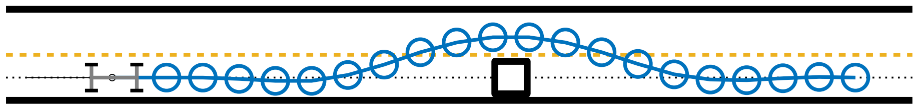

Figure 1.1: Example of motion planning around an obstacle. The centerline (dotted) is the reference trajectory, and the blue circles indicate the planned trajectory.

motion through a combination of steering and acceleration commands in terms of a given reference trajectory. The reference is a nominal trajectory defining desired positions and velocities for the autonomous vehicle, and it is not guaranteed to be obstacle-free. If an autonomous vehicle is to navigate a two-lane roadway, for example, and an obstruction exists or a pedestrian appears, then the autonomous vehicle needs a motion planner to negotiate between following the reference and avoiding collisions. For scenarios like these, particular techniques are conducive for motion planning. A good choice for a motion planner is model predictive control (MPC), which solves for a sequence of control inputs by optimizing a cost function according to a set of constraints along a prediction horizon in a receding fashion (see Figure 1.1) [3]. The designer of the MPC optimization problem must then determine the objectives for the cost function, the constraints, and the weights such that the vehicle realizes goals like mobility, safety, and legality. Other motion planning techniques may also be a good choice, and the designer of those algorithms will have to make similar considerations.

1.2.1 Mobility and Safety

In designing autonomous vehicle motion planning algorithms, engineers already make decisions that implicate values of mobility and safety. For example, Falcone, Borrelli, Asgari, Tseng, and Hrovat [4] explore the construction of two MPC optimization problems for executing a constant speed double-lane change maneuver along a collision-free reference trajectory at the limits of vehicle handling. Falcone *et al.* include heading deviation, lateral deviation, yaw rate deviation, and steering effort in the cost function, while the constraints consist of the vehicle model and actuator limits. The weight on heading deviation is relatively larger than the other states across implementations. For the nonlinear MPC formulation, the weight on steering effort is relatively lower than the weight on heading deviation but greater than the weight on lateral deviation. The linear time-varying formulation includes an additional constraint on slip angle, and the penalty on steering effort is two orders of magnitude higher than heading deviation. Because the authors do not explicitly mention how mobility and safety are captured in their algorithms, the formulations leads me to interpret that mobility arises by following the desired speed, while safety is realized by the vehicle models and system constraints because the reference trajectory is obstacle-free. Similarly, Ziegler, Bender, Dang, and Stiller [5] construct a constrained optimization problem and solve in a receding fashion. The cost function consists of maintaining equal distances from road edges and obstacles that appear, following the speed limit, smoothing acceleration and jerk, and attenuating high yaw rates. The constraints include steering geometry, the friction limits of the tires, and the driving corridor. The choice of weights for the experiments is not disclosed in the paper. Nor do the authors explain how their algorithm explicitly comprise mobility and safety. These formulations lead me, again, to interpret that mobility arises by following the speed limit when possible in addition to occupant comfort, while safety comes from avoiding obstacles and road edges.

Rather than solving a constrained optimization problem, Kuwata, Teo, Karaman, Fiore, Frazzoli, and How [6] use rapidly-exploring random trees (RRT) and sample from the control space to create reference trajectories given a target goal from a high-level route planner. A generative closed-loop vehicle model with constraints is used for randomly sampling feasible, smooth trajectories through a cluttered environment. Essentially, Kuwata *et al.* use the closed-loop RRT to bypass the need for a motion planner because the reference trajectory it generates already accounts for obstacles and for the closed-loop speed and steering controllers. The constraints consist of acceleration and steering limits. Although there is no cost function, the closed-loop controllers have gains to be chosen. Upon my interpretation, Kuwata *et al.* prioritize safety over mobility considerations with several safety mechanisms to override the planned trajectory. Bouton, Nakhaei, Fujimura, and Kochenderfer [7] explicitly account for mobility through efficiency as well as safety in the design of a partially observable Markov decision process (POMDP) for navigating occluded scenarios by penalizing collisions and rewarding the vehicle for completing maneuvers. Although constraints are not explicitly considered when optimizing the policy for a POMDP model, Bouton *et al.* use a discrete state space to limit the maximum speed for the maneuvers and a discrete action space to limit the change in acceleration commands. There are many more examples of engineers implicating values of safety and mobility in the design of motion planning algorithms.

1.2.2 Mobility, Safety, and Legality

Most autonomous vehicle motion planning algorithms suggest some consideration of mobility and safety, but legality is rarely or only implicitly addressed. For the 2007 DARPA Urban Challenge, participant vehicles were required to adhere to the California Vehicle Code [8], [9], but leading participants' reports provide only minimal details on how the algorithms in the decision structure (inclusive of mission planners, behavioral planners and motion planners) considered the traffic code [10]–[13]. All participants implemented some variant of a finite state machine to construct feasible paths that stayed within lane markers and considered stopping points at traffic lights and stop signs. The speed limit was an upper limit imposed on the motion planners. The value conflict between mobility, safety, and legality was implicitly mitigated by partitioning the decision problem with hierarchical decision structures, where the top level algorithms considered only the legal rules and bounded the feasible paths for the motion planners to consider only mobility and safety. Value conflicts with legality would arise only during error recovery, for example, due to sensor failure. If the legal requirements were too strict to obtain a feasible path, then the legal requirements were pruned away until an obstacle-free feasible path was determined. Essentially, the DARPA Urban Challenge vehicles would try to adhere to the vehicle code first unless a fail-safe mode was engaged. This paradigm of "try to be legal first" continued after the DARPA Urban Challenge as evidenced by the hybrid decision architecture that Wei, Snider, Gu, Dolan, and Litkouhi [14] propose, which considers traffic rules in the mission and reference planners, and considers safety, smoothness, and efficiency in the lower-level behavioral planner [15].

Rather than bisect the value conflict in the decision architecture, another approach is to construct a controller that guarantees adherence to a rule set such as the traffic code. Wongpiromsarn, Karaman, and Frazzoli [16] apply linear temporal logic (LTL) to the description of the traffic code to synthesize an obedient controller. In the simulations of an obstacle on the road and a double yellow line, the autonomous vehicle is able to navigate around the obstruction because the logic rules allow passing a double yellow after the vehicle comes to a stop first. Reyes Castro, Chaudhari, Tumova, Karaman, Frazzoli, and Rus [17] build upon this work by using a sampling-based algorithm known as optimal rapidly-exploring random trees (RRT*) to construct dynamically feasible trajectories that adhere to the rule set. Difficultly arises in defining a rule set that is reachable, meaning it will not overly constrain the vehicle motion. To address this difficulty, Reyes Castro *et al.* allow for (and chose the) prioritization of the rules.

On the opposite end of the spectrum are learning approaches that implicitly balance mobility, safety, and legality by learning from human drivers. Wulfmeier, Wang, and Posner [18] employ cost function learning from expert demonstrations. The planning problem is formulated as a Markov decision process (MDP) and learns a "perception-to-cost" mapping without specifying any features explicitly. This approach obfuscates the decision-making process and lacks the transparency needed to understand the implications of mobility, safety, and legality. Kuderer, Gulati, and Burgard [19] take a similar approach to learn a "perception-to-cost" mapping but they explicitly define its features. They use maximum entropy inverse reinforcement learning to parameterize driving style in a cost function. Taking this idea further, Lee and Seo [20] also define explicit features to implicitly learn how drivers navigate an obstructed roadway with a double yellow line. With these implementations, the autonomous vehicle embodies values of mobility, safety, and legality similar to those of the demonstrator.

It is evident in the above literature that there are many ways to formulate a motion planning algorithm. Some techniques focus on constraints, some on costs, and others allow for a combination of constraints and costs. In the general formulation of a motion planning algorithm, the choice of a cost, of a constraint, or of a weight may seem arbitrary. Even less clear is how these choices connect to human values. An investigation of ethics and human values may provide assistance to engineers in determining what should be a cost or constraint and even the weights when designing motion planning algorithms.

1.3 Autonomous Vehicles and Ethics

When the topic of autonomous vehicles and ethics arises, the infamous trolley problem may immediately come to mind. The trolley problem posits a scenario of an uncontrollable trolley loose on a set of railway tracks because of broken brakes. If the trolley continues on its way, it is guaranteed to kill five people. You as a bystander have the choice to intervene and divert the trolley to another set of tracks where there is an unassuming person that would surely perish [21]. Replacing the trolley with an autonomous vehicle, an algorithm may then need to determine whether to continue on a collision course with five pedestrians or swerve and kill an unassuming pedestrian. This would be an unfortunate scenario for an autonomous vehicle (or anyone) to encounter on the roadway. Hence, there is a prevalence of work that seeks to explore an understanding of and potential solutions to said problem.

Along the lines of understanding the trolley problem, some focus has been on learning the preferences of the public concerning what an autonomous vehicle should do in a crash scenario when it is confined to the choice of hitting one entity over another. It is assumed in these preference surveys and experiments that an autonomous vehicle will be able to affirmatively target an entity. The preference surveys and experiments largely conclude similar results, which are that a majority of the human subjects prefer the vehicle take a utilitarian approach to sacrifice the few for the many [22]–[25]. The social dilemma arises in the follow-up question about whether an autonomous vehicle should be programmed to adhere to utilitarian ethics. The authors conclude that the participants do not want the autonomous vehicle to be utilitarian or follow their own preferences, but want manufacturers and lawmakers to decide how the vehicle should behave in these crash scenarios. Philosophers and ethics researchers first posited the trolley problem to encourage engineers to think about the consequences of their design choices to avoid the unintentional development of targeting algorithms [26]. The focus on trolley problems has brought about solutions (and, in essence, targeting algorithms) which entail encoding an autonomous vehicle to accept a user-defined priority or “command” list for crash scenarios as posited by Fournier [27] and Jaiswal [28]. In particular, Jaiswal uses a priority list in conjunction with an object classification algorithm to inform the autonomous vehicle what action to take in a crash scenario such that it reflects the occupants’ personal ethics or use a default list provided by the manufacturer. The design implies that the outcome of the crash scenario is deterministic assuming accurate classification. Since the aforementioned preference surveys indicate users are less inclined to embed their ethics (i.e. their utilitarian choices) into the vehicle, Noothigattu *et al.* [29] propose the democratization of preference learning. In the case of a trolley scenario, should vehicle failure occur and legality not apply, then the aggregated voices of the public will decide who to target.

As an alternative to targeting solutions for trolley scenarios, some use the idea of social welfare to more equally distribute harm. In a trolley problem variant where the autonomous vehicle ferries a human occupant and it must suddenly swerve into a wall or kill a pedestrian, Kinjo and Ebina [30] propose evaluating an objective function that multiplies the utilities of the passenger and pedestrian such that non-binary steer angle solutions arise. Meaning, an autonomous vehicle equipped with this objective function could attempt to slightly swipe both the wall and pedestrian, thus distributing a small amount of harm to all involved. This starts moving trolley-like decisions from survey results to game theory approaches such as those further explored by Leben, who formulates social justice as a maximin algorithm [31]. The social justice approaches enforce neutrality with regards to the agents rather than assuming a particular target.

The proposed solutions to trolley scenarios solve a very different problem than a motion planning problem. In motion planning, the autonomous vehicle constantly evaluates how to move laterally and longitudinally such that it follows a given reference trajectory and avoids collisions. The above trolley scenario solutions focus on making last minute targeting decisions in crash scenarios. Wheeler demonstrates that there is a critical inflection point in terms of scenario risk well before an accident occurs [32], which suggests an autonomous vehicle equipped with an appropriate motion planner could avoid trolley scenarios altogether. This thesis takes a step back from trolley problems to consider philosophical frameworks and human values more broadly in order to help engineers design motion planning algorithms in socially acceptable and justifiable ways.

1.4 Dissertation Contributions

The above literature review indicates there are many approaches for autonomous vehicle motion planning and that the current take on autonomous vehicle ethics solves a very different problem. This thesis addresses the gap between motion planning and ethics through the following contributions: improving functionality by accounting for the actuation delays in a steering controller, incorporating traffic laws pertaining to lane dividers and crosswalks into problem formulations, mapping normative theories to mathematical concepts found in motion planning algorithms, and applying a modified version of a generic design methodology to connect human values to engineering specifications.

1.4.1 Delay compensation of steering actuation in model predictive control

It is well established that the models used in model predictive control (MPC) must be accurate because MPC is an open-loop planning technique. The steering system of an autonomous vehicle incurs some amount of time delay due to communication latency and may also have delay dynamics associated with actuation. If the combined delay is significant, then the path tracking performance and steer effort from MPC may suffer. To account for time delays in MPC, the technique of state-propagation estimates the initial state to a point in time beyond the effect of the delay but depends on having a highly accurate models of the delay. Additionally, the techniques for handling delays found in traditional controls cannot be applied to MPC without increasing the computational complexity of the problem formulation. The research presented here starts with a baseline MPC problem formulation without any pure time delays or actuation lags. Various efficient delay forecasting formulations are then appended to the problem formulation in a comparative analysis of path tracking and steering performance. With improved modeling of the steering system in the control formulation, system performance subsequently improves.

1.4.2 Incorporation of traffic laws: lane dividers in model predictive steering control and crosswalks in partially observable speed control

For model predictive steering control, a particularly relevant traffic law is California Vehicle Code §21460, which restricts drivers from driving on the left of a double yellow line. Lane and road dividers are implemented as slack variables on the environmental envelope of the MPC formulation. The type of lane or road divider contributes to the choice of weight on the slack variable. Alternatively, the type of divider is implemented as a secondary environmental envelope using the physical dimensions of the lanes rather than the full road width.

Traffic laws pertaining to the speed limit and unsignalized pedestrian crosswalks are embedded in a partially observable Markov decision process (POMDP) for speed control. The speed limit is implicitly accounted for by limiting the vehicle state space to the speed limit. The unsignalized pedestrian crosswalk is assumed to be a marked crosswalk or an unmarked crosswalk at an intersection as described by California Vehicle Code §21950. The vehicle distance to the crosswalk is used to account for the physical location of the crosswalk on the road.

The incorporation of these traffic laws demonstrates additional examples of how legality can be explicitly captured in the design of motion planning algorithms. In particular, §21460 is included in an MPC formulation, and §21950 is embedded in a POMDP.

1.4.3 Mapping of philosophical frameworks to mathematical frameworks in autonomous vehicle motion planning

There are mathematical frameworks that parallel philosophical frameworks, and can subsequently inform engineering design choices. In particular, the normative theory of deontological ethics maps to set theory and conditionals or constraints, while consequentialism maps to an optimization problem [33]. Commonalities amongst these frameworks motivate the design of many motion planning algorithms. In particular, this thesis considers the design of a model predictive steering controller. The design choices of the steering controller include accounting for path tracking, smooth steering, physical limits of the steering system, obstacle avoidance, and adherence to traffic laws. Path tracking and smooth steering are cast as consequential costs, while steering limits and obstacle avoidance are cast as deontological constraints. Adherence to traffic laws depends on whether a double yellow line should be treated deontologically or consequentially. Engineers already use rule-based and cost-based algorithms, so understanding the implications from philosophical principles provide engineers with further rationalization and justification for the design choices of a cost, constraint, and weight.

1.4.4 Demonstration of a modified value sensitive design process for autonomous vehicle motion planning

Value sensitive design (VSD) is a general design methodology that can be applied to any design task with value conflicts to explicitly account for human values with ethical import [34], [35]. The three phases of conceptualization, technical implementation, and empirical analysis iterate until the identified human values are captured by the technology. In order to apply VSD to autonomous vehicle motion planning algorithms, the conceptualization phase in the methodology is modified to focus on listing stakeholders, value assumptions and information needed to account for the values in subsequent phases. In line with VSD, the technology or algorithm chosen for implementation embodies certain values; the application of VSD to autonomous vehicle motion planning should acknowledge how the chosen algorithm upholds the outlined values. The empirical analysis phase investigates how well the technical implementation meets the conceptualization using qualitative and quantitative techniques. Engineers equipped with VSD can explicitly and formally account for values being embedded in the algorithm design.

1.5 Dissertation Outline

The contributions of this thesis are dispersed throughout the chapters. The theme of ethical considerations for autonomous vehicle motion planning is woven through each chapter as follows:

Chapter 2: Comparative Analysis of Steering System Modeling in Model Predictive Control

Consideration of ethics and values applies to situations where balancing various objectives may be difficult or not obvious. Certain engineering problems in autonomous vehicle motion planning do not suffer from value conflicts. Instead, those engineering problems concentrate on basic functionality. Chapter 2 focuses on proper account of the actuation level in the decision-layer in a computationally efficient manner. All problem formulations in Chapter 2 use model predictive control (MPC) for steering an autonomous vehicle because MPC can account for state and input constraints in the optimization problem. Variations in the problem formulation account for different system identification models of the actuation signal delay and dynamics found in a modified production vehicle with an electronic power assist steering system. The analysis comprises open-loop prediction mismatch with closed-loop performance and steering effort for the vehicle driving around an obstruction on the reference trajectory. The compensated system can be used for motion planning tasks in subsequent chapters.

Chapter 3: Costs, Constraints, and Weights from Philosophical Principles

In addition to safety, autonomous vehicle motion planning touches on the human values of mobility and legality. Translating these values to engineering specifications is not straightforward or often implicitly embedded by the design choices of the engineer. Balancing the values of mobility, safety, and legality is a human process. Philosophers are experts on the human process, and model the world through normative theories. The most common theories are a rule-based model known as deontological ethics, a cost-based model known as consequentialism, and a character-based model known as virtue ethics. In Chapter 3, these philosophical frameworks are mapped to mathematical frameworks. Deontological ethics parallels the mathematical concept of set theory or constraints and conditionals while consequentialism parallels an optimization problem. With this knowledge, design choices of how to formulate a motion planning algorithm for path tracking, obstacle avoidance, and adherence to traffic laws for the case of navigating a two-lane roadway with a double yellow line and an obstruction on the roadway manifests as a constrained optimization problem. The approach is implemented using MPC as a constrained finite-time optimal control problem thereby maintaining computational efficiency and tractability. Design choices are cast as either deontological constraints or consequential costs as a manner of design justification. The choice of weight values in the cost function of a constrained optimization algorithm still lead to different driving behavior, so virtue ethics in the form of role morality is used to help motivate the relative weightings between the design objectives of path tracking, obstacle avoidance, and adherence to traffic laws. Engineers equipped with this mapping can better reason through the design choice of a cost, constraint, or weight.

Chapter 4: Motion Planning with Human Values

Justifying the design of autonomous vehicle motion planning algorithms through normative theories is a useful discussion tool for engineers. Unfortunately, in design problems, value conflicts may not be resolved through discussion of normative theories because they cannot guarantee the realization of human values in the designed technology. Chapter 4 applies a modified version of the value sensitive design (VSD) methodology to explicitly connect human values to engineering specifications through iteration over conceptualization, technical implementation, and empirical analyses. The iterative nature of VSD challenges the engineering process to continually check back with the moral values at stake until all three phases align. Chapter 4 builds on the previous scenario of Chapter 3 by changing the obstruction to a large, occluding vehicle parked in front of a pedestrian crosswalk. For this scenario, the modified VSD is applied to the development of an autonomous vehicle speed control algorithm for safe navigation of an occluded pedestrian crosswalk. The conceptualization phase identifies the values with ethical import and stakeholders for the scenario. The choice of technical implementation upholds certain values, so the control algorithm takes the form of a POMDP to account for the uncertainty of the scenario in a closed-loop planning algorithm. The empirical analysis includes a comparison with a baseline deterministic proportional speed control design and a policy evaluation. By applying VSD, the implicit embedding of values by engineers becomes explicit in the design process.

Chapter 5: Conclusion

The contributions of this thesis and the journey therein are summarized in Chapter 5. This thesis comprises preliminary works that explicitly account for ethical considerations in autonomous vehicle motion planning algorithms, but there is still more to be done. Thus, additional attention is given to explaining what some of these next steps may look like. Furthermore, the conclusions emphasize that ethical considerations are not just at the decision-making layer in the autonomous vehicle stack but are at play throughout all aspects of the vehicle and system design such as the choice of sensors, the perception layer, or even how the vehicles are tested and deployed.

Chapter 2

Comparative Analysis of Steering System Modeling in Model Predictive Control

Autonomous vehicle motion planning requires ethical considerations in situations where value conflicts arise. Some engineering problems within motion planning, however, do not deal with value conflicts because they deal with having a functioning system. Chapter 2 focuses on accounting for the actuation level in the decision-layer in a computationally efficient manner in order to ensure the planned trajectories can actually be executed. The majority of this chapter appears in the Proceedings of the International Symposium on Advanced Vehicle Control of 2018 [36].

2.1 Introduction

Model predictive control (MPC) has been shown to perform well for automated driving applications because it can account for constraints on the system while optimizing multiple objectives. MPC also relies upon a model of the system to predict how control inputs affect future trajectories of the plant. For the motion planning problem defined in Chapter 1, one way to formulate an MPC problem is to track the reference trajectory by minimizing lateral deviation in the cost function and defining obstacle avoidance as a constraint. But an autonomous vehicle needs to steer properly in order to execute the planned trajectories.

Automated vehicles on the road today are modified production vehicles. An electric power assisted steering (EPAS) system involving a motor mounted to the steering column provides a straightforward external interface to the steering actuation. However, EPAS may introduce significant delay and compliance to the production steering system. If a system has a non-trivial time delay between the actuation requests and the fulfillment, it is important to model and embed the delay in the MPC formulation. Otherwise, path tracking and occupant comfort can be compromised.

Actuation modeling in MPC has been considered throughout the literature, typically by state propagation or additional delay states in the model. For systems with well known signal delays, Trimboli, Di Cairano, Bemporad, and Kolmanovsky [37] demonstrate state propagation in simulation, and Liniger, Domahidi, and Morari [38] address 20 ms signal delay through state propagation using a second-order Runge-Kutta method with validation on 1:43 scale electric cars. An alternative approach to state propagation for pure time delays entails modeling the signal delay directly in the model by appending the system dynamics with delay states as demonstrated by Di Cairano, Yanakiev, Bemporad, Kolmanovsky, and Hrovat [39]. These approaches demonstrate state propagation as a useful technique for well-characterized actuation dynamics and known signal delays in simulation and on small-scale electric cars.

Although the inclusion of actuation modeling in MPC is not new, to my knowledge there is no comparison of various approaches with MPC on a full-sized experimental vehicle. Nor do the actuation modeling methods address actuation dynamics in addition to pure time delays. As part of the research presented in this thesis, various actuation modeling problem formulations are evaluated using low-fidelity models as applied to Trudi, an automated Ford Fusion. In particular, the signal delay is captured by shifting the input vector by the delay time, thus reducing the computational complexity. The actuation dynamics are incorporated into the system dynamics using first- and second-order identified system models. All the models improve the system behavior. Yet, a simple first-order lag is demonstrated to be sufficient with Trudi.

The chapter is organized as follows: the steering system of Trudi is explained in

Section 2.2. In Section 2.3, a baseline MPC approach is benchmarked on Trudi and includes the experimental setup used for comparison of various actuation modeling approaches. A pure time delay approach is presented in Section 2.4, while approaches for addressing dynamic lag are demonstrated in Section 2.5. Section 2.5 also entails a discussion with further analysis of each approach, and a summary is provided in Section 2.6.

2.2 Steering System

Trudi is a fully automated hybrid Ford Fusion with drive-by-wire systems shown in Figure 2.1a. Because of the EPAS system, the hand wheel actuates via a motor mounted on the steering column. The relationship between the hand wheel angle (δ_{hw}) actuation and the road wheel angle (δ) can be approximated linearly as

$$\delta = K_{\mathrm{hw}} \delta_{\mathrm{hw}}, \tag{2.1}$$

where K_{hw} is the ratio between hand wheel angle and front road wheel angle.

In contrast X1, depicted in Figure 2.1c, has a steer-by-wire system meaning there is no mechanical linkage between the hand wheel and the road wheels. The corresponding step response of X1's steering system is shown in Figure 2.1d. The response is instantaneous and demonstrates the challenge of moving an MPC problem formulation designed for X1 onto a production vehicle with slower dynamics.

The dynamics of the steering system are captured through step response data at various speeds. Figure 2.1b displays commanded and measured hand wheel angles at speeds of 4, 6 and 8 m/s (gray dots). The Trudi step steers were performed with commands of 45° hand wheel angle, and the X1 step steers commanded 3° of road wheel angle. The figure also depicts the median and mean step response of the steering data to summarize the transients in one signal. The median transient is hereafter used to characterize the dynamics of the steering system because of the outliers in the measurements.

(a) Trudi, an automated Ford Fusion.

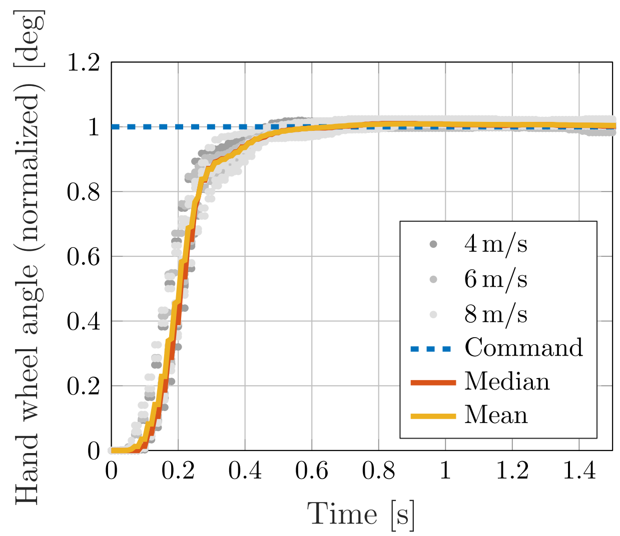

(b) Trudi steering step response.

(c) X1, an automated electric vehicle.

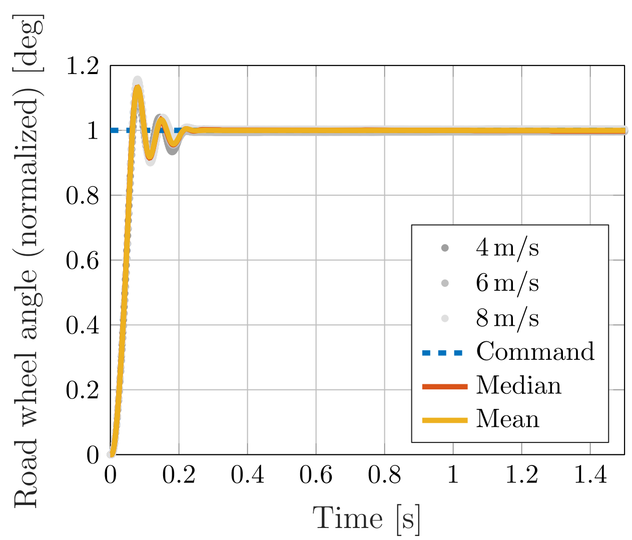

(d) X1 steering step response.

Figure 2.1: Experimental vehicles Trudi and X1, and their respective steering system step response. Trudi's steering system is controlled by EPAS, while X1 is a real steer-by-wire system.

Based on the step steer data, the response of Trudi's steering system is not instantaneous, and a pure time delay is evident. The steering system has a known communication delay of 40 ms. The pure time delay depicted in Figure 2.1b totals 80 ms: 40 ms for the command signal and 40 ms for the measure signal. Ignoring the measure signal communication delay, the steering system has a 40 ms pure time delay. The dynamic lag between the command and measure is characterized in Section 2.5.

While this step response data represents the steering system on Trudi, other automated vehicles with an EPAS system will also have dynamics in the steering actuation. By addressing the actuation dynamics, the following sections on modeling the delay can be generalized to other automated vehicle steering systems.

2.3 Model Predictive Control

Before accounting for the delays in the steering system, an MPC problem formulation without delay modeling as demonstrated by Funke, Brown, Erlien, and Gerdes [40] is benchmarked with Trudi's steering system.

2.3.1 Problem Formulation

The vehicle model used in the baseline MPC formulation is a four-state bicycle dynamic model with a constant acceleration assumption. The state vector (x) comprises vehicle lateral velocity (U_{y}), yaw rate (r), heading deviation ($\Delta\psi$), and lateral deviation (e):

$$x = [U_{\mathrm{y}} \quad r \quad \Delta\psi \quad e]^{\top}. \tag{2.2}$$

The control input to the vehicle model (u) is the front steering angle (δ) calculated from an affine vehicle model:

$$x^{(k+1)} = A^{(k)}x^{(k)} + B^{(k)}u^{(k)} + C^{(k)}, \quad k = 0, \ldots, n-1 \tag{2.3}$$

for each time step in the prediction horizon (k) up to a finite number of time steps (n). The original problem formulation by Funke *et al.* uses front lateral tire force as the steering input, but here the analysis uses steering angle (δ) similarly to Falcone, Borrelli, Asgari, Tseng, and Hrovat [4].

The relation between steering angle (δ) and front lateral tire force (F_{y}) is nonlinear. The Fiala brush tire model [41] as presented by Pacejka [42] describes the relation as

$$F_{\mathrm{y}} = \begin{cases} -C_\alpha \tan\alpha + \frac{C_\alpha^2}{3\mu F_{\mathrm{z}}}|\tan\alpha|\tan\alpha & \\ \qquad -\frac{C_\alpha^3}{27\mu^2 F_{\mathrm{z}}^2}\tan^3\alpha, & |\alpha| < \tan^{-1}\left(\frac{3\mu F_{\mathrm{z}}}{C_\alpha}\right) \\ -\mu F_{\mathrm{z}} \mathrm{sgn}\ \alpha, & \text{otherwise} \end{cases} \tag{2.4}$$

where slip angle (α) is the angle from the tire heading to the velocity vector of the tire, C_α is the cornering stiffness of the tire, μ is the coefficient of friction, and F_{z} is the normal load on the tire. For the front tires, the front slip angle (α_{f}) can be expressed as follows assuming small angles and the vehicle model in Appendix A:

$$\alpha_{\mathrm{f}} = \tan^{-1}\left(\frac{U_{\mathrm{y}} + ar}{U_{\mathrm{x}}}\right) - \delta \approx \frac{U_{\mathrm{y}} + ar}{U_{\mathrm{x}}} - \delta. \tag{2.5}$$

In order to capture accurate behavior while maintaining convexity of the bicycle model, Erlien, Funke, and Gerdes approximate the tire curve for rear tires with an affine, time-varying model that uses successive linearization points [43]. A similar method can be applied at the front tires:

$$F_{\mathrm{yf}} = \left.\frac{\partial F_{\mathrm{yf}}}{\partial \alpha_{\mathrm{f}}}\right|_{\alpha_{\mathrm{f},0}} (\alpha_{\mathrm{f}} - \alpha_{\mathrm{f},0}) + F_{\mathrm{yf}}(\alpha_{\mathrm{f},0}). \tag{2.6}$$

This is simply a Taylor expansion about an operating point ($\alpha_{\mathrm{f},0}$), where $\alpha_{\mathrm{f},0}$ at each time step in the prediction horizon is calculated from the previous optimization's solution. With Eq. (2.5) and using steer angle as the controller input, the system is represented by linear differential equations. Zhang, Thornton, and Gerdes [44] details the full derivation.

Path tracking is accomplished by associating a nonzero diagonal entry in the weighting matrix (Q) to lateral deviation and heading deviation as these states are defined relative to a nominal or desired path. The complete optimization problem is as follows:

$$\underset{\mathbf{u}}{\text{minimize}} \quad \sum_{k=0}^{n} v^{(k)\top} R^{(k)} v^{(k)} + \sum_{k=1}^{n} x^{(k)\top} Q^{(k)} x^{(k)} \tag{2.7a}$$

$$\text{subject to} \quad x^{(k+1)} = A^{(k)} x^{(k)} + B^{(k)} u^{(k)} + C^{(k)} \tag{2.7b}$$

$$|u^{(k)}| \leq u_{\max}^{(k)} \tag{2.7c}$$

$$|v^{(k)}| \leq v_{\max}^{(k)} \tag{2.7d}$$

where $v^{(k)} = u^{(k)} - u^{(k-1)}$ is the change in front steer angle, weighting matrix (R) penalizes changes in steer angle, and $u_{\max}^{(k)}$ and $v_{\max}^{(k)}$ are physical limits in the steering system.

Optimization problem (2.7) is a quadratic program with a significantly sparse structure that can be leveraged with an efficient solver for real-time implementation [45]. For this work, CVXGEN, developed by Mattingley and Boyd [46], is used to solve for the input vector, $\mathbf{u} = [u^{(0)} \dots u^{(n)}]^\top$, but only the first solution in the vector ($u^{(0)}$) actually commands to the steering system. The optimization problem runs on a single core of a ruggedized computer with an i7 CPU at a rate of 100 Hz.

2.3.2 Experimental Results

As a baseline, the performance of problem formulation (2.7) was benchmarked on Trudi to determine if delay compensation would be necessary. The experiment to validate all the problem formulations in the chapter entails Trudi driving down a straight path on a flat paddock area at Thunderhill Raceway in Willows, California at a speed of 8 m/s. The vehicle is at speed before encountering a disturbance protruding into the path forcing Trudi to deviate laterally from the path by approximately 1 m. Trudi is equipped with an Oxford Technical Solutions RT4001 GPS unit and receives RTK differential correction information from a Novatel basestation to measure the vehicle position states within 2 cm accuracy.

Although a delay exists in the steering system, Trudi successfully drives around the path perturbation as depicted in the top plot of Figure 2.2. The middle plot in Figure 2.2 includes the command and measure of the hand wheel angle as Trudi

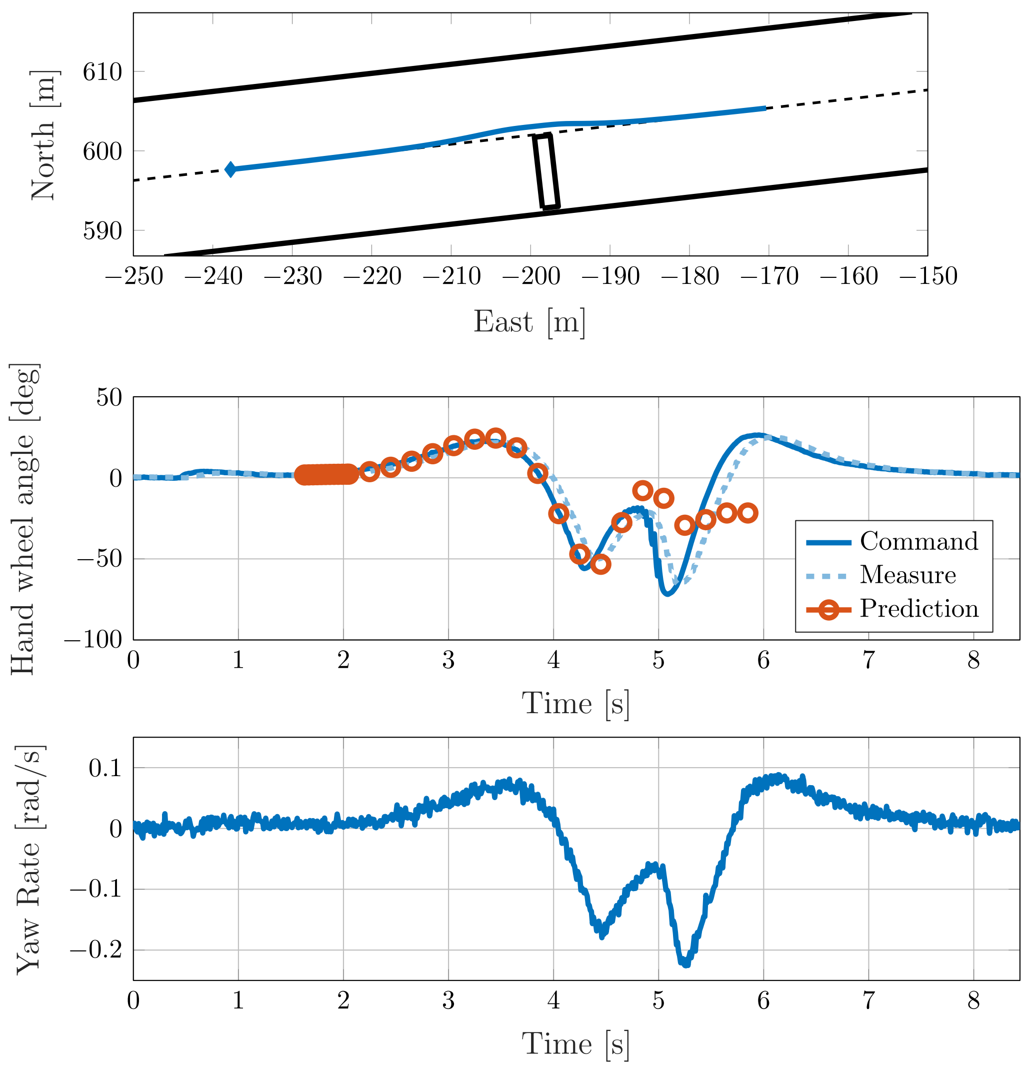

Figure 2.2: Baseline trajectory overhead, diamond indicates start of maneuver (top); command and measure hand wheel angle with open-loop prediction at one time step (middle); and yaw rate (bottom) due to lateral perturbation from the nominal path.

conducts the maneuver. It also portrays an open-loop prediction horizon for the steering command (scaled by the steering ratio) at time $t = 1.63\,\mathrm{s}$. At this instance in time, the optimal solution does not do well at predicting future steering behavior at the end of the prediction horizon, indicating a model mismatch. The root mean squared (RMS) error between the closed-loop command and the open-loop prediction is $16.13°$ along the prediction horizon, which indicates how well the open-loop prediction estimates the future behavior. The bottom of Figure 2.2 shows the vehicle yaw rate having large oscillations, signifying that the vehicle did not rotate sufficiently and had to adjust its rotation while completing the maneuver. This occurs because the delay prevents the correct amount of steer angle from commanding when needed, thus requiring more aggressive maneuvering to avoid the perturbation at a time step later than anticipated by the predictive controller. The RMS error of the yaw rate is $0.0638\,\mathrm{rad/s}$.

2.4 Modeling for Pure Delay

Based on the results in the previous sections, it is apparent that the delay in the steering system should be accounted for. In this section, only the pure time delay is incorporated in the MPC problem formulation.

2.4.1 Problem Formulation

Formally, the actuated steering command (u_{actual}) is delayed from the commanded steering command (u) by

$$u_{\text{actual}}(t) = u(t - T_{\text{delay}}), \tag{2.8}$$

resulting in the modeled state transition relationship

$$x(t + T_{\text{pred}}) = A(t)x(t) + B(t)u(t - T_{\text{delay}}), \tag{2.9}$$

where T_{delay} is the pure delay time and T_{pred} is the discretization time into the prediction horizon. Since T_{delay} is known to be approximately $40\,\mathrm{ms}$, problem formulation

(2.7) is modified to account for the pure delay time. Explicitly, I assume $T_{\text{delay}} \approx T_{\text{pred}}$ for the short time steps (first 10 time steps [47]) of the prediction horizon such that Eq. (2.9) can be equivalently represented in discretized form as

$$x^{(k+1)} = A^{(k)}x^{(k)} + B^{(k)}u^{(k-1)} + C^{(k)}, \quad k = 0, \dots, n-1. \tag{2.10}$$

The transition from $x^{(0)}$ to $x^{(1)}$ is a function of the optimal steering input solved for T_{pred} seconds prior.

The optimization problem accounting for $T_{\text{delay}} = T_{\text{pred}}$ can then be written as

$$\begin{aligned}
\underset{\mathbf{u}}{\text{minimize}} \quad & \sum_{k=1}^{n} v^{(k)\top}R^{(k)}v^{(k)} + x^{(k)\top}Q^{(k)}x^{(k)} && \text{(2.11a)}\\
\text{subject to} \quad & x^{(1)} = A^{(0)}x^{(0)} + B^{(0)}z^{-d}u^{*(0)} + C^{(0)} && \text{(2.11b)}\\
& x^{(k+1)} = A^{(k)}x^{(k)} + B^{(k)}u^{(k-1)} + C^{(k)} && \text{(2.11c)}\\
& |u^{(k)}| \leq u^{(k)}_{\text{max}} && \text{(2.11d)}\\
& |v^{(k)}| \leq v^{(k)}_{\text{max}} && \text{(2.11e)}
\end{aligned}$$

where $z^{-d}u^{*(0)}$ is the zero-th entry of the optimal solution vector calculated d control iterations prior. This solution should correspond to the solution T_{delay} seconds prior. Assuming the controller executes every T_{control} seconds, it follows that

$$d = \texttt{round}\left(\frac{T_{\text{delay}}}{T_{\text{control}}}\right). \tag{2.12}$$

Note that the decision variable vector $\mathbf{u}$ has been reduced from size $n+1$ to size n due to the delay formulation leveraging the previous control input in seeding Eq. (2.11b).

2.4.2 Experimental Results

By only accounting for the pure time delay in the steering system, the open-loop prediction improves, as shown in the middle plot of Figure 2.3. The hand wheel angle RMS over the prediction horizon for this implementation, also at time $t = 1.63\,\text{s}$, is reduced by about a factor of two to 7.87°. The yaw rate for this maneuver, shown

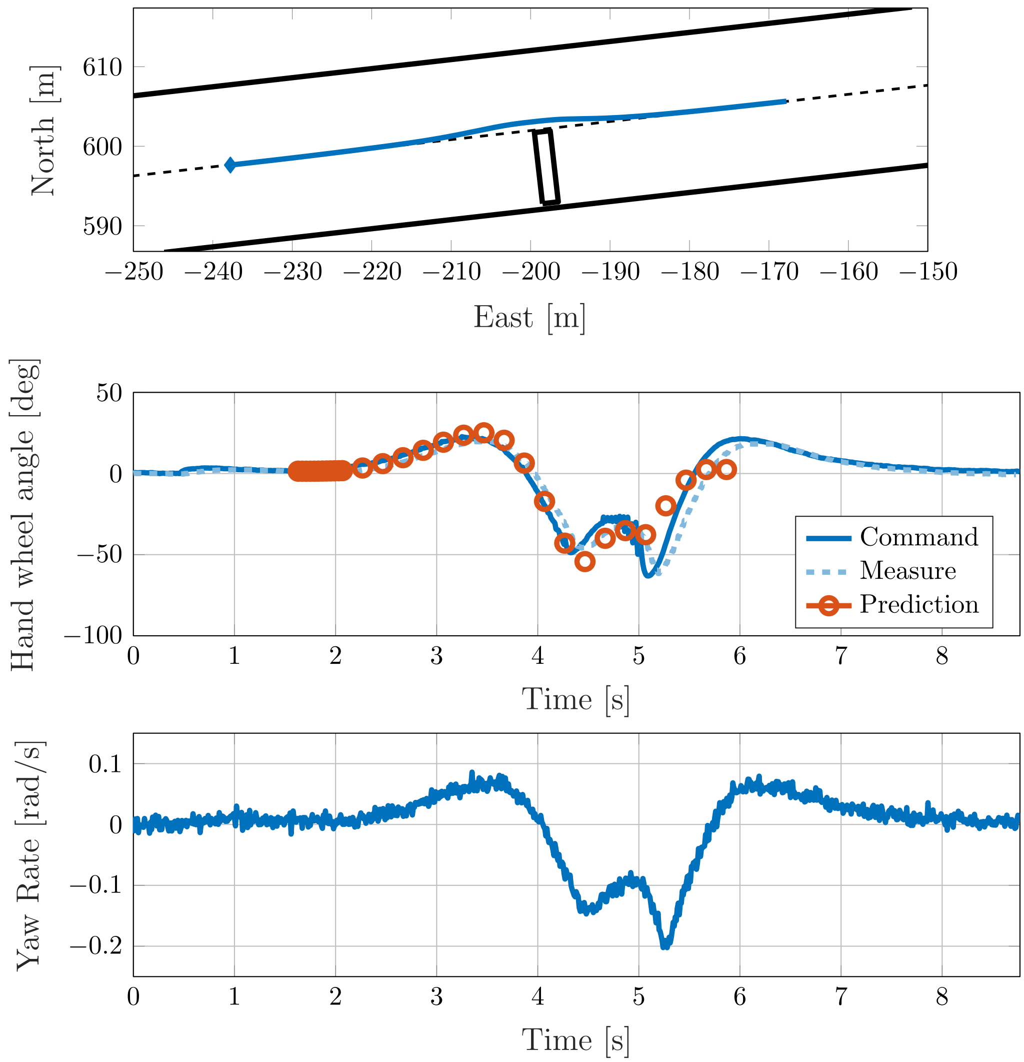

Figure 2.3: Pure time delay trajectory overhead, diamond indicates start of maneuver (top); command and measure hand wheel angle with open-loop prediction at one time step (middle); and yaw rate (bottom) due to lateral perturbation from the nominal path.

in the bottom plot of Figure 2.3, also has a reduced RMS error of 0.0574 rad/s. The trajectory taken with the pure time delay problem formulation is comparable to the baseline trajectory.

2.5 Modeling for Dynamic Lag

To further improve the open-loop prediction, the dynamic lag in the steering system as introduced in Figure 2.1b is incorporated into the MPC problem formulation. The following section compares three approaches: a first-order lag model in conjunction with the pure time delay, a first-order lag model tuned to lump in the pure time delay, and a second-order lag model also tuned to lump in the pure time delay.

2.5.1 First-Order Lag

To account for the first-order dynamics, the state vector is appended with a fifth state δ_{fol} to represent the steering input to the affine vehicle model after a first-order delay. The new state vector becomes

$$x = [U_{\text{y}} \quad r \quad \Delta\psi \quad e \quad \delta_{\text{fol}}]^\top. \tag{2.13}$$

The affine vehicle model is also appended with δ_{fol}, which can be expressed as

$$\dot{\delta}_{\text{fol}}(t) = -\frac{1}{\tau}\delta_{\text{fol}}(t) + \frac{1}{\tau}u(t), \tag{2.14}$$

where τ is the time constant.

The five state vector is used in both problem formulation (2.7) and (2.11), but with different values for the time constant. The step responses of the first-order models are shown in Figure 2.4.

2.5.2 Second-Order Lag

Since there is a pure time delay with seemingly first-order dynamics, the steering delay can be roughly approximated as a second-order system. Similarly to the first-order

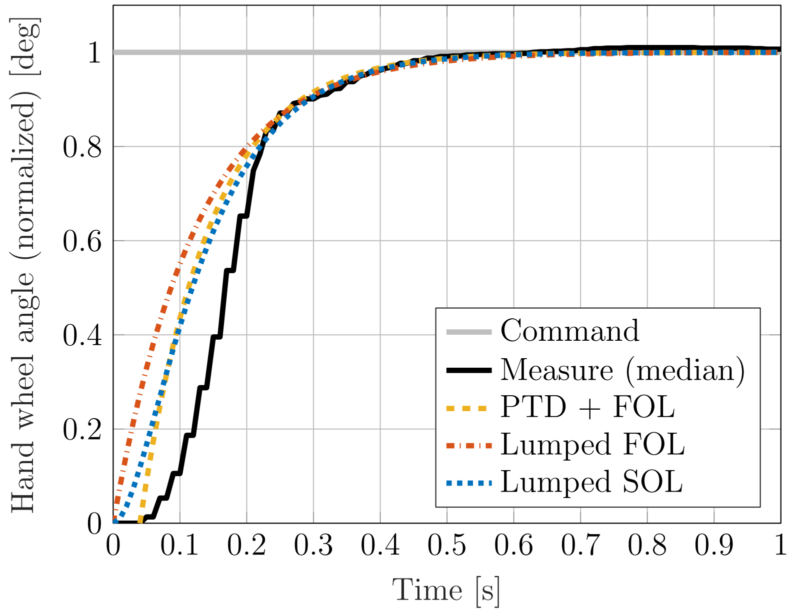

Figure 2.4: First- and second-order models on median step response data.

model, the original four state vector has two additional states appended:

$$x = [U_{\mathrm{y}} \quad r \quad \Delta\psi \quad e \quad \delta_{\mathrm{sol}} \quad \dot{\delta}_{\mathrm{sol}}]^{\top}. \tag{2.15}$$

The new states are related to the actual steering input as

$$\ddot{\delta}_{\mathrm{sol}}(t) = -\omega_n^2 \delta_{\mathrm{sol}}(t) - 2\zeta\omega_n\dot{\delta}_{\mathrm{sol}}(t) - \omega_n^2 u(t), \tag{2.16}$$

where ω_n is the undamped natural frequency and ζ is the damping ratio.

This delay model is only used with problem formulation (2.7) and thus is tuned to have the step response shown in Figure 2.4.

2.5.3 Experimental Results

Since the actual dynamics of the steering system are unknown, three additional experiments were performed to determine which of the problem formulations accounting for dynamic lag performs best on Trudi, as shown in Figures 2.5, 2.6 and 2.7. A summary

Table 2.1: RMS of hand wheel angle (HWA) along prediction horizon, RMS of yaw rate, and maximum absolute yaw rate

Experiment	Prediction horizon HWA RMS† [°]	Yaw rate RMS [rad/s]	Maximum yaw rate* [rad/s]
Baseline	16.13	0.0638	0.2259
Pure time delay	7.87	0.0574	0.2024
PTD + FOL	7.76	0.0536	0.1895
Lumped FOL	8.16	0.0555	0.1678
Lumped SOL	8.38	0.0562	0.1665

†at $t = 1.63\,\mathrm{s}$
*absolute value

of the RMS error for the hand wheel angle through the prediction horizon at time $t = 1.63\,\mathrm{s}$ along with RMS and maximum absolute yaw rate are in Table 2.1. The maximum absolute yaw rate is included because of its relation to lateral acceleration in the inertial frame:

$$a_{\mathrm{Y}} = \dot{U}_{\mathrm{y}} + rU_{\mathrm{x}}. \tag{2.17}$$

Lateral acceleration in the vehicle frame ($\dot{U}_{\mathrm{y}}$) is small compared to the rU_{x} term, and longitudinal velocity is consistent across all the experiments. Thus, the maximum absolute yaw rate provides insight into how much lateral acceleration vehicle occupants experience during each maneuver.

The open-loop prediction horizons in the middle plots of Figures 2.5, 2.6 and 2.7 qualitatively seem very similar because the respective models closely represent the steering actuation. Although the commanded hand wheel angles depict jitter, the measured response is quite smooth; the jitter is likely due to location of the steering measurement in the EPAS system. The overall smoothness of the maneuver is furthermore captured by the yaw rate response. Quantitatively, the RMS error and maximum of the yaw rate continues to improve compared to the baseline and pure time delay implementations.

A close look at the numbers in Table 2.1 suggests there is a trade-off in determining the "best" delay modeling problem formulation to use with Trudi. There is little

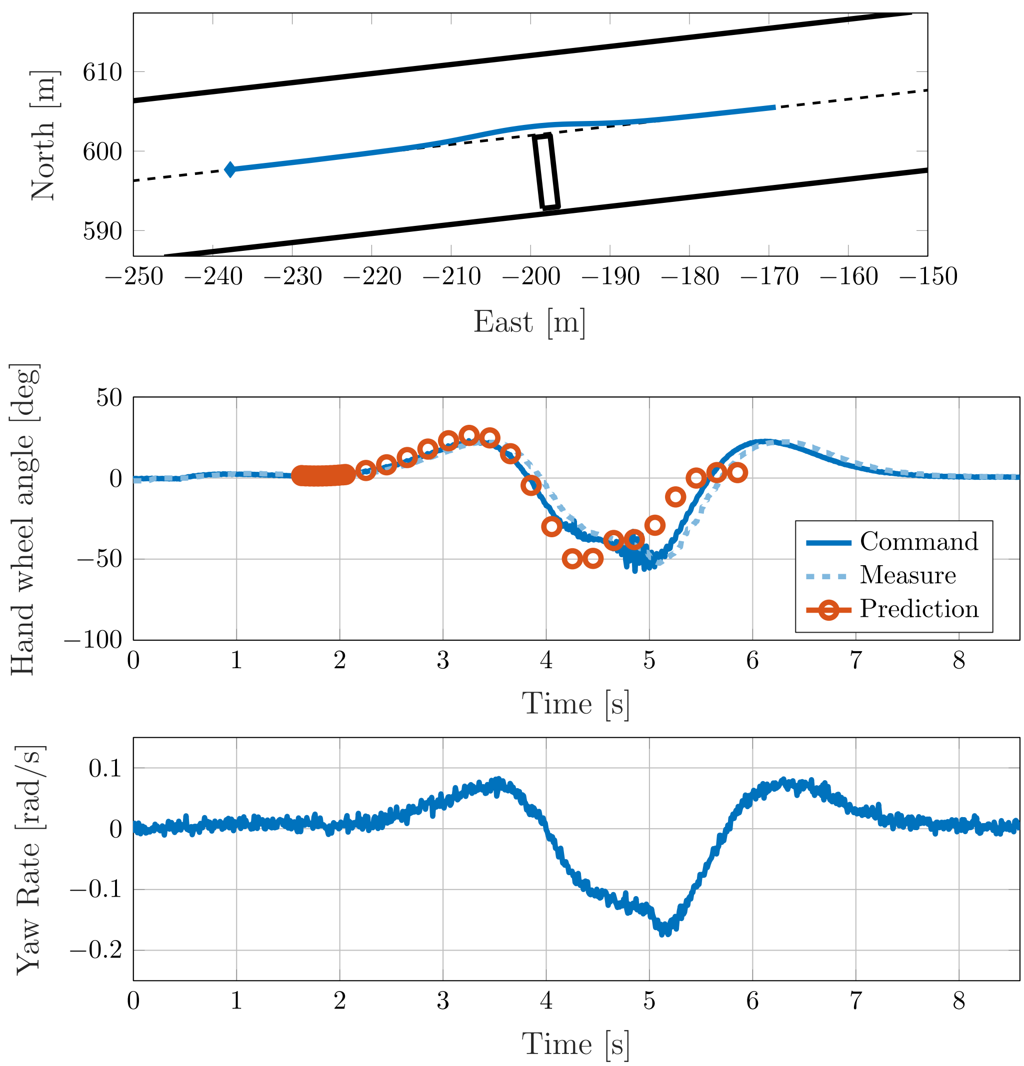

Figure 2.5: Pure time delay with first-order dynamics trajectory overhead, diamond indicates start of maneuver (top); command and measure hand wheel angle with open-loop prediction at one time step (middle); and yaw rate (bottom) due to lateral perturbation from the nominal path.

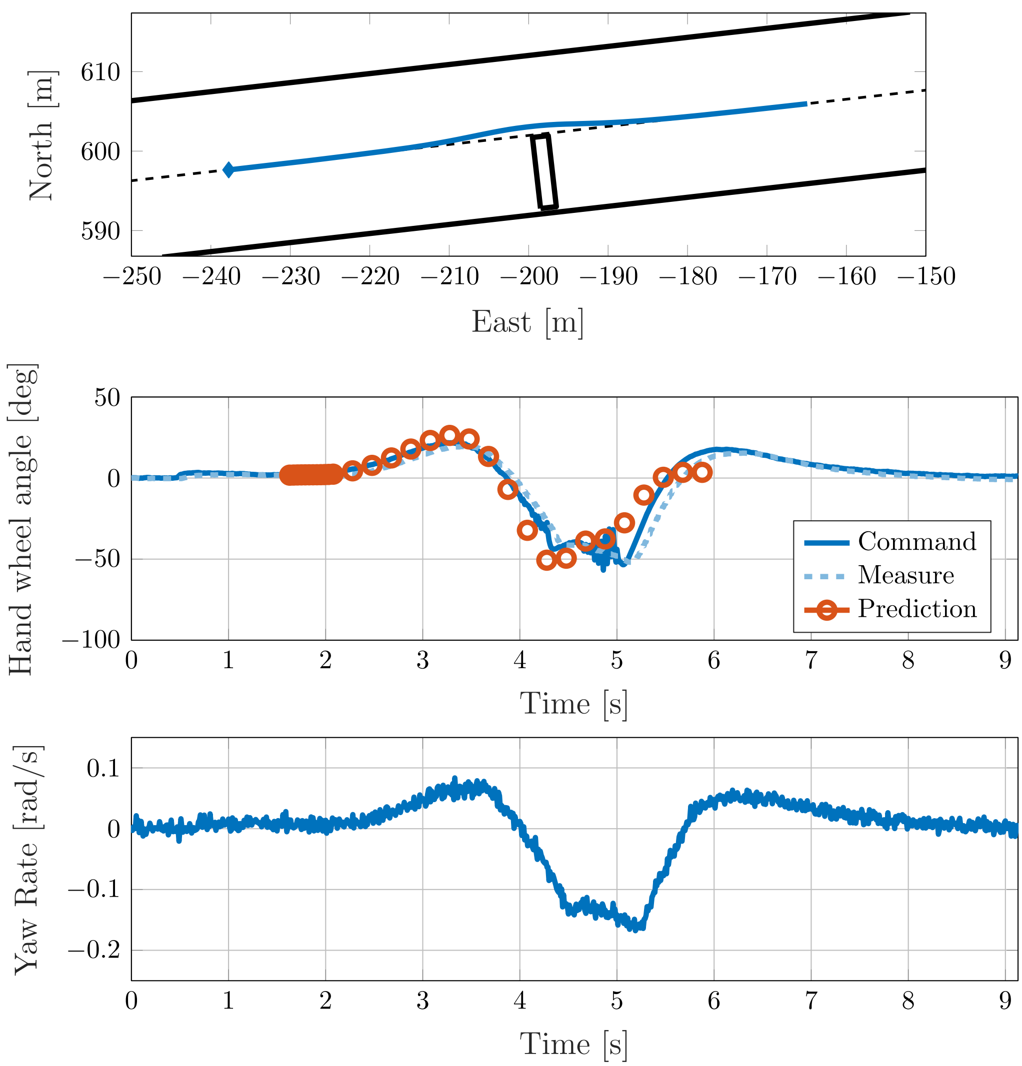

Figure 2.6: Lumped first-order lag trajectory overhead, diamond indicates start of maneuver (top); command and measure hand wheel angle with open-loop prediction at one time step (middle); and yaw rate (bottom) due to lateral perturbation from the nominal path.

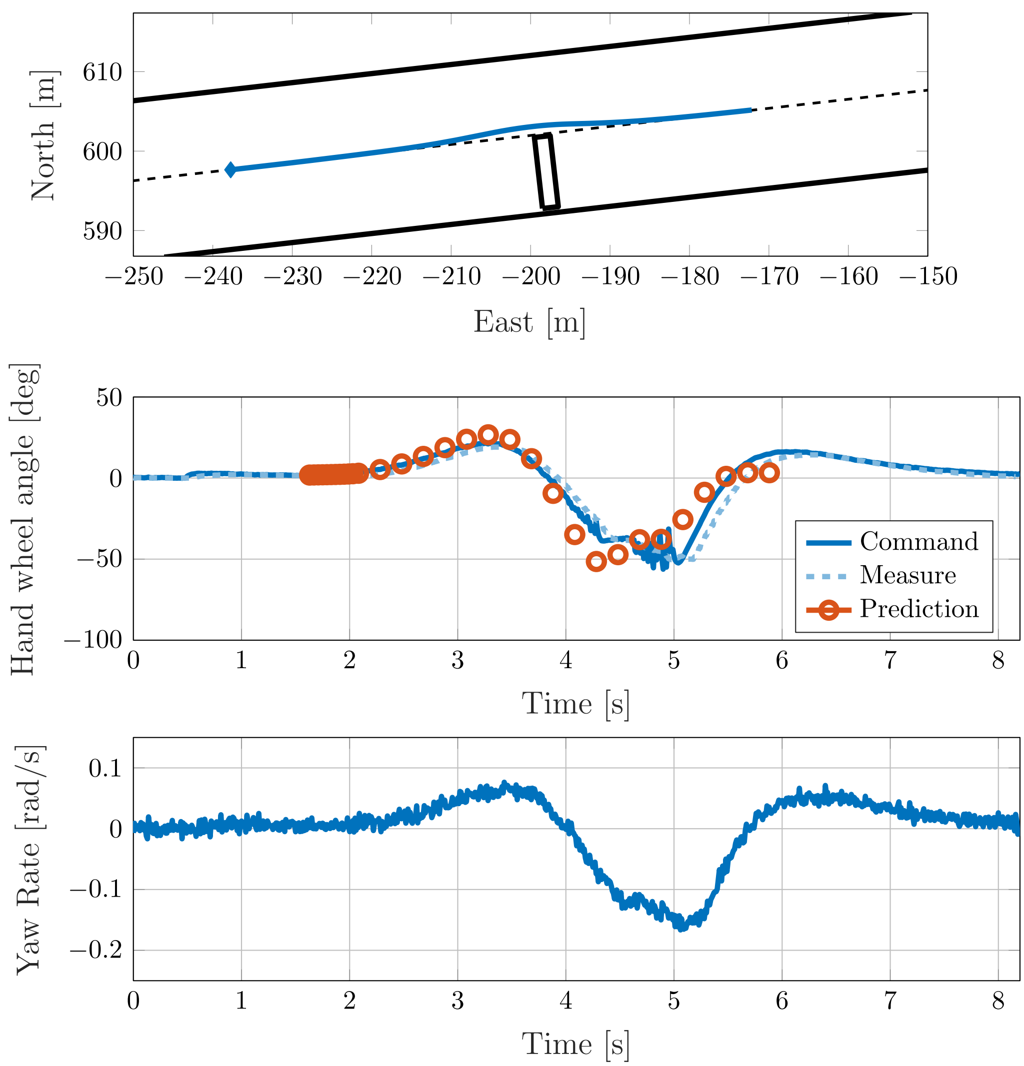

Figure 2.7: Lumped second-order lag trajectory overhead, diamond indicates start of maneuver (top); command and measure hand wheel angle with open-loop prediction at one time step (middle); and yaw rate (bottom) due to lateral perturbation from the nominal path.

difference between the lumped first-order and second-order RMS errors and maximum absolute yaw rate. At the computational cost of a larger state space, the lumped second-order approach produces a marginally smoother yaw rate response. If the pure time delay is to be modeled explicitly, then additional information needs to be maintained in-vehicle, such as how many control time steps are needed to delay the initial input to the optimization problem. For Trudi, I decide to proceed with the lumped first-order approach because of the comparable performance, smaller state space, and simpler implementation.

2.6 Summary

Four MPC delay modeling problem formulations are presented to address the pure time delay and dynamic lag in the steering actuation of an automated production vehicle. Each problem formulation is compared to a baseline approach. The first of the four modeled delay methods addresses just the pure time delay, while the last three also address the dynamic lag. In particular, the dynamic lag is modeled as a first-order lag and as a second-order lag, where the first-order lag is tuned both with and without the pure time delay.

The problem formulations are all validated in experiments: the vehicle travels down a straight path until a lateral perturbation forces the vehicle to deviate from the nominal path by approximately 1 m. Adding in the delay dynamics improves the open-loop prediction at a time prior to the maneuver compared to the baseline approach, although the closed-loop trajectory of the vehicle navigating around the perturbation is comparable across all the experiments. Ultimately, the lumped first-order approach meets the design objectives of comparable performance while maintaining low computational complexity and simple implementation. Now that the actuation delay is accounted for in the MPC formulation, it is safe to control the vehicle in more complicated scenarios.

Chapter 3

Costs, Constraints, and Weights from Philosophical Principles

The value of safety in autonomous vehicle motion planning goes beyond predictable control of the steering actuation system. Designing control algorithms for automated vehicles presents new challenges for engineers. Traditionally, control systems have desired specifications and performance measures to which programmers design the control algorithms. For fully automated vehicles, the ultimate desired performance outcome is the ability to drive safely and smoothly through traffic. Setting specifications for achieving this desired outcome is challenging because traffic situations involving humans are not easily quantified. Driving through traffic requires that the vehicle conform to societal expectations for roadway behavior. Expectations such as collision avoidance and observance of traffic laws go beyond mere technical specification to touch on moral issues that are long established and formally characterized in philosophy. In this chapter, value conflict resolution manifests through the connection of philosophical frameworks to mathematical frameworks used in programming autonomous vehicles. The majority of this chapter was originally published in the IEEE Transactions on Intelligent Transportation Systems of 2016 [48].

3.1 Introduction

The objective of accident avoidance is fundamentally motivated by the idea of caring for life and avoiding harm. Haidt *et al.* described care (and its opposite, harm) as one of the foundational principles for moral reasoning [49]–[52]. A vehicle's compliance with traffic laws involves another moral foundation—the degree of respect for authority. Interactions with other road users should, furthermore, be based on fairness and reciprocity, yet another category in the five (or sometimes six) moral foundations found in Haidt's work. The fact that these societal expectations of automated vehicles map so cleanly to ethical principles in philosophy suggests that philosophy can be a useful resource for translating such expectations to specifications. In the same realm of ethics in engineering but with different applications, Mladenovic and McPherson [53] draw from social justice in the design of a traffic control framework for automated vehicles. Miller, Wolf, and Grodzinsky [54] and Van den Hoven, Lokhorst, and Van de Poel [55] all advocate for the consideration of ethics throughout the entire engineering process. In particular, Miller *et al.* summarize ethics models used in operations research and extend it to ethical decision-making machines. Lin also makes the case that ethical issues are central to the design of automated vehicles [56].

But as both Lin *et al.* [57] and Wallach and Allen [58] point out, a single philosophical framework is unlikely to be sufficient for programming autonomous systems. As a result, researchers have proposed solutions that combine different concepts from philosophy. Deontology, a rule-based ethical framework, and consequentialism, a cost-based ethical framework, both contribute structured guidelines for vehicle behavior. Goodall presents a three-tiered system for ethical decision-making in automated vehicles [59]. The first tier is a rational approach in which the vehicle follows the ethical principles of deontology and consequentialism. The second and third tiers involve artificial intelligence and a combined rational-artificial intelligence approach. Gerdes and Thornton present the two ethical frameworks of deontology and consequentialism as parallel to constraints and cost, respectively, in an optimal control problem [33]. Since many semi-autonomous and autonomous vehicles are already designed based on this type of control formulation (Gray *et al.* [60], [61]; Gao *et al.* [62]; Erlien, Fujita, and Gerdes [47], [63]; Falcone, Borrelli, Asgari, Tseng, and Hrovat [4]), such an approach makes the links between philosophy and engineering quite direct. Thus, the work presented here uses multiple ethical frameworks, deontology and consequentialism, and parallels the concepts to a constrained optimization problem as a starting point.

The goal of this chapter is twofold. Firstly, the goal is to introduce engineers to ethical theories that parallel engineering paradigms in order to elucidate how such structures may implicate a particular ethical theory. Secondly, the goal is to use these ethical principles to motivate engineering decisions that result in reasonable, justifiable automated vehicle behavior. In particular, the approaches of normative ethical theories, deontology and consequentialism, are used to reason driving goals as rules or costs as appropriate. These goals can be translated into constraints and cost functions that can be used in motion planning algorithms, such as a model predictive control (MPC) formulation. This enables ethically motivated design decisions to be demonstrated and compared on a test vehicle for a simple traffic scenario. Assuming the construction of an optimization problem, there is a challenge in choosing appropriate weights for different objectives, so an ethical theory known as virtue ethics in the form of role morality may assist with this choice. Virtue ethics and role morality provide an ethical framework based on alignment with character (Hursthouse [64], Harman [65]). As applied to vehicles, this framework guides the algorithm design to achieve desired behavior for different types of vehicles. To the knowledge of the author, this is the first quantitative and in-vehicle experimental endeavor to incorporate ethical reasoning in the design of autonomous vehicle control.

This chapter is structured as follows: Section 3.2 presents a motivating driving scenario based on a lane blocked by an obstacle. Section 3.3 describes how the philosophical concepts of deontology and consequentialism relate to engineering choices. These ideas combine in Section 3.4, which suggests an ethical reasoning of how to address driving goals associated with the simple scenario through the philosophical concepts. The driving goals considered are path tracking, steering, obstacle avoidance, and traffic laws. While many of these are straightforward, traffic laws can

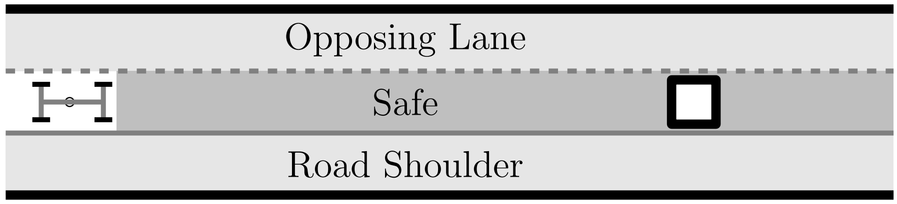

Figure 3.1: The shaded regions indicate driving regions. The safest region is the current lane of the vehicle excluding any obstacles. Driving regions decrease in safety as the vehicle departs the lane.

combine elements of cost and constraint. Section 3.5 formalizes an MPC problem and Section 3.6 shows in-vehicle experimental results, highlighting the impact of different formulations of traffic laws. This raises a larger question of virtue ethics and role morality as applied to automated vehicles, discussed with additional in-vehicle experiments in Section 3.8.

3.2 Scenario

To contextualize the relationship between ethics and engineering in autonomous vehicles, a simple, realistic driving scenario is constructed that involves a variety of factors, including collision avoidance, mobility considerations, and traffic laws. Figure 3.1 depicts the scenario of an autonomous vehicle traveling along a two-lane roadway at constant speed. The current lane is obstructed by an obstacle ahead of the vehicle. This simple scenario prompts a range of possibilities for engineering decisions. Section 3.5 details how the various parameters from this scenario (collision avoidance, mobility, traffic laws, speed) fit into an MPC formulation.

One engineering design option is to program the vehicle to prioritize the ability to continue moving. This would mean entering the opposing lane or the road shoulder to move around the obstacle to continue on its way. If the vehicle chooses the option of entering the opposing lane, it could cause the vehicle to briefly violate a traffic law, for example if the lane divider is a double yellow line. If the vehicle travels into the road shoulder to move around the obstruction, it obeys the double yellow line traffic law and continues to move, but the road shoulder would need to be accessible and safe and is not meant to be driven on ordinarily. Clearly, competing demands of mobility, safety, and legality must be weighed when considering these options.

With a potential sacrifice in mobility, another approach is to program the automated vehicle to adhere strictly to traffic laws. Here, the underlying assumption is that traffic laws are implemented as hard rules. However, such strict obedience to traffic laws has consequences: the vehicle could stop and remain stopped indefinitely if trapped between the dual goals of avoiding the obstruction and obeying the double yellow line. This action could negatively impact the mobility or safety of surrounding vehicles.

The scope of engineering decisions that must be determined does not merely lie in the type of action to take; the degree of the action needs to be assessed, too. For example, in the case of crossing the double yellow line to maneuver around the obstruction and encouraging smooth traffic flow, the amount of clearance between the vehicle and the obstruction is a design consideration. A narrow space between the vehicle and obstruction allows the vehicle to stay closer to its original designated lane should oncoming traffic emerge, but can increase the risk of brushing against the side of the obstruction. A wider berth ensures passage without hitting the obstruction, but the vehicle is then positioned further into the opposing lane and would take longer to return to its original designated lane. The degree to which a vehicle is tuned to disobey a traffic law is therefore another engineering decision that requires careful ethical consideration. In addition to the type and degree of action taken, another layer of engineering design decisions involves the fact that different types of vehicles could be given different traffic law allowances depending on their expected role in society. This is discussed later on through the example of an ambulance and a taxi, both of which ferry passengers but can exhibit very different behavior on roads due to their differing purposes.

Deconstructing this simple and common driving scenario highlights the many different vehicle behaviors that can result from making different engineering design decisions and underscores the need to make those design decisions in a systematic, reasoned, and justifiable manner. The decisions should be explainable not only to engineers but also to other road users that will share the streets with the automated vehicle and to regulators who ensure traffic safety. An understanding of philosophical frameworks in the context of engineering can provide engineers with a reasoning tool to consider the ethical implications of their engineering design choices.

3.3 Philosophical Frameworks

Ethical principles have been a topic of analysis among philosophers for centuries. This section examines those principles in the context of framing vehicle behavior. An autonomous vehicle is programmed by an engineer whose programming adheres to a system of decision-making and control logic. Although control logic and ethics are clearly not equivalent, there do exist ethical frameworks that provide applicable motivation to mathematical frameworks, which is examined in this section.

Deontology is one of the major normative ethical theories. Deontological ethics evaluates the morality of one's action according to some rules. So, in essence, to be moral, one must follow a set of rules that determine the correct, ethical action, and these rules are to be followed with no exception. Isaac Asimov's Three Laws of Robotics [66] are an example of deontological ethics, which state:

1. A robot may not injure a human being, or, through inaction, allow a human being to come to harm.

2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.

3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

The Three Laws of Robotics state a clear architecture of behavioral rules for the robots in Asimov's stories to follow, effectively serving as constraints on the robots' behavior. As long as a robot remains within the conditionals of the Three Laws, it can operate as necessary. Of course, Asimov also highlights the limits of such an approach as his robot stories often involve strange behavior arising from conflicting interpretations of these laws.

Deontology provides one type of motivating structure for the programming of automated vehicle algorithms: rules that can be defined and followed on the road. Such rules are analogous to conditionals and constraints used in decision-making and in control algorithms to bound and manipulate system behavior (for example, a conditional for actuation saturation or a constraint in an optimization problem). For an autonomous vehicle, examples of such rules can be found in constraints designed to prevent the vehicle from causing harm to human beings, from inducing property damage to itself or to other objects, or from violating traffic laws. A key feature of a deontological framework is that rules can be hierarchical, thus setting clear priorities. From a programming perspective, the ability to weave together dependencies and hierarchies provides the advantage of clarity in reasoning for the development of the algorithm. However, a strict deontological construction of an algorithm could define overly restrictive driving goals.

Another central normative ethical theory is *consequentialism*, which evaluates the moral validity of actions solely based on the consequences. Consequentialism is often posed as the exact opposite of deontological ethics in an attempt to address the limitations of deontology. There are many flavors of consequentialism, and the focus here is on a form of consequentialism known as *act utilitarianism*. Act utilitarianism analyzes the expected utility of a scenario and evaluates the consequences of actions based on what produces the most good [67]. Imagine Santa Claus is injured and an ambulance driver must rush him to the hospital because Santa brings a lot of good to the world with many gifts. The driver of the ambulance can justify breaking traffic laws and taking other actions as necessary in order to expedite Santa's recovery through consequentialism. Consequentialism also suffers from limitations because it can be difficult to define what is "good."

Consequentialism, through its more specific form of act utilitarianism, provides a basis for reasoning ethical behavior in an optimization problem. In control theory, optimal control uses an optimization problem to mathematically determine an optimal solution to be used as the control action. Specifically, the optimal control action (i.e., the ethically correct decision) is the feasible solution that minimizes the cost function (i.e., the desired outcome toward which one strives). An example of such a cost function for autonomous vehicles could be to minimize harm to the vehicle occupants. The optimal solution in consequentialist terms would be to maneuver the vehicle to achieve that goal of minimizing harm to the occupants, no matter what. This approach also has some limitations, such as the difficulty in actually forming or evaluating the cost function (as is the case with definitions such as "harm" [33]) or making that cost function comprehensive (by, for instance, considering road users other than the occupants in this case).

3.4 Design Choices

The frameworks of deontology and consequentialism are the result of extensive research in philosophy. Given their seminal presence in philosophy, these frameworks are used as reasoning tools to motivate and explain design decisions in the programming of an automated vehicle. There are many ways to program an automated vehicle as explored in Chapter 1. In this chapter, an MPC formulation of the problem is adopted since the explicit consideration of constraints and costs in MPC maps well to the concepts of deontology and consequentialism. With this philosophical reasoning, the vehicle actions are bound by deontological constraints and consequentialism is realized through the cost function. Thus, the philosophical frameworks are used to reason through the construction of this constrained optimization problem.

The following sections demonstrate how approaching the problem with these two philosophical frameworks in mind leads to a systematic treatment of different goals in the problem. Specifically, constraints are set that direct the vehicle to avoid collisions, follow dynamical equations, and stay within its steering capabilities. The cost function is designed to direct the vehicle toward the desired outcomes of tracking a prescribed path and providing acceptable occupant comfort. In contrast with these other objectives, it is less clear whether traffic laws represent a cost or a constraint so different representations are explored.

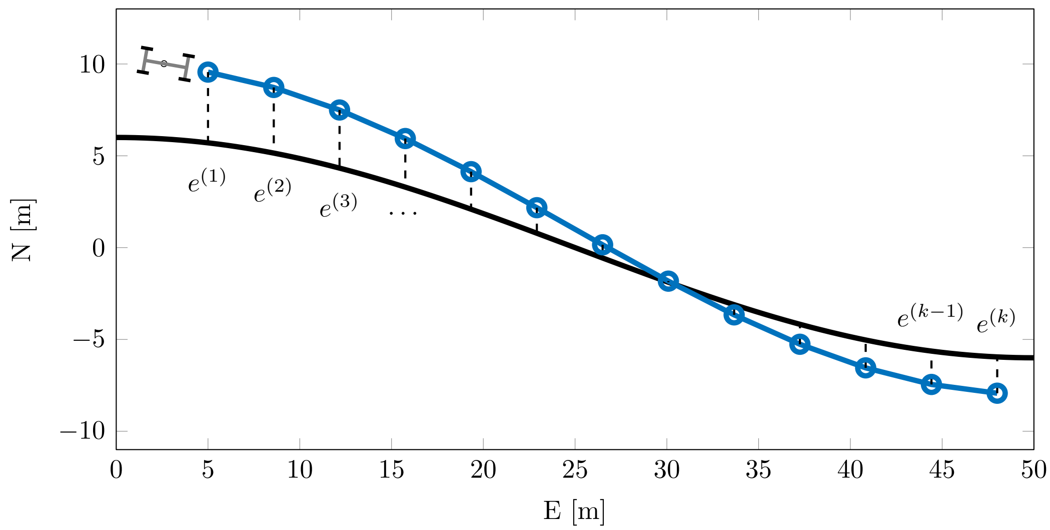

Figure 3.2: Generating a cost from the difference between a desired path (black) and the vehicle's actual path (blue with dots).

3.4.1 Path Tracking

A key objective of an autonomous vehicle is to follow a designated path. This objective assumes a higher level planner provides the reference path, which is not guaranteed to be obstacle-free. Following a path is a physical condition based on a measure of position difference, so a constraint constructed through deontological reasoning could be used to ensure path tracking. This could take the form of an equality constraint for the position of the vehicle to be equal to the desired position on the path. Upon further examination, however, following the path is not a strict requirement when it comes to maintaining safety; if an obstacle appears in the path, then the vehicle should have the option to deviate. Accounting for this line of reasoning, a cost function alternatively can serve as the instrument that accomplishes the goal of tracking a path via optimization, as depicted in Figure 3.2. Thus, in the choice of framing path tracking in a consequentialist manner, the vehicle is allowed freedom to deviate; if path tracking were denoted as a rule in a deontological framework, then a safety hazard would arise in rule conflict and problem feasibility. The fundamental idea behind deontology, which is that rules are to be followed without exception, is bypassed in favor of the more flexible principles from consequentialism for the specific goal of path tracking.

The objective of path tracking is translated from a consequentialist form into a mathematical structure. In order to follow the path, the vehicle must achieve the two goals of minimizing lateral deviation from the path (e) and heading error ($\Delta\psi$) using a cost function, such as

$$J_x = \sum_k x^{(k)\top} Q^{(k)} x^{(k)}, \tag{3.1}$$

where x is the vehicle state vector encompassing e and $\Delta\psi$ (explained in further detail in Section 3.5.1 and Appendix A), k is the discrete time step in the prediction horizon, and the weight matrix (Q) only has diagonal, non-zero entries corresponding to e and $\Delta\psi$.

3.4.2 Steering

The vehicle steering encompasses a few different design goals. The steering must operate within the actuator limits, should contribute toward path tracking and obstacle avoidance, and should perform smoothly as an aspect of mobility. The first of those goals, operating within the actuator limits, can be cast in the form of a constraint on a maximum slew rate. The reasoning behind this design choice is that this cap is a physical limit on an actuator. Physical limits must be highly prioritized in the control scheme and therefore are enforced as constraints in the deontological sense of strict rule obedience. A solution requiring control inputs that violate physical limits is simply not feasible. Thus, categorizing the slew rate limit in a deontological manner is most appropriate. Further, the physical limits of the vehicle are akin to acknowledging the laws of nature, which can be considered to be governing rules to follow. This limit can be represented mathematically:

$$\left| F_{\text{yf}}^{(k)} - F_{\text{yf}}^{(k-1)} \right| \le F_{\text{yf,max slew}}, \tag{3.2}$$

where F_{yf} is the lateral front tire force and $F_{\text{yf,max slew}}$ is the maximum slew rate of the steering system.

When obeying the constraint of maximum slew rate, additional design goals regarding the steering arise. One key goal is the smoothness of the steering, which is affected by the change in input from time step to time step. Steering smoothness is a reasonable consideration to include in the control algorithm because most occupants expect a level of comfort as they ride in a vehicle. Occupant comfort is a desirable feature, but, similarly to path tracking, if encoded as a hard and fast rule of matching a specific rate via an equality constraint or remaining under a rate via an inequality constraint could result in safety compromises if an emergency maneuver is required. Specifically, if the vehicle needs to suddenly swerve to avoid an obstacle, it will be constrained by a requirement to maintain smooth steering and may not be able to steer and maneuver sufficiently to avoid a collision. However, when viewed from a consequentialist standpoint, a cost associated with steering smoothness that the algorithm will minimize can be included in the cost function. I choose to account for smoothness in the cost function as long as it is subservient to more highly prioritized rules associated with safety. Therefore, steering smoothness for occupant comfort is cast as a cost through consequential reasoning. Occupant comfort level is incorporated into the objective function by associating a cost to the change in steering, or effectively, the lateral front tire force:

$$J_{F_{\text{yf}}} = R \sum_k \left\| F_{\text{yf}}^{(k)} - F_{\text{yf}}^{(k-1)} \right\|_2^2, \tag{3.3}$$

where R is the associated cost. By minimizing the differential in lateral front tire force throughout the prediction horizon, the differential in front steer angle is also minimized. Thus, this term in the cost function results in smooth steering enforcement, which can factor into occupant comfort.

3.4.3 Obstacle Avoidance

Obstacle avoidance is a high priority in navigating roadways. As mentioned within the context of Sections 3.4.1 and 3.4.2, the possibility of collisions and preserving the ability to avoid them are the basis for choosing to associate path tracking and steering smoothness with consequentialist costs rather than deontological rules. Collision avoidance is arguably the highest priority of the automated vehicle, so I choose to consider it through the lens of a deontology and implement collision avoidance as a constraint.

The deontological rules that govern obstacle avoidance arise from dividing the environment into tubes in which the vehicle can safely travel and describing the envelope in which each tube lies. In the example scenario depicted in Figure 3.1, the vehicle can choose from one of three tubes—pass the vehicle by entering the lane on the left, pass the vehicle by moving onto the shoulder to the right, or stay in the lane. The boundaries of each tube in the environment can be constructed in reference to a nominal path which, in the scenario considered here, is simply the centerline of the lane in which the vehicle travels. These envelopes are described by a set of time-varying constraints on the maximum and minimum lateral offset from the nominal path (e) necessary to remain in the tube. The vehicle's trajectory over the prediction horizon is constrained to remain within this envelope to ensure the trajectory is collision-free. The environmental envelope may require the vehicle to deviate from the nominal path due to the fact that the nominal path need not be obstacle-free.

Figure 3.3 illustrates the methodology for the generation of the environmental envelope, beginning with the scenario illustrated in Figure 3.3a as an example. The environment is sampled at discrete points along the nominal path based on the maintained longitudinal vehicle speed (U_{x}), and the assumption that the distance along the path is a function of only U_{x}. Figure 3.3b shows the future position of the vehicle k steps into the prediction horizon. In order to align with the discrete sampling, the objects are extended to match the sampling, shown in Figure 3.3c. This extension allows for the identification of feasible gaps (defined as distances greater than vehicle width) between the objects. A graph search algorithm links adjacent feasible gaps, forming tubes like those shown in Figure 3.3d. One tube must contain the full prediction horizon for the vehicle to avoid collisions. This tube methodology is similar

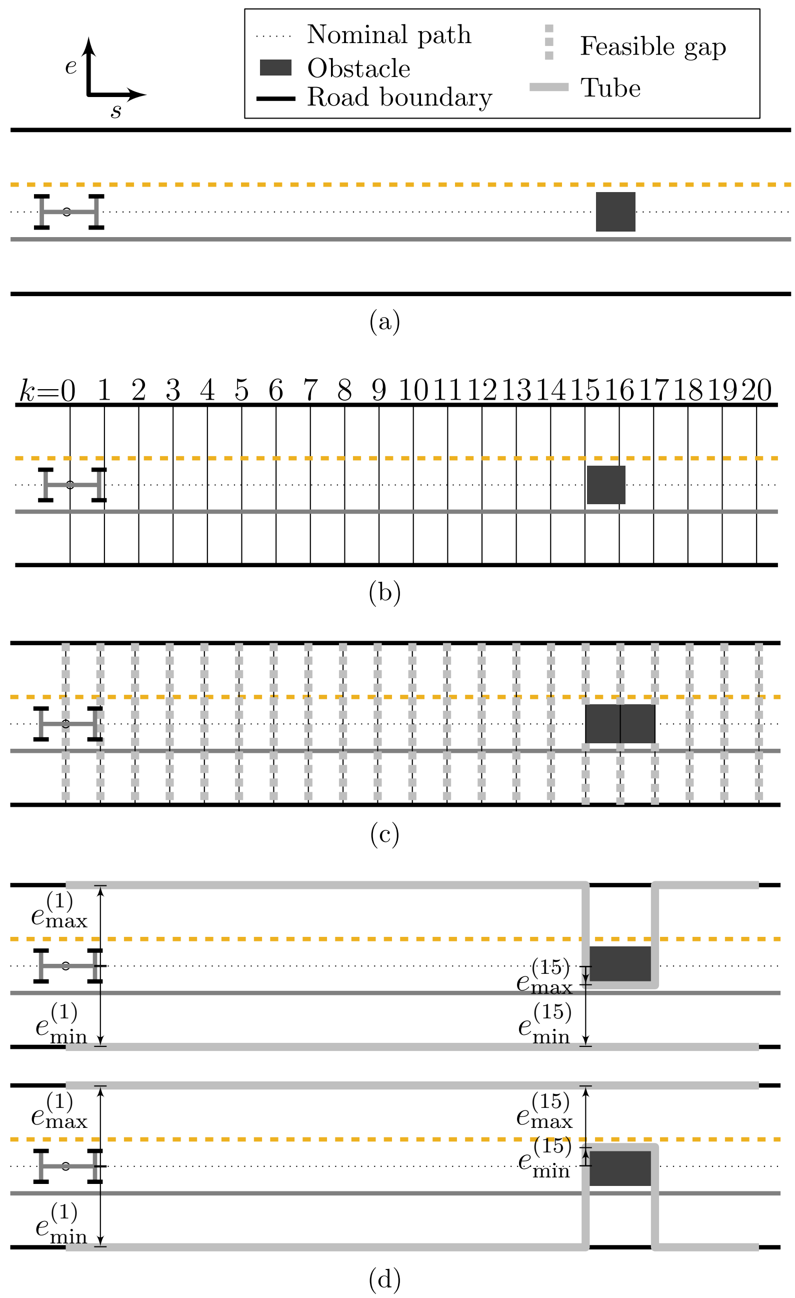

Figure 3.3: Environmental envelope generation. The process consists of (a) starting with a set of obstacles along the nominal path, (b) discretization along the s direction, (c) extension of objects along that same s direction, which creates alignment with the discretization and from which feasible gaps between objects are identified, and (d) connecting adjacent gaps into tubes which define maximum ($e_{\text{max}}^{(k)}$) and minimum ($e_{\text{min}}^{(k)}$) lateral deviation from the nominal path at each time step (k). Here, two tubes are given as examples.

to LaValle's vertical cell decomposition in [68], tube feasibility for robotic arm motion planning, as in Suh and Bishop [69], and driving corridors presented by Ziegler, Bender, Dang, and Stiller [5].

The set of collision-free trajectories corresponding to a single tube is a convex set, due to the property that any linear combination of generated trajectories that is contained within a tube will also be contained within that tube. This property opens up the ability to use fast optimization techniques for optimal trajectories to be identified quickly.

The following linear inequality represents the bound on lateral deviation (e) at each time step (k) represented by each tube:

$$H_{\text{env}} x^{(k)} \leq G_{\text{env}}^{(k)}, \tag{3.4}$$

where

$$\begin{aligned} H_{\text{env}} &= \begin{bmatrix} H_{\text{env,left}} \\ H_{\text{env,right}} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & -1 \end{bmatrix} \\ G_{\text{env}} &= \begin{bmatrix} G_{\text{env,left}} \\ G_{\text{env,right}} \end{bmatrix} = \begin{bmatrix} e_{\text{max}}^{(k)} - \frac{1}{2}d - d_{\text{buffer}} \\ -e_{\text{min}}^{(k)} - \frac{1}{2}d - d_{\text{buffer}} \end{bmatrix}. \end{aligned}$$

The environmental envelope is denoted with the subscript env, the lateral deviation bounds for time step k are given as $e_{\text{max}}^{(k)}$ and $e_{\text{min}}^{(k)}$, and the vehicle width is d. Occupant comfort can be further specified via d_{buffer}, which represents a preferred minimum distance between obstacles and the vehicle and can account for vehicle orientation changes in determining minimum gaps from obstacles as well.

3.4.4 Traffic Laws

Traffic laws present the most ambiguous choice between rule- and cost-based design. The scenario posed in Section 3.2 serves as a clear example of this conflict. Traffic laws by definition enforce structure and rules, and thus naturally can be framed as deontological. However, humans do not always treat traffic laws as deontological.

Realistically, drivers in Section 3.2's scenario often judge factors such as clearance from the obstacle, traffic in the opposing lane, and speed for overtaking. The driver then makes a decision on whether or not to cross the double yellow line. Since humans will often opt to cross the line, particularly when overtaking a bicyclist, human compliance with traffic laws seems less deontological and more like a consequentialist weighting of safety, mobility, and legality. Thus, in moving from human driver actions to programming automated driving, the decision to treat traffic laws as deontological or consequentialist is a fundamental design choice.

If defined as a rule, obedience to traffic laws can easily result in congestion, as described in Section 3.2. If defined as a cost, the implication is that laws are *a priori* programmed to be broken. Thus, given the dilemma, a "soft" constraint is employed by incorporating a slack variable to encode traffic laws in the MPC formulation. The slack variable creates a scalable cost on constraint violation. When that cost is comparably lower than other objectives, the constraint is treated in a consequentialist manner because the slack variable augments the constraint to make the constraint less strict. In contrast, a very high weight on the slack variable causes the constraint to dominate all other objectives in a deontological manner because of the large penalty associated with making the slack variable value non-zero. Continuing with the two-lane roadway example and translating the ethically motivated design decisions into the algorithm, a cost on the slack variable is included corresponding to either crossing the road divider (S_{left}) or entering the road shoulder (S_{right}). Treating traffic law obedience as deontological or consequentialist results in dramatically different vehicle behavior, which is demonstrated in Section 3.6. These different driving outcomes underscore the importance of matching societal expectations to programming decisions, and these philosophical frameworks assist with reasoning and justifying what driving behavior to deploy.

3.5 Model Predictive Control Formulation

The control algorithm attempts to determine the optimal path for the vehicle in light of the costs and constraints placed on its motion. The vehicle divides the

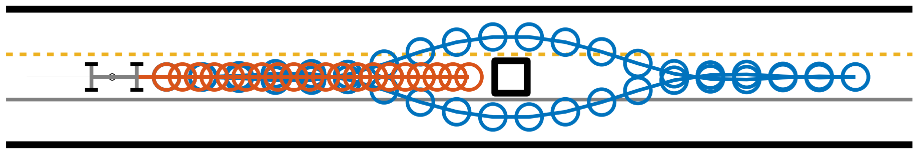

Figure 3.4: The three tubes define the generic maneuver options to avoid an obstacle. The left and right tubes are depicted in blue while stopping is depicted in red.

world into a number of feasible tubes, with each tube representing a separate convex optimization problem. The vehicle then calculates the optimal path in each tube as illustrated in Figure 3.4, choosing the option with the lowest cost. The following sections describe the mathematics behind solving for the optimal trajectory in each tube and its associated cost.

3.5.1 Vehicle Model

The vehicle model used in the online MPC controller is a bicycle model with four states. There are two velocity states, vehicle sideslip (β) and yaw rate (r), and two position states, heading deviation ($\Delta\psi$) and lateral deviation (e), the details of which can be found in Appendix A. Thus, the vehicle state vector is

$$x = [\beta \quad r \quad \Delta\psi \quad e]^\top. \tag{3.5}$$

In this chapter, the actuator is considered to be front steering, and the vehicle is assumed to be equipped with steer-by-wire technology. This enables the computer algorithm to command the desired lateral front tire force (F_{yf}). For simplicity, the results here consider a constant longitudinal speed maintained by a PD cruise controller unless the vehicle needs to brake, though this is not strictly necessary.

3.5.2 Optimization

For each tube in the environment, the optimal path and control inputs are the solution to the following optimization problem:

$$\begin{aligned}
\text{minimize} \quad & \sum_k x^{(k)\top} Q^{(k)} x^{(k)} && \text{(3.6a)}\\
& + R \sum_k \left|\left| F_{\text{yf,opt}}^{(k)} - F_{\text{yf,opt}}^{(k-1)} \right|\right|_2^2 && \text{(3.6b)}\\
& + \sum_l \left[\sigma_{\text{env}} \quad \sigma_{\text{env}}\right] S_{\text{env,opt}}^{(l)} && \text{(3.6c)}\\
& + \sum_l \left[\sigma_{\text{tra}} \quad \sigma_{\text{tra}}\right] S_{\text{tra,opt}}^{(l)} && \text{(3.6d)}\\
\text{subject to} \quad & x^{(k+1)} = A_d^{(k)} x^{(k)} + B_d^{(k)} F_{\text{yf,opt}}^{(k)} + d_d^{(k)} && \text{(3.6e)}\\
& \left| F_{\text{yf,opt}}{}^{(k)} \right| \leq F_{\text{yf,max}} && \text{(3.6f)}\\
& \qquad k = 0 \ldots (T-1) &&\\
& H_{\text{env}} x^{(l)} \leq G_{\text{env}}^{(l)} + S_{\text{env,opt}}^{(l)} + S_{\text{tra,opt}}^{(l)} && \text{(3.6g)}\\
& \qquad l = (T_{\text{split}} + 1) \ldots T &&\\
& \left| F_{\text{yf,opt}}{}^{(i)} - F_{\text{yf,opt}}{}^{(i-1)} \right| \leq F_{\text{yf,max slew}} && \text{(3.6h)}\\
& \qquad i = 0 \ldots T &&
\end{aligned}$$

where $T_{\text{split}} + 1$ is the instance when the time steps are longer for the environmental envelope as detailed by Erlien, Fujita, and Gerdes [47]. The variables to be optimized are the lateral front tire forces ($F_{\text{yf,opt}}$), the vehicle states (x), the slack variable on the environmental constraint ($S_{\text{env,opt}}$), and the slack variable on the traffic laws ($S_{\text{tra,opt}}$). The tunable parameters in this optimization problem are the costs on the vehicle states (Q), the cost on the change between inputs (R), and the costs on the slack variables (σ_{env} and σ_{tra}).

The slack variables are not only used to represent traffic laws but also to set a hierarchy of deontological constraints and ensure that the problem always returns a feasible solution. Hence, the slack variables are unbounded. Higher weights on a slack variable give it a higher precedence in a deontological framework. Here the slack variables are used to weight the constraint version of traffic laws below that of obstacle avoidance. In a more comprehensive version of a vehicle control system, they could also be used to incorporate vehicle stability constraints such as those in Beal and Gerdes [70] and Bobier and Gerdes [71]. The slack variable weights for such constraints should be placed below that of collision avoidance in a deontological sense as demonstrated in Funke, Brown, Erlien, and Gerdes [40]. Additional cost terms could also be incorporated in the cost function such as those presented by Wei, Dolan, and Litkouhi [15].

Optimization problem (3.6) is a quadratic program with a significantly sparse structure that can be leveraged with an efficient solver for real-time implementation [45]. For this work, CVXGEN, developed by Mattingley and Boyd [46], is used to solve for the vector of optimal lateral front tire forces ($F_{\text{yf,opt}}$). The control input to the autonomous vehicle at the next time step is the first solution ($F^{(0)}_{\text{yf,opt}}$) which, in turn, is translated to a desired steering angle as described in Appendix A. An alternative problem formulation, where the road lane dividers and shoulder are implemented as additional constraints with a slack variable can be seen in Appendix B.

3.6 Experimental Results

The example from Section 3.2 sets the stage for demonstrating ethically motivated algorithm design in real vehicle experiments. Depending upon the driving situation and the engineering design choices made in the algorithms, the vehicle chooses a different tube or a different trajectory within that tube. In these experiments, the weights are changed to reflect different interpretations of the traffic laws. The philosophical treatment of different objectives can therefore be translated through mathematics to the actual motion of the vehicle. The weights were first chosen in simulation and validated in-vehicle for the experiments. The actual numerical values for the weights are not crucial; instead, the relative values affect the solution to the optimization problem. In practice, they can be derived as a function of geographic indicators or determined by perception algorithms.

Figure 3.5: X1, an all electric, steer- and drive-by-wire research testbed.

Table 3.1: X1 vehicle parameters

Parameter	Symbol	Value	Unit
Vehicle mass	m	2009	kg
Yaw moment of inertia	I_{zz}	3000	$\mathrm{kg\,m^2}$
Distance from front axle to CG	a	1.53	m
Distance from rear axle to CG	b	1.23	m
Track width	d	1.63	m
Front cornering stiffness	$C_{\alpha\mathrm{f}}$	140	kN/rad
Rear cornering stiffness	$C_{\alpha\mathrm{r}}$	170	kN/rad

The test vehicle used for these experiments is X1, an all-electric, drive- and steer-by-wire vehicle shown in Figure 3.5. The parameters for this vehicle are specified in Table 3.1. The vehicle is equipped with an integrated GPS/INS system that provides real-time estimates of the vehicle states. In these experiments, obstacles and road boundary locations are assumed to be known. All of the following experiments took place on a cement surface which the controller models with constant friction.

3.6.1 Traffic Laws as Consequentialist Costs

In Section 3.4, the design choice to implement traffic laws as slack variables on the environmental constraints in the MPC formulation is discussed. Changing the weights on the slack variables can influence the vehicle's behavior as it navigates the scenario described in Section 3.2. A consequentialist approach allows for flexibility in weighting the road bounds to reflect the severity of costs associated with crossing them. Table 3.2 shows a set of weights where the shoulder is treated effectively as a hard constraint by placing a high cost on the slack variable (representing, for instance, a curb on the side of the road) while the double yellow lane divider is treated more as a cost with lower weight. As seen in Figure 3.6, the vehicle maneuvers to the left of the obstacle and crosses the divider. Because of the relatively high weight put on the divider relative to the path tracking weight, the vehicle does not cross far into the opposing lane but rather stays fairly close to the obstacle.

Upon interchanging the costs on road shoulder and road divider (representing a stricter adherence to the law but an open space adjacent to the lane), Figure 3.7 shows that the vehicle maneuvers to the right of the obstacle. Table 3.3 shows the weights chosen for this particular situation. Since the weights describe essentially a mirror image of the previous experiment, both the trajectory and the resulting steering angle are reflections of the original case.

3.6.2 Traffic Laws as Deontological Constraints

Another philosophical approach to accounting for traffic laws is to describe them deontologically as rigid rules. As the slack variable weights increase relative to the other weights, the road boundary constraints begin to resemble hard or deontological constraints. With this choice of weights, shown in Table 3.4, the tubes to the left and right of the vehicle have unacceptably high costs and therefore the vehicle must remain in the tube corresponding to the lane. Since that tube is blocked by the obstacle, the vehicle must brake to a complete stop. Figure 3.8 shows the trajectory as an independent longitudinal PD controller commands a brake force to bring the vehicle to a stop before the end of the tube.

The MPC formulation is flexible enough to incorporate different treatments of the environmental boundaries and the corresponding traffic laws. As this simple example demonstrates, however, treating the double yellow lane boundary as a hard, deontological constraint removes much of the flexibility from the vehicle path and

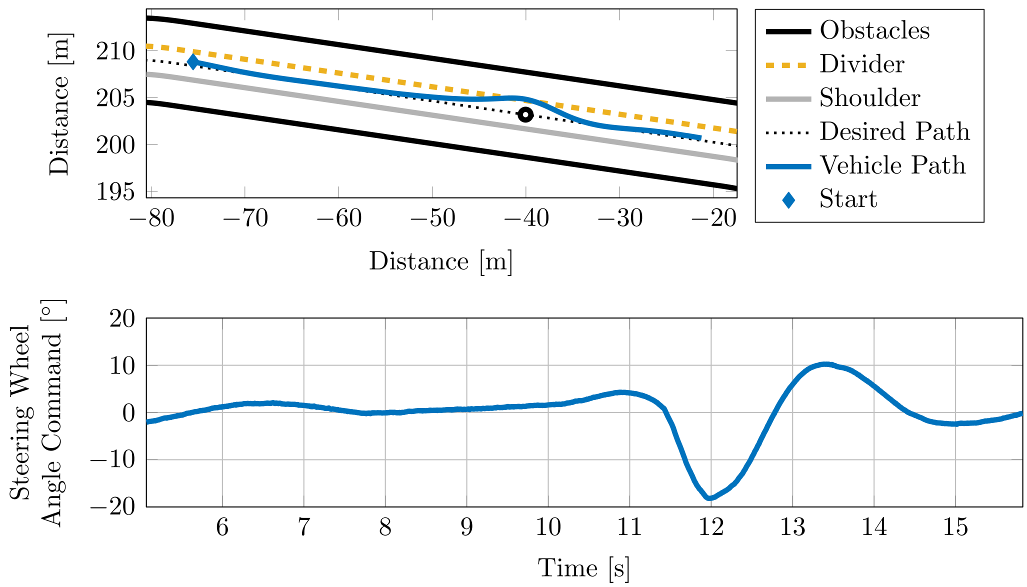

Figure 3.6: The left tube is chosen because the traffic lane divider is considered safe to cross.

Table 3.2: Weights resulting in a pass on the left

Parameter	Symbol	Value	Unit
Lateral error	Q_e	0.7	m^{-1}
Heading error	$Q_{\Delta\psi}$	0.5	rad^{-1}
Smoothness	R	0.1	kN^{-1}
Environmental slack	σ_{env}	500	m^{-1}
Road divider slack	σ_{left}	10	m^{-1}
Road shoulder slack	σ_{right}	150	m^{-1}

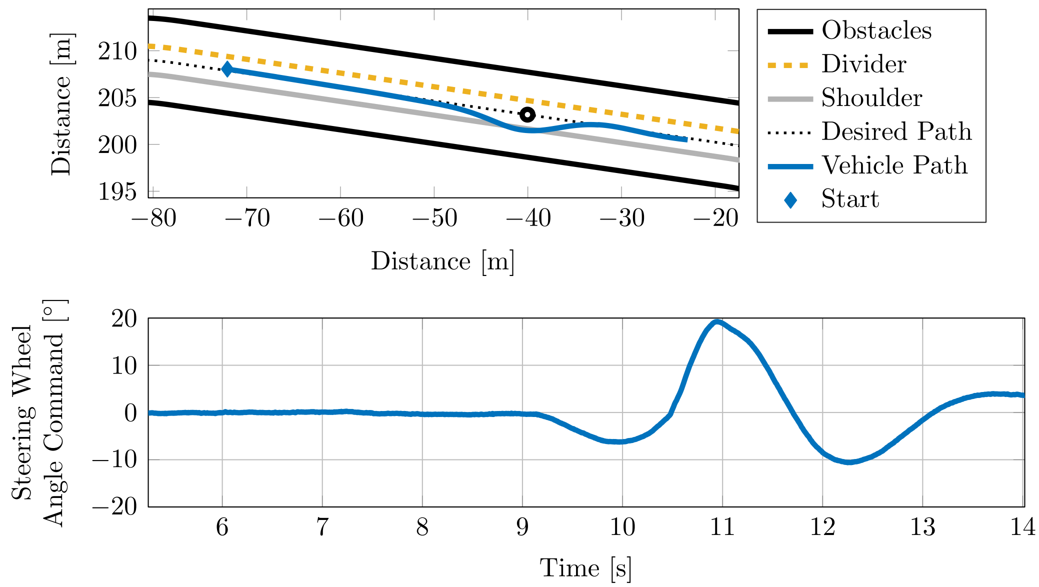

Figure 3.7: The right tube is chosen because evaluation of the scenario determined it is safer to pass around the obstacle via the road shoulder.

Table 3.3: Weights resulting in a pass on the right

Parameter	Symbol	Value	Unit
Lateral error	Q_e	0.7	m^{-1}
Heading error	$Q_{\Delta\psi}$	0.5	rad^{-1}
Smoothness	R	0.1	kN^{-1}
Environmental slack	σ_{env}	500	m^{-1}
Road divider slack	σ_{left}	150	m^{-1}
Road shoulder slack	σ_{right}	10	m^{-1}

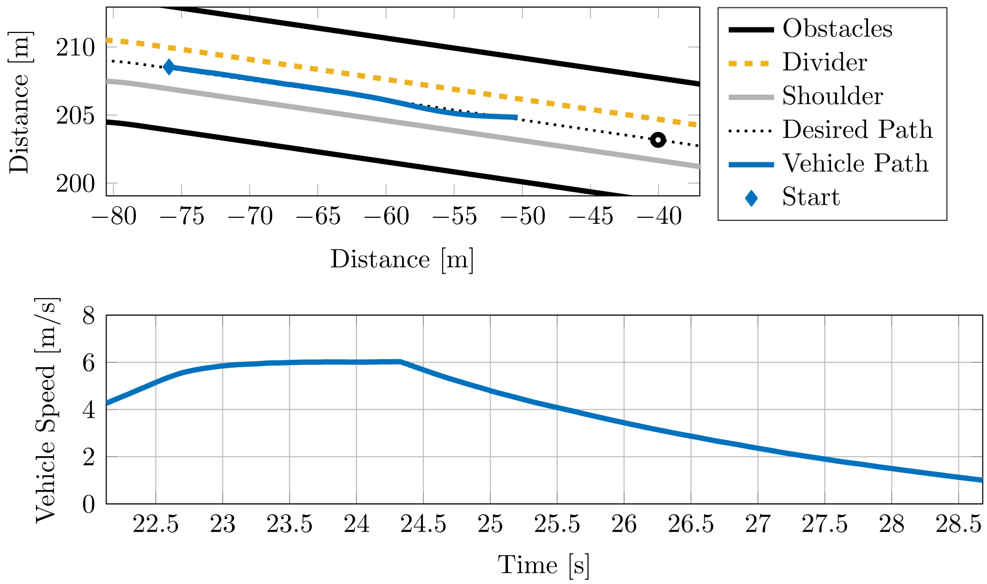

Figure 3.8: Since left and right path options are weighted equivalently, the tube is considered to be blocked and the vehicle brakes to a complete stop.

Table 3.4: Weights resulting in a full stop

Parameter	Symbol	Value	Unit
Lateral error	Q_e	10	m^{-1}
Heading error	$Q_{\Delta\psi}$	1	rad^{-1}
Smoothness	R	0.1	kN^{-1}
Environmental slack	σ_{env}	500	m^{-1}
Road divider slack	σ_{left}	150	m^{-1}
Road shoulder slack	σ_{right}	150	m^{-1}

fails to reflect the way humans drive. This suggests that traffic laws may have to be modified to give programmers the same flexibility demonstrated by their human driver counterparts. Without that option, programmers need to choose how much to weight the traffic laws in a consequentialist approach. In this MPC formulation, the relative values of the weights chosen for the traffic slack variables determine how the vehicle maneuvers around obstacles. These values can be a function of geographic indicators, such as when a lane divider is a single, dashed yellow line or when it is a double yellow line. Additionally, the weights can be determined by a perception algorithm using information from lidars and cameras that determine the safety of the terrain in each tube. These results demonstrate the flexibility of the MPC formulation to account for responsible decision-making.

3.7 Discussion of Costs, Constraints, and Weights

The above examples illustrate the autonomous vehicle can exhibit very different behavior when the traffic laws are treated either deontologically or consequentially. The benefit of mapping costs and constraints to deontology and consequentialism is two-fold. Firstly, as demonstrated in this chapter, it can be used to justify a design goal to be implemented as a cost or constraint by using either deontological or consequential arguments. Secondly, if an engineer implements an algorithm with costs or constraints, then this mapping can provide insight into what type of behavior the implementation may implicate.

To further understand the implications of this mapping, moral psychologist Greene suggests deontological rationalization arises during sensitive situations while consequential reasoning comes forth in less severe scenarios [72]. This provides an explanation as to why adhering to the double yellow line might be best represented as a cost, unless the severity of the situation increases (i.e. oncoming traffic or erratic behavior with the obstruction). When the scenario is of low risk, then treating certain traffic laws as rules may be cumbersome, if not, inherently restrictive.

Beyond the choice of costs and constraints, the choice of weights in the optimization problem ultimately influenced the vehicle behavior with regards to the traffic laws (σ_{left} and σ_{right}). The other weights in the cost function also influence vehicle behavior, such as path deviation (Q_e and $Q_{\Delta\psi}$), obstacle avoidance (σ_{env}), and even occupant comfort (R). Some of these weights are purely in the cost function (i.e. not a slack variable) so they can only influence the vehicle behavior consequentially. Yet, the principles of deontology and consequentialism do not provide insight into the choice of weights which could affect the other aspects of the vehicle behavior.

3.8 Vehicle Character

Thus far, the variations in the MPC formulation have been motivated by traffic situations for one particular vehicle. The experimental results indicate the principles behind the ethical theories of consequentialism and deontology and how the theories can implicate varying vehicle behavior. However, these theories do not explicitly guide the choice of relative numerical values for the weights, which have a huge impact on design goals beyond safety for automated vehicles. Thus, this section is motivated by introducing a third normative ethical theory known as *virtue ethics*. Virtue ethics points the focus of ethical behavior toward character rather than correct actions or outcomes as emphasized by deontology and consequentialism. Under the framework of virtue ethics, a decision is ethical as long as it adheres to the disposition of the moral being. In other words, moral beings operate virtuously provided they always perform the correct action in the correct situation that is in alignment with their character [64].

The idea of the character of an agent naturally leads to a more specific concept which will be utilized in this work called *role morality*. Role morality is the idea that behavior within the context of a specific professional role and situation is acceptable but may not be acceptable outside that setting [73]. Role morality is cited in fields such as law and medicine to justify behavior that might be deemed unacceptable were the behavior to take place outside a professional situation specific to that field. An extreme example is Sanson, the executioner of Paris, as noted by Applbaum [74]. A less extreme example is that of a physician prescribing medication to someone who is not officially his or her patient. While writing prescriptions is acceptable within the professional bounds of a doctor-patient relationship, it is illegal outside of this role. These acceptable roles and codes of conduct are based upon the societal expectation of the service provided by professionals in that particular field; therefore, role morality is derived not from any individual in charge of making the rules but from a collective decision on what is best for society [73].

An important issue in the development of automated vehicles is therefore the type of role or character different vehicles should have. Specifically, the role of a vehicle affects the level of strictness by which a vehicle should adhere to a traffic law. Previous experiments in Figures 3.6, 3.7, and 3.8 showcase the ability of an autonomous vehicle to choose whether or not to violate the traffic law of obeying a double yellow line boundary due to safety considerations. The level of fidelity of either adherence or violation begs for a guiding principle, which can be motivated through role morality. Different types of vehicles have different roles to play in society. An ambulance running a red light while carrying a passenger with a life-threatening condition to the hospital can acceptably break traffic laws: the role of an ambulance in society is to transport people to the hospital as quickly as possible in order to save lives. In contrast, a taxi carrying a passenger in a hurry may not run a red light to save time because its societal role does not merit that behavior.

Thus, while deontology and consequentialism enable justification of vehicle goals to be described as either constraints or costs, role morality can help determine the strength of the applied rules and costs for different vehicles. For example, role morality establishes the context for why an ambulance would be programmed to consider breaking traffic laws more freely than a taxi.

Using the MPC formulation presented in Section 3.5, the character of a vehicle that can acceptably break laws, such as an ambulance, can be modeled by varying the weights on different objectives. As an example, the weights on tracking error can be relatively reduced to imitate the desired behavior of an emergency response vehicle. The relatively lower costs on lateral and heading error allow the vehicle more freedom to deviate from the path. The experimental results using the weights in Tables 3.5 and 3.6 are shown in Figures 3.9 and 3.10, respectively. These relative weights were determined using the process described in Section 3.6. Because the

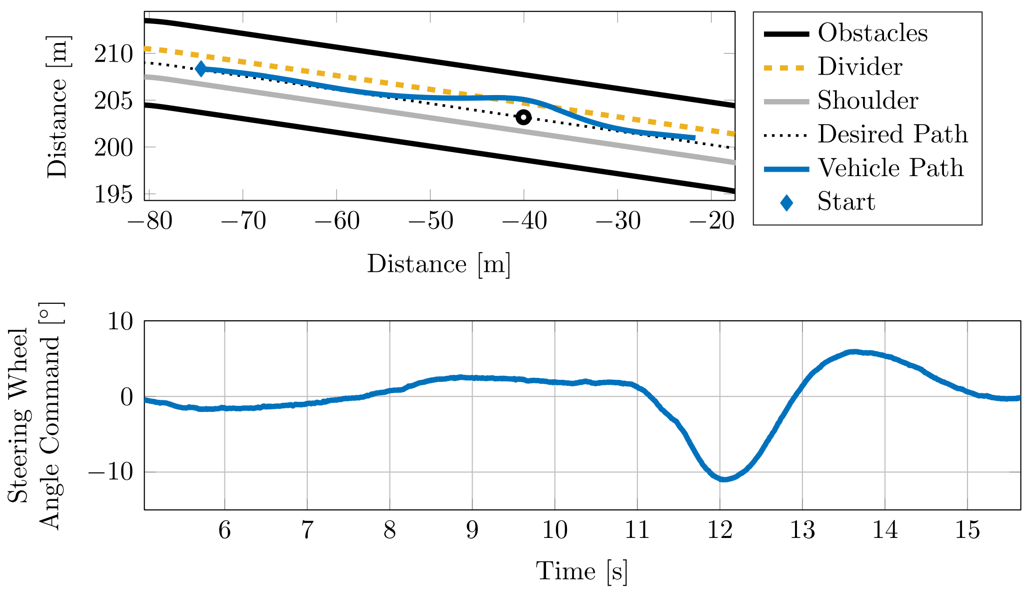

Figure 3.9: Reduction of the cost on path following allows the vehicle behavior to model emergency response vehicle character.

Table 3.5: Weights causing an ambulance vehicle character to pass on the left

Parameter	Symbol	Value	Unit
Lateral error	Q_e	0.3	m^{-1}
Heading error	$Q_{\Delta\psi}$	0.25	rad^{-1}
Smoothness	R	0.1	kN^{-1}
Environmental slack	σ_{env}	500	m^{-1}
Road divider slack	σ_{left}	5	m^{-1}
Road shoulder slack	σ_{right}	200	m^{-1}

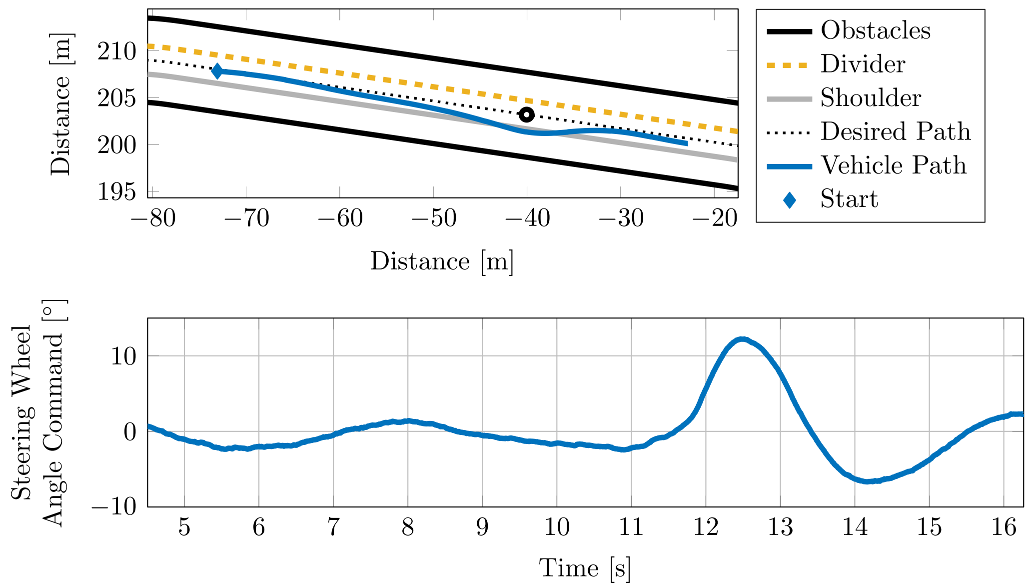

Figure 3.10: Reduction of the cost on path following allows the vehicle behavior to model emergency response vehicle character.

Table 3.6: Weights causing an ambulance vehicle character to pass on the right

Parameter	Symbol	Value	Unit
Lateral error	Q_e	0.3	m^{-1}
Heading error	$Q_{\Delta\psi}$	0.25	rad^{-1}
Smoothness	R	0.1	kN^{-1}
Environmental slack	σ_{env}	500	m^{-1}
Road divider slack	σ_{left}	200	m^{-1}
Road shoulder slack	σ_{right}	5	m^{-1}

lower relative costs on lateral and heading error with respect to smoothness allow for greater path deviation, the emergency vehicle begins executing the maneuver earlier, as shown in Figures 3.9 and 3.10. In this particular scenario, the smoothness of the maneuver is potentially advantageous for an injured passenger.

While a vehicle can indeed be programmed to weight traffic laws and maneuver objectives according to its societal role, developing an engineered system which may not always observe the law is an uncomfortable proposition. This raises another question of role morality: should the engineers designing such a system be responsible for setting weights on the laws? Or, is a better solution to adapt the law to the programming realities of automated systems? While not producing a simple answer, the translation of philosophical frameworks to engineering terms does help to raise the appropriate questions that must be resolved to field automated vehicles.

3.9 Limitations

An understanding of these philosophical frameworks provides engineers a way to better judge the ethical implications of a motion planning algorithm. However, focusing on philosophical frameworks may give the impression that the goal here is to program an ethics module, when the actual goal is to program ethically, as envisioned by philosopher Vallor [75]. Even though philosophers are the most trained in ethical analysis and thus more equipped to address deeper ethical issues that may be present in the technology or algorithm [76], it can be helpful to the engineer to connect their design decisions to normative ethical theories as a way of realizing the impact of those choices. An engineer may also recognize the importance of getting further input on those choices from experts in ethics. Thus, the work here is a first step at ethics capacity-building for engineers to better understand how deontological and consequential reasoning can bring about different behavior.

Another limitation of this approach is the inability to guarantee the realization of human values in the algorithm design. The focus has been on using philosophical frameworks to reason through the construction of a motion planning algorithm. Because this does not explicitly address human values, a successful integration of human values can be difficult to determine. For example, when the traffic laws were considered through consequential reasoning, mobility and safety were better achieved over legality compared to when the traffic laws were considered deontologically. Even though the choice of a consequentialist treatment of traffic law led to behavior likely emulating what most human drivers do, this justification alone does not explicitly address the value conflict. This will be explored further in the next chapter.

3.10 Summary

The normative ethical theories of deontology and consequentialism provide insight for engineers to understand the implications of various design decisions when programming autonomous vehicles. In particular, these concepts map to rule- and cost-based engineering paradigms such as constraints and cost functions. Making these connections can enable engineers working at the deepest levels of programming automated vehicles to connect their design choices with broader issues of societal acceptance. The mapping of philosophical frameworks to mathematical frameworks can be used as a tool by engineers as a way of thinking and reasoning through autonomous vehicle motion planning designs. This chapter has examined how to incorporate objectives such as path tracking, vehicle occupant comfort, and traffic laws as priorities in the cost function while obstacle avoidance and vehicle slew rate limits enter as constraints through an MPC formulation. The concept of role morality provides a further basis for different weighting schemes within the control formulation depending on the vehicle type and function.

More complex scenarios may require the consideration of additional objectives for the vehicle. Nevertheless, the simple obstacle avoidance maneuver in this chapter already points out key challenges for balancing a respect for the law with the desire to integrate smoothly with human drivers in traffic. Realizing the benefits of automated vehicles will require further integration of legal and ethical considerations with the underlying control code. In particular, a formal realization of human values could help society better understand and trust autonomous vehicle technology.

Chapter 4

Motion Planning with Human Values

Recognizing the parallelism between normative philosophical frameworks and mathematical frameworks helps with framing and discussing the development of particular algorithms for autonomous vehicle motion planning. The parallelism also helps with understanding the human value implications for various objective trade-offs. But how do engineers know what values to account for in the algorithms? Human values go beyond just mobility, safety, and legality. In this chapter, ethical considerations that map values to engineered technology are achieved through the formal, iterative process of value sensitive design (VSD). The first iteration of VSD, which is a large part of this chapter, appears in the Proceedings of the IEEE Intelligent Vehicles Symposium of 2018 [77].

4.1 Introduction

The roadways are populated by many stakeholders, such as pedestrians, bicyclists and vehicle occupants, and these stakeholders have values that predicate their expectations. Public expectations regarding autonomous vehicle driving behavior will likely be based on similar human values. Mobility, safety, and legality are some reasonable values to consider [78], [79]. Designers of autonomous vehicle technologies have the challenge of connecting these human values to engineering specifications. One way to address this challenge is to integrate the stakeholders and their values into the design process of algorithms for autonomous vehicle motion planning.

There are many design processes that account for human values and needs as well as various stakeholders. One popular design process is known as human-centered design (HCD) [80], [81]. The methodologies of HCD involve engaging with a group of stakeholders, largely direct users of the technology, in order for the designers to get feedback on how to make design improvements. Both Millar [76] and Niemelä *et al.* [82] suggest current HCD approaches lack a view point from ethics in terms of justifying the designs or understanding a design's potential ethical implications. Another design process that includes an ethics evaluation is known as life-based design (LBD) [83]. LBD takes a broad and holistic approach to the design task by investigating the needs of stakeholders through their quality of life. It starts by identifying the human requirements in the design activity, then the users and the technology requirements, and then determines if the human's quality of life improves based on the designed technology. The analysis on the quality of life improvement explicitly includes an ethical evaluation. LBD has this inherent justification that the technology must improve quality of life for the users. HCD and LBD are both iterative design processes that attempt to improve the design of technology for the respective users. HCD focuses on values relating to usability and LBD focuses on values relating to quality of life. With many stakeholders impacted by a potential technology, many more values need to be considered, especially since these values may conflict.

Value sensitive design (VSD) [34], [35] is another human design process that addresses ethical considerations, but it does so earlier on in the design cycle by explicitly investigating the values (usually emphasizing those with ethical import [1]) throughout the entire design process. VSD is a tripartite methodology that can be applied to any general design task by iterating over conceptual, technical, and empirical investigations. At every stage of the design process, human values are connected to the designed technology. VSD is most applicable to a design task in which value conflicts exist for ethical issues. Friedman, Kahn Jr., and Borning indicate VSD has been found to be widely useful in the human-computer interaction community to help balance privacy concerns with usability for end-users [35], [84]. In the design of an office space with a virtual window viewing a public plaza, Friedman *et al.* also demonstrate that recognizing indirect stakeholders (a component of the conceptualization phase) uncovered privacy concerns for passersby. Denning, Borning, Friedman, Gill, Kohno, and Maisel use VSD to construct a list of specifications to guide future designs of a security system in implantable medical devices [85]. Furthermore, the generality of VSD allows for modification to be in line with certain design tasks. Wynsberghe appends to VSD the moral theory of care ethics in the design of health care robots to ensure the robots reflect stakeholder values [86].

In designing algorithms for autonomous vehicle motion planning, engineers already account for some human values. Many algorithm designs focus on the values of safety and efficiency as is demonstrated by Chen, Zhao, and Peng's evaluation framework of an autonomous vehicle approaching an unsignalized pedestrian crosswalk [87]. Bandyopadhyay, Jie, Hsu, Ang, Rus, and Frazzoli [88] and Bandyopadhyay, Won, Frazzoli, Hsu, Lee, and Rus [89] also focus on safety and efficiency in the construction of the reward function of a partially observable Markov decision process (POMDP) for speed control in pedestrian environments. As do Brechtel, Gindele, and Dillmann in the construction of the reward function of a speed control POMDP for entering occluded intersections [90]. These examples demonstrate engineers want to connect human values to an engineered technology. The focus on safety and efficiency in the evaluation framework and motion planning policies highlights the difficulty of designing for two conflicting values. Brechtel *et al.* additionally consider occupant comfort in the reward function and suggest traffic rules can likewise be included in future iterations of the POMDP design. Their discussion indicates a desire to account for the various human values at stake in the design of a motion planning policy. To account for these values, it would be useful to have a methodology that can help with determining which values to include because humans value more than just safety and efficiency. Having a list of identified values is also useful to determine conflicts amongst stakeholders and values. Establishing value conflicts early in the design process can help engineers design technology that explicitly resolves them early on rather than requiring patchwork solutions to dissolve value tensions after a system deploys.

Here, it is proposed that VSD can help fill the gaps in the design process of motion planning algorithms for autonomous vehicles. VSD is used to formalize the connection of human values to engineering specifications by identifying a more complete list of human values that are at stake in the design problem and by resolving value conflicts through justification of design choices. This chapter demonstrates a modified application of VSD for autonomous vehicle motion planning with the design task of a speed controller for the scenario of a pedestrian crosswalk. The only constraint on the speed along the path is an upper bound set by the speed limit. In order for the autonomous vehicle to navigate the scenario safely, it likely needs to reduce its speed. The speed is controlled by acceleration commands from POMDP policies designed using VSD. The VSD speed controllers are the first two iterations from the design process and are not final products. The iterative process of VSD helps document how values are incorporated into the speed control design as well as how tensions amongst the values are resolved.

The subsequent sections proceed as follows: Section 4.2 details the three phases of the VSD methodology: conceptualization, technical implementation, and empirical analysis. For the scenario of an occluded pedestrian crosswalk, Section 4.3 presents the first iteration of VSD for the speed control design task. Section 4.4 describes the second VSD iteration. An important part of designing a motion planning algorithm with human values is ensuring the realization of said values, so Section 4.5 elaborates on one technique that can assist with closing the loop on the integration of human values into the technology. A summary is provided in Section 4.6.

4.2 Value Sensitive Design

The methodology of value sensitive design (VSD) consists of three phases: conceptual, technical, and empirical [34], [35]. During the conceptual phase, the methodology involves identifying the values encompassed by the designed technology. Additionally, the conceptualization phase determines the direct and indirect stakeholders of the technology. A feature of VSD holds that some technological implementations are

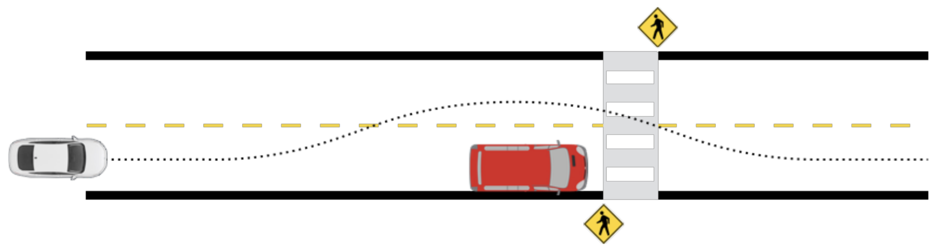

Figure 4.1: Experimental scenario of occluded pedestrian crosswalk.

better suited to uphold certain values than other implementations. For the technical phase, the technical solutions most in line with the identified values (from the conceptual phase) are used to develop the technology being designed. Finally, the empirical phase allows for quantitative and qualitative analyses of the developed design, such as data analysis or observations from human-user studies. This period allows for inspection of how successfully the designed technology meets the conceptualization. Throughout the design development, the designer iterates over the various phases until all three align. Engineers already implicitly iterate over conceptual, technical and empirical phases as they design new technology. VSD provides a tool to help formalize the engineering process to explicitly account for values embedded in the technology by identifying the values and tracking these values throughout the iterations.

4.3 The First Iteration

Designing an autonomous vehicle motion planning algorithm is a broad design task because of the many situations the vehicle may encounter on the roadways. With such potentially broad impact, the list of stakeholders and values may be untenable to design for as a first iteration. To simplify the design task, this chapter will focus on a particular scenario as a form of case-study in order to confine the stakeholder and value consideration space. The scenario in Figure 4.1 depicts a two-lane roadway with a single, dashed yellow line. The roadway also comprises a marked pedestrian crosswalk. In front of the crosswalk is a large, illegally parked van. From the perspective of the autonomous vehicle approaching the crosswalk, the crosswalk is partially occluded because of the obstructing van. The steering controller from Chapter 2 will continue to laterally control the vehicle around the van along an obstacle-free path. The design task is to develop a speed control algorithm along the given path such that the autonomous vehicle safely navigates through the scenario.

4.3.1 Conceptualization

To start the VSD process, the direct and indirect stakeholders involved with the scenario and design task are identified in addition to the human values. Identifying both the direct and indirect stakeholders forces engineers and programmers to think more deeply about the consequences and who is affected by the designed technology. For this scenario, the direct stakeholders are the autonomous vehicle, occupants in the autonomous vehicle, the pedestrian potentially crossing the street, and the authority of traffic laws. An indirect stakeholder is the obstructing vehicle parked on the road because it is assumed that the autonomous vehicle can follow an obstacle-free path around the occlusion. For this first iteration, the focus is on these stakeholders, but there are many more stakeholders, such as bicyclists or bystanders.

Determining the human values involved in the scenario and design task is critical to VSD and the engineering process. Traffic scenarios, in general, relate to balancing the human values of mobility, safety, and legality. By considering the stakeholders, more values at stake can be uncovered. In lieu of engaging with actual stakeholders, the human values to consider in this conceptualization stem from Haidt [91] and Choi and Ji [92]. Haidt suggests there is a set of values (or moral foundations) inherent to human beings, such as care and respect for others, fairness and reciprocity, respect for authority, and individual autonomy, while Choi *et al.* indicates trust and transparency are important for the acceptance of autonomous vehicles. How these moral values are defined could lead to different technology designs, so more thorough definitions are provided by considering the stakeholders to clearly define what is meant by each value for this particular scenario:

- *Care and respect for others* manifests by the desire to not harm other persons.
- *Fairness and reciprocity* affect both the vehicle occupants and pedestrian stakeholders in that the autonomous vehicle should not take biased or discriminatory

actions based on information about the stakeholders. The autonomous vehicle should treat all individuals involved as equal agents.

- *Respect for authority* engages the relationship between the autonomous vehicle and its adherence to traffic laws.

- *Trust* emerges when the pedestrian assumes an oncoming vehicle yields to his or her right-of-way while crossing within the crosswalk. *Transparency* occurs when the autonomous vehicle's actions facilitate this trust.

- *Individual autonomy* of the vehicle occupants acknowledges the desire to get from one destination to another with little impedance.

Since these human values have been identified as necessary to incorporate in the motion planning algorithm, I relate them all to an engineering specification for use in the technical implementation phase. A summary is shown in Table 4.1.

Safety and Legality

When navigating an occluded pedestrian crosswalk, the relationship between legality and safety are closely tied. According to the California Driver's Handbook [93], pedestrians have the right-of-way in a crosswalk. When a pedestrian may be present, a driver is encouraged to reduce the vehicle speed when approaching a crosswalk and be prepared to stop as specified by California Vehicle Code §21950:

> (a) The driver of a vehicle shall yield the right-of-way to a pedestrian crossing the roadway within any marked crosswalk or within any unmarked crosswalk at an intersection, except as otherwise provided in this chapter.
>
> (b) This section does not relieve a pedestrian from the duty of using due care for his or her safety. No pedestrian may suddenly leave a curb or other place of safety and walk or run into the path of a vehicle that is so close as to constitute an immediate hazard. No pedestrian may unnecessarily stop or delay traffic while in a marked or unmarked crosswalk.

(c) The driver of a vehicle approaching a pedestrian within any marked or unmarked crosswalk shall exercise all due care and shall reduce the speed of the vehicle or take any other action relating to the operation of the vehicle as necessary to safeguard the safety of the pedestrian.

(d) Subdivision (b) does not relieve a driver of a vehicle from the duty of exercising due care for the safety of any pedestrian within any marked crosswalk or within any unmarked crosswalk at an intersection.

As the vehicle code alludes, following the law and driving safely strongly correlate. For this iteration, I assume safety and legality to be the same engineering specification for this scenario: if the autonomous vehicle adheres to the legal requirement, then it also takes safe actions. The key pieces of information necessary for safe and legal decision-making are vehicle speed (v_t), distance to crosswalk (d_t), and whether or not a pedestrian is crossing the street (c_t). Again, this connects to the ethical values of care and respect for others, respect for authority, and fairness and reciprocity.

Efficiency and Mobility

The metric of time efficiency captures the human value of mobility. Time efficiency is directly related to the speed of the vehicle (v_t) for the given path. This objective relates to the moral value of individual autonomy.

Smoothness

Smooth driving affects occupant comfort and interjects trust and transparency between the stakeholders. For longitudinal control, smoothness can be captured through the change in vehicle speed, which is equivalent to knowing the acceleration command (a_t) and change in time (Δt).

4.3.2 Technical Implementation

Many motion planning approaches can be tailored to solve the pedestrian occlusion scenario. Rather than just choosing an arbitrary approach, VSD maintains that

Table 4.1: Summary of human values mapping to engineering specifications for the first VSD iteration.

Human Value	Engineering Specification	Information
Safety Legality Care and respect for others Respect for authority Fairness and reciprocity	Safety and Legality	v_t d_t c_t
Mobility Individual autonomy	Efficiency	v_t
Trust Transparency	Smoothness	a_t Δt

the choice of technology or algorithm implicates certain values. Here, the lateral motion continues to be controlled by the formulation in Chapter 2 and the longitudinal motion is chosen to be controlled by a stochastic optimization problem. A stochastic optimization problem can account for modeled uncertainty present in the driving scenario while balancing the identified values through the objective function. Furthermore, the problem can be formulated as an open-loop or a closed-loop planning problem. Closed-loop planning accounts for future state information because it breaks the planning problem into simpler sub-problems (i.e. dynamic programming), while open-loop does not (like stochastic MPC [94]) because it requires constructing a static sequence of actions. To obtain an offline policy to inspect and verify before putting on an autonomous vehicle, a closed-loop planning approach is chosen and the problem is constructed as a partially observable Markov decision process (POMDP) [95]. Throughout the design of the POMDP, every design choice is connected back to values from the conceptualization phase in order to justify the engineering and explicitly record the embedding of said values.

Partially Observation Markov Decision Process

In a POMDP, an agent makes decisions based on its history of observations $o_1, \ldots, o_t$. To reduce the information stored, the history is summarized in a belief state b, which is a distribution over the states. The optimal policy is represented as a set of alpha vectors, which convert the belief state to a control input or action. Given the values of safety and legality, efficiency and mobility, and smoothness, the information necessary to address each value in the objective function is captured by the state vector

$$x = [v_t \quad d_t \quad c_t]^\top \tag{4.1}$$

and the control input

$$u_t = a_t, \tag{4.2}$$

where v_t is the vehicle speed, d_t is the vehicle distance to the crosswalk, c_t is the pedestrian detection and a_t is the longitudinal acceleration. An alternative formulation with the actions as desired speed is presented in Appendix C. The top speed of the roadway is assumed to be $10\,\mathrm{m/s}$, so the vehicle speed is upper bounded by the speed limit to coincide with the safety and legality objective. The pedestrian detection is a Boolean value because the pedestrian is either crossing or not, and, in order to uphold the values of fairness and reciprocity, the detection does not rely on other information about the pedestrian that may be discriminatory. The control input, or action, is limited to $\pm 3\,\mathrm{m/s^2}$ to provide comfortable acceleration and deceleration values, which further addresses the objective of smoothness for occupant comfort.

For the dynamics (or state transitions), a point mass model of the vehicle is used to calculate the distance to the crosswalk and vehicle speed. The detection of a pedestrian crossing maintains some uncertainty over time. When the pedestrian is detected, there is a $90\,\%$ probability the pedestrian will continue to be detected at the next time step. This probability is chosen in order to capture the high likelihood the pedestrian will remain in the crosswalk as he or she crosses the street while acknowledging he or she will not remain in the crosswalk for all time. When the pedestrian is not detected, then there is a $50\,\%$ chance he or she will continue to not be detected, which captures the uncertainty due to the occlusion. This probability is chosen to capture the randomness of a pedestrian potentially appearing. These pedestrian state transition probabilities are rather arbitrarily chosen in the first iteration to demonstrate the process. In practice, the state transition probabilities could come from event-based statistics or another model. The control loop assumes perfect information for the distance to the crosswalk and vehicle speed. However, there is observation uncertainty for the pedestrian crossing with a false positive of 5 % for detecting and a false positive of 5 % for not detecting the pedestrian, which captures the uncertainty due to sensor noise. These false positive rates were chosen arbitrarily small but, in practice, would come from the perception system's capability of detecting pedestrians.

The goal is for the autonomous vehicle to smoothly drive safely and efficiently through the crosswalk while adhering to the relevant traffic laws. The reward function defines the stage cost $g(x_t, u_t)$ for every state and action, which further connects the conceptualization values to the technical implementation. The reward for a state-action pair involves adding stage costs (4.4), (4.5), and (4.6) for that state and action.

The stage cost for safety and legality is partly derived from physical properties so that the computed reward at a given state is then a function of the amount of deceleration needed to stop the vehicle. The constant acceleration point mass equation used to relate the constant deceleration needed to come to a complete stop given the distance to the crosswalk and vehicle speed is

$$a_t = -\frac{v_t^2}{2d_t}. \tag{4.3}$$

This leads to the following stage cost for safety and legality:

$$g_{\text{safe}}(x_t, u_t) = -\left(\zeta \frac{v_t^2}{d_t + \epsilon} + \eta \mathbf{1}(d_t = 0)\right) \mathbf{1}(c_t), \tag{4.4}$$

where $\epsilon > 0$ is a buffer in the denominator to soften the constraint, $\zeta > 0$ is a weight on the penalty incurred by driving quickly as the vehicle gets closer to the crosswalk, $\eta > 0$ is a terminal penalty independent of velocity to encourage the vehicle to stop when the pedestrian is crossing, and $\mathbf{1}(\cdot)$ is a function that evaluates to 1 if the Boolean logic is true and 0 if false.

For efficiency and mobility, the stage cost takes the form of

$$g_{\text{efficient}}(x_t, u_t) = \lambda v_t \mathbf{1}(\neg c_t), \tag{4.5}$$

where $\lambda > 0$ is a reward weight to encourage higher speed when the pedestrian is not crossing.

To achieve smoothness for occupant comfort, the objective is realized through a penalty stage cost on the change in velocity:

$$g_{\text{smooth}}(x_t, u_t) = -\xi(v_t - v_{t+1})^2 = -\xi(a_t \Delta t)^2, \tag{4.6}$$

where ξ is the weight penalizing large changes in velocity. This stage cost only depends on the current input and the time step.

In order to solve the POMDP, the method of QMDP is used to approximate an optimal solution [95]. Although QMDP assumes that at the next time step the state will be fully observable, it is well suited for this problem because the actions are not information gathering, meaning the actions do not directly reduce the uncertainty of the scenario. Another motivation to use QMDP is that it is an offline solver, which means in the next phase (empirical analysis), the policy can be inspected ahead of time before deploying on a vehicle. To solve the POMDP with this method, the state and action spaces are discretized, but continuousness is maintained using multilinear grid interpolations [96] for the state transitions. Vehicle speed increments in steps of 0.5 m/s, vehicle distance to crosswalk increments by 1 m, and accelerations are quantized by 0.1 m/s^2 intervals. These discretizations were chosen such that the sizes of the state and action spaces in the POMDP were kept small: 2,563 total states (includes terminal state) and 61 possible actions.

Throughout the technical implementation, every design choice in the POMDP is connected back to a value from the conceptualization phase as a way to document and justify the engineering of this technology. The next section shows how well this implementation realizes the conceptualization.

4.3.3 Empirical Analysis

The third phase of the VSD methodology entails qualitative and quantitative analyses. In order to discuss how the VSD process impacts the design of a speed control algorithm, it is first compared to a deterministic proportional speed control as a baseline. A qualitative discussion compares the resulting policies of the baseline and POMDP. Quantitative analysis comes from real-time in-vehicle experiments on Trudi, an automated Ford Fusion.

Baseline

The baseline is a deterministic proportional speed control for comparison. Once the pedestrian is detected, a constant deceleration is commanded based on the current vehicle velocity and distance to the crosswalk as in Eq. (4.7a), which also assumes a point mass model with constant acceleration. When no pedestrian is detected, then the vehicle resumes a proportional cruise control with gain k_p and known desired velocity v_{des} as in Eq. (4.7b). The logic is as follows:

$$\texttt{if}\ c_t$$

$$a_t = -\frac{v_t^2}{2d_t} \tag{4.7a}$$

$$\texttt{else}$$

$$a_t = k_p(v_{\text{des}} - v_t) \tag{4.7b}$$

The baseline is intentionally simple as it allows for examination of the design characteristics due to the limited number of design choices considered in this baseline implementation.

Policy Comparison

The baseline controller (4.7) is a closed-loop policy, mapping every state to an action, and is graphically represented in Figure 4.2. The values along the horizontal-axis are the vehicle speed, and the values along the vertical-axis are the distance to the crosswalk. The colors indicate the action. The policy figure is bisected by pedestrian detection. The baseline policy seems to be safe when the pedestrian is detected as evidenced by coasting (zero acceleration) when the vehicle is further from the crosswalk and increasing braking commands as it gets closer. But this part of the policy lacks efficiency when the pedestrian is crossing. In contrast, when the pedestrian is not detected, it is efficient but not safe. This dichotomous behavior highlights the need to anticipate transitioning from one set of logic to the other because of the uncertainty about a pedestrian crossing. The baseline policy lacks an explicit resolution of the value conflict between safety and efficiency for the scenario.

The closed-loop policy of the POMDP designed in the technical implementation phase is represented as a set of alpha vectors, where each alpha vector corresponds to an action. Using the optimal expected utility for each state, the corresponding action is shown in Figure 4.3. When the pedestrian is crossing, the policy indicates a balance between efficiency further away from the crosswalk and safety as the vehicle approaches. There is significant improvement in terms of safety while the pedestrian is not crossing while maintaining some efficiency. The policy also indicates smooth actions across the state space.

The choice of weights in the reward function for this policy are summarized in Table 4.2. These weights were originally hand-tuned to produce a nice blending of positive accelerations away from the crosswalk and negative accelerations closer to the crosswalk. The weights were then more finely tuned when conducting in-vehicle experiments in order to resolve the value tension between the pedestrian and the vehicle as presented in the next section. With these weights, at the extreme state of $v_t = 10\,\mathrm{m/s}$, $d_t = 0\,\mathrm{m}$, and c_t, the term for constant deceleration in g_{safe} gives a penalty of -2.5, while $g_{\text{efficient}}$ is 2.5 when $\neg c_t$. The additional penalty by η in g_{safe} means safety and legality are prioritized in this first implementation. The numerical value for the buffer is the most arbitrary decision; it was chosen such that the numerator does not evaluate to zero and to limit the magnitude of the constant deceleration term. Lastly, ξ is chosen to try to produce smooth acceleration commands.

The goal of this first iteration is to determine how value tensions can be addressed in the implementation. The specific weights that tie directly to the human values are

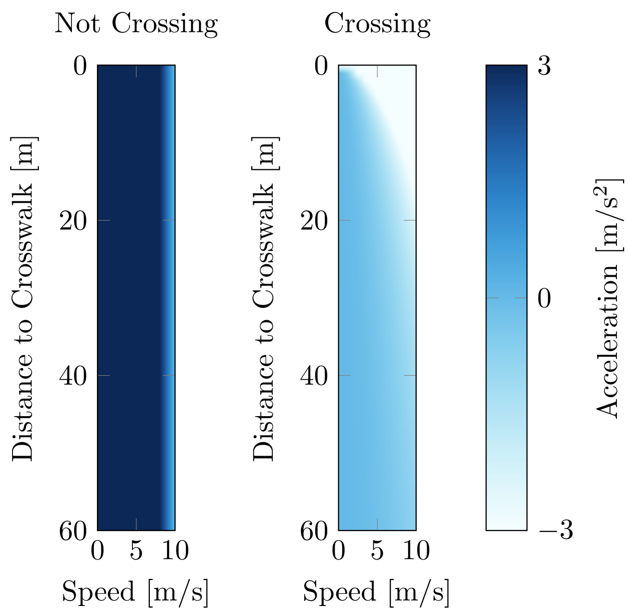

Figure 4.2: Baseline closed-loop policy mapping each state to an action.

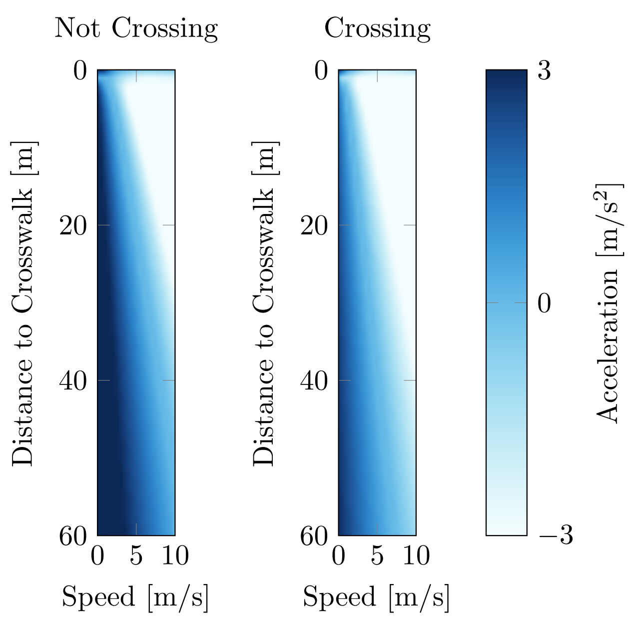

Figure 4.3: Closed-loop policy depicting optimal action at that state assuming perfect state information.

Table 4.2: Weights of the reward function

Variable	Weight	Unit
Safety and legality (ζ)	0.2	s^2/m
Safety and legality (η)	0.2	–
Buffer (ϵ)	8	m
Efficiency (λ)	0.25	s/m
Smoothness (ξ)	1	s^2/m^2

ζ and η for safety and legality, λ for efficiency, and ξ for smoothness. If the overall design seems satisfactory, then exact gains can be further tuned through a Pareto, or multi-objective, optimization over these weights to better determine the value trade-off to implement. This is demonstrated later in Section 4.5. In this case, however, further tuning of the weights is forgone until additional analysis determines this design is satisfactory.

Experimental Results

To further illustrate how well the VSD speed controller realizes the human values, an in-vehicle experiment was conducted using a fully automated hybrid Ford Fusion known as Trudi. Trudi is equipped with an Oxford Technical Solutions RT4000 differential GPS/INS unit, which obtains pose information of 2 cm accuracy with the addition of a Novatel base station to communicate corrections. Vehicle state information is calculated using the pose information as well as velocities and accelerations from the INS. Additionally, Trudi has four Velodyne HDL-32E lidars, which return a 3D point cloud with intensity values. The point cloud data is used to construct a 2D occupancy grid to parse an obstacle detections list. The intensity values are used as a simple classifier for the pedestrian, which is a large person-shaped cardboard cutout (Figure 4.4) covered in retro-reflective material. The vehicle is tasked with following an obstacle-free path around the occluding vehicle using a deterministic model predictive steering control as in Brown, Funke, Erlien, and Gerdes [97] and Chapter 2.

For the policy execution of the POMDP, an observation of the vehicle speed,

Figure 4.4: Experimental setup of occluded pedestrian crosswalk using an inflatable van for the occluding vehicle and a retro-reflective cardboard cutout for the pedestrian that moves along a track.

vehicle distance to crosswalk, and detection of the pedestrian are used to update the belief with a Bayesian filter. The approximate optimal action taken is then

$$\arg\max_a \boldsymbol{\alpha}_a^\top \mathbf{b}, \tag{4.8}$$

where $\boldsymbol{\alpha}_a$ is an alpha vector for each action a and $\mathbf{b}$ is the belief state as a vector. Both the policy solver and policy execution are implemented using the POMDPs.jl library [98].

The experimental scenario involves an occluded pedestrian crosswalk on a two-lane roadway. The vehicle starts from a stopped position at the beginning of the road. As the vehicle approaches the crosswalk, the pedestrian suddenly appears in the crosswalk from behind the occluding vehicle. The control algorithms have no prior knowledge as to when the pedestrian will appear. Figures 4.5 and 4.6 depict the overhead driven trajectory, acceleration commands and speed profile for the baseline and POMDP

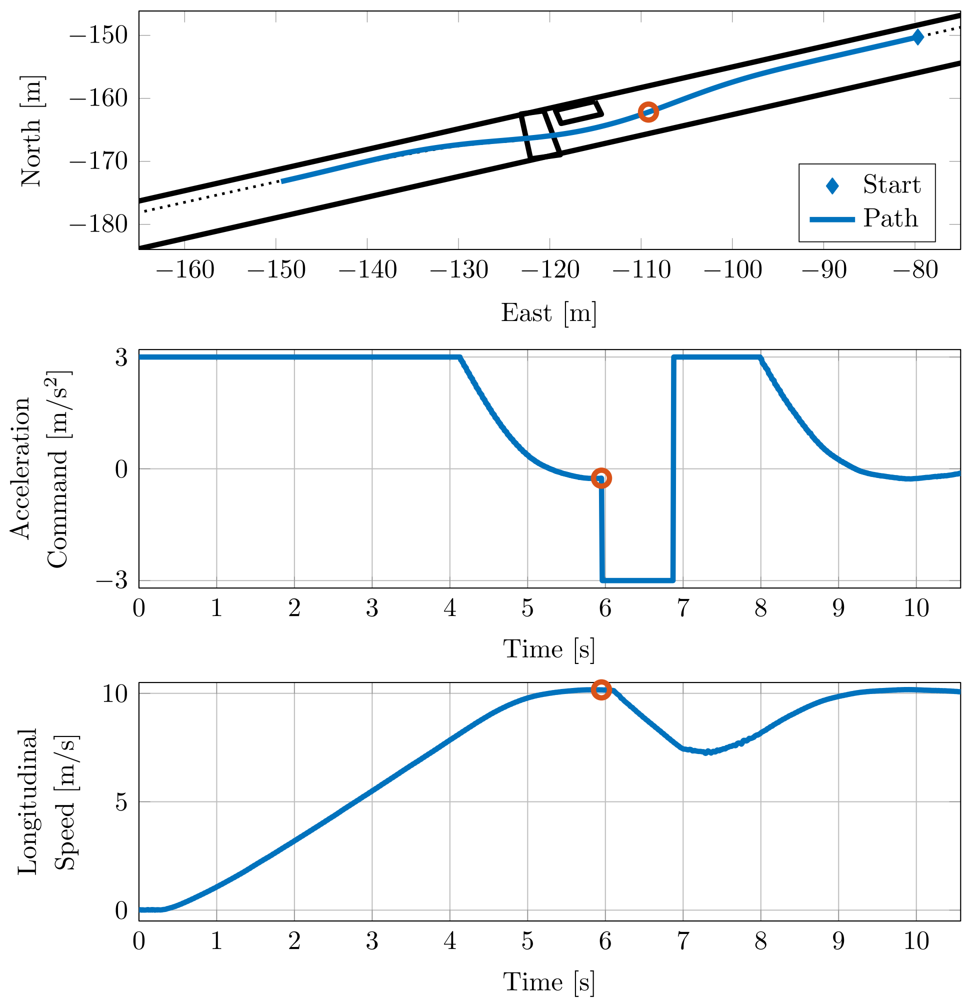

Figure 4.5: Baseline trajectory overhead, acceleration command, and speed profile using deterministic speed control (circle indicates when the pedestrian was detected). The vehicle decelerates upon detection of the pedestrian, but does not yield.

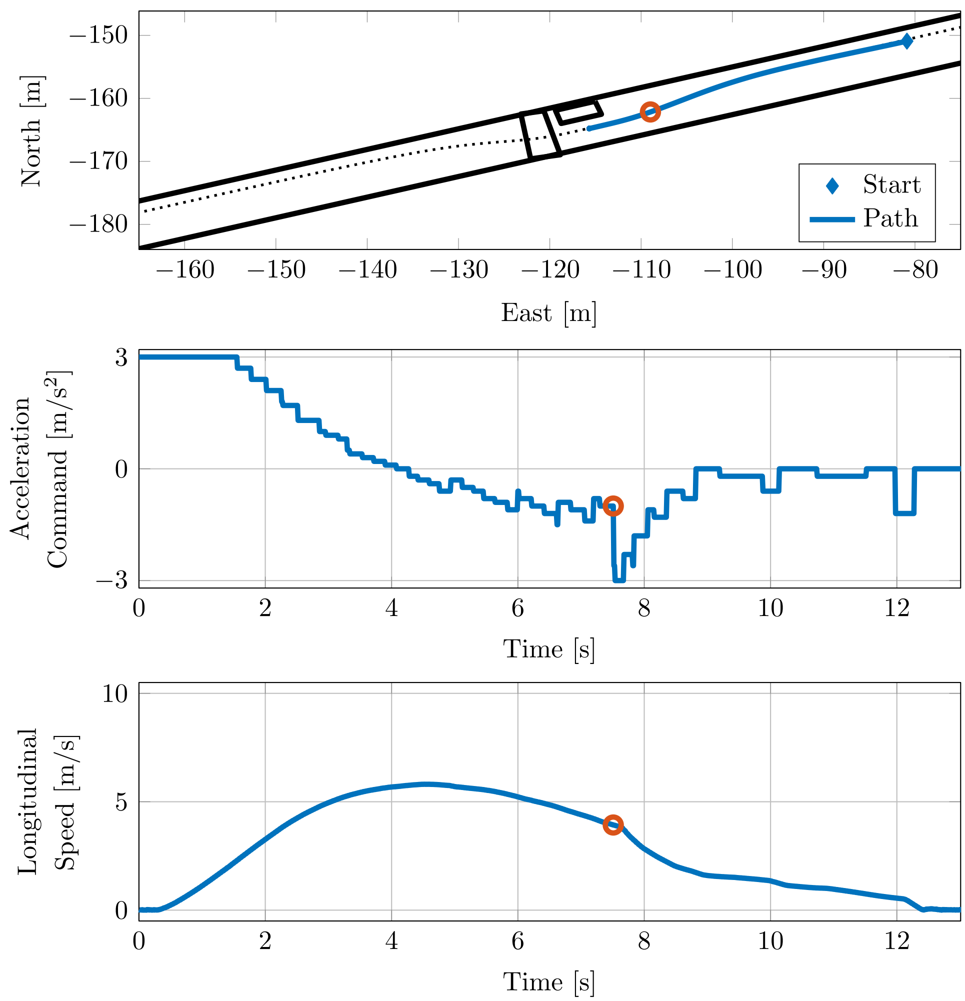

Figure 4.6: POMDP trajectory overhead, acceleration command, and speed profile using belief about pedestrian detection (circle indicates when the pedestrian was detected).

policies, respectively. The circles in Figure 4.5 and 4.6 indicate when the pedestrian was detected by the intensity filter. For both approaches, the immediate time step after detection commands a large deceleration. Because the baseline control is at full speed when the pedestrian appears, it is unable to legally yield to the pedestrian. In contrast, the POMDP policy has the vehicle decelerate much earlier and has it reach a lower maximum speed. Consequently, Trudi successfully stops for the sudden pedestrian.

4.3.4 Lessons Learned

This first iteration of the speed control design demonstrates the difficulty of designing an algorithm in the midst of competing values. There are some components of the implementation to highlight and some things to improve.

Positive Outcomes

- Accounting for the pedestrian uncertainty allowed the vehicle to successfully yield to the pedestrian. This effect largely came from the vehicle approaching the crosswalk at a "reasonable speed" because the POMDP anticipated future state information.

- The only information with regard to the pedestrian used was whether he or she was detected. This largely upheld the values of fairness and reciprocity but should be more explicit.

- With proper choice of weights, the tension between safety/legality and efficiency/mobility can be balanced.

- The design decision to model the problem as a POMDP and solve for an offline policy helped with investigating and balancing some of the value tensions in this design task.

Things to Improve

- Remove the limitation on braking authority, which led to a prioritization of occupant comfort above safety.
- Although the POMDP formulation is intentionally designed for occupant comfort, it optimized only for smoothness in velocity and did not account for the jerk vehicle occupants experience due to choppy acceleration commands. When considering merely the closed-loop policies, the value of smoothness seems to be achieved, but the in-vehicle experiments indicate that smoothness may not have been properly accounted for with this first iteration.
- This particular scenario is not generalizable to non-occluded crosswalks.
- Pedestrian modeling is key to the value tension, so more focus is needed there.

The next iteration will explore how to maintain these positive attributes while addressing some of the downsides of this implementation.

4.4 The Second Iteration

The iterative process of value sensitive design is not only helpful to identify how to improve the technical implementation but can also be used to re-evaluate the design task. In the second iteration, the scenario is revised to focus on the uncertainty of pedestrian behavior. By eliminating the occluding vehicle, the design task can investigate how pedestrian intent affects the vehicle behavior and vice versa. The pedestrian behavior introduces a lot of uncertainty in crosswalk scenarios, and the inclusion of an occluding vehicle obfuscates the pedestrian-vehicle interaction. Hence, the occlusion is removed and a pedestrian is positioned at the side of the road as shown in Figure 4.7. As the iterations progress, it is reasonable to assume the designer will gain a better understanding of the pedestrian-vehicle interaction. At which point, the occlusion can be re-introduced to the design task or used as a test case in the analysis phase.

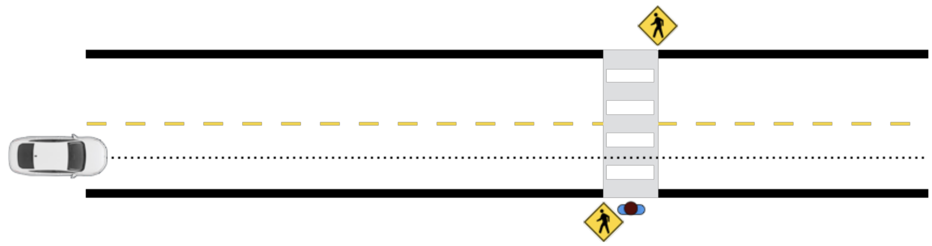

Figure 4.7: Experimental scenario of pedestrian crosswalk.

4.4.1 Conceptualization

The design task still involves various stakeholders and touches upon many human values. The direct stakeholders are now the occupants in the autonomous vehicle, the pedestrian potentially crossing the street, and the authority of traffic laws. Even with the removal of the occluding vehicle, the values at stake in the scenario are still mobility, safety, legality, care and respect for others, fairness and reciprocity, respect for authority, trust and transparency, and individual autonomy. The values take on the same definition as in Section 4.3.1, but how they translate to an engineering specification is going to be refined.

In the last iteration, the only human values explicitly considered were those that related to an engineering objective. This iteration serves to clarify how each identified value is to be captured in the technology. In particular, the value of fairness and reciprocity does not translate directly to an engineering objective. Instead, it becomes a higher-level design constraint that limits the information to be non-discriminatory, e.g. no age or gender information.

The other values are addressed by relating them to specifications that can be captured by engineering terms. These are all summarized in Table 4.3

Legality and Respect for Authority

Upon closer inspection of the California Vehicle Code, safety and legality are not strictly the same requirement. The vehicle code only requires drivers to exhibit “due care” to be safe, which is not the same requirement to actually be safe. The vehicle code further specifies reducing the vehicle speed and taking actions as necessary to safeguard the safety of the pedestrian. The key pieces of information necessary for legal decision-making are vehicle speed (v_t), vehicle distance to crosswalk (d_t), and pedestrian behavior. In order to safeguard the pedestrian, the autonomous vehicle must have information about whether the pedestrian is going to transition from the sidewalk to the crosswalk. The pedestrian behavior is assumed to be captured by the pedestrian position (c_t) and pedestrian posture (p_t), which are non-discriminatory.

Safety, Care and Respect for Others

The value of safety is a more strict interpretation of the vehicle code. Safety focuses on harm and injury reduction. To achieve this value, the same information as for legality is needed: vehicle speed (v_t), vehicle distance to crosswalk (d_t), pedestrian position (c_t) and pedestrian posture (p_t).

Mobility and Efficiency

The metric of time efficiency is still captured by the value of mobility, and is directly related to the speed of the vehicle (v_t) for a straight path.

Mobility and Smoothness

An additional aspect of mobility is smooth driving, which still affects occupant comfort and interjects trust and transparency between the stakeholders. In this iteration, the value of mobility is intended to improve by using both the previous acceleration (a_{t-1}) and current acceleration command (a_t) for smooth change in actions.

4.4.2 Technical Implementation

A new iteration provides an opportunity to choose a different technique or algorithm that better aligns with the defined values. Since the POMDP helped to illuminate value tensions in the previous iteration, the POMDP is kept in this iteration since it seems to offer potential resolution. Dynamic programming is used again to compute an optimal policy to control the longitudinal acceleration of the vehicle based on the belief of a pedestrian crossing.

Table 4.3: Summary of human values mapping to engineering specifications for the second VSD iteration.

Human Value	Engineering Specification	Information
Fairness and reciprocity	Do not use discriminatory information.	
Legality Respect for authority	Legality	v_t d_t c_t p_t
Safety Care and respect for others	Safety	v_t d_t c_t p_t
Mobility Individual autonomy Trust Transparency	Mobility	v_t a_{t-1} a_t

Given the engineering specifications of legality, safety, and mobility, the information necessary to address their respective values in the objective function is captured by the state vector

$$x = [v_t \quad d_t \quad c_t \quad p_t \quad a_{t-1}]^\top \tag{4.9}$$

and the control input

$$u_t = a_t, \tag{4.10}$$

where v_t is the vehicle speed, d_t is the vehicle distance to the crosswalk, c_t is the pedestrian position, p_t is the pedestrian posture, and a_{t-1} and a_t are the previous and current longitudinal acceleration, respectively. The top speed of the roadway is assumed to be 10 m/s, so the vehicle speed is upper bounded by the speed limit to coincide with both the legality and safety objectives. The pedestrian position is either in the crosswalk or on the sidewalk, and the pedestrian posture is either stopped while the pedestrian makes eye contact with the vehicle, is distracted, or is in motion. In order to continue to uphold the values of fairness and reciprocity, the pedestrian states do not rely on other information about the pedestrian that may be discriminatory. Previously, the control input was limited to $\pm 3\,\mathrm{m/s^2}$ to provide comfortable acceleration values, but this impeded the vehicle's ability to be safe. Here, the control algorithm allows the vehicle to use its full braking authority by allowing deceleration up to $-10\,\mathrm{m/s^2}$.

The dynamics (or state transitions) still use a point mass model of the vehicle to calculate the distance to the crosswalk and vehicle speed. A new model for the pedestrian is developed in order to further investigate the value tensions for the design task (Table 4.4). The likelihood of the pedestrian transitioning from the sidewalk to the crosswalk is a function of the pedestrian posture. The likelihood is 50 % when the pedestrian is distracted and 86.7 % [99] when the pedestrian is in motion. (The probability of 86.7 % is calculated from Schroeder and Rouphail's [99] statistics on yield and non-yield events for an assertive pedestrian at site B.) When the pedestrian is stopped while making eye contact with the vehicle, the probability of transitioning is a function of the vehicle's distance to the crosswalk

$$\Pr(c_t \mid \neg c_t; p_t = \texttt{STOPPED}) = (p_{\mathrm{xing}}/d_{\mathrm{max}})d_t, \tag{4.11}$$

where p_{xing} is 52.3 % [99] likelihood and d_{max} is the maximum distance the vehicle is defined to be away from the crosswalk. (The probability of 52.3 % is calculated from Schroeder and Rouphail's [99] statistics on yield and non-yield events for a pedestrian waiting on the near side at site B.) Once within the crosswalk, the pedestrian is assumed to stay in the crosswalk for the next time step. The control loop assumes perfect information for the vehicle distance to the crosswalk, vehicle speed, and, for simplicity, the pedestrian posture. However, there is observation uncertainty for the pedestrian position with a false positive of 5 %, which captures sensor uncertainty. These false positive rates were again chosen arbitrarily small but, in practice, would come from the perception system's capability of detecting pedestrians.

The goal is still for the autonomous vehicle to smoothly drive safely and efficiently through the crosswalk while adhering to the relevant traffic laws. The reward function defines the stage cost $g(x_t, u_t)$ for every state and action, which again further connects

Table 4.4: Pedestrian transition model for the second VSD iteration.

Pedestrian posture	Transition probability $\Pr(c_t \mid \neg c_t)$
Distracted	0.5
Stopped	$0.523^{\dagger}(d_t/d_{\max})$
Moving	$0.867^{\dagger}$

†calculated from yield event statistics [99]

the conceptualization values to the technical implementation. The reward for a state-action pair involves adding stage costs (4.12), (4.13), and (4.16) for that state and action.

The stage cost for legality derives from the constant acceleration point mass equations relating the constant deceleration needed to come to a complete stop given the distance to the crosswalk and vehicle speed, and is as follows

$$g_{\text{legality}}(x_t, u_t) = -\zeta \frac{v_t^2}{d_t + \epsilon} \mathbf{1}(c_t), \tag{4.12}$$

where $\epsilon > 0$ is a buffer in the denominator to soften the constraint, and $\zeta > 0$ is a weight on the penalty incurred by driving quickly as the vehicle gets closer to the crosswalk.

The stage cost for safety is

$$g_{\text{safety}}(x_t, u_t) = -\eta \mathbf{1}(c_t \wedge d_t < 0), \tag{4.13}$$

where $\eta > 0$ is a terminal penalty independent of velocity to encourage the vehicle to stop when the pedestrian is crossing.

For mobility, the stage cost consists of two terms:

$$g_{\text{efficient}}(x_t, u_t) = \lambda v_t \mathbf{1}(\neg c_t) \tag{4.14}$$

and

$$g_{\text{smooth}}(x_t, u_t) = -\xi (a_{t-1} - a_t)^2. \tag{4.15}$$

The total stage cost for mobility then becomes:

$$g_{\text{mobility}}(x_t, u_t) = g_{\text{efficient}}(x_t, u_t) + g_{\text{smooth}}(x_t, u_t) = \lambda v_t \mathbf{1}(\neg c_t) - \xi(a_{t-1} - a_t)^2, \quad (4.16)$$

where $\lambda > 0$ is a reward weight to encourage higher speed when the pedestrian is not crossing, and $\xi > 0$ is a penalty on large changes in acceleration.

To solve the POMDP, the method of QMDP is used again to approximate an optimal solution. In this iteration, vehicle speed increments in steps of 0.5 m/s, vehicle distance to crosswalk increments by 1 m, and accelerations are quantized by 0.5 m/s^2 intervals. These discretizations were chosen such that the sizes of the state and action spaces in the POMDP were kept small: 142,884 total states (includes terminal states) and 27 possible actions.

4.4.3 Empirical Analysis

The empirical analysis for the second iteration focuses on experimental results. A policy comparison is not included in this analysis because, unlike the baseline, the state space of this POMDP cannot be fully represented in three dimensions. This is because the POMDP policies are conditioned on the previous acceleration command.

Experimental Results

Once again, in-vehicle experiments are conducted using Trudi. However, instead of using the lidar sensors and retro-reflective material for pedestrian detections, these experiments use computer vision to identify the presence of a pedestrian [100]. With the definition of a static polygon for the shape of the road, the pedestrian detection is used to determine whether the pedestrian is in the crosswalk influence area (i.e. the sidewalk) or in the crosswalk. The vehicle is tasked with following a straight line path down the road still using a deterministic model predictive steering control as in Brown *et al.* and Chapter 2. The experimental scenario involves a pedestrian crosswalk on a two-lane roadway (Figure 4.8). The vehicle is at speed when the pedestrian enters the crosswalk influence area, at which point the policy starts executing. As the vehicle approaches the crosswalk, the pedestrian may or may not transition into the

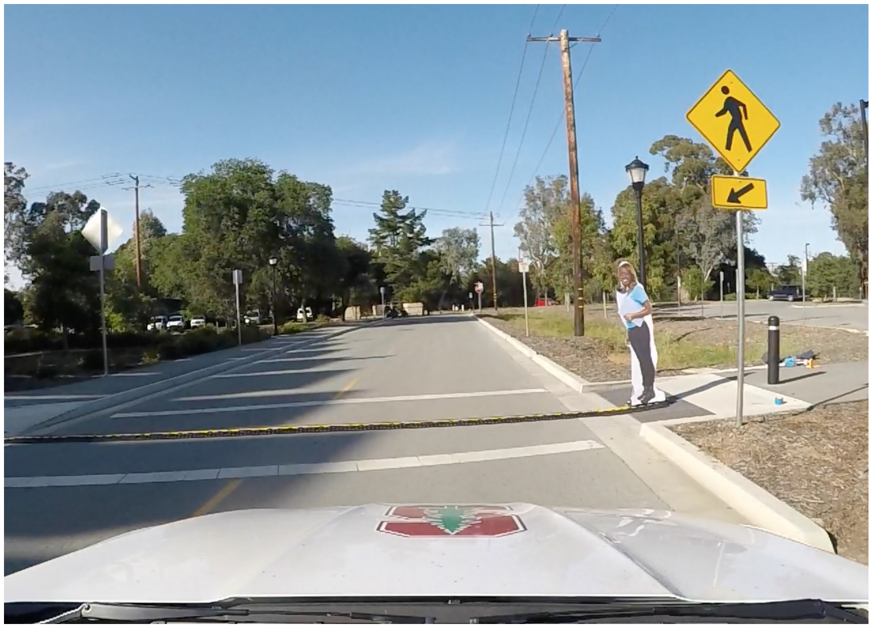

Figure 4.8: Experimental setup of pedestrian crosswalk using a cardboard cutout for the pedestrian that moves along a track. Depicts the pedestrian posture of stopped.

crosswalk. The control algorithms have no prior knowledge as to whether or when the pedestrian will transition.

The baseline from the first iteration is used again for comparison against the new POMDP policies. It is considered to be an aggressive baseline since it will not yield to the pedestrian until he or she has entered into the crosswalk. As an alternative, a conservative baseline is also considered, according to which the vehicle starts to yield to the pedestrian once he or she enters the crosswalk influence area. The policies are the same, except for what is considered to be the crosswalk influence area which determines when to switch between cruise control and braking. Yet, they do not account for the pedestrian posture. Figures 4.9 and 4.10 depict the overhead driven trajectory, acceleration commands, and speed profile for the aggressive and conservative baselines, respectively. The circles indicate when the pedestrian was detected to be in the crosswalk by the computer vision algorithm. Since there is no circle in the aggressive baseline, this means the pedestrian never entered the crosswalk. The vehicle continued at the speed limit, never yielding to the pedestrian, because the

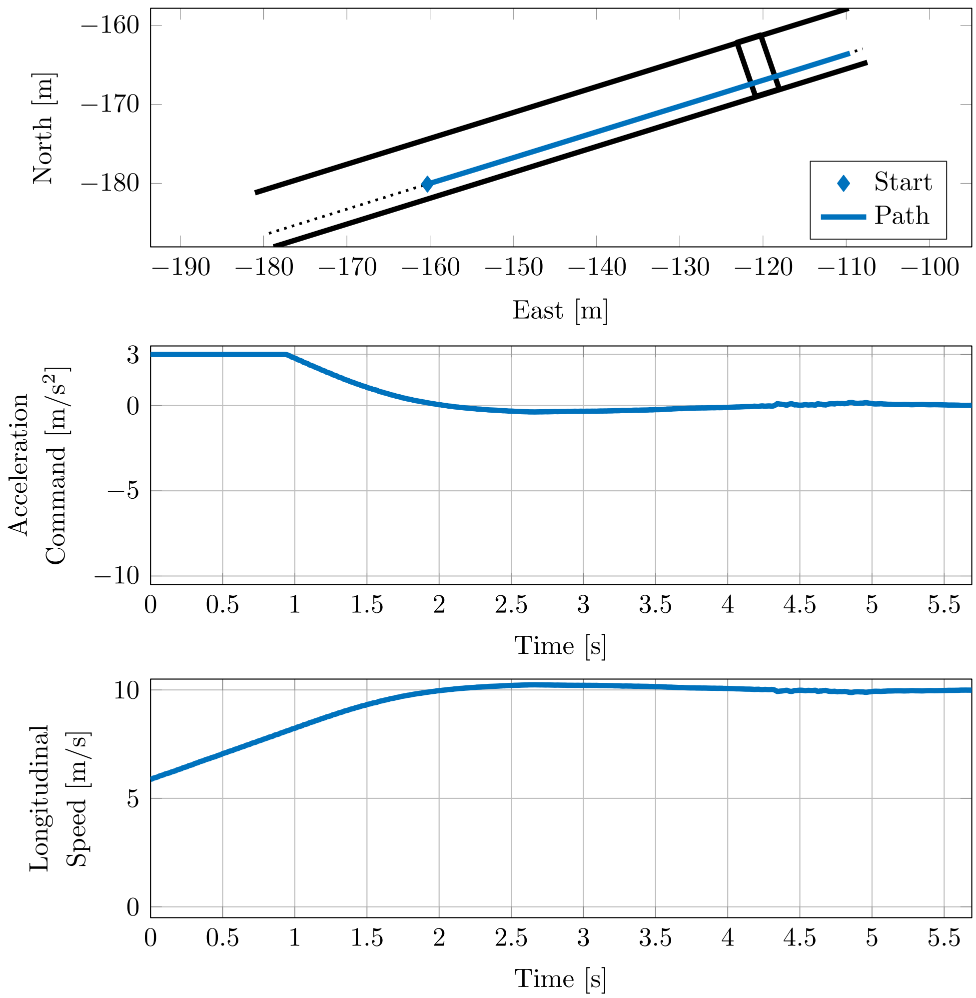

Figure 4.9: Aggressive baseline trajectory overhead, acceleration command, and speed profile using deterministic speed control. There is no red circle because the pedestrian does not enter the crosswalk.

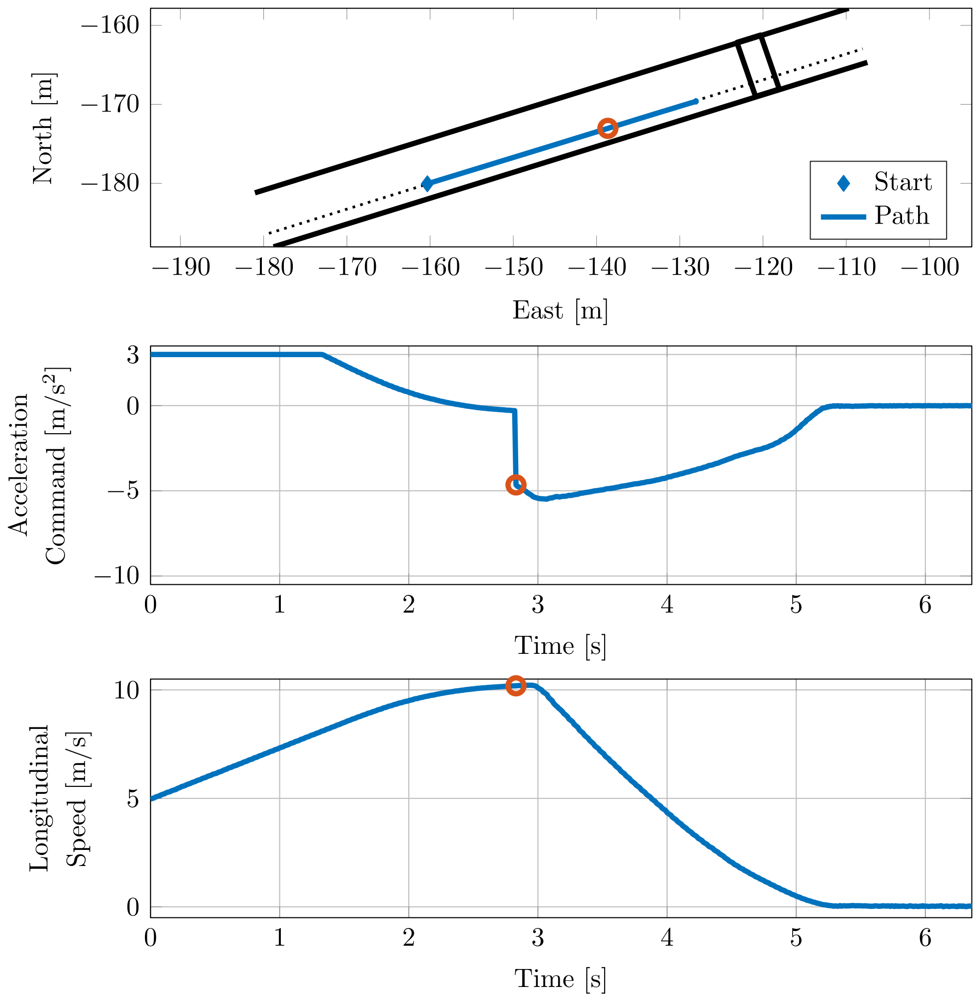

Figure 4.10: Conservative baseline trajectory overhead, acceleration command, and speed profile using deterministic speed control (circle indicates when the pedestrian was detected).

pedestrian did not enter the crosswalk. For the conservative baseline, the perception system detected the pedestrian to be in the crosswalk influence area when the vehicle was 12.99 m away from the crosswalk, and the vehicle successfully yielded to the pedestrian.

For the policy execution of the POMDP, an observation of the vehicle speed, vehicle distance to crosswalk, pedestrian posture, and pedestrian location are used to update the belief with a Bayesian filter just like the previous iteration. Figures 4.11-4.14 depict the overhead driven trajectory, acceleration commands, and speed profile for the POMDP policies. The circles indicate when the pedestrian was detected in the crosswalk by the computer vision algorithm.

In this second POMDP implementation, the pedestrian postures are independent from each other. Hence, a different set of weights are used in the reward function for each posture (Table 4.5). This is reasonable because each pedestrian posture is a unique sub-scenario that requires a different vehicle response. The numerical value for the buffer is still rather arbitrary; it is chosen such that the numerator does not evaluate to zero and to limit the magnitude of the constant deceleration term. Again, the other weights can be further tuned with additional analysis through Pareto optimization (see Section 4.5) but are chosen preliminarily here to see if this design is satisfactory.

For the scenario of the distracted pedestrian (Figure 4.11), the weights are chosen such that safety, efficiency, and smoothness are prioritized to similar normalized values: $\eta_{\mathrm{n}} = 0.5$, $\lambda_{\mathrm{n}} = 0.5$, and $\xi_{\mathrm{n}} = 0.507$, respectively, at the extreme states and actions when $v_t = 10\,\mathrm{m/s}$, $d_t < 0\,\mathrm{m}$, and $a_{t-1} - a_t = 13\,\mathrm{m/s^2}$. The legality term is normalized to $\zeta_{\mathrm{n}} = 0.125$ when $v_t = 10\,\mathrm{m/s}$ and $d_t = 0\,\mathrm{m}$, suggesting lower prioritization. Figure 4.11 shows that once the pedestrian enters the crosswalk influence area, the policy executes small negative accelerations to slow the vehicle down to around 1.5 m/s and to make it coast until the pedestrian enters the crosswalk. Once the pedestrian enters the crosswalk, the vehicle comes to a complete stop.

When the pedestrian is walking (Figure 4.12), the terms for efficiency and smoothness increase to normalized values of $\lambda_{\mathrm{n}} = 1$ and $\xi_{\mathrm{n}} = 1.69$. With higher efficiency, the vehicle drives at a faster speed down the road and more smoothness is consequently

Table 4.5: Weights of the reward function with respect to pedestrian posture (p_t)

Variable	$p_t = \texttt{DISTRACTED}$ Weight	$p_t = \texttt{WALKING}$ Weight	$p_t = \texttt{STOPPED}$ Weight	Unit
Legality (ζ)	0.01	0	0.01	s^2/m
Buffer (ϵ)	8	8	8	m
Safety (η)	0.5	0.5	0.5	–
Mobility (λ)	0.05	0.1	0.03	s/m
Mobility (ξ)	0.003	0.01	0.003	$\mathrm{s}^2/\mathrm{m}^2$

needed to smoothly decelerate the vehicle from the faster speed in case the pedestrian enters the crosswalk. Because of the high probability the pedestrian transitions to the crosswalk, there is consequently a high belief of 0.36 that the pedestrian is crossing. This means the impact of the safety and legality terms largely influence the vehicle behavior well before the pedestrian is physically within the crosswalk. To reduce the large impact of the safety and legality terms, one of the terms is reduced to 0 (legality, in this instance) to allow the vehicle to progress towards the crosswalk. With these weights, the vehicle again coasts until the pedestrian is detected. As the vehicle gets closer to the crosswalk, it smoothly decelerates to a complete stop.

For the stopped pedestrian, the normalized legality term increases back to $\zeta_{\mathrm{n}} = 0.125$ while efficiency and smoothness decrease to $\lambda_{\mathrm{n}} = 0.3$ and $\xi_{\mathrm{n}} = 0.507$. Smoothness is really important in this scenario for the autonomous vehicle to demonstrate transparency about its intentions to move through the environment, and hence is chosen to have the highest prioritization. Using these weights, two different scenarios were tested. The scenario depicted in Figure 4.14 shows the vehicle gradually accelerating as the vehicle gets closer to the crosswalk because of decreasing belief that the pedestrian will cross the street. Since the pedestrian does not want to cause an immediate hazard, it does not enter the crosswalk. In the second stopped pedestrian scenario, Figure 4.14 portrays the pedestrian crossing the street before the vehicle accelerates too much. Once the pedestrian is in the crosswalk, the vehicle comes to a complete stop.

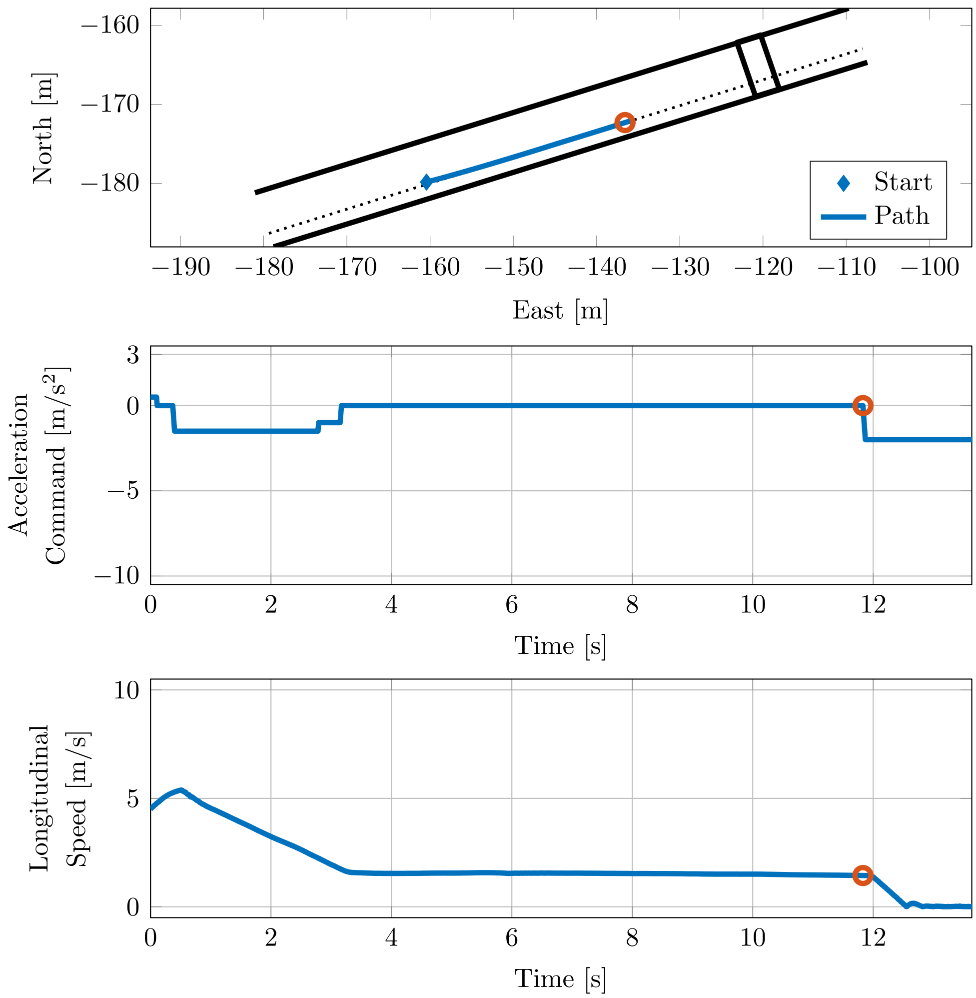

Figure 4.11: Distracted pedestrian POMDP trajectory overhead, acceleration command, and speed profile using the belief of the pedestrian crossing (circle indicates when the pedestrian was detected).

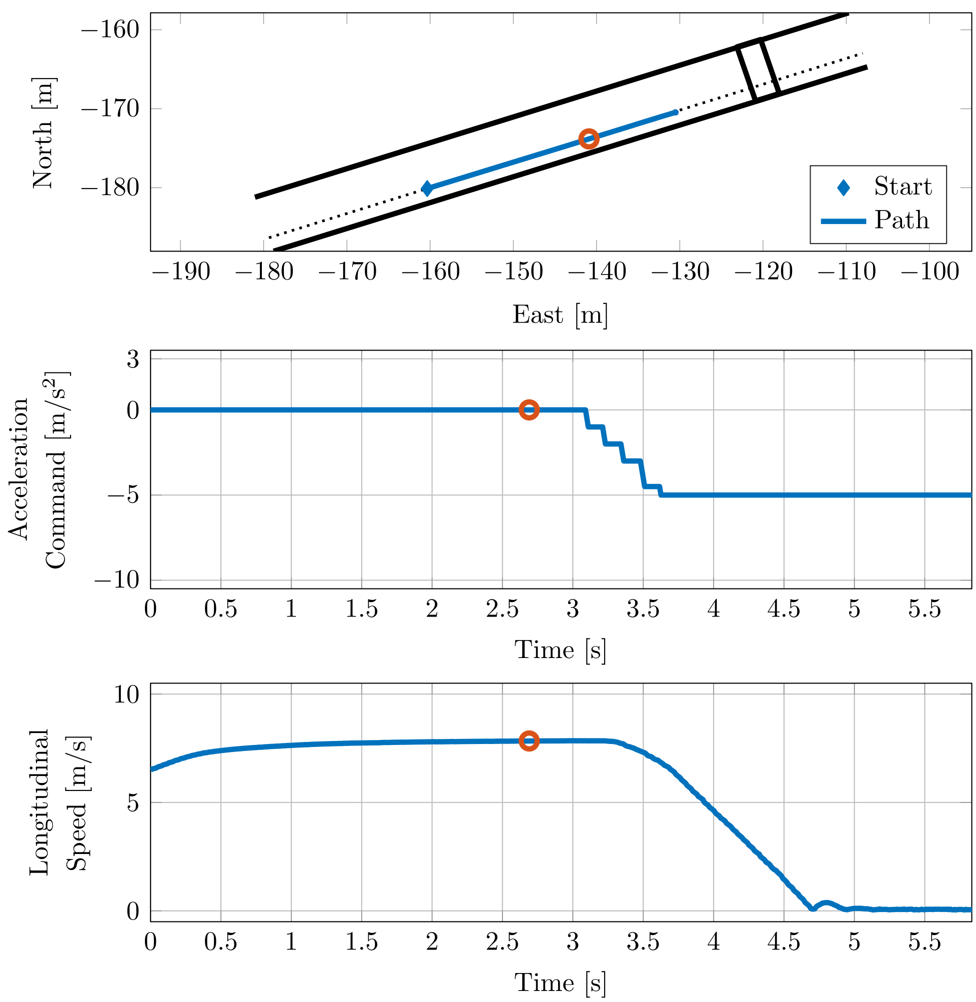

Figure 4.12: Walking pedestrian POMDP trajectory overhead, acceleration command, and speed profile using the belief of the pedestrian crossing (circle indicates when the pedestrian was detected).

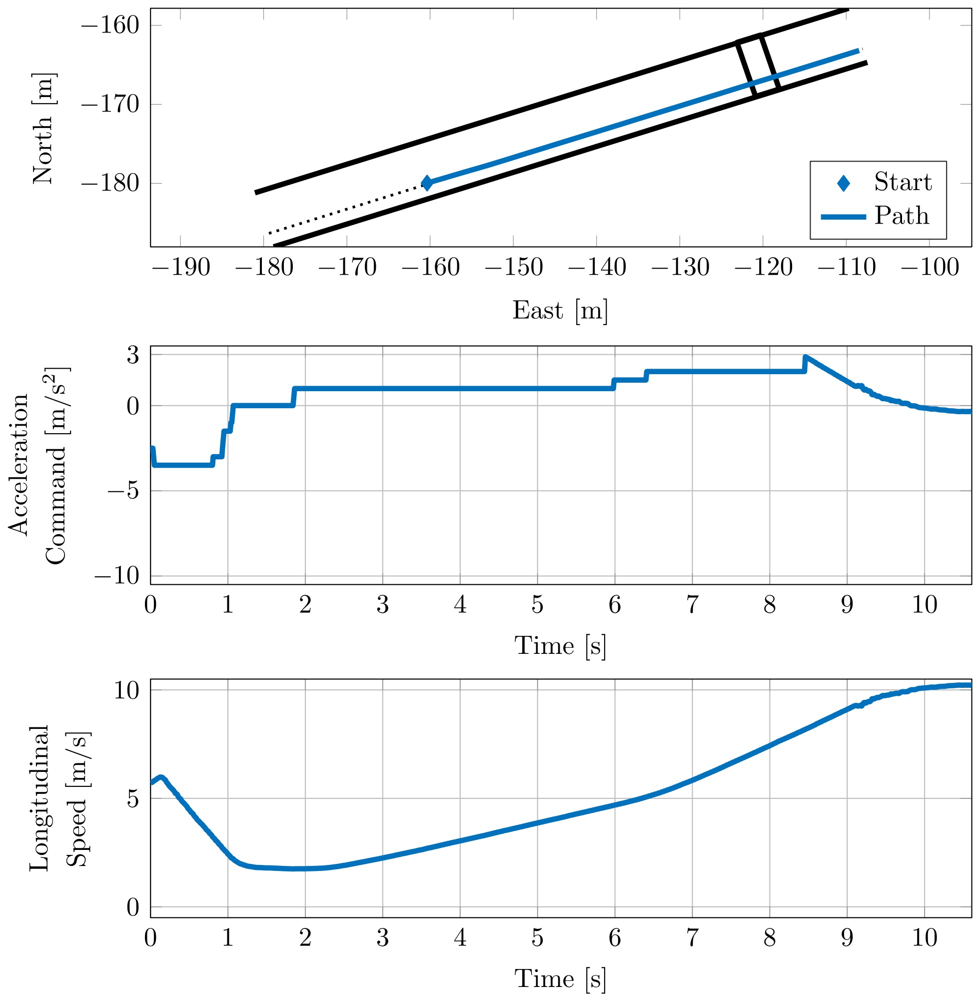

Figure 4.13: Stopped pedestrian POMDP trajectory overhead, acceleration command, and speed profile using the belief of the pedestrian crossing. There is no red circle because the pedestrian does not enter the crosswalk.

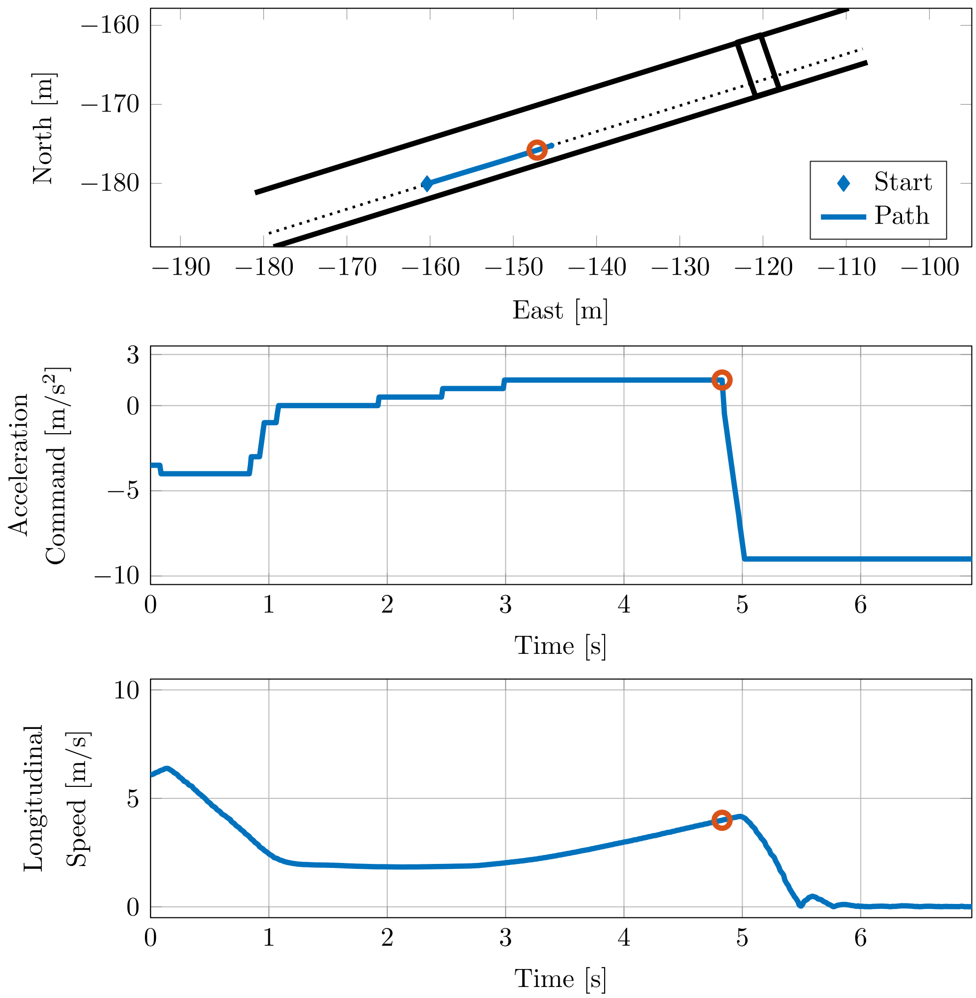

Figure 4.14: Stopped pedestrian POMDP trajectory overhead, acceleration command, and speed profile using the belief of the pedestrian crossing (circle indicates when the pedestrian was detected). Pedestrian takes right of way.

With the focus on pedestrian behavior in this second iteration, the scenario inherently gains complexity around the value tension between the pedestrian's intent and the autonomous vehicle's desire to travel down the road. The weights chosen here are still rather arbitrary trade-offs over the value statements. Hence, a deeper analysis is likely appropriate at this point to better determine if a particular design point can resolve the value tensions. For example, a Pareto optimization would be useful to simulate over many scenarios as the choice of weights change. This is demonstrated in the next section (Section 4.5).

4.4.4 Lessons Learned

This second iteration of the speed control design demonstrates improvements in handling the value conflicts as the engineering specifications refine in terms of the identified values. There are still components of the implementation to highlight and some things to improve because this is not the final product.

Positive Outcomes

- In all scenarios, accounting for the pedestrian uncertainty allowed the vehicle to successfully yield to the pedestrian. The design choice of modeling the problem as a POMDP meant dynamic programming could be used to account for future state information in the policy.

- The only information about the pedestrian used was whether he or she was in the crosswalk and what posture he or she composed. This continued to uphold the values of fairness and reciprocity.

- The continued design decision to model the problem as a POMDP and solve for an offline policy helped to investigate and balance some of the value tensions in this design task. Although the actual choice of weights were rather arbitrary, it demonstrated the potential for the value tensions to be resolved.

- Smoothness improved with the penalty on change in acceleration and the influence corresponded directly with ξ.

- Efficiency continued to correspond to the term λ.
- Modeling the pedestrian as a function of posture gave insight into pedestrian intent and crossing the street.
  - For the stopped pedestrian, the vehicle slowly increased its speed as it approached the crosswalk to indicate to the pedestrian it will yield if he or she enters the crosswalk but also wants to travel down the road.
  - The random probability for the distracted pedestrian allowed the vehicle to approach the crosswalk at a very cautious speed.

Things to Improve

- Pedestrian modeling needs to continue to improve
  - Pedestrian posture and position did not seem to fully grasp the pedestrian's intent to cross the street. Other parameters could be considered while keeping in mind the values of fairness and reciprocity to mitigate use of biased information or discriminatory actions by the vehicle.
  - There is likely some correlation between distraction and motion for the pedestrian which may indicate a higher likelihood of transitioning for the distracted posture. Other pedestrian models could be considered to further study nuances in pedestrian posture and motion.
  - The pedestrian transitions assume the pedestrian will stay within the crosswalk once they enter the crosswalk. This is not true in reality. More exploration into modeling the transition from the crosswalk to the sidewalk (or safe distance from the autonomous vehicle's traveling lane) should be considered.
- The vehicle tended to stop short of the crosswalk. This could be a result of poor choice of weights or the reward function could adjust, i.e. add a slight penalty on large decelerations so the vehicle only comes to a full stop when necessary. This will help with the engineering specification of mobility and efficiency.

- In the situation with the moving pedestrian, the high likelihood of transitioning greatly increased the influence of the safety and legality terms on the policy. These weights either need to be tuned down significantly or maybe an alternative formulation should be considered to better isolate the impact of safety and legality.

If another iteration were to occur (this thesis only presents two iterations), then it would investigate how to maintain these positive attributes while addressing some of the downsides of this second implementation. Further investigation into the choice of weights in the reward function is also needed in order to determine how well mobility, safety, and legality can be realized with this implementation. This could be done with a Pareto, or multi-objective, optimization over the weights, for example, and is demonstrated in Section 4.5 for the first iteration.

4.5 Closing the Loop on Human Values

The policy comparison and experimental results demonstrate a potentially reasonable speed control algorithm design, but only for a particular set of weights. The behavior of the vehicle can greatly vary depending on the choice of weights in the reward function. To more directly analyze how well the designed technology aligns with the stakeholder values, an analysis technique is needed. One way to perform this analysis is with the technique of Pareto (or multi-objective) optimization in order to determine which set of weights best align with the values. A design is Pareto optimal if one objective cannot improve without worsening at least one other objective. To construct a frontier of Pareto optimal points, the design objectives map to a criterion space using evaluation criteria. The determination of Pareto optima serves to close the loop on the design process where human values map to engineering objectives, engineering objectives map to evaluation criteria, and evaluation criteria map to human values. Thus, engineers can focus on Pareto optimal designs without committing to a particular prioritization between objectives ahead of time. The Pareto frontier can then be taken back to a larger group of stakeholders to determine the final design to deploy.

An example of a Pareto frontier for the first VSD iteration is constructed by varying the weights in the reward function that correspond to the engineering objectives. For each combination of weights in the reward function, a different optimal policy is generated. For each given optimal policy, Monte Carlo simulations are run and the simulation results are calculated against evaluation criteria. The objective of safety and legality maps to the criterion of the vehicle velocity at the crosswalk. The objective of efficiency maps to the criterion of average time to complete maneuver. The objective of smoothness maps to the criterion of average maximum change in acceleration. The Monte Carlo simulations entail the pedestrian suddenly appearing from behind the occluding vehicle at random times whenever the autonomous vehicle is within 20 m of the crosswalk. The simulations include the assumption that the pedestrian takes about 4 s to cross the street.

The resulting Pareto frontier can then be brought back to a larger group of stakeholders, such as policymakers, lawyers, and public interest groups, to determine which set of weights to deploy on the autonomous vehicle. Figure 4.15 shows an example of a slice of the Pareto frontier for safety and legality vs. mobility. It is additionally colorized by the yield rate for the simulated suddenly appearing pedestrian scenarios. The larger group of stakeholders can confer the Pareto frontiers to additional information, such as injury curves [101], user studies, emissions curves [102], and congestion studies [103].

Pareto optimization is not the only tool that can help close the loop on human values. Another utilitarian-like analysis tool could be a risk management or cost-benefit analysis for a set of outcomes [104]. Or maybe a deontological-like analysis is more desirable where thresholds or conditions are determined by policymakers or by re-engaging with stakeholders. This Pareto analysis is a first step on demonstrating how appropriate analysis tools can help determine how successfully a technical implementation embodies the human values identified in the conceptualization phase.

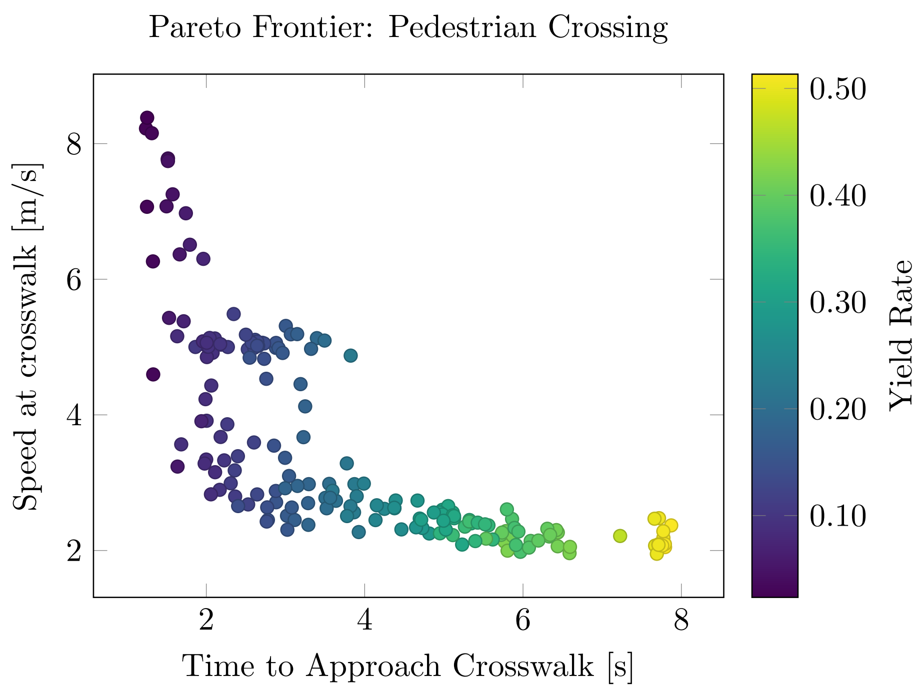

Figure 4.15: Pareto frontier of POMDP for various weights mapped to evaluation criteria.

4.6 Summary

This chapter demonstrates the formal connection of human values into the design of a speed control algorithm through the conceptualization and technical implementation phases. The empirical analysis phase helps identify areas of improvement for subsequent iterations. In the first iteration, a POMDP is chosen to help realize the values of safety and legality, efficiency and mobility, and smoothness for a scenario with a large vehicle parked in front of a pedestrian crosswalk. The POMDP helped with capturing the uncertainty in the situation and allowed the vehicle to be proactive by approaching the crosswalk at a reasonable speed, which led to a successful yield to a suddenly appearing pedestrian. The pedestrian model in the first iteration is very simple, so the second iteration considers a slight change in the scenario by removing the occlusion and improves the pedestrian model to look more closely at the value tension between the pedestrian and autonomous vehicle. In the second iteration, the values of legality, safety, efficiency, and smoothness were refined in terms of the technical implementation. Additional analysis with Pareto optimization provides further insight into how well an implementation aligns with the identified values. Iterating through VSD helps engineers think more deeply about how human values are implicated in the technology as it develops.

The focus here has been on engineers and programmers as designers, but VSD allows for other stakeholders to be involved in the design process, such as policymakers and civil servants. VSD is not only a valuable tool for engineers, but can widely encompass other contributors at a company or from third party groups to help ensure autonomous vehicles behave in socially acceptable ways.

Chapter 5

Conclusions

Engineers make design decisions every day and may not realize that these choices can have ethical implications, especially with regards to autonomous vehicle motion planning algorithms. The following distance for platooning autonomous vehicles seems like a straightforward parameter to choose. However, it can influence fuel efficiency, reaction time, and potentially impact surrounding traffic. This thesis serves to provide engineers with tools to gain a better understanding of how the choice of these parameters and various algorithm implementations can influence vehicle behavior and society at large.

Of course, not all engineering decisions deal with conflicting values; often times they deal with the one value of safety. Safety is an important value to treat judiciously, so Chapter 2 centers on modeling steering system delays in a motion planning algorithm to safely control the vehicle laterally. Within this task are several engineering trade-offs, such as algorithm complexity and implementation overhead.

For autonomous vehicle motion planning with ethical considerations, the two approaches suggested in this thesis are steps towards better engineers and better autonomous vehicle motion planning algorithms. The first approach turns directly to philosophy and uses parallels between philosophical and mathematical frameworks to justify the decisions in the design of an autonomous vehicle steering control algorithm. Deontological reasoning, a rule-based philosophy, can be used to justify the use of rule-based mathematical techniques, such as set theory and constraints.

Consequential reasoning, a cost-based philosophy, can be used to justify the use of cost-based mathematical paradigms, like optimization. There are many ways to program an autonomous vehicle, and Chapter 3 suggests that model predictive control is one way to harness the positive attributes of both deontology and consequentialism, given that MPC solves a constrained optimization problem. The choice of weights can lead to different vehicle behavior, so further insight is sought from a third philosophy known as virtue ethics. This mapping can equip engineers with better reasons for the design choice of a cost, constraint, or weight when formulating motion planning algorithms.

If a wide range of stakeholders (ideally, representative of a society) and explicit consideration of human values are involved in the design process, then engineers may be better equipped to ethically program autonomous vehicle motion planning algorithms. The stakeholders and the identified values provide guidance to the engineers and bring it perspectives they may not have considered at the start in isolation. Chapter 4 explores such an approach by applying a modified value sensitive design process to connect human values into engineering specifications. Various analysis techniques can be incorporated to help ensure the realization of the identified human values. An example design task of an autonomous vehicle speed control algorithm for navigating through a pedestrian crosswalk demonstrates such a process. Even with an ethically designed implementation, there could be an array of design points to choose from. Additional analysis, such as Pareto optimization, can help facilitate communication back to the larger group of stakeholders, along with other justifying resources, to determine which design point to deploy.

5.1 Contributions

This thesis makes the following contributions:

- **An MPC formulation to compensate for the delay due to the steering actuation on an autonomous vehicle platform.** The formulation used simple, computationally efficient models to improve the functionality of an autonomous vehicle (Chapter 2).

- **An incorporation of the traffic laws pertaining to lane dividers in model predictive steering control and crosswalks in partially observable speed control.** The traffic law §21460 is included in the formulation of the model predictive steering control in order to investigate adherence to double yellow lines when steering an autonomous vehicle around an obstruction (Chapter 3). The traffic law §21950 is embedded in the model of a POMDP for controlling the speed of an autonomous vehicle through a crosswalk (Chapter 4).

- **A mapping of philosophical principles to mathematical principles.** The philosophical frameworks of deontology, consequentialism, and virtue ethics map to the mathematical concepts of constraints, costs, and choice of weights (Chapter 3). This helps engineers to understand the implications of choosing when to use a cost and constraint or even the weights.

- **An application of a modified value sensitive design approach to the design of motion planning algorithms.** VSD is applied to the design of a speed controller for crosswalk scenarios in order to identify human values that could be implicated in the design task. Iterating through the design task teases out tensions amongst the values until the designed technology aligns with the identified human values.

5.2 Further Work

Although this thesis focuses on the decision layer of the autonomous vehicle stack, human values can be impacted at all levels in the design of an autonomous vehicle: how sensors are used, how perception algorithms are developed, and even how the vehicles are tested and deployed can have ethical implications.

5.2.1 Generalizability

The efficacy of the approaches in the thesis were demonstrated by scenarios. It is then unclear how well these approaches may scale or generalize to the real-world because the real-world has an unbounded number of scenarios to encounter. This is likely an important consideration for those developing autonomous vehicles, so including scalability or generalizability in the design process could help clarify the feasibility of such consideration.

Even if scalability was not accounted for at the beginning of the design, there are some techniques for scaling problem formulations with single-users. For example, Bandyopadhyay *et al.* [88], [89] formulate a speed control POMDP for an autonomous vehicle to drive down a street with a single pedestrian nearby with unknown intent. In the experiment, they started a new instance of the POMDP for each individual pedestrian encountered. Each instance output an action. The ultimate action taken was the most conservative of the set. Another approach is utility fusion [7], where the utilities of each encountered pedestrian are combined additively or minimized. Once the utilities reconcile, the ultimate action is taken from the policy.

5.2.2 Quantify Engineering Improvement with VSD

Value sensitive design has the potential to change the way engineers engineer. It forces the engineers to consider whether the conditional statement or reward function they are coding will further the values identified in the conceptualization. This is a purely qualitative observation. In order to quantify the usefulness of such an approach, a user-study with engineers should be conducted. There could be two groups of engineers that go through the same design task: one group learns about VSD and the other does not. At the end of the design task, the engineering implementation can be compared but also interviews of the engineers to document their thought and justification process of the designs.

5.2.3 Philosophical Frameworks and VSD

Although there are limitations of the philosophical frameworks approach in Chapter 3, they can still serve a purpose in the engineering process. The focus of Chapter 4 was on engineering analysis, but a more rigorous philosophical analysis could also be included. An engineer equipped with an understanding of the relationship between these philosophical and mathematical frameworks could conduct some preliminary analysis or gage simpler implications in the design process. However, the philosophers are those truly trained to make moral judgements, so VSD could be a great way to further engage them in the design process as they can do a philosophical analysis in parallel to the engineering analysis. (A similar argument can be made for a legal analysis.)

5.2.4 Policy

Government regulators recognize the importance of ethical considerations in the design of autonomous vehicle technology [105]–[107]. Claims have been made in this thesis that a process like VSD can help policymakers and regulators. This is largely due to a conversation centered around human values, which everyone can understand. However, it is important to actually engage policymakers to determine if they find this process useful.

5.3 Outlook

Before society can enjoy the benefits of autonomous vehicle technology, society should have a way to contribute to the conversation of its development. By engaging a larger group of stakeholders, a wider of range of perspectives can be included in the design process. This means better technology for everyone.

Appendix A

Vehicle Model

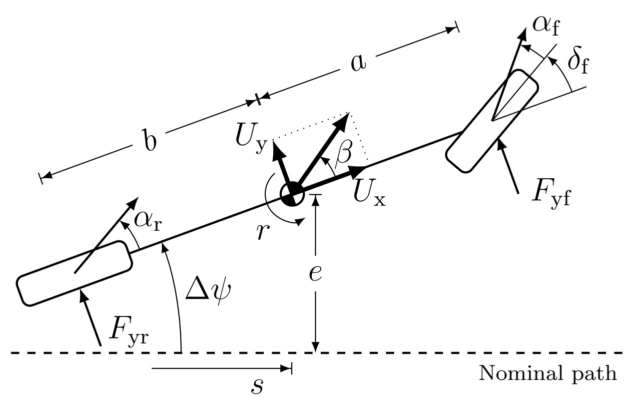

Figure A.1: Schematic of bicycle model.

A.1 Velocity States

The vehicle sideslip angle, β, and yaw rate, r, are the velocity states in the vehicle model. The sideslip angle is

$$\beta = \tan\left(\frac{U_y}{U_x}\right) \approx \frac{U_y}{U_x}, \tag{A.1}$$

where the lateral and longitudinal velocities are denoted in the body-fixed frame as U_y and U_x, respectively. As simplifications, we assume $U_x >> U_y$ and that U_x is constant.

The equations of motion for the sideslip and yaw rate are

$$\dot{\beta} = \frac{F_{\mathrm{yf}} + F_{\mathrm{yr}}}{mU_{\mathrm{x}}} - r \tag{A.2}$$

$$\dot{r} = \frac{aF_{\mathrm{yf}} - bF_{\mathrm{yr}}}{I_{\mathrm{zz}}}. \tag{A.3}$$

Here, the lateral tire force of the [front, rear] axle is denoted as $F_{\mathrm{y[f,r]}}$, the vehicle mass is denoted as m, the yaw inertia is denoted as I_{zz}, and the distances from the center of gravity of the vehicle to the front and rear axles are denoted a and b, respectively. The front tire slip angle, α_{f}, and rear tire slip angle, α_{r}, can be expressed as:

$$\begin{aligned} \alpha_{\mathrm{f}} &= \tan^{-1}\left(\beta + \frac{ar}{U_{\mathrm{x}}}\right) - \delta \\ &\approx \beta + \frac{ar}{U_{\mathrm{x}}} - \delta \end{aligned} \tag{A.4}$$

$$\begin{aligned} \alpha_{\mathrm{r}} &= \tan^{-1}\left(\beta - \frac{br}{U_{\mathrm{x}}}\right) \\ &\approx \beta - \frac{br}{U_{\mathrm{x}}} \end{aligned} \tag{A.5}$$

The linear expressions result from small angle approximations.

The brush tire model proposed by Fiala [41] and presented in the following form by Pacejka [42] provides a model of the relationship between the lateral tire forces and tire slip angles:

$$\begin{aligned} F_{\mathrm{y}} &= \begin{cases} -C_\alpha \tan\alpha + \frac{C_\alpha^2}{3\mu F_{\mathrm{z}}}|\tan\alpha|\tan\alpha \\ \qquad -\frac{C_\alpha^3}{27\mu^2 F_{\mathrm{z}}^2}\tan^3\alpha, & |\alpha| < \arctan\left(\frac{3\mu F_{\mathrm{z}}}{C_\alpha}\right) \\ -\mu F_{\mathrm{z}} \mathrm{sgn}\ \alpha, & \text{otherwise} \end{cases} \\ &= f_{\mathrm{tire}}(\alpha) \end{aligned} \tag{A.6}$$

Here, the surface coefficient of friction is given as μ, the normal load is given as $F_{\mathrm{z[f,r]}}$, and the tire cornering stiffness is given as C_α.

The vehicle model used by the online MPC controller utilizes the front tire force to keep the problem linear with regards to the input. The resulting steer angle, δ, follows from (A.4) and (A.6):

$$\delta \;=\; \beta + \frac{ar}{U_{\mathrm{x}}} - f_{\mathrm{tire}}^{-1}\left(F_{\mathrm{yf}}\right) \tag{A.7}$$

To address the nonlinearity of the rear tires, the brush tire model is linearized at a given rear tire slip angle ($\bar{\alpha}_{\mathrm{r}}$), and the rear tire force (F_{yr}) is thus modeled as an affine function of α_{r}:

$$F_{\mathrm{yr}} = \bar{F}_{\mathrm{yr}} - \bar{C}_{\bar{\alpha}_{\mathrm{r}}}(\alpha_{\mathrm{r}} - \bar{\alpha}_{\mathrm{r}}) \tag{A.8}$$

where $\bar{F}_{\mathrm{yr}} = f_{\mathrm{tire}}\left(\bar{\alpha}_{\mathrm{r}}\right)$ and $\bar{C}_{\bar{\alpha}_{\mathrm{r}}}$ is the equivalent cornering stiffness at $\bar{\alpha}_{\mathrm{r}}$. In the initial time steps of the prediction horizon, the current rear slip angle, α_{r}, is chosen to be $\bar{\alpha}_{\mathrm{r}}$. This allows the MPC controller to explicitly consider rear tire saturation in the near term prediction [70].

The equations of motion of the velocity states can now be formulated as affine functions of the states and input, F_{yf}:

$$\dot{\beta} \;=\; \frac{F_{\mathrm{yf}} + \bar{F}_{\mathrm{yr}} - \bar{C}_{\bar{\alpha}_{\mathrm{r}}}\left(\beta - \frac{br}{U_{\mathrm{x}}} - \bar{\alpha}_{\mathrm{r}}\right)}{mU_{\mathrm{x}}} - r \tag{A.9}$$

$$\dot{r} \;=\; \frac{aF_{\mathrm{yf}} - b\left[\bar{F}_{\mathrm{yr}} - \bar{C}_{\bar{\alpha}_{\mathrm{r}}}\left(\beta - \frac{br}{U_{\mathrm{x}}} - \bar{\alpha}_{\mathrm{r}}\right)\right]}{I_{\mathrm{zz}}} \tag{A.10}$$

A.2 Position States

The position states of the vehicle, the heading deviation ($\Delta\psi$) and lateral deviation (e), are in reference to a nominal path that need not be obstacle-free.

The equations of motion of the heading deviation and lateral deviation are:

$$\dot{\Delta\psi} \;=\; r \tag{A.11}$$

$$\dot{e} \;=\; U_{\mathrm{x}}\sin\left(\Delta\psi\right) + U_{\mathrm{y}}\cos\left(\Delta\psi\right) \tag{A.12}$$

To approximate the above nonlinear equations as linear functions of the vehicle states, small angle assumptions for $\Delta\psi$ and β are employed to yield:

$$\dot{e} \approx U_{\mathrm{x}}\Delta\psi + U_{\mathrm{x}}\beta \tag{A.13}$$

Thus, (A.9), (A.10), (A.11), and (A.13) combine to produce a continuous state-space representation of the vehicle model:

$$\dot{x} = A_{\mathrm{c}}\left(\bar{\alpha}_{\mathrm{r}}\right)x + B_{\mathrm{c}}F_{\mathrm{yf}} + d_{\mathrm{c}}\left(\bar{\alpha}_{\mathrm{r}}\right) \tag{A.14}$$

with

$$\begin{aligned}
x &= \begin{bmatrix}\beta & r & \Delta\psi & e\end{bmatrix}^{\top} \\
A_{\mathrm{c}}\left(\bar{\alpha}_{r}\right) &= \begin{bmatrix} -\frac{\bar{C}_{\bar{\alpha}_{\mathrm{r}}}}{mU_{\mathrm{x}}} & \frac{b\bar{C}_{\bar{\alpha}_{\mathrm{r}}}}{m{U_{\mathrm{x}}}^{2}} - 1 & 0 & 0 \\ \frac{b\bar{C}_{\bar{\alpha}_{\mathrm{r}}}}{I_{\mathrm{zz}}} & -\frac{b^{2}\bar{C}_{\bar{\alpha}_{\mathrm{r}}}}{I_{\mathrm{zz}}U_{\mathrm{x}}} & 0 & 0 \\ 0 & 1 & 0 & 0 \\ U_{\mathrm{x}} & 0 & U_{\mathrm{x}} & 0 \end{bmatrix} \\
B_{\mathrm{c}} &= \begin{bmatrix}\frac{1}{mU_{\mathrm{x}}} & \frac{a}{I_{\mathrm{zz}}} & 0 & 0\end{bmatrix}^{\top} \\
d_{\mathrm{c}}\left(\bar{\alpha}_{\mathrm{r}}\right) &= \begin{bmatrix}\frac{\bar{F}_{\mathrm{yf}}-\bar{\alpha}_{\mathrm{r}}\bar{C}_{\bar{\alpha}_{\mathrm{r}}}}{mU_{\mathrm{x}}} & -\frac{b\left(\bar{F}_{\mathrm{yf}}-\bar{\alpha}_{\mathrm{r}}\bar{C}_{\bar{\alpha}_{\mathrm{r}}}\right)}{I_{\mathrm{zz}}} & 0 & 0\end{bmatrix}^{\top}
\end{aligned}$$

Here, c denotes a continuous-time model. $A_{\mathrm{c}}\left(\bar{\alpha}_{\mathrm{r}}\right)$ indicates a linearization of A_{c} about $\bar{\alpha}_{\mathrm{r}}$.

Appendix B

Alternative Lane Divider MPC Formulation

The formulation below was presented at the IEEE International Symposium on Ethics in Engineering, Science and Technology of 2016 [108].

For control of lateral tire forces, a model predictive control (MPC) algorithm calculates the vehicle input from an affine vehicle model. The vehicle model used in the MPC formulation is a 4-state bicycle dynamic model with a constant speed assumption. The state vector is comprised of vehicle sideslip angle β, yaw rate r, heading deviation $\Delta\psi$, and lateral deviation e:

$$x = [\beta \quad r \quad \Delta\psi \quad e]^\top.$$

The control input to the vehicle model u, is the lateral tire force which has a nonlinear relationship to the steering angle. However, representing the vehicle input as a lateral tire force results in a linear relationship to the vehicle states:

$$x^{k+1} = A^k x^k + B^k u^k + C^k, \quad k = 1 \dots T \tag{B.1}$$

for each time step in the prediction horizon k up to finite time T.

Path tracking is accomplished by associating a nonzero diagonal entry in the weighting matrix Q to lateral deviation and heading deviation as these states are defined relative to a nominal or desired path.

The safe driving space a vehicle may traverse in the environment is represented via the following constraint

$$H^k_{\text{env}} x^k \leq G^k_{\text{env}}, \quad k = 1 \dots T \tag{B.2}$$

where the vehicle states are bound to stay within an environment that is obstacle-free.

Traffic lane boundaries for a roadway are encoded as

$$H^k_{\text{tra}} x^k \leq G^k_{\text{tra}}, \quad k = 1 \dots T \tag{B.3}$$

where the subscript "tra" represents terms associated with traffic laws.

The complete optimization problem is as follows:

$$\begin{aligned}
\underset{u}{\text{minimize}} \quad & \sum_{k=1}^{T} (x^k)^\top Q^k x^k + (v^k)^\top R^k v^k + \dots \\
& W^k_{\text{env}} \sigma^k_{\text{env}} + W^k_{\text{tra}} (\sigma^k_{\text{tra}})^2 && \text{(B.4a)} \\
\text{subject to} \quad & x^{k+1} = A^k x^k + B^k u^k + C^k && \text{(B.4b)} \\
& H^k_{\text{env}} x^k \leq G^k_{\text{env}} + \sigma^k_{\text{env}} && \text{(B.4c)} \\
& H^k_{\text{tra}} x^k \leq G^k_{\text{tra}} + \sigma^k_{\text{tra}} && \text{(B.4d)} \\
& |u^k| \leq u_{\max} && \text{(B.4e)} \\
& |v^k| \leq v^k_{\max} && \text{(B.4f)}
\end{aligned}$$

where $v^k = u^k - u^{k-1}$ is the change in lateral tire force, σ_{env} and σ_{tra} are slack variables on the constraints enforcing the safe environment and traffic lanes, and $v^k_{\max}$ and $u^k_{\max}$ are physical limits in the steering system and tire forces.

Cost R determines how much weight is placed on change in steering inputs. Cost W_{env} determines how much weight is placed on remaining within the safe environmental envelope; i.e. avoiding collisions with obstacles. Cost W_{tra} determines the weight placed on obeying traffic laws, such as not crossing a double yellow line. The more weight that is placed on the cost terms containing a slack variable, the more the constraint of environmental envelope violation and/or traffic law violation are prioritized in the list of constraints. The traffic law slack variable σ_{tra} is quadratic in the cost function to prevent penalizing relatively small violations of the lane boundary, while the environmental slack variable σ_{env} is linear because even small violations of the environment is undesirable.

Appendix C

Speed Scale POMDP Formulation

The speed control partially observable Markov decision process (POMDP) formulation of Chapter 4 directly commands acceleration values to the vehicle. Because the acceleration commands change every time step, the steering MPC cannot predict how the speed will change throughout the prediction horizon. Thus, Chapter 4 assumes constant speed throughout the prediction horizon. An alternative speed prediction for steering MPC would be to use a speed profile as presented by Funke *et al.* In order to marry POMDP speed control with MPC speed profile predictions, an alternative POMDP formulation would be to scale the desired speed profile rather than command acceleration directly. This appendix demonstrates the speed scaling POMDP formulation through the scenario of a pedestrian crosswalk on a two-lane roadway with a large vehicle occluding the event of a pedestrian crossing as shown in Figure 4.1. Simulation results with this implementation are available on arXiv [109].

The state space is represented in a low dimensional subspace that captures pose and motion of the vehicle as well as perception information. The components of the state considered in this work are velocity of the vehicle (v_t), vehicle distance along the path (d_t), and the event of a pedestrian crossing (c_t). Speed and distance along the path are continuous states. To further reduce the problem size, states v_t and d_t are discretized. The max speed considered for the scenario is 10 m/s with discretization intervals of 1 m/s. The distance along the path is discretized into 0.5 m intervals for a path that is 60 m long. State c_t is already discretized as a binary occurrence.

The vehicle actuation considered here is longitudinal acceleration. Commanded longitudinal acceleration is determined by proportional speed control. Thus, the POMDP action space is a speed scaling factor applied to the desired speed in the longitudinal control. After discretization of the action space, the actions are $\mathcal{A} = \{0\,\%, 10\,\%, 20\,\%, 30\,\%, 40\,\%, 50\,\%, 60\,\%, 70\,\%, 80\,\%, 90\,\%, 100\,\%\}$.

The observation space captures information the agent observes after taking an action. In this work, the observations are provided just from the lidar sensors. Two types of observations are considered number of unobservable tiles (n_t) and the detection of a pedestrian crossing (c_t). To simplify the problem size, the number of unobservable tiles is reduced to 10 discretized bins linearly spaced between 0 and 1800 unobservable tiles. The detection of a pedestrian crossing is handled by a different perception algorithm specifically designed to detect pedestrians. For example, it could be an image recognition algorithm using cameras.

The reward function in this POMDP formulation is designed with the following objectives in mind:

- Encourage the vehicle to drive to the end of the path.
- If a pedestrian is detected, then the vehicle should yield to the pedestrian. Thus, non-zero scale factors are penalized when a pedestrian crossing event is true.
- Additionally, the vehicle should not drive fast when it cannot see the pedestrian.

To achieve the goals outlined, the reward function is specified using two costs:

- **Complete path reward:** The reward for the vehicle to drive to the end of the path +100.
- **Not yielding cost:** The cost when the vehicle does not yield to the crosswalk -50.
- **Too fast cost:** The cost to deter the vehicle from speeding around the occlusion because it is close to a pedestrian crosswalk is set to -5, which is orders of magnitude lower than the collision cost. For this work, it is simply implemented as penalizing the vehicle for going faster than 6 m/s when it cannot see a pedestrian crossing.

The reward is assumed to be zero for all other states.

The dynamics of the system are not actually stochastic, but rather uncertainty is introduced from the crude discretization of the state space. Also, the event of a pedestrian crossing is modeled as a random process. The following parameters characterize the state transition model:

- Speed scaling and changes in speed are not immediately realized in the state space discretization because the vehicle simulation is closer to continuous time.
- Vehicle is assumed to be stopped or moving forward (does not reverse direction).
- If a pedestrian is present within the crosswalk, then the person is assumed to be standing still.

Even though the state in a POMDP is a belief state, the state transition function for a POMDP is the same as for an MDP (assumes no state uncertainty). Even though the state includes truth about a pedestrian crossing event, the problem only maintains a belief about whether there is a pedestrian crossing event based on observations.

In a typical POMDP problem, the observation model is defined as the conditional probability of observing each observation o given the current state s and the action a taken to get there: $\Pr(o \mid s, a)$. For this work, it is assumed the action does not contribute to the observation o given s. Thus, the dependence on a is dropped and the observation model need only specify $\Pr(o \mid s)$. Using the observation space described above, the observation model is the probability of having a number of unobservable tiles and detecting a pedestrian given the current state $\Pr(o_n, o_c \mid s)$. A simple observation model of uniform distribution is implemented if the state where a pedestrian is not crossing is observed. Given a state where the pedestrian is crossing, the observation distribution increases to favor detecting a pedestrian crossing by 30 %.

Bibliography

[1] A. Borning and M. Muller, "Next steps for value sensitive design," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2012, pp. 1125–1134.

[2] "CACI No. 707," in *Judicial Council of California Civil Jury Instructions*, LexisNexis Matthew Bender, 2018.

[3] J. B. Rawlings, "Tutorial overview of model predictive control," *IEEE Control Systems Magazine*, vol. 20, no. 3, pp. 38–52, 2000.

[4] P. Falcone, F. Borrelli, J. Asgari, H. E. Tseng, and D. Hrovat, "Predictive active steering control for autonomous vehicle systems," *IEEE Transactions on Control Systems Technology*, vol. 15, no. 3, pp. 566–580, 2007.

[5] J. Ziegler, P. Bender, T. Dang, and C. Stiller, "Trajectory planning for Bertha–A local, continuous method," in *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2014, pp. 450–457.

[6] Y. Kuwata, J. Teo, S. Karaman, G. Fiore, E. Frazzoli, and J. How, "Motion planning in complex environments using closed-loop prediction," in *AIAA Guidance, Navigation and Control Conference and Exhibit*, 2008.

[7] M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, "Scalable decision making with sensor occlusions for autonomous driving," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2018.

[8] *Urban Challenge Rules*, Defense Advanced Research Projects Agency, 2007.

[9] *Urban Challenge Technical Evaluation Criteria*, Defense Advanced Research Projects Agency, 2006.

[10] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, M. Gittleman, S. Harbaugh, M. Hebert, T. M. Howard, S. Kolski, A. Kelly, M. Likhachev, M. McNaughton, N. Miller, K. Peterson, B. Pilnick, R. Rajkumar, P. Rybski, B. Salesky, Y.-W. Seo, S. Singh, J. Snider, A. Stentz, W. Whittaker, Z. Wolkowicki, J. Ziglar, H. Bae, T. Brown, D. Demitrish, B. Litkouhi, J. Nickolaou, V. Sadekar, W. Zhang, J. Struble, M. Taylor, M. Darms, and D. Ferguson, "Autonomous driving in urban environments: Boss and the Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.

[11] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhnke, D. Johnston, S. Klumpp, D. Langer, A. Levandowski, J. Levinson, J. Marcil, D. Orenstein, J. Paefgen, I. Penny, A. Petrovskaya, M. Pflueger, G. Stanek, D. Stavens, A. Vogt, and S. Thrun, "Junior: The Stanford entry in the Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 9, pp. 569–597, 2008.

[12] A. Bacha, C. Bauman, R. Faruque, M. Fleming, C. Terwelp, C. Reinholtz, D. Hong, A. Wicks, T. Alberi, D. Anderson, S. Cacciola, P. Currier, A. Dalton, J. Farmer, J. Hurdus, S. Kimmel, P. King, A. Taylor, D. V. Covern, and M. Webster, "Odin: Team VictorTango's entry in the DARPA Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 467–492, 2008.

[13] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, O. Koch, Y. Kuwata, D. Moore, E. Olson, S. Peters, J. Teo, R. Truax, M. Walter, D. Barrett, A. Epstein, K. Maheloni, K. Moyer, T. Jones, R. Buckley, M. Antone, R. Galejs, S. Krishnamurthy, and J. Williams, "A perception-driven autonomous urban vehicle," *Journal of Field Robotics*, vol. 25, no. 10, pp. 727–774, 2008.

[14] J. Wei, J. M. Snider, T. Gu, J. M. Dolan, and B. Litkouhi, "A behavioral planning framework for autonomous driving," in *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2014, pp. 458–464.

[15] J. Wei, J. M. Dolan, and B. Litkouhi, "A prediction- and cost function-based algorithm for robust autonomous freeway driving," in *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2010, pp. 512–517.

[16] T. Wongpiromsarn, S. Karaman, and E. Frazzoli, "Synthesis of provably correct controllers for autonomous vehicles in urban environments," in *Proceedings of the IEEE Conference on Intelligent Transportation Systems*, 2011, pp. 1168–1173.

[17] L. I. Reyes Castro, P. Chaudhari, J. Tumova, S. Karaman, E. Frazzoli, and D. Rus, "Incremental sampling-based algorithm for minimum-violation motion planning," in *Proceedings of the IEEE Conference on Decision and Control*, 2013, pp. 3217–3224.

[18] M. Wulfmeier, D. Z. Wang, and I. Posner, "Watch this: Scalable cost-function learning for path planning in urban environments," in *Proceedings of the IEEE and RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 2089–2095.

[19] M. Kuderer, S. Gulati, and W. Burgard, "Learning driving styles for autonomous vehicles from demonstration," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2015, pp. 2641–2646.

[20] S.-H. Lee and S.-W. Seo, "A learning-based framework for handling dilemmas in urban automated driving," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2017, pp. 1436–1442.

[21] J. J. Thomson, "Killing, letting die, and the trolley problem," *The Monist*, vol. 59, no. 2, pp. 204–217, 1976.

[22] J. Li, X. Zhao, M.-J. Cho, W. Ju, and B. F. Malle, "From trolley to autonomous vehicle: Perceptions of responsibility and moral norms in traffic accidents with self-driving cars," SAE, Tech. Rep., 2016.

[23] J.-F. Bonnefon, A. Shariff, and I. Rahwan, "The social dilemma of autonomous vehicles," *Science*, vol. 352, no. 6293, pp. 1573–1576, 2016.

[24] M. A. Wächter, A. Faulhaber, F. Blind, S. Timm, A. Dittmer, L. R. Sütfeld, A. Stephan, G. Pipa, and P. König, "Human decisions in moral dilemmas are largely described by utilitarianism: Virtual car driving study provides guidelines for ADVs," *ArXiv e-prints*, 2017. arXiv: `1706.07332 [cs.CY]`.

[25] P. Wintersberger, A.-K. Prison, A. Riener, and S. Hasirlioglu, "The experience of ethics: Evaluation of self harm risks in automated vehicles," in *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2017, pp. 385–391.

[26] P. Lin. (2014). The robot car of tomorrow may just be programmed to hit you, [Online]. Available: `http://www.wired.com/2014/05/the-robot-car-of-tomorrow-might-just-be-programmed-to-hit-you/`.

[27] T. Fournier, "Will my next car be a libertarian or a utilitarian?: Who will decide?" *IEEE Technology and Society Magazine*, vol. 35, no. 2, pp. 40–45, 2016.

[28] S. Jaiswal, "Autonomous cars: An adaptable feedback mechanism for customised ethics," Master's thesis, Dalhousie University, 2017.

[29] R. Noothigattu, S. S. Gaikwad, E. Awad, S. Dsouza, I. Rahwan, P. Ravikumar, and A. D. Procaccia, "A voting-based system for ethical decision making," *ArXiv e-prints*, 2017. arXiv: `1709.06692 [cs.AI]`.

[30] K. Kinjo and T. Ebina, "Optimal program for autonomous driving under Bentham- and Nash-type social welfare functions," *Procedia Computer Science*, vol. 112, pp. 61–70, 2017.

[31] D. Leben, "A Rawlsian algorithm for autonomous vehicles," *Ethics and Information Technology*, vol. 19, no. 2, pp. 107–115, 2017.

[32] T. Wheeler, "Automotive safety validation in simulation," PhD thesis, Stanford University, 2018.

[33] J. C. Gerdes and S. M. Thornton, "Implementable ethics for autonomous vehicles," in *Autonomes Fahren: Technische, rechtliche und gesellschaftliche Aspekte*, M. Maurer, J. C. Gerdes, B. Lenz, and H. Winner, Eds., Berlin, Heidelberg: Springer, 2015, pp. 87–102.

[34] B. Friedman and P. H. Kahn Jr., "Human values, ethics, and design," in *The Human-Computer Interaction Handbook*, J. A. Jacko and A. Sears, Eds., Mahwah, NJ: Lawrence Erlbaum Associates, 2003, pp. 1177–1201.

[35] B. Friedman, P. H. Kahn Jr., and A. Borning, "Value sensitive design and information systems," in *Human-Computer Interaction and Management Information Systems*, P. Zhang and D. Galletta, Eds., vol. 5, Armonk, NY: M.E. Sharpe, 2006, pp. 348–372.

[36] S. M. Thornton, V. Zhang, S. Varnhagen, and J. C. Gerdes, "Comparative analysis of steering system models in model predictive control of automated vehicles," in *Proceedings of the International Symposium on Advanced Vehicle Control*, 2018.

[37] S. Trimboli, S. Di Cairano, A. Bemporad, and I. V. Kolmanovsky, "Model predictive control with delay compensation for air-to-fuel ratio control," in *Time Delay Systems: Methods, Applications and New Trends*, R. Sipahi, T. Vyhlídal, S.-I. Niculescu, and P. Pepe, Eds., Berlin, Heidelberg: Springer, 2012, pp. 319–330.

[38] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1:43 scale RC cars," *Optimal Control Applications and Methods*, vol. 36, no. 5, pp. 628–647, 2015.

[39] S. Di Cairano, D. Yanakiev, A. Bemporad, I. V. Kolmanovsky, and D. Hrovat, "An MPC design flow for automotive control and applications to idle speed regulation," in *Proceedings of the IEEE Conference on Decision and Control*, 2008, pp. 5686–5691.

[40] J. Funke, M. Brown, S. M. Erlien, and J. C. Gerdes, "Prioritizing collision avoidance and vehicle stabilization for autonomous vehicles," in *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2015, pp. 1134–1139.

[41] E. Fiala, "Lateral forces on rolling pneumatic tires," *Zeitschrift VDI*, vol. 96, no. 29, pp. 973–979, 1954.

[42] H. Pacejka, *Tire and Vehicle Dynamics*. Elsevier, 2012.

[43] S. M. Erlien, J. Funke, and J. C. Gerdes, "Incorporating non-linear tire dynamics into a convex approach to shared steering control," in *Proceedings of the American Control Conference*, 2014, pp. 3468–3473.

[44] V. Zhang, S. M. Thornton, and J. C. Gerdes, "Tire modeling to enable model predictive control of automated vehicles from standstill to the limits of handling," in *Proceedings of the International Symposium on Advanced Vehicle Control*, 2018.

[45] J. Mattingley, Y. Wang, and S. Boyd, "Code generation for receding horizon control," in *Proceedings of the IEEE International Symposium on Computer-Aided Control System Design*, 2010, pp. 985–992.

[46] J. Mattingley and S. Boyd, "CVXGEN: A code generator for embedded convex optimization," *Optimization and Engineering*, vol. 13, no. 1, pp. 1–27, 2012.

[47] S. M. Erlien, S. Fujita, and J. C. Gerdes, "Shared steering control using safe envelopes for obstacle avoidance and vehicle stability," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 2, pp. 441–451, 2016.

[48] S. M. Thornton, S. Pan, S. M. Erlien, and J. C. Gerdes, "Incorporating ethical considerations into automated vehicle control," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 6, pp. 193–241, 2017.

[49] J. Haidt and C. Joseph, "Intuitive ethics: How innately prepared intuitions generate culturally variable virtues," *Daedalus*, vol. 133, no. 4, pp. 55–66, 2004.

[50] J. Haidt, J. Graham, and C. Joseph, "Above and below left–right: Ideological narratives and moral foundations," *Psychological Inquiry*, vol. 20, no. 2-3, pp. 110–119, 2009.

[51] R. Shweder, N. Much, M. Mahapatra, and L. Park, “The ‘big three’ of morality (autonomy, community, divinity) and the ‘big three’ explanations of suffering,” in *Morality and Health*, A. M. Brandt and P. Rozin, Eds., New York, London: Routledge, 1997, pp. 119–169.

[52] J. Graham, J. Haidt, and B. A. Nosek, “Liberals and conservatives rely on different sets of moral foundations,” *Journal of Personality and Social Psychology*, vol. 96, no. 5, pp. 1029–1046, 2009.

[53] M. N. Mladenovic and T. McPherson, “Engineering social justice into traffic control for self-driving vehicles?” *Science and Engineering Ethics*, vol. 22, no. 4, pp. 1131–1149, 2016.

[54] K. W. Miller, M. J. Wolf, and F. Grodzinsky, “This “ethical trap” is for roboticists, not robots: On the issue of artificial agent ethical decision-making,” *Science and Engineering Ethics*, vol. 23, no. 2, pp. 389–401, 2017.

[55] J. Van den Hoven, G.-J. Lokhorst, and I. Van de Poel, “Engineering and the problem of moral overload,” *Science and Engineering Ethics*, vol. 18, no. 1, pp. 143–155, 2012.

[56] P. Lin, “Why ethics matters for autonomous cars,” in *Autonomes Fahren: Technische, rechtliche und gesellschaftliche Aspekte*, M. Maurer, J. C. Gerdes, B. Lenz, and H. Winner, Eds., Berlin, Heidelberg: Springer, 2015, pp. 69–85.

[57] P. Lin, K. Abney, and G. A. Bekey, *Robot Ethics: The Ethical and Social Implications of Robotics*. MIT Press, 2011.

[58] W. Wallach and C. Allen, *Moral Machines: Teaching Robots Right From Wrong*. Oxford University Press, 2008.

[59] N. J. Goodall, “Machine ethics and automated vehicles,” in *Road Vehicle Automation*, G. Meyer and S. Beiker, Eds., Cham: Springer, 2014, pp. 93–102.

[60] A. Gray, Y. Gao, T. Lin, J. K. Hedrick, H. E. Tseng, and F. Borrelli, “Predictive control for agile semi-autonomous ground vehicles using motion primitives,” in *Proceedings of the American Control Conference*, 2012, pp. 4239–4244.

[61] A. Gray, Y. Gao, T. Lin, J. K. Hedrick, and F. Borrelli, "Stochastic predictive control for semi-autonomous vehicles with an uncertain driver model," in *Proceedings of the IEEE Conference on Intelligent Transportation Systems*, 2013, pp. 2329–2334.

[62] Y. Gao, A. Gray, J. V. Frasch, T. Lin, E. Tseng, J. K. Hedrick, and F. Borrelli, "Spatial predictive control for agile semi-autonomous ground vehicles," in *Proceedings of the International Symposium on Advanced Vehicle Control*, 2012.

[63] S. M. Erlien, S. Fujita, and J. C. Gerdes, "Safe driving envelopes for shared control of ground vehicles," in *Proceedings of the IFAC Symposium on Advances in Automotive Control*, 2013, pp. 831–836.

[64] R. Hursthouse, "Virtue ethics," in *The Stanford Encyclopedia of Philosophy*, E. N. Zalta, Ed., Fall 2013, Metaphysics Research Lab, Stanford University, 2013.

[65] G. Harman, "Skepticism about character traits," *The Journal of Ethics*, vol. 13, no. 2-3, pp. 235–242, 2009.

[66] I. Asimov, *I, Robot*. Gnome Press, 1950.

[67] J. Driver, "The history of utilitarianism," in *The Stanford Encyclopedia of Philosophy*, E. N. Zalta, Ed., Winter 2014, Metaphysics Research Lab, Stanford University, 2014.

[68] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.

[69] S.-H. Suh and A. B. Bishop, "Collision-avoidance trajectory planning using tube concept: Analysis and simulation," *Journal of Robotic Systems*, vol. 5, no. 6, pp. 497–525, 1988.

[70] C. E. Beal and J. C. Gerdes, "Model predictive control for vehicle stabilization at the limits of handling," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 4, pp. 1258–1269, 2013.

[71] C. G. Bobier and J. C. Gerdes, "Staying within the nullcline boundary for vehicle envelope control using a sliding surface," *Vehicle System Dynamics*, vol. 51, no. 2, pp. 199–217, 2013.

[72] J. D. Greene, "The secret joke of kant's soul," in *Moral Psychology*, W. Sinnott-Armstrong, Ed., vol. 3, MIT Press, 2008, pp. 35–79.

[73] R. R. Radtke, "Role morality in the accounting profession–how do we compare to physicians and attorneys?" *Journal of Business Ethics*, vol. 79, no. 3, pp. 279–297, 2008.

[74] A. I. Applbaum, *Ethics for Adversaries: The Morality of Roles in Public and Professional Life*. Princeton University Press, 2000.

[75] S. Vallor, *Philosophy and law panel*, Remarks by Shannon Vallor at the "Towards Programming Ethics in Automated Vehicles" Workshop at Stanford University, Jun. 2015.

[76] J. Millar, "An ethics evaluation tool for automating ethical decision-making in robots and self-driving cars," *Applied Artificial Intelligence*, vol. 30, no. 8, pp. 787–809, 2016.

[77] S. M. Thornton, F. E. Lewis, V. Zhang, M. J. Kochenderfer, and J. C. Gerdes, "Value sensitive design for autonomous vehicle motion planning," in *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2018.

[78] NHTSA's National Center for Statistics and Analysis, "How vehicle age and model year relate to driver injury severity in fatal crashes," U.S. Department of Transportation, Tech. Rep., 2013.

[79] *Federal Automated Vehicles Policy*, U.S. Department of Transportation, 2016.

[80] M. Maguire, "Methods to support human-centred design," *International Journal of Human-Computer Studies*, vol. 55, no. 4, pp. 587–634, 2001.

[81] J. Giacomin, "What is human centred design?" *Design Journal*, vol. 17, no. 4, pp. 606–623, 2014.

[82] M. Niemelä, V. Ikonen, J. Leikas, K. Kantola, M. Kulju, A. Tammela, and M. Ylikauppila, "Human-driven design: A human-driven approach to the design of technology," in *Proceedings of the IFIP International Conference on Human Choice and Computers*, 2014, pp. 78–91.

[83] J. Leikas, P. Saariluoma, J. Heinilä, and M. Ylikauppila, "A methodological model for life-based design," *International Review of Social Sciences and Humanities*, vol. 4, no. 2, pp. 118–136, 2013.

[84] B. Friedman, P. H. Kahn Jr., and A. Borning, "Value sensitive design: Theory and methods," University of Washington, Tech. Rep., 2002.

[85] T. Denning, A. Borning, B. Friedman, B. T. Gill, T. Kohno, and W. H. Maisel, "Patients, pacemakers, and implantable defibrillators," in *Proceedings of the International Conference on Human Factors in Computing Systems*, 2010, pp. 917–926.

[86] A. van Wynsberghe, "Designing robots for care: Care centered value-sensitive design," *Science and Engineering Ethics*, vol. 19, no. 2, pp. 407–433, 2013.

[87] B. Chen, D. Zhao, and H. Peng, "Evaluation of automated vehicles encountering pedestrians at unsignalized crossings," in *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2017, pp. 1679–1685.

[88] T. Bandyopadhyay, C. Z. Jie, D. Hsu, M. H. Ang, D. Rus, and E. Frazzoli, "Intention-aware pedestrian avoidance," in *Experimental Robotics: The 13th International Symposium on Experimental Robotics*, J. P. Desai, G. Dudek, O. Khatib, and V. Kumar, Eds., Heidelberg: Springer, 2013, pp. 963–977.

[89] T. Bandyopadhyay, K. Won, E. Frazzoli, D. Hsu, W. Lee, and D. Rus, "Intention-aware motion planning," in *Algorithmic Foundations of Robotics X*, 2013, pp. 475–491.

[90] S. Brechtel, T. Gindele, and R. Dillmann, "Probabilistic decision-making under uncertainty for autonomous driving using continuous pomdps," in *Proceedings of the IEEE Conference on Intelligent Transportation Systems*, 2014, pp. 392–399.

[91] J. Haidt, *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. Vintage, 2012.

[92] J. K. Choi and Y. G. Ji, "Investigating the importance of trust on adopting an autonomous vehicle," *International Journal of Human-Computer Interaction*, vol. 31, no. 10, pp. 692–702, 2015.

[93] *California Driver Handbook*, State of California Department of Motor Vehicles, 2017.

[94] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.

[95] M. J. Kochenderfer, *Decision Making Under Uncertainty: Theory and Application*. MIT Press, 2015.

[96] S. Davies, "Multidimensional triangulation and interpolation for reinforcement learning," in *Proceedings of the International Conference on Neural Information Processing Systems*, 1997, pp. 1005–1011.

[97] M. Brown, J. Funke, S. Erlien, and J. C. Gerdes, "Safe driving envelopes for path tracking in autonomous vehicles," *Control Engineering Practice*, vol. 61, pp. 307–316, 2017.

[98] M. Egorov, Z. N. Sunberg, E. Balaban, T. A. Wheeler, J. K. Gupta, and M. J. Kochenderfer, "POMDPs.jl: A framework for sequential decision making under uncertainty," *Journal of Machine Learning Research*, vol. 18, no. 26, pp. 1–5, 2017.

[99] B. J. Schroeder and N. M. Rouphail, "Event-based modeling of driver yielding behavior at unsignalized crosswalks," *Journal of Transportation Engineering*, vol. 137, no. 7, pp. 455–465, 2011.

[100] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6517–6525.

[101] B. C. Tefft, "Impact speed and a pedestrian's risk of severe injury or death," AAA Foundation for Traffic Safety, Tech. Rep., 2011.

[102] P. S. Bokare and A. K. Maurya, "Study of effect of speed, acceleration and deceleration of small petrol car on its tail pipe," *International Journal for Traffic and Transport Engineering*, vol. 3, no. 4, pp. 465–478, 2013.

[103] F. Soriguera, I. Martínez, M. Sala, and M. Menéndez, "Effects of low speed limits on freeway traffic flow," *Transportation Research Part C: Emerging Technologies*, vol. 77, pp. 257–274, 2017.

[104] N. J. Goodall, "Away from trolley problems and toward risk management," *Applied Artificial Intelligence*, vol. 30, no. 8, pp. 810–821, 2016.

[105] *Automated Driving Systems 2.0: A Vision for Safety*, National Highway Traffic Safety Administration, 2017.

[106] National Highway Traffic Safety Administration. (2017). What is NHTSA's approach to ethical considerations? [Online]. Available: `https://www.nhtsa.gov/manufacturers/automated-driving-systems` (visited on 11/22/2017).

[107] *Ethics Commission: Automated and Connected Driving*, Bundesministerium für Verkehr und digitale Infrastruktur (BMVi), 2017.

[108] S. Pan, S. M. Thornton, and J. C. Gerdes, "Prescriptive and proscriptive moral regulation for autonomous vehicles in approach and avoidance," in *Proceedings of the IEEE International Symposium on Ethics in Engineering, Science and Technology*, 2016.

[109] S. M. Thornton, "Autonomous vehicle speed control for safe navigation of occluded pedestrian crosswalk," *ArXiv e-prints*, 2018. arXiv: `1802.06314 [cs.RO]`.