

Formalizing Human-Machine Communication in the Context of Autonomous Vehicles

May 2020 | Final Report



Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by a grant from the U.S. Department of Transportation's University Transportation Centers Program. However, the U.S. Government assumes no liability for the contents or use thereof.

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. 02-014	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Formalizing Human Machine Communication in the Context of Autonomous Vehicles		5. Report Date May 2020	
		6. Performing Organization Code:	
7. Author(s) Swaminathan Gopalswamy Srikanth Saripalli Dylan Shell Jeff Hickman Ya-Chuan Hsu		8. Performing Organization Report No. Report 02-014	
		9. Performing Organization Name and Address: Safe-D National UTC Texas A&M University; Virginia Tech Transportation Institute	
12. Sponsoring Agency Name and Address Office of the Secretary of Transportation (OST) U.S. Department of Transportation (US DOT) State of Texas		10. Work Unit No.	
		11. Contract or Grant No. 69A3551747115/Project 02-014	
15. Supplementary Notes This project was funded by the Safety through Disruption (Safe-D) National University Transportation Center, a grant from the U.S. Department of Transportation – Office of the Assistant Secretary for Research and Technology, University Transportation Centers Program, and, in part, with general revenue funds from the State of Texas.		13. Type of Report and Period Final Research Report	
		14. Sponsoring Agency Code	
16. Abstract There are many situations where tacit communication between drivers and pedestrians governs and enhances safety. The goal of this study was to formalize this communication and apply it to the driving strategy of an autonomous vehicle. Toward this, we performed a field study of the interaction between drivers and pedestrians. Vehicles were instrumented to capture behavioral information on a driver as well as passengers and the traffic scenario in general. The data captured were reduced by data analysts to provide insights into the communication and driving patterns. The categorical reduction on driver, pedestrian, and environmental variables was captured. A domain specific language (DSL) was developed to precisely describe the driver-pedestrian behavior, toward the development of a behavioral model for generating autonomous vehicle controls. Specifically, interaction was formalized through a probabilistic model, namely a partially observable Markov decision process (POMDP). This enabled study of what-if scenarios with different risk averseness characteristics. One particular strategy was implemented on an autonomous vehicle and experimental observations were made. Future work will consider (i) richer DSLs to better quantify the driver-human communication, (ii) faster POMDP solvers for real-time operation, and (iii) further applications.			
17. Key Words Driver Pedestrian Communication, Vehicle Pedestrian Communication, Human-Machine-Interaction, Human-Robot-Interaction, MDP, POMDP, DSL, Naturalistic Driving Data		18. Distribution Statement No restrictions. This document is available to the public through the Safe-D National UTC website , as well as the following repositories: VTechWorks , The National Transportation Library , The Transportation Library , Volpe National Transportation Systems Center , Federal Highway Administration Research Library , and the National Technical Reports Library .	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 27	22. Price \$0

Abstract

There are many situations where tacit communication between drivers and pedestrians governs and enhances safety. The goal of this study was to formalize this communication and apply it to the driving strategy of an autonomous vehicle. Toward this, we performed a field study of the interaction between drivers and pedestrians. Vehicles were instrumented to capture behavioral information on a driver as well as passengers and the traffic scenario in general. The data captured were reduced by data analysts to provide insights into the communication and driving patterns. The categorical reduction on driver, pedestrian, and environmental variables was captured. A domain specific language (DSL) was developed to precisely describe the driver-pedestrian behavior, toward the development of a behavioral model for generating autonomous vehicle controls. Specifically, interaction was formalized through a probabilistic model, namely a partially observable Markov decision process (POMDP). This enabled study of what-if scenarios with different risk averseness characteristics. One particular strategy was implemented on an autonomous vehicle and experimental observations were made. Future work will consider (i) richer DSLs to better quantify the driver-human communication, (ii) faster POMDP solvers for real-time operation, and (iii) further applications.

Acknowledgements

This project was funded by the Safety through Disruption (Safe-D) National University Transportation Center, a grant from the U.S. Department of Transportation – Office of the Assistant Secretary for Research and Technology, University Transportation Centers Program, and, in part, with general revenue funds from the State of Texas. Additionally we acknowledge the detailed comments from SME reviewers that have tremendously helped shape the final report.

Table of Contents

TABLE OF CONTENTS	III
LIST OF FIGURES	V
LIST OF TABLES	V
INTRODUCTION	1
Original Research Proposal	1
Research Focus	2
FIELD STUDY ON DRIVER-PEDESTRIAN INTERACTIONS	2
Participants	2
Equipment	2
Study Design	2
Results	3
Descriptive Variables	3
Predictors of Driver Yield Behavior	4
Discussion	4
Summary	5
FORMALIZING DRIVER-PEDESTRIAN COMMUNICATIONS TO SUPPORT AUTONOMOUS DRIVING.....	5
Literature Review	5
Studies of Pedestrian Crossing Behavior: Definitions and Known Determinants	6
Prior Studies of Pedestrian Interaction and Cooperation with AVs	6
Domain Specific Language (DSL) for Capturing Communication Events	7
DSL Framework	7
Key Components of DSL for Describing the Communication Scenario	7
Sample DSL Implementation	8
Code Generation from DSL to Implementation in Vehicle	9

Crossing Intent Model	9
Formal Problem Definition	10
Vehicle-Pedestrian Interaction Model	11
Optimal Action Policy for Autonomous Driving	13
Vehicle Dynamics.....	13
Rewards	13
Evaluation	14
Simulation Setup.....	14
Results.....	15
 CONCLUSIONS AND RECOMMENDATIONS	 18
Recommendations for Future Work.....	18
 ADDITIONAL PRODUCTS	 19
Education and Workforce Development Products	19
Technology Transfer Products	20
Data Products.....	20
 REFERENCES	 21
 APPENDIX A: FIELD STUDY RESULTS	 24
 APPENDIX B: EXPLICIT COMMUNICATION SIMULATION RESULTS:	 27

List of Figures

Figure 1. Screen capture. Example of communication scenario DSL editor 8

Figure 2. Screen capture. Projection of communication DSL as CNL 9

Figure 3. Diagram. Bird’s eye view of crossing scenario..... 10

Figure 4. Graphs. Simulation results showing (a) vehicle executing a policy, interacting with a cautious pedestrian, and (b) vehicle executing a policy, interacting with a reckless pedestrian. . 16

Figure 5. Chart. Strategies of the vehicle in ambiguous situations with a reckless pedestrian (based on [13] but redrawn with modifications) 17

Figure 6. Reckless pedestrian interacts with vehicle equipped to flash its lights as a form of explicit communication..... 27

List of Tables

Table 1. Example of Comparison of Rewards between Policies 14

Table 2. Results of Chi-Square Tests Between the Driving Context Condition and Behavioral Variables 24

Table 3. Results of Chi-Square Tests Between the Narration Condition and Behavioral Variables 25

Table 4: Model 2 Results–Naturalistic Driving 26

Introduction

Recent studies, such as Lavasani et al. [1], predict that more than 80% of the vehicles sold by 2050 will be autonomous. Additionally, studies suggest that the transformation of transportation from manually operated to autonomous vehicles (AVs) promises significant safety benefits. For example, Fagnant et al. [2] predict that a 90% market penetration of AVs will lead to 21,700 lives saved per year and 4.2 million fewer crashes.

One crucial safety concern in the development of AVs is their interactions with pedestrians. Current AV designs gather contextual information on pedestrians, such as their location, speed, and direction of motion from a variety of sensors (see Gandhi and Trivedi [2]). However, such information does not explicitly factor in any communication between the pedestrian and vehicle, such as when drivers and pedestrians make eye contact, or the driver waves the pedestrian through. This report presents a series of efforts to study pedestrian-driver communication and translate that data into a model that can be used by AVs.

Original Research Proposal

As originally proposed to the Safety Through Disruption (Safe-D) University Transportation Center (UTC), this research involved three phases:

1. Human Behavior Analysis – Using an instrumented vehicle, collect data on how human drivers and pedestrians use tacit communication at crosswalks.
2. Domain Specific Language – Abstract and codify the data collected during the field study in a domain specific language (DSL).
3. Autonomous Vehicle as a Technology Demonstrator Platform – Use an autonomous Polaris GEM e4 electric golf cart to integrate and evaluate the code derived from the DSL developed in the previous phase.

However, it quickly became apparent that the data collection and human behavior analysis proposed in Phase 1 would take significant effort and time. Thus, performing Phases 2 and 3 in series to directly incorporate the results of Phase 1 into the AV experiments would result in the project exceeding its planned duration. Additionally, Phase 2 was addressing an open research question that required exploration of ideas, which could be done independent of Phase 1. This led to a revised research plan that involved conducting the field study in parallel with the other efforts. As a result, the data from the field study were not available to inform the development of the DSL.

For Phase 3, while we started the AV technology demonstrator platform on the Polaris golf cart (including initial experimentation with explicit communication devices such as programmable LEDs), our research informed us that it was more important to consider implicit rather than explicit communication. A majority of the effort was shifted to addressing this question of implicit

communication. Such implicit communication requires a relatively precise understanding of the state of the pedestrian, which is a complex task, still subject to open research, and outside the scope of our original proposal. Such a capability was more readily available from sensors mounted on the roadside. In order to leverage this capability, we moved to a Lincoln MKZ as a demonstration platform. However this capability could not be realized before the completion of the project, and therefore Phase 3 was performed through simulations. (A demonstration of the capabilities using the MKZ platform was performed after completion of the contracted work.)

Research Focus

The objective of this research, as revised, was to explore the communication that happens between drivers and pedestrians, and develop a method to apply that to autonomous decision-making that will mimic human driver behavior. Prior work has explored communication with pedestrians using a technology focus. For example, Anaya et al. [6] developed a Wi-Fi warning system for pedestrians and determined the minimum communication distance that would provide adequate reaction time. Our focus is to infer the communication intent of the pedestrian through probabilistic models defining the interaction. This was accomplished through a field study and a three-step process to formalize driver-pedestrian communication to support autonomous driving.

Field Study on Driver-Pedestrian Interactions

The field study collected data on driver-pedestrian interactions at crosswalks.

Participants

The field study collected data from 11 drivers: six females and five males. Due to an equipment failure, one female driver had incomplete data and therefore the final sample consisted of five females and five males. The drivers were recruited by gender to account for possible gender differences in driving behavior. Participants received research credit for their participation. One hour of research credit was awarded for each hour of participation.

Equipment

The vehicle for the field study was a 2008 Cadillac STS. Data were collected from the Controller Area Network (CAN) bus and a MiniDAS data acquisition system (DAS). The MiniDAS is a compact data collection instrument that attaches easily to the vehicle windshield, providing continuous forward video of the roadway, along with video of the driver via an internal facing camera. In addition, the MiniDAS collects audio and accelerometer data.

Study Design

The study followed a 3-way factorial design to explore tacit communication between vehicles and pedestrians in high-density pedestrian traffic. The three variables from the 3-way factorial each had two levels, providing a $2 \times 2 \times 2$ full factorial study design. The three independent variables were (1) driving context, (2) driving route, and (3) narration. The two levels of driving context

were driving the experimental vehicle with and without a sign reading “self-driving vehicle.” Driving route included one route primarily covering unsignalized crosswalks and a second route primarily covering signalized crosswalks. The narration condition included driving the vehicle with no special instructions and driving the vehicle while performing a think-aloud protocol.

During the think-aloud protocol, participants were instructed to narrate whatever they were thinking about while reacting to external stimuli. For example, if the driver altered their behavior in response to a pedestrian entering the roadway, they were requested to talk through the decision. We hypothesized that the think-aloud protocol would provide a small window into the participant’s decision-making during pedestrian interactions. Unfortunately, most participants failed to convey their thoughts during the narration condition (we suspect due to poor training/comprehension).

We attempted to minimize ordering effects by having all participants drive both routes under each driving context without performing the think-aloud protocol. The first four drives involved no narration, with the routes and driving context randomized. Subsequently, the driver performed the same order for the routes and driving context while performing the narration protocol. This selection was made because performing without the think-aloud protocol is unlikely to affect performance in the think-aloud protocol. However, if the think-aloud protocol had been introduced first, the participants may have experienced a carryover into the other drives.

Results

An event of interest was defined as occurring when a pedestrian was less than or equal to 10 feet from a crosswalk while the experimental vehicle was less than or equal to 50 feet from the crosswalk. From the video, a total of 1,738 cases were appropriate for analysis. Results are reported for the descriptive analysis and regression models. The descriptive analysis details the pedestrian demographics and the influences of the experimental conditions on pedestrian and driver behavior. The regression analyses identify predictors of driver yield behavior.

Descriptive Variables

For the driving context, just over half of the events occurred when the sign was placed on the car (51.2%). A total of nine variables were investigated using the chi-square test for independence. Of the variables investigated, only the position of the driver’s hands on the wheel, $\chi^2(2, n = 1,738) = 41.7, p < .05$, and the driver waving pedestrians through, $\chi^2(1, n = 1,738) = 4.2, p < .05$, were significant. No dependence was detected between the driving context and yield behavior, $\chi^2(1, n = 1,738) = 1.4, p > .05$. All other behavioral variables investigated, including all pedestrian behavior variables, were non-significant, indicating no interdependence in the data. Full results are reported in Table 1 in Appendix A.

Events within the narration condition were split, with just over half of the events occurring when the participant was thinking aloud (50.6%). Driver yield behavior was independent from the narration condition, $\chi^2(1, n = 1,738) = 0.1, p > .05$. Driver eye contact had a significant chi-square test, indicating dependence with the narration condition, $\chi^2(1, n = 1,738) = 7.8, p < .05$. All other

behavioral variables were found to be non-significant. Full results are reported in Table 2 in Appendix A.

Predictors of Driver Yield Behavior

A multilevel logistic regression analysis was conducted to investigate variables of interest that predict driver yield behavior. Only odds ratio estimates that approximate large and medium effect sizes were considered for interpretation. Based on previous research by Chen et al. (2010), the odds ratio estimates for large and medium effect size for the prevalence rate of the dependent variable are 4.1387 and 2.4972, respectively. The intraclass correlation was calculated from an unconditional means model and demonstrated that 15.5% of the variance is due to driver effects. The current analysis focuses on the remaining 84.5% of the variance attributed to event effects. Predictors were selected for inclusion if every category included at least 5% of the overall data. All continuous predictors were person-mean centered to control for driver effects.

The odds ratio estimates for six categorical predictors approximated large effects: (1) stop sign compared to no traffic control, (2) the driver making eye contact compared to no eye contact, (3) a factor affecting the driver's path compared to no path obstruction, (4) pedestrian distance to crossing 0 feet compared to being in the crossing, (5) pedestrian distance to crossing 0 to 5 feet compared to being in the crossing, and (6) pedestrian distance to crossing 5 to 10 feet compared to being in the crossing. The odds ratio estimate for an additional three predictors approximated a medium effect: (1) pedestrian making eye contact compared to no eye contact, (2) a vehicle traveling in the opposing lane compared to no vehicle, and (3) a factor affecting the pedestrian's path compared to no path impairment. The full results are presented in Table 3 in the Appendix A.

Discussion

The comparison of pedestrian and driver behavior related to the driving context provided some interesting results. Here, the drivers were aware that the vehicle was marked with a sign reading "self-driving vehicle" and altered their visible behavior. Changes in steering wheel hand position toward the bottom and a reduction in waving pedestrians through can be interpreted as attempts by the participants to conform to the study expectations. Despite the changes in visible behavior, the safety criterion, yielding behavior, was not impacted by the sign condition.

Another interesting finding was that pedestrians did not change any of their behavior when interacting with the vehicle when it was marked as a "self-driving vehicle." The overall statistical test did not indicate any mean differences. However, during driving sessions the experimenter noticed some pedestrians had extreme reactions to the vehicle when it was equipped with the "self-driving vehicle" sign. These included a pedestrian starting to walk into the crosswalk, followed by quickly returning to the curb before circling behind the vehicle to cross. However, an AV may or may not have an individual seated in the "driver's seat"; thus, these results may not characterize a scenario when the driver's seat is unattended.

Driver eye glance toward the pedestrian was influenced by the narration condition. The decrease in driver eye glance toward pedestrians is likely due to an increase in cognitive load. This finding can also be connected to past research on roadway eye glance for drivers using a hands-free phone. The general finding indicates that drivers using a hands-free phone spend more time looking at the roadway compared to a baseline measure. The overall conclusion is usually that this presents a protective effect on crash risk by reducing eye glance away from the roadway. In this case, eye glance away from the roadway may be a crash risk factor.

The predictive model demonstrated that pedestrian distance and eye glance are key indicators for SAE Level 5 AV programming to consider. In addition, driver eye-glance behavior indicates the importance of considering how SAE Level 4 and 5 AVs indicate their intentions. The remaining predictors are situational factors. The path obstruction of the vehicle and pedestrian should also be considered for programming.

Summary

To maximize safety and transportation efficiency, it is essential for Level 5 AVs to be programmed to understand and interact with pedestrians during roadway crossing situations. Ten participants each completed eight drives in a 3-way factorial designed study. The research provided nearly 90 hours of real-world video data. Differences in pedestrian and driver behavior for the independent variables indicated that pedestrians do not demonstrate significant behavior changes in response to seeing a vehicle marked as a “self-driving vehicle.” In addition, the predictive model identified variables for future research to improve SAE Level 5 AV programming by including consideration for pedestrian-vehicle interactions.

Formalizing Driver-Pedestrian Communications to Support Autonomous Driving

The second and third phases of this research involved formalizing driver-pedestrian communications to support autonomous driving. This involved the development of a DSL, a crossing intent model, and an action policy for an AV. The crossing intent model and action model were subsequently tested through simulation.

Literature Review

In the original proposal for this project, the mathematical formalization of driver-pedestrian communications was intended to be informed by the results of the field study. However, the field study was performed in parallel to the rest of the project, including the development of the DSL, and therefore there were no results available in time to use the DSL to abstract the field study data. Consequently, the formalization effort was based on a review of the literature, including studies of pedestrian crossing behavior and studies of pedestrian-AV interactions.

Studies of Pedestrian Crossing Behavior: Definitions and Known Determinants

Early studies that analyzed pedestrian road-crossing behaviors observed that pedestrians are primarily concerned with time gaps instead of distance gaps when negotiating approaching traffic [13]. According to the *Highway Capacity Manual* [15], a *critical gap* is defined to be the time in seconds below which a pedestrian will not attempt to begin crossing the street. R. L. Moore [13] identified that each pedestrian has their own critical distance gap, but this gap changes according to an oncoming vehicle's speed, and the pedestrian will not cross if the vehicle is nearer than this critical distance. Cohen et al. [16] also proposed a similar idea called maximum risk-taking. They found that 92% of pedestrians crossed a 7.0-m wide road when the vehicle was 7 s away, and 0% crossed when the time difference was shorter than 1.5 s.

Several researchers followed up on these studies with a thorough evaluation of the factors that may influence the critical gap. Brewer et al. [17] categorized pedestrians' crossing maneuvers based on different traffic flow and road geometric conditions: (1) a *single-stage crossing* is when pedestrians cross the road in one crossing maneuver; (2) *two-stage* is when pedestrians cross to the median first and then cross to the far side; (3) *rolling* is when pedestrians search for gaps between a continuous flow of vehicles by adjusting the speed and direction of their movements. W. A. Harrell [18] analyzed crossing behaviors with variables for multiple parameters, including traffic volume, temperature, and the width of the roadway. He concluded that traffic volume has an inverse relationship with cautiousness.

In addition, it has been observed that personal characteristics influence pedestrians' critical gap. Studies have shown that gender affects pedestrian behavior, most results indicating that males tend to take riskier actions, whereas females often cross with greater caution [19, 20]. Age is another influence on pedestrian behavior. Oxley et al. [21] suggest that age-related perceptual and cognitive deficits affect crossing behavior. Moreover, the complexity of the traffic has a larger effect on the behavior of older pedestrians. For example, on two-way undivided roads, elderly people are frequently found crossing even when the traffic is already closing up [23].

Prior Studies of Pedestrian Interaction and Cooperation with AVs

Researchers have proposed different approaches to predict pedestrian behavior to enable planning and action-selection for AVs. Schneemann and Heinemann [23] classified pedestrian intention using a combination of a support vector machine (SVM) along with a context-based feature descriptor. Bandyopadhyay [24] used a mixed observable Markov decision process to maintain uncertainty over pedestrians' headings to generate a conservative avoidance policy. Most of these studies focus on avoiding a collision by relying on collecting pedestrian data.

The result is that the vehicle's action choices come to be dominated by the pedestrian's behavior which, while guarded, can undermine social resolution of contention. For example, those approaches have the potential to make rolling crossing dangerous or entirely infeasible. Moreover, such an approach can be counterproductive, reducing overall efficiency and encouraging (or, if not exactly rewarding, then certainly never penalizing) interposition by pedestrians.

There are other studies that focus on communicating the vehicle's intentions to nearby pedestrians. No single solution dominates, and a variety of different communication methods for AVs and pedestrians have been explored. Some studies state that showing physical information such as gap distance dominates the communication [5]. Others suggest that designing external interaction devices can help reassure the pedestrian that it is safe to cross. Lagstrom and Lundgren [25] conducted experiments with an LED strip mounted on the windshield of a car and provided evidence that it boosts confidence in pedestrian decision-making. Companies such as Google [26], Mercedes [27], and Nissan [28] have also proposed their own external hardware (such as programmable LED displays and speakers) to interact with pedestrians. In contrast, in this study we deliberately opted to use only standard production vehicle features.

Domain Specific Language (DSL) for Capturing Communication Events

DSL Framework

The focus of this task was to explore the use of a controlled natural language (CNL) representation of events relevant to driver-pedestrian communication. There has been a significant interest in the use of CNL for knowledge representation in general (e.g., Schwitter [9]). The primary reason to look at CNL is to provide a mechanism to describe communication events precisely yet in a human-readable form.

A CNL is often developed to not just describe knowledge but also enable programmatic access to this knowledge for the purpose of automation. In this case, the description evolves into a domain-specific language (DSL). A DSL is a more general concept than a CNL, with well-defined syntax and semantics, and well-defined transformations that can eventually enable analysis and execution directly from the DSL descriptions (for example, see Volter [10]).

An editor is required to describe the communication events in the DSL, along with associated tools that can analyze the syntax and semantic constraints of (i.e., compile) the DSL. This is normally a huge task, and therefore we decided to use a language work bench (LWB) [11] to aid in the development of the DSL. We chose MPS (Meta Programming System) [12] as a suitable LWB because it allows relatively easy construction of a CNL-type editor without compromising the ability to define the underlying semantics of the language.

Key Components of DSL for Describing the Communication Scenario

Based on the literature review and engineering intuition, we selected key components to develop the first iteration of the DSL. Thus, it was not possible to comprehensively identify all the key components of the DSL within the scope of this project. However, we attempted to characterize some of the critical components, with the intent that this will be expanded in future studies to make it more comprehensive.

- Scenario description: Before describing the specifics of the communication, it is important to describe the scenario in a uniform fashion. Typically, this consists of defining the infrastructural elements near the scene (e.g., presence and position of a crosswalk [or the absence of a crosswalk], stop sign, traffic signal, number of vehicles, number of passengers, etc.), and the state of motion of the dynamic objects (i.e., velocity and position of the vehicles and passengers).
- Communication description: Based on the type of information that was going to be available from the field study, the key components of the communication that we identified were change of state of motion (e.g., speed up, slow down, stop, etc.), unilateral actions such as waving or flashing lights, and bilateral actions such as eye contact.

Sample DSL Implementation

While the DSL will continue to evolve in the future as we iterate on the key components and the user interface, we will describe a sample implementation of the DSL using MPS for this project. The communication scenario is described in three steps: definition of the infrastructure, description of the dynamic objects, and description of the communication.

Definition of the Infrastructure

Since the description of the infrastructure is usually static in nature, we felt that it was efficient to have a tabular user interface to define this instead of a CNL. This is seen in the table at the top part of Figure 1.

Traffic Scene T1

Table Infrastructures

	Existance	Location	Volumn	Visibility
crosswalk	false	(90 , 78)	(width: 90 , length: 66 , height: 0)	false
stop sign	false	(12 , 6)	(width: 8 , length: 88 , height: 67)	false
traffic lights	false	(0 , 0)	(width: 6 , length: 78 , height: 67)	false

Pedestrian Sophie is walking at a pace of 4 m/s and is alert
 Vehicle Ford is moving at a speed of 10 m/s and is 100 m away from the crosswalk at acceleration of 0 m/s^2
 Driver Tim and is alert

While approaching crosswalk:
 Tim is moving towards the crosswalk at constant speed and Sophie is moving

Driver and Pedestrian Communicating:
 Select a communication expression

- Ⓜ CommunicationReference (Communication in iteration_0)
- Ⓜ Only action (Communication in iteration_0)
- Ⓜ Pre-condition and action and end condition (Communication in iteration_0)
- Ⓜ Pre-condition with action (Communication in iteration_0)
- Ⓜ Select a communication expression (BaseConcept in jetbrains.mps.lang.core)

End script

Figure 1. Screen capture. Example of communication scenario DSL editor.

Description of the State of Motion of the Dynamic Objects (Vehicles and Pedestrians)

For the dynamic scenario description, we used CNL as the editing interface for the DSL. The resultant CNL description is seen in the middle part of Figure 1.

The example shows that it is possible to specify explicit values for the state of motion (e.g., 10 m/s). It is also possible to specify the states of mind of the pedestrian and driver (in this example, “alert”), to be chosen from a predefined set of possible states.

Even though the description appears in the form of “text,” the editor allows a more nuanced interface, where the available options for typing change depending on the context.

Description of the Communication

The interface for describing the communication is very similar to that of the dynamic scenario description. The editor provides nuanced support, as shown in the bottom part of Figure 1. A drop-down list of available element types is provided depending on the context, and this list of options gets expanded or whittled down depending on what the user types.

MPS provides the ability to define an underlying DSL that can be projected in multiple editors, such as a table or text. We leveraged this to generate the CNL for the communication as shown in Figure 2.

```
Pedestrian Sophie is walking at a pace of 3 m/s and is alert  
Vehicle Ford is moving at a speed of 10 m/s and is 100 m away from the crosswalk at acceleration of 0 m/s^2  
Driver Tim and is alert  
  
While approaching crosswalk:  
Tim is moving towards the crosswalk at constant speed and Sophie is moving  
  
Driver and Pedestrian Communicating:  
When vehicle is 10 m away from crosswalk and detects the pedestrian waiting beside the crosswalk  
the vehicle stops until detect a clear crosswalk
```

Figure 2. Screen capture. Projection of communication DSL as CNL.

Code Generation from DSL to Implementation in Vehicle

While the eventual goal of the DSL is to enable the synthesis of an optimal action strategy for the vehicle, we could also use the DSL as way of describing the “typical” human driver, and therefore make the AV mimic the behavior of such a human driver. Since the DSL is defined precisely, with well-defined underlying semantics, we can use the LWB to auto-generate the code that will drive the behavior of the AV. We implemented a proof-of-concept of this by taking a scenario description defined using the DSL and auto-generating code (in Python) that would then interface with code that was already running in our AV, ARV_003, a Polaris GEM golf cart.

Crossing Intent Model

The core part of our work was to develop a method to capture the communication between the human driver or the vehicle and the pedestrian at an unsignalized intersection. Since the field study had a broader scope, and data were not readily available at the time of development of the models, we focused on using knowledge distilled from studies of pedestrian crossing behavior (see the literature review above) in order to first design and then investigate a simple decision-theoretic model. The model is instantiated in a prototypical road crossing setting, where it forms the basis for a planning problem that an AV might use to interact and communicate with pedestrians.

Formal Problem Definition

The model concerns a pedestrian and a vehicle both approaching the same segment of a roadway. As they start, it is ambiguous as to who will cross the intersection first. The laws usually leave it open-ended (e.g., see the *California Driver Handbook* [35]). As the scenario unfolds, both vehicle and pedestrian interact via their respective actions to efficiently and smoothly resolve this question. A graphic representation of this scenario appears in Figure 3.

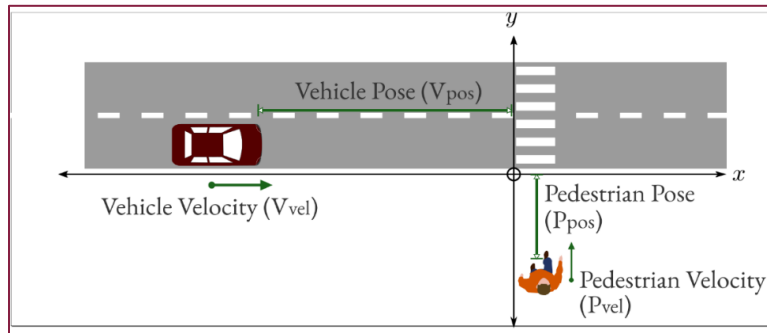


Figure 3. Diagram. Bird's eye view of crossing scenario.

Our goal is to understand the crossing behavior of both the vehicle and the pedestrian as a dynamic process. The basic assumption that motivates our modeling approach is to presume that both agents resolve the question of who will cross first as a form of uncertainty reduction. The most concise representation is a single binary variable whose value encodes who will cross the intersection first. We consider this variable as representing the (changing) crossing intent of the pedestrian. The crossing intent that the pedestrian has in mind cannot be observed directly by the vehicle. Instead, the vehicle maintains a “belief” (i.e., a probability distribution) over this binary variable. As the scenario evolves, the vehicle gains information regarding the pedestrian, integrating observations and using its model of the pedestrian's progress to learn more about the pedestrian's state.

An “observation model” is created that describes the dependence of a vehicle's observation on the actual behavior of the pedestrian. When the pedestrian's behavior, as sensed under the observation model, depends markedly on the pedestrian's sense of the crossing order, then the vehicle can leverage this dependence to “learn” the crossing order implicitly.

We formulate an instance of a partially observable Markov decision process (POMDP) [29] that describes the impact of actions for the vehicle. The solutions to such decision problems balance actions that gain information with those that attain valuable reward. In our case, the former turn out to be actions that the vehicle takes to better ascertain the pedestrian's understanding of the crossing order, *when that is valuable*. Further, it must be emphasized that “better ascertain” does not merely mean observing but, potentially, also influencing.

Vehicle-Pedestrian Interaction Model

In the sections that follow, we will examine the dynamics of the vehicle's and the pedestrian's beliefs about the intersection crossing order in order to analyze how they resolve the ambiguity and reach an agreement.

Non-observable States: Crossing Order

Let us denote the binary variable encoding crossing order as ξ . We define $\xi \in \{0,1\}$, where $\xi = 0$ means that the pedestrian crosses first and $\xi = 1$ means the vehicle crosses first. We can describe the dynamics of ξ based on some domain knowledge (i.e., a decision based on the time gap).

Pedestrian Dynamics

We express the dynamics of a pedestrian with a collection of Markov chains. Each state in the chain contains variables that describe the pedestrian's physical states and status. We also think of ξ as being associated with the pedestrian.

In this study, we restrict ourselves to a very basic motion of the pedestrian: the pedestrian can either move along the crosswalk or not. We assume that the pedestrian can move at any reasonable speed, but we treat speed in a particular way. In each state of the Markov chain, the physical state is a representation of the distance (discretized) from the crosswalk. The transition probability between each physical state is calculated based on the speed the pedestrian is traveling. To define the speed, we first need to know the crossing order (recall, ξ is seen as the pedestrian's belief).

We can break down the influences on pedestrian crossing decision-making into two main factors: contextual and habitual. Contextual factors include the position and velocity of the vehicle and the location of the crosswalk. Habitual factors include the pedestrian's traits and personal characteristics, such as age and gender. We condense contextual factors into a notion of the "level of perilousness" of the current world state. The level of perilousness is computed based on the time difference between the time remaining until the vehicle's arrival at the crossing point and the pedestrian's arrival at the crossing. The sooner the vehicle arrives compared to the pedestrian, the higher the level of perilousness and vice versa. The habitual factors determine how the pedestrian will act according to its sense of the level of perilousness. For modeling purposes, consider the extremes: a reckless and a cautious pedestrian. Let us denote the behavior by γ , with $\gamma = 1$ implying that the pedestrian's behavior is reckless and $\gamma = 0$ implying the pedestrian is cautious.

The final result is the making of a decision, which we consider as having ξ take a value. Based on the factors described above, the dynamics of ξ can be expressed as below:

$$P(\xi_{t+1}^0 | \xi_t^1, S_t^p, S_t^v) =$$

- 0.9 if vehicle slows down near the crosswalk and the pedestrian is reckless
- 0.7 if the vehicle slows down dramatically
- 0.5 if pedestrian is reckless
- 0.3 if vehicle slows down from far
- 1.0 otherwise

$$P(\xi_{t+1}^1 | \xi_t^1, S_t^p, S_t^v) = \begin{cases} 0.9 & \text{if vehicle speeds up near the crosswalk and pedestrian is cautious} \\ 0.5 & \text{if vehicle speeds up from far} \\ 1.0 & \text{otherwise} \end{cases}$$

Here the number in the superscript is the value for which the probability is being calculated. For example with $P(\xi_{t+1}^1 | \xi_t^1, S_t^p, S_t^v)$, S_t^p represents $P(\xi_{t+1} = 1 | \xi_t = 1)$. (The S_t^p and S_t^v serve as indicators that pedestrian and vehicle state information is needed to compute the expressions).

Crossing Order

The pedestrian's motion depends on their intent to cross first or second. This is, of course, precisely the information in ξ_t . Hence, the motion can be clearly defined in two cases.

Pedestrian Intends to Cross First ($\xi_t = 0$)

When the pedestrian decides to cross the intersection before the vehicle does, the pedestrian will attempt to travel at some speed to ensure this. We will assume that if the pedestrian decides to cross, they would do so at a minimum speed corresponding to nominal walking speed (1.4 m/s). However, if they perceive that the vehicle might arrive at the road crossing before they leave the crossing, they will tend to speed up, up to a maximum walking speed (2.5 m/s). Quantitatively, we capture it as below, with f_0 representing the walking speed:

$$f_0(\cdot) = \begin{cases} 1.4 & \text{if } o_{\Delta t} > 2 \text{ s} \\ 2.5 \times e^{\alpha o_{\Delta t}} & \text{if } o_{\Delta t} \leq 2 \text{ s} \end{cases}$$

Here $o_{\Delta t}$ is the difference between the remaining time for the vehicle to arrive at the intersection and the remaining time for the pedestrian to finish crossing the intersection. α is a negative constant that represents the incline of the pedestrian's speed. (We might write S_t^p and S_t^v again as a reminder but omit them for simplicity.)

Pedestrian Expects Vehicle to Cross First ($\xi_t = 1$)

In the case of the pedestrian expecting the vehicle to cross first, they will first determine if they can reach the other side of the road while walking at nominal speed. If they believe that that is not possible, they will stop at the curb and wait for the vehicle to pass before proceeding. This behavior is captured by the following equation for their velocity:

$$f_1(\cdot) = \begin{cases} 1.4 & \text{if } o_{\beta} < -1\text{m or } o_{\beta} > 1\text{m} \\ 0 & \text{if } -1\text{m} \leq o_{\beta} \leq 1\text{m and can't cross at 1.4m/s} \\ 1.4 & \text{if } -1\text{m} \leq o_{\beta} \leq 1\text{m and can cross at 1.4m/s} \end{cases}$$

where o_{β} represents the distance between the pedestrian and crosswalk. Once the pedestrian starts crossing the crosswalk, o_{β} becomes a negative distance in our representation.

Optimal Action Policy for Autonomous Driving

Vehicle Dynamics

The vehicle, unlike the pedestrian, has actions that we wish to determine. Hence, we model the vehicle's controls as actions of a decision process. The vehicle needs to avoid collision with the pedestrian, whose crossing behavior is not perfectly known. The vehicle must deal with two forms of uncertainty: partial observability and stochasticity. By choosing actions, the vehicle seeks an optimal strategy through reasoning about the pedestrian's behavior as expressed in the stochastic model.

Vehicle's Motion Model

Let V_{pos} be the state that represents the vehicle's distance from the crosswalk and state V_{vel} represent the vehicle's velocity. The evolving physical state of the vehicle is specified as (V_{pos}, V_{vel}) . The vehicle is constrained to move in a fixed direction towards the crosswalk and its control is based on acceleration rate $a \in \{a_{dec}, 0, a_{inc}\}$, where $a_{dec} < 0$ and $a_{inc} > 0$. Given a , the new state of the vehicle is calculated as

$$\begin{aligned}V'_{pos} &= V_{pos} + V_{vel}, \\V'_{vel} &= V_{vel} + a.\end{aligned}$$

Vehicle-Pedestrian Interaction

The interactions between vehicle and pedestrian near the crossing point are embedded into transition functions. When the vehicle and the pedestrian are far from the crossing, they transition to their next state based on their individual dynamics. However, the pedestrian's crossing behavior considers the vehicle position and velocity. Once the pedestrian is near the crosswalk, both the crossing behavior of the vehicle and the pedestrian are now tightly coupled: both their state transition probabilities are influenced by not only the vehicle's state but its actions as well.

Sensors and Observations

We assume that the vehicle is equipped with sensors capable of detecting the pedestrian and reporting their position and velocity. These sensors produce data that has an error range, which decreases as the vehicle gets closer to the pedestrian.

Additionally, the vehicle is assumed to have sensors that return an accurate value of its velocity and pose (the latter is merely the distance from the crosswalk). Taken together, this sensing equipment generates observations for the POMDP. We can represent these observations as a 4-tuple: $(P_{pos}, P_{vel}, V_{pos}, V_{vel})$.

Rewards

The decisions for the vehicle shall be synthesized based on maximizing *rewards* that will reflect the objectives of the pedestrian and the vehicle during a crossing. The reward model is simple: the primary objective of the vehicle is to minimize the risk of colliding with the pedestrian. Consequently, we assign a large penalty when both the vehicle and the pedestrian are on the

crosswalk simultaneously. Additionally, to incentivize efficiency, the vehicle receives rewards for those states with a higher velocity.

We emphasize that the vehicle is not specifically rewarded for knowing things about the pedestrian; any information of value is valuable because it has implications for safe, efficient motion indirectly.

The Vehicle’s Perspective on the Crossing Order

Unlike the pedestrian, who has a state ξ to represent whether they intend to cross first or not, the vehicle has no such explicit state. Instead, the POMDP maintains a distribution over the entire state space, i.e., a belief state. When all dimensions of the state other than ξ are marginalized out, what remains is a probability that represents the vehicle’s estimate of the pedestrian’s conception.

Quantifying the Value of Knowing the Crossing Order

Given that the vehicle has a distribution over ξ , the POMDP permits us to quantify the value of that single bit of information. We present this by way of an example. Assume there are three policies generated by the POMDP solver denoted π_1, π_2 , and π_3 . Assume that π_1 has a belief distribution for ξ as shown in Table 1. Further, suppose that π_2 and π_3 are policies with a 100% belief on $\xi = 0$ and $\xi = 1$, respectively. The value of knowing (or communicating), denoted V_ξ , at each time step can be calculated by comparing the expected reward of π_1 with that of π_2 and π_3 :

$$E[\pi_1] \leq E[\pi_2] \times 0.7 + E[\pi_3] \times 0.3$$

$$V_\xi = E[\pi_2] \times 0.7 + E[\pi_3] \times 0.3 - E[\pi_1].$$

Table 1. Example of Comparison of Rewards Between Policies

	$\xi = 0$	$\xi = 1$	Expected Reward
Policy π_1	70%	30%	50
Policy π_2	100%	0%	≥ 50
Policy π_3	0%	100%	≥ 50

Evaluation

We evaluated the methodology developed above using a mathematical simulation of the vehicle and pedestrian behavior.

Simulation Setup

Our simulator emulates a continuous world describing a crossing scenario and employs DESPOT [30, 31] as a POMDP solver to create a safe and efficient crossing policy for the vehicle. (DESPOT is an online solver that uses a belief-tree approach in which sampled scenarios produce nodes that are connected via edges to produce approximate policies.) Both the simulator and solver are connected through the Robot Operating System (ROS) [32]. We implemented them as ROS nodes

with inter-process communication handled by having them subscribed to one another, with messages sent via publication functions.

We developed a custom simulator to model the crossing scenario depicted above in Figure 3. The pedestrian motion is simulated using the pedestrian crossing behavior model described earlier. The vehicle was simulated to move towards the crossing point. The vehicle proceeds at the speed of 3 m/s initially, with its speed changing according to the acceleration rates produced as actions. These actions are themselves generated directly by DESPOT.

One round of simulation occurs when the vehicle or pedestrian finishes crossing from one side to the other. Each round begins with both the vehicle and the pedestrian in the simulator moving steadily towards the shared crossing. When the vehicle is 14 m or less away from the crosswalk, the simulator node sends a message to start the POMDP solver. For every execution step, the solver reads the current state of the simulator, including both the vehicle and the pedestrian information, and outputs an acceleration rate. The simulator will transition to the next world state as the vehicle transitions based on the generated acceleration rate as input and the pedestrian transitions based on its crossing behavior model. The simulator continues to subscribe for new acceleration rates until the round finishes.


For the DESPOT solver, we used 500 sampled scenarios with the maximum depth of the belief tree as 100, and the discount factor set to 0.98. The solver was given 1 s to construct the search tree and choose an action. The simulations were executed on an Intel Core i7-6670HQ 2.6 GHz processor with 32 GB of RAM running Ubuntu 16.04.

We constructed a crosswalk 4-m wide and 5-m long. The vehicle was set to start 14 m away from the crosswalk, and the pedestrian was set 4.2 m away from the crosswalk. The vehicle's and the pedestrian's position state space were generated by discretizing the continuous space by an interval of 1.75 and 0.75, respectively. The velocity state for the vehicle contained values from 0 to 3 with an interval of 1. The pedestrian had three velocity states: $\{0.0, 1.4, 2.5\}$. We defined the action space of the POMDP model to be $\{-1.0, 0.0, 0.5\}$, where each item is an acceleration rate that can be feasibly executed by the vehicle.

Results

In this section, we discuss the performance of the simulated AV by analyzing the resulting behavior based on overall safety (i.e., occurrence of a collision) and the quality of the vehicle's belief of the pedestrian's intent (which reflects the quality of implicit communication). We also briefly discuss results where the vehicle has an option to take an action that generates explicit communication. We present the simulation via detailed plots of a variety of variables as they evolve in time. Figure 4(a) and (b) depict simulations of a cautious and a reckless pedestrian, respectively. The positions of both pedestrian and vehicle are shown in the graph entitled "continuous position." The vertical axis of the graph is the distance from the crossing, where negative values represent positions that

are before the crossing point. The horizontal axis of the graphs is the number of iterations the belief states are updated. The time interval between each update is 0.04 s.



(a) (b) Figure 4. Graphs. Simulation results showing (a) vehicle executing a policy, interacting with a cautious pedestrian, and (b) vehicle executing a policy, interacting with a reckless pedestrian.

Crossing Safety

We can see that regardless of whether interacting with a reckless or a cautious pedestrian, the lines for the vehicle and the pedestrian positions are never both greater than 0 simultaneously. This indicates adequate spatial separation, i.e., no collision occurs.

In our simulation, we decelerate at the rate of -1.0 m/s^2 , which is within the Institution of Transportation Engineers (2009) recommended average maximum deceleration rate of 3.0 m/s^2 and the comfortable deceleration rate of 3.4 m/s^2 defined by AASHTO (2004). This assures the safety and comfort of the speed change.

Beliefs Over Non-observable States

In our scenario, the key to communication is the inference of behavior, which itself is resolved as a question about the pedestrian's crossing decision (ξ). Recall that ξ is calculated based on the perilousness of the crossing for the pedestrian and also their habitual characteristics. Since neither this characteristic nor ξ are observable, the vehicle's knowledge of these two elements is understood in terms of the belief state (or distribution) over both variables. The third plot in the figures shows how the belief of the characteristic converges to the correct trait. As for the belief distribution of the pedestrian's ξ , it appears (along with the actual pedestrian's ξ value for comparison) in the fourth plot of the figures. Notice that ξ changes, but the vehicle's belief distribution is shown to align with the changes in the pedestrian's ξ .

Implicit Communication: An Interpretation

To help assimilate the results, we chose to compare the behavior of the simulated AV with the behavior of human drivers under circumstances where they are uncertain of the pedestrian's sense of who should cross.




Figure 5. Chart. Strategies of the vehicle in ambiguous situations with a reckless pedestrian (based on [13] but redrawn with modifications).

Figure 5 is a summary of the behavior of our simulated vehicle; it is redrawn from Schneemann and Gohl [13] using their style of summarization, along with some modifications for clarity, and with numbers reporting the data from our simulations. The percentages and speed values are based on the results of running simulations using a reckless pedestrian as input. (One example of the strategy is shown in Figure 4). Schneemann and Gohl [13] report that human drivers resolve ambiguous situations by initially reducing their speed, and then deciding whether to speed up or come to a stop depending on the pedestrian’s response to their speed reduction. We see that the vehicle’s strategy is less conservative than that of human drivers. Figure 5 can be interpreted as the vehicle trying to gain efficiency (i.e., directly modelled via rewards) while also balancing uncertainty. Instead of slowing down to learn the pedestrian’s crossing order decision passively, the vehicle remains at moderately high speeds, seemingly to express its intention to cross first. This communicates with the pedestrian, and the pedestrian’s subsequent motion can be explained as a reply to the crossing arrangement. In Figure 4(a), the cautious pedestrian is shown to slow down, acceding to the vehicle crossing first. In Figure 4(b), the reckless pedestrian accelerated to express disagreement on the vehicle crossing order. Both the vehicle and the pedestrian continue to adapt their maneuvers thereafter in order to reach an agreement on the crossing order.

Explicit Communication

We also conducted a simple simulation to analyze the value of communicating crossing order by creating an action that communicates ξ explicitly. In Figure 6 of Appendix B, the vehicle can flash its headlights. We model the pedestrian as understanding this action as an indication that the vehicle intends to let the pedestrian cross first. Additionally, to have the vehicle’s policy be deliberate in choosing to communicate intent and establish ξ , we assign a negative reward as penalty for opting to flash the lights.

As seen from the results in Figure 6, the vehicle chooses to flash its lights (quite frequently in fact) despite a penalty being incurred. Moreover, when we compare the third graph in Figure 4(b) to that in Figure 6, it is clear that knowledge of the pedestrian’s characteristic is recognized faster with explicit communication. The fourth graph in both figures also shows that ξ stabilizes sooner

too. As the ambiguity is resolved, the result is that both the vehicle and pedestrian cross the crosswalk more efficiently.

Conclusions and Recommendations

This project had the objective of understanding the interaction between a vehicle with a human driver and pedestrians so that this understanding could be leveraged to synthesize driving policies for AVs that are derived from a rational trade-off between traffic efficiency and safety. There were two primary activities towards this understanding: (i) a field study that captured and documented the behavior of drivers and pedestrians under a variety of realistic traffic scenarios, and (ii) a mathematical formalization of the vehicle-human communication leading to an optimization-based decision strategy for autonomous driving.

The field study resulted in a large amount of knowledge captured through both raw video and other sensory data, as well as through curated information captured in Excel files.

The formalization of communication effort resulted in (1) the development of a DSL methodology for precisely capturing the driving behavior, which can be used for both translating human behavior directly to AVs (which was demonstrated on an actual AV), as well as for parameterizing a more formal communication model that can be used to generate optimal control strategies for AVs; (2) a communication intent model that recognizes the fundamental nondeterminism associated with human behavior and addresses it through a POMDP; and (3) a practically realizable decision support system that leverages state-of-the-art solvers for POMDP, and demonstrating this on a real AV.

By modeling the pedestrian behavior as a POMDP, a vehicle maintains a distribution over a binary variable that encodes whether a pedestrian aims to cross first and generates a sequence of actions that manage uncertainty, including some actions that seem to elicit information—bearing the hallmarks of implicit communication. Our simulation results show a vehicle capable of resolving uncertainty in order to achieve efficiency. An examination of the vehicle’s crossing behavior suggests that the strategy is less conservative than some driving behavior, including some humans, while trying to resolve ambiguous situations. Motivated by studies conducted on and patents issued for external hardware designs for AVs, we also briefly discussed vehicle strategies that include explicit communication actions.

Recommendations for Future Work

Here we make a few recommendations on future work:

1. The development of the DSL is not a one-time activity, but something that needs to be evolved as more and more scenarios are considered for description using the DSL. A direct way to achieve this is by applying the DSL to field data. While this had been one of the paths identified, it was not possible to do this during the course of this project as the field

study was happening in parallel with the formalization activity. A future study should take the large amounts of field study data that have been generated and use the DSL to abstract the descriptions of the communication events. During this process, the DSL itself needs to be improved.

2. As a corollary to the above, if sufficient data are processed through the DSL, it becomes a rich knowledge base that can be programmatically analyzed, so the formal communication models can be more accurately parameterized to represent typical driver-pedestrian behavior.
3. While we have had very remarkable results eliciting the tacit communication using just one behavioral state—the binary variable capturing crossing intent—we recognize that there are multiple behavioral states that are relevant and need to be considered. This would be a natural area for future work.
4. The ultimate value in a keener understanding of communication between vehicles and pedestrians is being able to use that understanding for decision support. Using more variables to describe the behavior results in exponential growth in complexity, making the implementation of optimal decision support a challenge. Thus, activity is needed to look at efficient optimization and implementation techniques that can be implemented in real-time.

Additional Products

The Education and Workforce Development (EWD) and Technology Transfer (T2) products created as part of this project can be downloaded from the [project page on the Safe-D website](#). The final project dataset is located on the Safe-D Collection of the VTTI [Dataverse](#).

Education and Workforce Development Products

The research project supported an M.S. student, Ya-Chuan Hsu. She presented a poster at the 3rd Annual Texas A&M Transportation Technology Conference, fitting the theme of “Preparing for Connected Automation.” A preliminary report on the research she coauthored, “An MDP Model of Vehicle-Pedestrian Interaction at an Unsignalized Intersection,” was accepted to a workshop held at the IEEE Vehicular Technology Conference in Chicago, Illinois. She has a paper titled “Implicit Coordination via Uncertainty-aware Plans: A POMDP Treatment of Vehicle-pedestrian Interaction” under submission to IEEE/RSJ International Conference on Intelligent Robots and Systems 2019. She is well on her way to successfully defending her M.S. thesis and has expressed an interest in continuing her graduate studies to pursue a Ph.D. in Computer Science.

Dylan Shell, co-principal investigator, used the topic of uncertainty-aware planning as a means to realize implicit communication as a basis for a topical discussion on September 3, 2019, in the Fall 2019 offering of CSCE-691, section 646. Aspects of POMDPs were introduced and reviewed, then information gathering in standard formulations was examined. Finally, discussion with application to social settings where the planner realizes communication for interaction was described. This

constituted a topic for 1.5 hours of time, and culminated in a student present (not affiliated with the original project) expressing that both the application and the non-standard formulation were interesting and novel to him.

Co-principal investigator Srikanth Saripalli led a group of senior undergraduate students as part of a capstone project. Five students teamed up as part of MEEN 401/402 over two semesters and built a communication system consisting of LEDs, speakers, and GPS that can communicate the intentions of an AV to pedestrians. The project ended in December 2019.

Technology Transfer Products

The methodology that we have developed has broader implications in the context of Human Robot Interaction (HRI). The generality of the approach, specifically in treating actions as informative and hence expressing the subtle intercouplings of agents in the world, has raised the question of how widely state-of-the-art planning techniques may be used to realize social competency in robots. In particular, we have already engaged a new project that extends the methodologies developed here to the AGV-pedestrian interaction scenario in a factory setting. The belief is that the separation of a specification language (for recording observations and posing scenarios) from the realization via probabilistic models (e.g., particular discrete Markov decision processes) is a methodology applicable to a range of settings of practical and potential importance.

Data Products

The data collected as part of the field study are available via the Safe-D collection on the VTTI Dataverse at <https://doi.org/10.15787/VTTI/IC4KCCQ>. The CommAutoExport dataset contains the data analysts' annotations for the 1,808 pedestrian interactions with detailed descriptions of each variable in the dataset, including data type, range of values, coding for categorical variables, and identifiers used to denote missing values or errant data. The Tacit Communication Video Reduction Directory includes the values for the annotations in the CommAutoExport dataset. The CommAuto Driver Dataset includes information on the driving context, driving route, narration, and participant demographics.

References

1. Lavasani, M., X. Jin, and Y. Du. Market Penetration Model for Autonomous Vehicles Based on Previous Technology Adoption Experiences. In 95th Annual Meeting of the Transportation Research Board, Washington, DC, January 2016.
2. Fagnant, D. J., and K. M. Kockelman. Preparing a Nation for Autonomous Vehicle: Opportunities, Barriers and Policy Recommendations. William P. Eno Paper, October 2013.
3. Gandhi, T., and M. M. Trivedi. Pedestrian Protection Systems: Issues, Survey, and Challenges. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 8, September 2007.
4. Rothenbucher, D., J. Li, D. Sirkin, B. Mok, and W. Ju. Ghost Driver: A Field Study Investigating the Interaction Between Pedestrians and Driverless Vehicles. 25th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2016, pp. 795–802, 2016.
5. B. Schoettle and M. Sivak. *A Survey of Public Opinion About Autonomous and Self-Driving Vehicles in the US, the UK, and Australia*. University of Michigan Transportation Research Institute, July 2014.
6. Anaya, J. J, P. Merdrignac, O. Shagdar, F. Nashashibi, and J. E. Naranjo Vehicle to Pedestrian Communications for Protection of Vulnerable Road Users. 2014 IEEE Intelligent Vehicles Symposium, Dearborn, Michigan, United States, June 2014.
7. Matthews, M., G. Chowdhary, and E. Kieson. Intent Communication Between Autonomous Vehicles and Pedestrians. 2017. arXiv preprint arXiv:1708.07123.
8. Clamann, M., M. Aubert, and M. L. Cummings. Evaluation of Vehicle-to-Pedestrian Communication Displays for Autonomous Vehicles. In 96th Annual Meeting of the Transportation Research Board (paper No. 17-02119), Washington, DC, 2017.
9. Schwitter, R. Controlled Natural Languages for Knowledge Representation. In Proceedings of the 23rd International Conference on Computational Linguistics: Posters (COLING '10). Association for Computational Linguistics, Stroudsburg, PA, USA, 2010, pp. 1113-1121.
10. Voelter, M. *DSL Engineering: Designing, Implementing and Using Domain Specific Languages*. 2013. dslbook.org
11. Erwig, M., R. F. Paige, and E. Van Wyck. The State of the Art in Language Workbenches - Conclusions from the Language Workbench Challenge. International Conference on Software Language Engineering. 2013.

12. MPS. www.jetbrains.com/mps
13. Schneemann, F., and I. Gohl. Analyzing Driver-Pedestrian Interaction at Crosswalks: A Contribution to Autonomous Driving in Urban Environments. In *Proceedings of the IEEE Intelligent Vehicles Symposium*, 2016, pp. 38-43.
14. Moore, R. L. Pedestrian Choice and Judgment. *Operational Research Society*, Vol. 4, 1953, No. 1, pp. 3-10. <https://www.jstor.org/stable/3006905>
15. National Research Council, *Highway Capacity Manual*, 2000, No. 9.
16. Cohen, J., E. J. Dearnaley, and C. E. M. Hansel. The Risk Taken in Crossing a Road. *Operational Research Society*, Vol. 6, 1955, No. 3, pp. 120-128. <https://www.jstor.org/stable/3007711>
17. Brewer, M., K. Fitzpatrick, J. Whitacre, and D. Lord. Exploration of Pedestrian Gap-Acceptance Behavior at Selected Locations. *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 1982, 2006, pp. 132-140. <http://trjournalonline.trb.org/doi/10.3141/1982-18>
18. Harrell, W. A. Factors Influencing Pedestrian Cautiousness in Crossing Streets. *Journal of Social Psychology*, Vol. 131, 1991, No. 3, pp. 367-372.
19. Khan, F.M., M. Jawaid, H. Chotani, and S. Luby. Pedestrian Environment and Behavior in Karachi, Pakistan. *Accident Analysis and Prevention*, Vol. 31, 1999, No. 4, pp. 335-339.
20. Tiwari, G., S. Bangdiwala, A. Saraswat, and S. Gaurav. Survival Analysis: Pedestrian Risk Exposure at Signalized Intersections. *Transportation Research Part F: Traffic Psychology and Behaviour*, Vol. 10, 2007, No. 2, pp. 77-89.
21. Oxley, J. A., E. Ihsen, B. N. Fildes, J. L. Charlton, and R. H. Day. Crossing Roads Safely: An Experimental Study of Age Differences in Gap Selection by Pedestrians. *Accident Analysis and Prevention*, Vol. 37, 2005, pp. 962-971.
22. Oxley, J., B. Fildes, E. Ihsen, J. Charlton, and R. Day. Differences in Traffic Judgements Between Young and Old Adult Pedestrians. *Accident Analysis and Prevention*, Vol. 29, 1997, No. 6, pp. 839-847.
23. Schneemann, F., and P. Heinemann. Context-based Detection of Pedestrian Crossing Intention for Autonomous Driving in Urban Environments. In *IEEE International Conference on Intelligent Robots and Systems*, 2016, pp. 2243-2248.
24. Bandyopadhyay, T., C. Z. Jie, D. Hsu, M. H. Ang Jr., D. Rus, and E. Frazzoli. Intention-Aware Pedestrian Avoidance. In *Experimental Robotics. Springer Tracts in Advanced Robotics*, Vol. 88, 2013, pp. 963-977. <http://link.springer.com/10.1007/978-3-540-77457-0>

25. Lagstrom, T., and V. Lundgren. AVIP-Autonomous Vehicles' Interaction with Pedestrians. Ph.D. dissertation, 2015.
26. Urmson, C. P., I. J. Mahon, D. A. Dolgov, and J. Zhu. Pedestrian Notifications. 2012.
27. Mercedes-Benz. Mercedes-Benz F 015 Luxury in Motion. https://www.mbusa.com/en/future-vehicles/f-015?sd_campaign_type=Search&sd_digadprov=Resolution&sd_campaign=Brand%7CCorp%7CFuture+Models%7CConcept+Cars%7CGeneral%7CExact&sd_channel=GOOGLE&sd_adid=F+015+Concept&sd_digadkeyword=mercedes+benz+f+015+luxury+in+motion&gclid=Cj0KCQjwzunmBRDsARIsAGrt4msAvxuinpPDhWcRwxyH3TOrGw3Wp7C0RNj1qH1F9fiovvpX2VD4OwaArfUEALw_wcB&gclsrc=aw.ds
28. Nissan Motor Corporation. Nissan IDS Concept: Nissan's Vision for the Future of EVs and Autonomous Driving. 2015. <http://www.nissan-global.com/EN/NEWS/2015/STORY/151028-01-e.html%5Cnhttp://bit.ly/1PGMJmw>
29. Smallwood, R. D., and E. J. Sondik. The Optimal Control of Partially Observable Markov Processes over a Finite Horizon. *Operations Research*, Vol. 21, 1973, No. 5, pp. 1071-1088.
30. Somani, A., N. Ye, D. Hsu, and W. S. Lee. DESPOT: Online POMDP Planning with Regularization. In *Advances in Neural Information Processing Systems*, 2013, pp. 1772–1780.
31. Ye, N., A. Somani, D. Hsu, and W. S. Lee. DESPOT: Online POMDP Planning with Regularization. *Journal of Artificial Intelligence Research*, Vol. 58, 2017, pp. 231-266.
32. Quigley, M., K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng. ROS: An Open-source Robot Operating System. In *ICRA Workshop on Open Source Software*, 2009.
33. Drive.ai. <https://www.drive.ai>
34. Pennycooke N. AEVITA: Designing Biomimetic Vehicle-to-Pedestrian Communication Protocols for Autonomously Operating & Parking On-Road Electric Vehicles. Master's Thesis, Media Arts and Science, MIT, 2012.
35. *California Driver Handbook – Laws and Rules of the Road Online*: https://www.dmv.ca.gov/portal/dmv/detail/pubs/hdbk/right_of_way

Appendices

Appendix A: Field Study Results

Table 2. Results of Chi-Square Tests Between the Driving Context Condition and Behavioral Variables

		Driving context		χ^2 (df)
		Vehicle without sign	Vehicle with sign	
Driver yield behavior	Yield	763	783	1.4(1)
Driver hand position	Failure to yield	86	106	41.7(2)*
	Top of wheel	446	339	
Driver eye contact	Side of wheel	36	28	0.0(1)
	Bottom of wheel	367	522	
Driver wave through	Yes	769	805	4.2(1)*
	No	80	84	
Driver wave palm outward	Yes	79	59	.1(1)
	No	770	830	
Driver reaction	Yes	34	38	6.6(4)
	No	815	851	
	No reaction	88	113	
	Continue, but accelerate	34	45	
	Continue, but decelerate	367	359	
	Comes to complete stop	214	199	
Pedestrian eye contact	Interrupted driving, then continued	146	173	0.0(1)
	Yes	482	505	
Pedestrian wave palm outward	No	367	384	0.0(1)
	Yes	72	77	
Pedestrian reaction	No	777	812	7.3(6)
	No reaction	601	635	
	Continue, but accelerate	54	52	
	Continue, but decelerate	20	28	
	Begin walking from stationary	59	42	
	Interrupted walking, then continued	91	99	
	Interrupted walking and aborted	14	24	
	Other	10	9	

Table 3. Results of Chi-Square Tests Between the Narration Condition and Behavioral Variables

		Narration		χ^2 (df)
		No think-aloud	Think-aloud	
Driver yield behavior	Yield	762	784	0.1(1)
	Failure to yield	97	95	
Driver hand position	Top of wheel	377	408	1.1(2)
	Side of wheel	33	31	
	Bottom of wheel	859	879	
Driver eye contact	Yes	795	779	7.8(1)*
	No	64	100	
Driver wave through	Yes	63	75	0.9(1)
	No	796	804	
Driver wave palm outward	Yes	30	42	1.8(1)
	No	829	837	
Driver reaction	No reaction	93	108	3.7(4)
	Continue, but accelerate	45	34	
	Continue, but decelerate	367	359	
	Comes to complete stop	204	209	
	Interrupted driving, then continued	150	169	
Pedestrian eye contact	Yes	496	491	0.6(1)
	No	363	388	
Pedestrian wave palm outward	Yes	69	80	0.6(1)
	No	790	799	
Pedestrian reaction	No reaction	630	606	6.1(6)
	Continue, but accelerate	50	56	
	Continue, but decelerate	23	25	
	Begin walking from stationary	45	56	
	Interrupted walking, then continued	81	109	
	Interrupted walking and aborted	20	18	
	Other	10	9	

Table 4: Model 2 Results–Naturalistic Driving

Fixed effects	Odds ratio	Estimate	Standard error	Z-score	p-value	Effect size
Intercept	0.15	-1.91	0.69	-2.77	0.01	*
Vehicle speed	1.01	0.01	0.01	0.93	0.35	
Event history - stopped for infrastructure	0.47	-0.75	0.40	-1.88	0.06	Small
Event history - stopped for vehicle	0.72	-0.33	0.44	-0.75	0.45	
Event history - stopped for pedestrian	0.49	-0.71	0.46	-1.55	0.12	Small
Visual obstruction - Obstructed	0.80	-0.22	0.24	-0.92	0.36	
Traffic control - yield to pedestrian	0.97	-0.03	0.25	-0.12	0.90	
Traffic control - stop sign	0.06	-2.81	0.87	-3.21	0.00	Large
Right of way - pedestrian	2.28	0.82	0.39	2.12	0.03	Small
Vehicle distance - 15+ feet away	0.66	-0.42	0.26	-1.62	0.11	Small
Number of pedestrians in path	0.98	-0.02	0.09	-0.27	0.79	
Number of pedestrians near path	1.02	0.02	0.08	0.31	0.76	
Pedestrian distance - 0 feet away	16.32	2.79	0.55	5.09	0.00	Large
Pedestrian distance - 0-5 feet away	16.03	2.77	0.52	5.38	0.00	Large
Pedestrian distance - 5-10 feet away	76.02	4.33	0.55	7.82	0.00	Large
Phone use - using	1.34	0.30	0.39	0.75	0.45	
Phone use - unable to determine	2.18	0.78	0.50	1.57	0.12	*
Pedestrian facial expression - eye contact	3.31	1.20	0.36	3.35	0.00	Medium
Pedestrian facial expression - unable to determine	0.49	-0.72	0.71	-1.01	0.31	*
Pedestrian assertiveness	1.10	0.09	0.07	1.40	0.16	
Driver facial expression - eye contact	0.02	-3.97	0.34	-11.64	0.00	Large
Vehicle in opposing lane - yes	0.35	-1.06	0.38	-2.77	0.01	Medium
Factor affecting pedestrian path - path obstruction	3.85	1.35	0.38	3.57	0.00	Medium
Factor affecting driver path - path obstruction	0.16	-1.82	0.32	-5.77	0.00	Large
Random effect	Variance	SD				
Driver	1.45	1.21				

*Indicates an effect that is not a meaningful comparison

Appendix B: Explicit Communication Simulation Results:

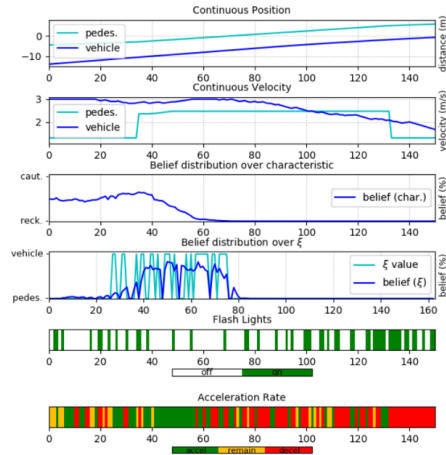


Figure 6. Reckless pedestrian interacts with vehicle equipped to flash its lights as a form of explicit communication.




Figure 7. Illustration. Capstone project MEEN 401/402 for explicit communications.