Modeling Collaborative Decisions in Car-Pedestrian Encounters

Extended Abstract

***Abstract*—In this work we build a model for pedestrian encounters parametrized by aggressiveness and attentiveness based on the SHRP2 dataset, which includes dashcam and driver facing video footage from several different regions within the US. From this dataset we use inverse reinforcement learning to extract a reward function to model pedestrian actions parametrized by attentiveness/distractedness and passivity/aggressiveness. The dataset is parsed and labeled by a video analytics toolkit we develop. Finally, we use these models to design an autonomous driver that makes optimal decisions according to a tunable parameter of desired aggressiveness.**

Introduction

Imagine you are driving your car in a residential area and you notice a pedestrian standing by the edge of the road with their body oriented to face the street. As you slow down, but probably before you reach a complete stop, the pedestrian notices your willingness to stop and begins to cross the street. The two of you have collectively decided to let the pedestrian go first. Now imagine you're driving on a larger road, perhaps two or three lanes wide on each side, and you encounter the same scenario. In this case, the pedestrian is probably planning to jaywalk, violating right of way protocols. However, they are probably planning to do so in a conscientious way, prioritizing your goals and plans above their own as to cause minimal disturbance to the flow of traffic. The two of you have the same decision to make, who do you let go first and how do you communicate intent?

Pedestrian encounters like this happen on a continuous basis, and while most of them happen without incident, such encounters result in an average of 15 fatalities in the US every day[1]. Of course the hope is that with the introduction of autonomous vehicles this number will be reduced substantially, however 2018 witnessed the first pedestrian fatality caused by an autonomous vehicle[7], which very closely resembled the jaywalking scenario described

There are a few interesting things to notice about these two scenarios. In both cases, the driver and the pedestrian took into account the plans of the other in deciding what to do; they planned assuming their actions would affect the plans of the other. The pedestrian may spend some time observing the situation, perhaps even making a small movement to see if the driver yields, then makes the final decision of whether or not to cross.

In such an encounter what does each observe and what features do they use to make the decision of whether to proceed or not? We hypothesize that in such a scenario both agents are looking to see whether the other agent is a myopic or a planning agent. By myopic agent, we mean one who ignores the other agent in the process of making its decisions, either as a result of being distracted or as a result of simply not caring. And by planning, we mean one who makes predictions of the other agent's actions in order to inform their future plans.

Another aspect of this encounter is the efficiency with which each agent reaches their goal. In both of these scenarios, both agents have an opportunity to aggressively and selfishly prioritize their goal over the other agent's goal, while still maintaining a basic level of safety. Although a non-aggressive driving strategy may perform well in a rural setting, an unaggressive driver could very conceivably freeze into inaction because of the aggressive pedestrians in a dense urban setting like New York City.

The most important consideration of this encounter is the need to preserve safety, which in this the ultimate responsibility belongs to the driver since the car is capable of killing/injuring the pedestrian, but with high probability the pedestrian is incapable of killing/injuring the pedestrian. This means that even the most aggressive strategy of the driver must still ensure the safety of pedestrians where physical laws of motion permit.

In this work we focus on two tasks: the first is to build a model of pedestrian behavior that takes into account the driver behavior, the second is to use this model to develop a controller for an autonomous vehicle that has the optimal tradeoff of courtesy vs efficiency while maintaining guarantees on safety. The first task requires using off-policy reinforcement learning of the pedestrian's actions. The second task can use a variety of known methods in control and optimization to determine the optimal policy.

Related Work

Pedestrian movement has traditionally been modeled by the civil and traffic engineering communities for the purpose of modeling aggregate movement of crowds of pedestrians through transportation hubs, shopping malls, or corridors in order to consider flow dynamics and bottlenecks. This type of work typically uses very simple representations of the pedestrians as perhaps a single point and generates aggregate interactions based on a cellular automaton, guided random walks, a sum of forces [2], [8]. Even these simple models are able to demonstrate emergent crowd behavior such as the formation of lanes in crowded hallways, but they are less capable when it comes to modeling fewer numbers of pedestrians in less structured environments.

There are a number of works that specifically consider car-pedestrian encounters at unsignalized crossings, but many of these works build models based of a single instant in the encounter, rather than the evolution of the state as the encounter unfolds and the agents react to one another. Some works focus specifically on using video analytics to predict whether or not the pedestrian will cross[9], [18], while others consider also the driver's decision to yield. A recent approach by Chen *et. al.* studies a dataset of 2973 encounters of pedestrians with a bus at a single unsignalized crossing. They use a Gaussian-mixture model to capture the interaction between bus and pedestrian velocity, distance, and time advantage and use this to simulate pedestrians for comparing their autonomous strategy[4]. Similarly, multiple works have successfully modeled pedestrian's decision of whether or not to cross (probability of gap acceptance) and the motorist's decision of whether or not to yield using logistic regression on a number of factors including speed, time until collision, age, gender, location, and several other characteristics[15], [16]. In these setups the pedestrian and the car are both myopic, choosing a fixed strategy based on their initial state observation and not reacting to the changes in state of the other.

In other parts of Human-Robot Interaction literature, however, there is significant work done in modeling different aspects of the internal human state in situations where the two interact which captures the interaction and intent signaling that happens in multi-agent exchanges. For instance, Nikolaidis et. al. has successfully modeled the internal human states of adaptability[12] and trust[5] in collaborative tasks such as moving a table or manipulating objects, and Majumdar et. al. have provided a formal structure for representing risk in both human and robot agents[10]. Sadigh, et. al. even demonstrate an ability to actively probe the internal state of the human agent[14], something that can be used to make the information gathering phase more effective. They also demonstrate that many humans are not myopic planners, and that when good models of human reward are known, the robot can use this to influence the path taken by the human in a driving scenario[13]. A similar goal of helping human agents understand robot intent can be found in works on legibility for arm manipulation[6] and micro UAV path planning[17], and specifically in car-pedestrian encounters with the use of a lights and LED text to communicate intent[11].

This work benefits from and builds heavily on the previous work, but makes the following unique contributions:

- Many of the related works consider the state of the pedestrian or the state of the motorist, but none of them consider the interaction between them and how it evolves over the course of the encounter. In this work we aim to build internal models for both the motorist and the pedestrian.
- We specifically consider the awareness/attentiveness of each agent and how that impacts their ability to plan and predict the actions of the other agent.
- The system we propose to build this model performs video analytics on dashcam footage. By providing this tool to researchers, this model can be continuously refined to learn pedestrian dynamics specific to certain regions.
- We learn an optimal controller for an autonomous car, parametrized by passiveness/aggressiveness, or how much the car wishes to prioritize its own goal over that of the pedestrian.

Problem Outline

The key to our approach is the wealth of data in the Strategic Highway Research Program (SHRP2) naturalistic driving survey dataset[3]. This dataset is comprised of trip data acquired from 2400 drivers in five different states, distributed uniformly distributed across age and gender. The trip data consists of multiple video feeds, including a dashcam and one aimed at the driver, GPS, accelerometer, radar, vehicle control inputs and state (speed, gas pedal, brake, horn, etc..) and even a passive alcohol sensor. A very important part of this work consists of parsing through the more than two petabytes of data that resulted from the SHRP2 study to identify and label the scenarios involving pedestrian encounters, and extract a time-series of states for every encounter.

Thus, the primary aim of this first step is to produce the trimmed dataset

$$\mathcal{D} = \{\boldsymbol{x}^{(j)}\}_{j=1}^{J},$$

where x is a single state from a single encounter, $\boldsymbol{x}$ represents the entire time-series of states from that encounter and the dataset $\mathcal{D}$ is composed of a total of J encounters. The state x at each time-step is composed of some things that are observed (position and velocity of car), some things that are estimated from sensor data (position and velocity of pedestrian), and some internal state variables that are not observed. One of these latent features is the attentiveness, responsiveness, or degree to which each agent plans ahead. Video footage of the direction the driver/pedestrian is looking may provide a prior probability for this. Another unobserved part of the state is agressiveness, which characterizes the urgency with which the agent pursues their reward.

In learning the reward function of the pedestrian we take the approach of off-policy inverse reinforcement learning (IRL). This assumes that the pedestrian is acting optimally according to some unknown reward function. The assumption of optimality is the typical drawback in IRL, and the hope is that with the inclusion of responsiveness and aggressiveness as part of the internal state, this will introduce enough realistic suboptimality to make the IRL successful.

Lastly, with the full pedestrian model we develop an autonomous vehicle controller parametrized by aggressiveness to act optimally according to our human model and study the strategies that emerge at each level of urgency. We validate our model on human test subjects in a simulated environment and, pending IRB approval, real life encounters using a human driver that takes instructions on optimal strategy from the control algorithm.

References

[1] N. H. T. S. Administration. Traffic safety facts 2015 data–pedestrians. Technical report, U.S. Department of Transportation, 2017.

[2] M. Batty. Agent-based pedestrian modelling. *Advanced spatial analysis: The CASA book of GIS*, 81:81–106, 2003.

[3] K. L. Campbell. The shrp 2 naturalistic driving study: Addressing driver performance and behavior in traffic safety. *TR News*, (282), 2012.

[4] B. Chen, D. Zhao, and H. Peng. Evaluation of automated vehicles encountering pedestrians at unsignalized crossings. In *Intelligent Vehicles Symposium (IV), 2017 IEEE*, pages 1679–1685. IEEE, 2017.

[5] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa. Planning with trust for human-robot collaboration. *arXiv preprint arXiv:1801.04099*, 2018.

[6] A. Dragan and S. Srinivasa. Generating legible motion. 2013.

[7] T. Griggs and D. Wakabayashi. How a self-driving uber killed a pedestrian in arizona, 2018.

[8] D. Harney. Pedestrian modelling: current methods and future directions. *Road & Transport Research*, 11(4):38, 2002.

[9] C. G. Keller and D. M. Gavrila. Will the pedestrian cross? a study on pedestrian path prediction. *IEEE Transactions on Intelligent Transportation Systems*, 15(2):494–506, 2014.

[10] A. Majumdar, S. Singh, A. Mandlekar, and M. Pavone. Risk-sensitive inverse reinforcement learning via coherent risk models. In *Robotics: Science and Systems*, 2017.

[11] M. Matthews, G. Chowdhary, and E. Kieson. Intent communication between autonomous vehicles and pedestrians. *arXiv preprint arXiv:1708.07123*, 2017.

[12] S. Nikolaidis, A. Kuznetsov, D. Hsu, and S. Srinivasa. Formalizing human-robot mutual adaptation: A bounded memory model. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, pages 75–82. IEEE Press, 2016.

[13] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan. Planning for autonomous cars that leverage effects on human actions. In *Robotics: Science and Systems*, 2016.

[14] D. Sadigh, S. S. Sastry, S. A. Seshia, and A. Dragan. Information gathering actions over human internal state. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 66–73. IEEE, 2016.

[15] B. Schroeder, L. Elefteriadou, V. Sisiopiku, N. Rouphail, K. Salamati, E. Hunter, B. Phillips, T. Chase, Y. Zheng, and S. Mamidipalli. Empirically-based performance assessment and simulation of pedestrian behavior at unsignalized crossings. *Southeastern Transportation Research, Innovation, Development and Education Center (STRIDE) Project 2012-016S*, 2014.

[16] D. Sun, S. Ukkusuri, R. F. Benekohal, and S. T. Waller. Modeling of motorist-pedestrian interaction at uncontrolled mid-block crosswalks. *Transportation Research Board Annual Meeting 2003*, 51:61801, 2003.

[17] D. Szafir, B. Mutlu, and T. Fong. Communication of intent in assistive free flyers. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pages 358–365. ACM, 2014.

[18] R. Tian, L. Li, K. Yang, S. Chien, Y. Chen, and R. Sherony. Estimation of the vehicle-pedestrian encounter/conflict risk on the road based on tasi 110-car naturalistic driving data collection. In *Intelligent Vehicles Symposium Proceedings, 2014 IEEE*, pages 623–629. IEEE, 2014.