A POMDP Treatment of Vehicle-Pedestrian Interaction: Implicit Coordination via Uncertainty-Aware Planning

Ya-Chuan Hsu[1], Swaminathan Gopalswamy[2], Srikanth Saripalli[2], and Dylan A. Shell[1]

***Abstract*— Drivers and other road users often encounter situations (e.g., arriving at an intersection simultaneously) where priority is ambiguous or unclear but must be resolved via communication to reach agreement. This poses a challenge for autonomous vehicles, for which no direct means for expressing intent and acknowledgment has yet been established. This paper contributes a minimal model to manage ambiguity and produce actions that are expressive and encode aspects of intent. Specifically, intent is treated as a latent variable, communicated implicitly through a partially observable Markov decision process (POMDP). We validate the model in a simple setting: a simulation of a prototypical crossing with a vehicle and one pedestrian at an unsignalized intersection. We further report use of our self-driving Ford Lincoln MKZ platform, through which we conducted experimental trials of the method involving real-time interaction. The experiment shows the method achieves safe and efficient navigation.**

I. Introduction

Field studies have suggested that current-generation autonomous vehicles lack social competence [1], [2]. Humans resolve ambiguities in traffic via social interaction, including expressing *intent*. In some cases, these interactions are so effective that they are almost transparent. Examples involve acknowledging/asserting the right of way, or communicating the intention to yield [3]. Drivers use actions, such as approaching at speed, as signals to implicitly communicate the intention of not giving way to pedestrians at unsignalized crossings [4]. This paper explores specific means to treat factors that cannot be sensed directly, such as a pedestrian's intent to cross, but for which it is important to reason over.

One place where informal interaction protocols are particularly important, and where autonomous vehicles must be competent participants, is when pedestrians wish to cross the road on which the vehicle is traveling. We examine pedestrian-vehicle interaction by focusing on a scenario with an *unsignalized intersection*[3], as this is a representative circumstance in which communication is crucial [6], [7].

Our focus is on approaches that are directly practicable: though current technologies encourage the installation of ever richer sensors on robots, factors such as age, gender, culture, faith, and past experiences—factors known to influence pedestrian crossing [8]—are unlikely to be precisely sensed any time soon. Moreover, our emphasis is to preserve *social* aspects so that existing understanding of pedestrians, including the role of implicit coordination and indirect signals, remains applicable without further presumptions. This study uses actions already commonplace, such as speeding up/slowing down. The present study is appropriate also in those cases where pedestrians cannot determine whether the approaching vehicle is autonomous or not.

We give a decision-theoretic treatment that considers those aspects of pedestrian behavior which are not directly observable to be a form of uncertainty that must be modeled. A plan is constructed that reasons over and manipulates this uncertainty. When the vehicle executes the plan, the result exhibits hallmarks of social competence, at least as applies to the simple, small-scale scenario studied. Our broader philosophy is that several aspects of social interaction cope with uncertainty, thus representing uncertainty explicitly and dealing with it efficiently, can yield socially effective robots.

We consider the paper's main contributions to be:

- *A minimal model of social ambiguity:* Despite social interaction being far from trivial, we boil several fairly complex and abstract concepts down to a single source of uncertainty, formulating a 'lumped parameter' model with a variable we dub the pedestrian's crossing intention. Although a collection of factors exist, our approach gives a single expression that can be interpreted in terms of probability.
- *Framing a practically solvable partially observable decision problem:* The conciseness of the representation can be exploited by maintaining a low-dimensional distribution, making it practical for the vehicle to solve for a sequence of actions. Those actions manage uncertainty, including some which seem to elicit information—bearing the hallmarks of implicit communication. The model we present expresses this satisfactorily (albeit indirectly) via transition dynamics.

Beyond those two primary contributions, we also report on an implementation and demonstration with a real autonomous vehicle. The brief description suffices to illustrate the feasibility of the approach. On further examining our vehicle's behavior, it appears to be less conservative than some human drivers, potentially indicating that that planner helps resolve ambiguous situations quite efficiently.

II. Literature review

Resolving ambiguities in traffic poses three main challenges for autonomous vehicles: 1) sensing is hard, especially for subtle gestures and cues; 2) human-to-human social protocols are informal and only serve to establish expectations;

[1]Ya-Chuan Hsu and Dylan A. Shell are with the the Department of Computer Science and Engineering at Texas A&M University, College Station, TX, USA.{yachuan815, dshell}@tamu.edu

[2]Swaminathan Gopalswamy and Srikanth Saripalli are with the Department of Mechanical Engineering at Texas A&M University, College Station, TX, USA.{sgopalswamy, ssaripalli}@tamu.edu

[3]An unsignalized intersection "is defined as any at-grade junction of two or more public roads at which the right-of-way for motorists, bicyclists, and pedestrians is not controlled by a highway traffic signal." [5]

3) human-to-autonomous vehicle protocols have yet to be forged along with associated expectations.

The studies we describe next provide evidence that a great deal of useful behavior can follow from quite minimal information about pedestrians. This gives the basis for our approach which only measures pedestrian position and velocity; these data, the daunting challenges of perception notwithstanding, can be sensed tractably today. We approach both challenges (2) and (3) by capturing a range of human-to-x interaction via active management of uncertainty.

A. Pedestrians crossing: definitions and known determinants

Early studies of pedestrians crossing roads observed that pedestrians are primarily concerned with time-gaps [9]. The Highway Capacity Manual [10] defines the *critical gap* as the time below which a pedestrian will not attempt to begin crossing. Detailed studies [9], [11] identified that each person has their own critical gap, which depends on the oncoming vehicle's speed; people do not cross when the vehicle violates their threshold.

Subsequent studies thoroughly evaluated factors, including traffic flow conditions, road geometry, temperature, *etc.*, that influence peoples' critical gap [12], [13]. They show that the critical gap correlates positively with cautiousness. Our model includes a factor representing one notion of caution.

Besides external factors, personal characteristics pertaining to specific pedestrians influence their critical gap. Some studies show that gender affects pedestrian behavior [14], [15]; others point out that age is an influencer of pedestrian behavior [16], [17].

B. Cooperative pedestrian interaction with autonomous cars

There are many successful systems for autonomous driving [18], [19], but the study of close interactions between autonomous vehicles and pedestrians is more recent. However, most studies [20]–[24] focus on examining pedestrian-vehicle interaction to provide reference data for future implementations of autonomous vehicles. These studies do not include practical approaches for actual autonomous vehicle implementations that consider interaction.

A critical difficulty for autonomous vehicles driving amid pedestrians is to incorporate pedestrian intentions and behaviors into their decision making. Among different approaches, the simplest approach is to create a reactive system [25]. However, this ignores uncertainty inherent in making predictions, resulting in fast computation but sub-optimal solutions over time. Thus, different methods to reason about the prediction uncertainty during decision making have been proposed [26], [27]. The POMDP approach is general, assuming neither linear dynamics nor Gaussian noise [28]. Though POMDPs are widely known to be demanding computationally, steady improvements in efficiency have seen them being implemented in reasonably-scaled experiments [29]. This work leverages POMDPs to balance the uncertainties the pedestrian intention while having the vehicle operating safely and efficiently. Whereas [28] and [29] seek to reduce obstruction by inferring pedestrians' navigation goals, we instead express the interplay between vehicle and pedestrian. Under our model, information-gathering actions influence the pedestrian's behavior.

Various methods have been proposed for autonomous vehicles to communicate their intended actions to pedestrians nearby. Though no single solution dominates, one approach involves displaying physical information [30]–[32] and several companies [33]–[35] have developed specialized external hardware. In contrast, we deliberately opted to use features found on a standard vehicle.

III. Problem Definition

Consider a pedestrian and vehicle both approaching the same segment of a roadway and, initially, it is ambiguous as to who will cross the intersection first. (See Fig. 1.) Both vehicle and pedestrian interact, via their respective choices of actions, to efficiently and smoothly resolve this question as the situation unfolds. Treating the (coupled) behavior of both participants as a dynamic process, the basic assumption underlying our model is that both agents resolve the question of who crosses first as a form of uncertainty reduction.

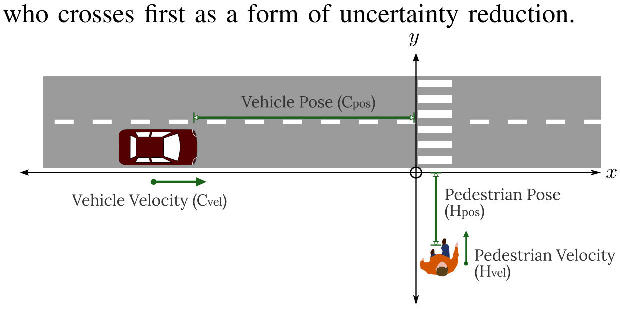

Fig. 1: Bird's eye view of the unsignalized crossing.

IV. Approach

The crossing order that the pedestrian has in mind cannot be observed directly by the vehicle. However, as the vehicle drives, it gains information regarding the pedestrian by integrating observations and using its model of the pedestrian's progress to learn more about the pedestrian's state. We propose the most concise representation possible, a single binary variable, whose values encode who will cross the intersection first. This variable represents the pedestrian's evolving understanding of the crossing order question. The vehicle maintains a belief/distribution over this variable.

From the point of view of the vehicle, the pedestrian's states evolve stochastically. But, at the same time, they are influenced by the vehicle's state, which is itself altered via the vehicle's actions. An appropriate choice of action helps ensure that it will make a sequence of observations that are informative. When the pedestrian's behavior, as sensed under the observation model, depends markedly on the pedestrian's conviction regarding the crossing order, then the vehicle learns about this hidden variable indirectly. Thus by choosing actions the vehicle has an interactive mechanism, which it can initiate, to actively manage its belief.

The preceding complexities are included in our model, along with one further, important nuance. The vehicle is only disposed to gain information that is valuable for driving safely and efficiently through the intersection. We formulate an instance of a POMDP [36] that describes the effect of actions for the vehicle. As is well known, POMDP solutions balance actions that gain information with ones that attain valuable reward. For us, the former are actions that the vehicle takes to better ascertain the pedestrian's understanding of the crossing order, *whenever valuable*. Furthermore, here 'better ascertain' does not merely mean observing but potentially also influencing.

Below, we carefully examine the dynamics of the vehicle and pedestrian's beliefs about the crossing order, interpreting how they resolve initial ambiguity and reach agreement. Sec. VII compares the vehicle's performance against the strategies human drivers are known to use in resolving ambiguous situations, as mentioned in [21].

V. Model Design

At a high level, our POMDP model considers a state space S comprising states $S = \{S^H, S^C\}$, where S^H represents information about the pedestrian (mnemonic: $\square^H$ for human), including their position (H_{pos}), velocity (H_{vel}), characteristic (H_{chr}) and crossing intention (ξ); S^C represents information of the vehicle ($\square^C$ for car), including vehicle position (C_{pos}) and velocity (C_{vel}). The subsections that follow detail how we approach the abstract concept of human intention and capture the notion of beliefs and agreement over crossing order; how we model the pedestrian and vehicle's physical transitions in a decoupled fashion; and then how we connect the pieces to construct the final POMDP model.

A. Mental states: Crossing order

Let us denote the binary variable encoding crossing order at time t as ξ_t. We define $\xi_t \in \{0, 1\}$, where $\xi_t = 0$ means the pedestrian crosses first and $\xi_t = 1$ means the vehicle crosses first. The dynamics of ξ_t are based on domain knowledge (i.e., a time-gap–based decision), detailed in Sec. V-B.2.

B. Pedestrian dynamics

The dynamics of a pedestrian is expressed with a collection of Markov chains. Each state in the chain contains variables that describe the pedestrian's physical states. We also think of ξ_t as being associated with the pedestrian.

1) Overview: We restrict ourselves to a consideration of very basic motion: the pedestrian can either move along the crosswalk or pause. We will assume that the pedestrian can move at any reasonable speed, but, as clarified shortly, 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_t is seen as the pedestrian's crossing intention).

2) Deciding to cross or not: Summing up the studies in Sec. II-A, the influences on pedestrian crossing decision-making comprise two main factors: (i) Contextual factors include the position/velocity of the vehicle and the location of the crosswalk; (ii) Habitual factors include the pedestrian's traits and personal characteristics, like age and gender.

For contextual factors, we condense them 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 pedestrian's arrival at the crossing and the remaining time of the vehicle's arrival at the crossing point. If the vehicle has not finished crossing the intersection while the pedestrian starts to cross, 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, we consider the extremes: a reckless ($H_{\text{chr}}^{\text{rkl}}$) and a cautious ($H_{\text{chr}}^{\text{cts}}$) pedestrian.

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

$$P(\xi_{t+1} = 0 | \, \xi_t = 0, S_t^H, S_t^C) = \begin{cases} 0.3 \text{ if critical gap is small, and pedestrian is reckless,} \\ 0.9 \text{ if critical gap is large, and pedestrian is reckless,} \\ 0.14 \text{ if critical gap is small, and pedestrian is cautious,} \\ 0.74 \text{ if critical gap is large, and pedestrian is cautious,} \\ 1.0 \text{ otherwise.} \end{cases} \tag{1}$$

For $P(\xi_{t+1} = 0 | \, \xi_t = 1, S_t^H, S_t^C)$ the values for the five cases are 0.26, 0.86, 0.1, 0.7, and 1.0, respectively. For $P(\xi_{t+1} = 1 | \, \xi_t = 0, S_t^H, S_t^C)$ they were 0.7, 0.1, 0.86, 0.26, and 1.0; similarly for $P(\xi_{t+1} = 1 | \, \xi_t = 1, S_t^H, S_t^C)$ the probabilities are 0.74, 0.14, 0.3, 0.9, and 1.0. Here a reckless (or cautious) pedestrian is just one that has $H_{\text{chr}} = H_{\text{chr}}^{\text{rkl}}$ (or $H_{\text{chr}}^{\text{cts}}$, respectively).

(The numbers above are the values used for the trials we report below; more precise values could be obtained via psychological experiments and data collection.)

3) Pedestrian locomotion: The pedestrian's motion depends on whether they currently intend to cross first or second. This is, of course, precisely the information in ξ_t. Hence the motion can be clearly defined into two cases expressed with functions $f_{\xi_t} : S \to \mathbb{R}$ yielding velocities.

a) Pedestrian crosses first ($\xi_t = 0$): When the pedestrian decides to cross the intersection before the vehicle does, it will attempt to travel at some speed to ensure it crosses first. We enforce some basic constraints: should the pedestrian reach the fastest walking speed, $2.5\,\mathrm{m/s}$ [37], it remains moving at the highest speed it is capable of maintaining. If the speed needed to arrive in time is below average walking speed $1.4\,\mathrm{m/s}$ [37], the pedestrian continues at a normal pace. Quantitatively, this is

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

where $o_{\Delta t}$, computed from $S_t \in S$, is the time difference between the remaining time for the vehicle to arrive at the intersection and for the pedestrian to finish crossing the intersection. Here α is a negative constant that represents the incline of the pedestrian's speed.

b) Vehicle crosses first ($\xi_t = 1$): In this case, the pedestrian will stop at the curb and wait for the vehicle to pass when it determines that, continuing at its current speed, it cannot reach the other side of the road before the vehicle arrives. The pedestrian's velocity is given as

$$f_1(\cdot) = \begin{cases} 1.4 & \text{if } o_{\Delta p} < -1\,\mathrm{m} \text{ or } o_{\Delta p} > 1\,\mathrm{m}, \\ 0 & \text{if } -1\,\mathrm{m} \leq o_{\Delta p} \leq 1\,\mathrm{m} \text{ and can't cross at } 1.4\,\mathrm{m/s}, \\ 1.4 & \text{if } -1\,\mathrm{m} \leq o_{\Delta p} \leq 1\,\mathrm{m} \text{ and can cross at } 1.4\,\mathrm{m/s} \end{cases} \quad (3)$$

where $o_{\Delta p}$, computed from $S_t \in S$, represents the relative position the pedestrian is from the crosswalk. Once the pedestrian starts crossing the crosswalk, $o_{\Delta p}$ becomes a negative value in our representation.

C. 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.

1) Vehicle's motion model: Let C_{pos} be the state that represents the vehicle's distance from the crosswalk and state C_{vel} represent the vehicle's velocity. The evolving physical state of the vehicle is specified as $(C_{\text{pos}}(t), C_{\text{vel}}(t))$ at time t. The vehicle is constrained to move in a fixed direction towards the crosswalk and its control is based on acceleration $a(t) \in \{a_{\text{dec}}, 0, a_{\text{inc}}\}$, where $a_{\text{dec}} < 0$ and $a_{\text{inc}} > 0$. Given $a(t)$, the new state of the vehicle is calculated as

$$\begin{aligned} C_{\text{vel}}(t + \Delta t) &= C_{\text{vel}}(t) + a(t), \\ C_{\text{pos}}(t + \Delta t) &= C_{\text{pos}}(t) + \Delta t \cdot C_{\text{vel}}(t). \end{aligned} \quad (4)$$

2) 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, as modeled in Sec. V-B.2, the pedestrian's crossing behavior considers the vehicle position and velocity. Once the pedestrian is near the crosswalk, the behavior of both 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 action as well.

3) Sensors and observations: We assume that the vehicle is equipped with sensors capable of detecting the pedestrian and reporting his/her 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 the vehicle's velocity and pose (the latter is merely the distance from the crosswalk). Taken together, this sensing equipment generates observations that we represent as a 4-tuple: $(H_{\text{pos}}, H_{\text{vel}}, C_{\text{pos}}, C_{\text{vel}})$.

4) Rewards: The vehicle's primary objective is to minimize the risk of colliding with the pedestrian. Consequently, a large penalty is assigned 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.

5) The vehicle's perspective on the crossing order: Unlike the pedestrian, who has a state ξ_t to represent who he/she considers to be crossing first, 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_t are marginalized out, what remains is a probability that represents the vehicle's estimate of the pedestrian's conception.

VI. Experimental Setup and Details

We construct a continuous world describing a crossing scenario and employ DESPOT [38], [39] 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–based approach in which sampled scenarios produce nodes that are connected via edges to produce approximate policies.) Both the simulator (or autonomous vehicle) and solver are connected through the Robot Operating System (ROS) [40]. We implemented them as ROS nodes with inter-process communication handled by them subscribing to one another.

A. Experimental flow

We treat the crossing scenario depicted in Fig. 1. Each experimental trial begins with both the vehicle and the pedestrian in the simulator moving steadily towards the crossing. A trial concludes when the vehicle finishes crossing from one side to the other. As our approach is designed for the purpose of resolving crossing ambiguity, the POMDP solver is only triggered once a pedestrian is detected heading towards the crossing. The POMDP solver solves for vehicle-pedestrian interaction when the vehicle is $14\,\mathrm{m}$ or less away from the crosswalk. For every execution step, the solver considers the current state of the world, including both the vehicle and the pedestrian information, and outputs an acceleration value. The vehicle enacts the new acceleration and the aspects pertaining to the pedestrian evolve as per the transition model detailed earlier in Sec. V-B.2. The solver continues to output new acceleration values based on new inputs until the trial finishes.

B. Experiment Parameters

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 $1\,\mathrm{s}$ to

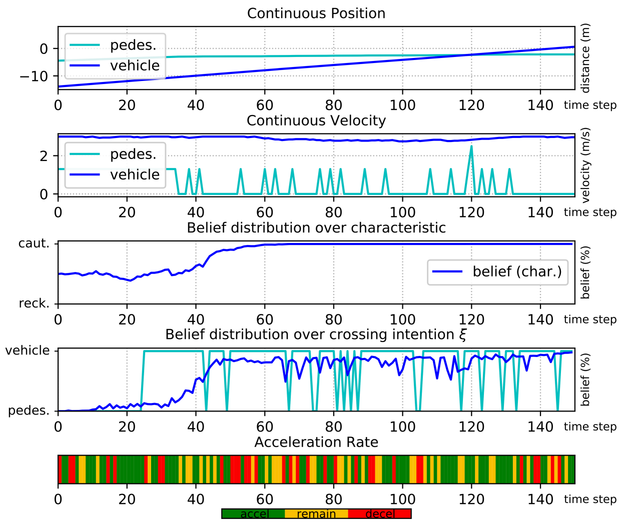

(a) Cautious pedestrian

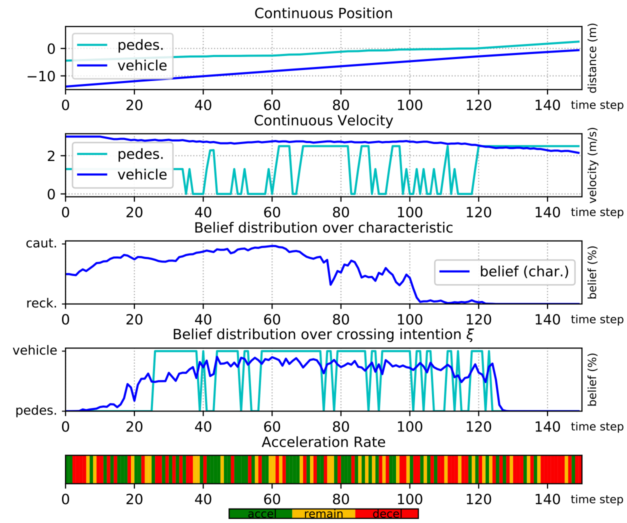

(b) Reckless pedestrian

Fig. 2: Simulation results showing a vehicle executing a policy, interacting with a cautious pedestrian (a) and a reckless pedestrian (b). The plotted variables evolve in each belief update time step.

construct the search tree and choose an action. The experiments were executed on an Intel Core i7-6670HQ $2.6\,\mathrm{GHz}$ processor with $32\,\mathrm{GB}$ of RAM running Ubuntu 16.04.

VII. Implementation Details and Experiments

In this section, we first summarize an extensive and carefully controlled evaluation conducted in simulation. The performance of the simulated autonomous vehicle will be discussed by analyzing the resulting behavior in terms of the overall safety and further, for non-observable states, the dynamics of the vehicle's belief. We then show that with the minimal model constructed, the results of our implementation of the solver running on an autonomous vehicle possesses some hallmarks of social competence. Finally, we briefly discuss experiments where the vehicle may also generate actions that are explicitly communicative.

A. Simulation setup

We developed a custom simulator to model the crossing setting of Fig. 1. It simulates pedestrian motions using the pedestrian crossing behavior model described in Sec. V-B. Pedestrian transitions at each step consider the vehicle's state as well as the pedestrian's characteristic. (Recall, these represent contextual and habitual factors respectively; the latter is set manually at the start of the simulation.) Both are needed to simulate the dynamics, which is achieved by sampling in proportion to the associated probabilities.

The vehicle and pedestrian's position state space is generated by discretizing continuous space into intervals of $1.75\,\mathrm{m}$ and $0.75\,\mathrm{m}$. The velocity state for the vehicle contains values from $0\,\mathrm{m/s}$ to $3\,\mathrm{m/s}$ with an interval of $1\,\mathrm{m/s}$. The pedestrian has three velocity states: $\{0.0, 1.4, 2.5\}$. We define the action space of the POMDP model as $\{-1.0, 0.0, 0.5\}$, where each is an acceleration value that can be executed by the vehicle.

B. Analysis simulation results

We present the simulation results via detailed plots of a variety of variables as they evolve in time. Figs. 2a and 2b depict runs of a cautious and a reckless pedestrian, respectively. The position of both pedestrian and vehicle are shown in the graph with the title 'continuous position.' The vertical axis of the graph is the relative position from the crossing, where negative values represent positions that are before the crossing point.

1) Crossing safely: We can see that whether interacting with a reckless or a cautious pedestrian, the lines for the vehicle and the pedestrian positions are never seen to be between the crossing region, 0 to 4, simultaneously. This indicates that no collision occurs.

2) Beliefs over non-observable states: In our scenario, the key to communication is the inference of behavior, which resolves to a question about the pedestrian's crossing decision (ξ_t). Where ξ_t is calculated based on the perilousness of the crossing for the pedestrian and also their habitual characteristics (H_{chr}). Since neither this characteristic nor ξ_t 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 results figures shows how the belief of the characteristic converges to the correct trait. As for the belief distribution of ξ_t, it appears (along with the actual pedestrian's ξ_t value for comparison) as in the fourth plot of the result figures. Notice that ξ_t changes, but the vehicle's belief distribution is shown to track the changes in the pedestrian's ξ_t.

3) Implicit communication — an interpretation: To help comprehend the results, we chose to compare the behavior of simulated autonomous vehicle with human drivers under circumstances where they are uncertain of the pedestrian's sense of who should cross. Fig. 3 is a summary of the behavior of our simulated vehicle; it is redrawn from [21], using their style of summarization, along with some modifications to improve clarity, but with numbers reporting results from our experiments. The percentages and speed values are based on data from simulations using a reckless pedestrian. (One example of the strategy is shown in Fig. 2b.)

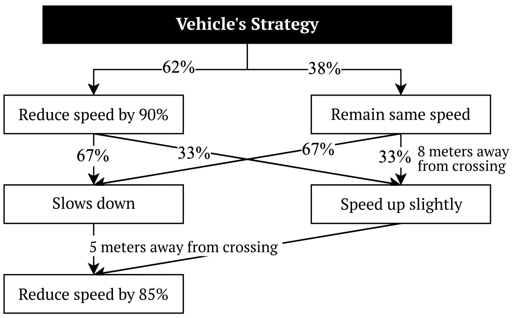

Fig. 3: Strategies of the vehicle in ambiguous situations with a reckless pedestrian simulated for crossing. (Based on [21].)

Schneemann and Gohl [21] report that human drivers resolve ambiguous situations by initially reducing their speed, and then decide to 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 compared to human drivers. Fig. 3 can be interpreted as the vehicle trying to gain efficiency rewards but also balancing uncertainty. Instead of slowing down to passively learn the pedestrian's crossing order decision, the vehicle remains at moderately high speeds, seemingly expressing its desire to cross first. This communicates with the pedestrian and the pedestrian's following movement can be explained as a reply to the crossing arrangement. In Fig. 2a, the cautious pedestrian is shown to slow down, giving the vehicle permission to cross first. In Fig. 2b, the reckless pedestrian accelerated to express disagreement on the vehicle crossing arrangement. Both the vehicle and the pedestrian continue to adapt their maneuvers thereafter in order to reach agreement on crossing order.

Readers interested in a planner that always yields to the pedestrian are referred to [41], where no uncertainty in the pedestrian's intention was modelled. Observations of that work show the autonomous vehicle to be more conservative when facing non-communicative pedestrians who insist on crossing ahead of the vehicle.

C. Autonomous vehicle experiment setup

We carried out experiments at our university autonomous vehicle testing ground. A virtual crossing was setup in the system that is $5\,\mathrm{m}{\times}4\,\mathrm{m}$ overseen by the camera mounted on a lamp post on the side of a roadway. (See Fig. 4.)

Our test vehicle is a Ford Lincoln MKZ with auto-driving enabled. The auto-driving system is capable of following a pre-recorded path with GPS and vehicle orientation included by sending command through ROS to the low-level vehicle controller to control the throttle, brake and steering.

Our pedestrian is a manikin mounted on a pole installed on a remote control car. It is operated by people with the intent to either cross before the vehicle or after the vehicle has passed the crosswalk.

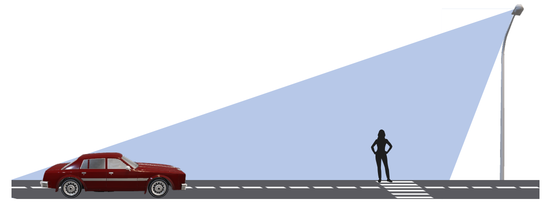

Fig. 4: Cartoon depicting the experimental infrastructure. The drawn crosswalk is just an representation of all the different types of unsignalized intersection.

Our sensors include a camera mounted on a lamp pole overseeing the crossing area and an RTK GNSS receiver, Piksi. With infrastructure enabled autonomy [42], we coordinated both sensors to provide the location of the vehicle and pedestrian relative from the base station. The information gathered from the sensors becomes observations for the planner to locate the current state in the belief state space and output an acceleration value as the action to the vehicle (as shown in Fig. 5).

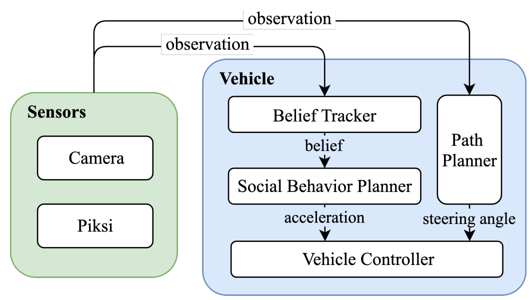

Fig. 5: System architecture employed for the Lincoln MKZ.

Both pedestrian and vehicle's motion, and sensor readings, are imperfect owing to factors such as friction, bumps on the road, wind, sunlight, *etc.*, in the world affecting both agents. Seeking to balance between decision quality and computational expediency, observations of the vehicle and pedestrian information are discretized. The prior has a resolution of $1.5\,\mathrm{m}$ for position and $1.0\,\mathrm{m/s}$ for speed and the later has a resolution of $1.0\,\mathrm{m}$ for position and $1.0\,\mathrm{m/s}$ for speed and The vehicle actions are sparse: $0.5\,\mathrm{m/s^2}$, $-1.0\,\mathrm{m/s^2}$, and $-2.0\,\mathrm{m/s^2}$. The maximum planning time per step is $0.2\,\mathrm{s}$, and the planning horizon is 100 steps.

D. Results from the autonomous vehicle experiment

Though extensive tests (46 trials) were conducted in this environment, we can only report some illustrative instances here due to the limited space. Once the vehicle is as close as $70\,\mathrm{m}$, it receives input from the camera sensors for pedestrian detection. If the pedestrian is detected to be approaching the crossing, the vehicle slows down to $3\,\mathrm{m/s}$; Otherwise, it continues at the same speed, proceeding to approach and cross the crossing. Once the vehicle slows down and is within $14\,\mathrm{m}$ before the crossing, it activates the behavior planner to interact with the pedestrian and begins to maintain a belief distribution over the pedestrian's crossing intention. In the case of interacting with a reckless pedestrian, the vehicle slows down in advance (Fig. 6a) with an initial belief that the pedestrian will cross first. While slowing down (Fig. 6b),

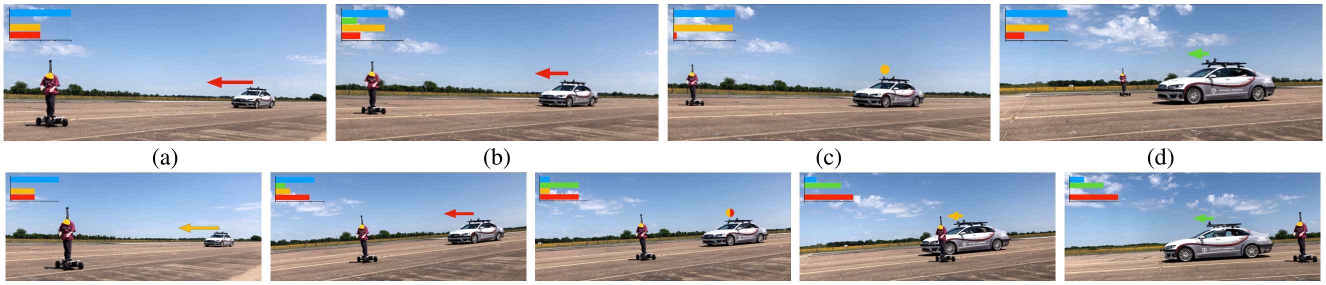

(a) (b) (c) (d)

(e) (f) (g) (h) (i)

Fig. 6: In figures (a)–(d), the vehicle encounters a reckless pedestrian who speeds up to start crossing before vehicle arrives. In figures (e)–(i), the vehicle encounters a cautious pedestrian who stops to wait for vehicle to cross. Histogram indicates beliefs over pedestrian crossing intentions and characteristic: blue for *pedestrian crosses first*, green for *vehicle crosses first*, yellow for a pedestrian who is characteristically *reckless*, red for a *cautious* one. The length of the arrow above the vehicle expresses vehicle's velocity, in which, a circle indicates that the velocity is approximately zero. The color of the arrow describes the acceleration value: green for *accelerate*, yellow for *maintain*, and red for *decelerate*.

the vehicle's belief in the pedestrian being reckless increases which leads to it later picking up speed (Fig. 6c) and crosses (Fig. 6d). As for the cautious pedestrian case, the vehicle slows down due to the initial belief the pedestrian intends to cross first (Fig. 6e). Before it comes to a stop, the vehicle changes its belief distribution over the pedestrian's characteristic and intention as it now observes the pedestrian to be slowing down before the crossing (Fig. 6f). The vehicle's acceleration changes from decelerating to maintaining speed as the majority of the weighting of belief shifts toward the pedestrian being cautious and who, thus, intends to let the vehicle cross first (Fig. 6g). Then, passing by the cautious pedestrian slowly, the vehicle gradually gains speed (Fig. 6h). Finally, the vehicle accelerates once past the pedestrian (Fig. 6i)

These examples from our autonomous vehicle trials indicate that the planner produces effective crossing behavior and that it performs well at managing uncertainty for states of the pedestrian that are not directly observable.

E. Explicit communication

We also conducted a simple experiment to analyze the value of communicating crossing order by creating an action that communicates ξ_t explicitly. In Fig. 7, an action is added where the vehicle may flash its headlights. We model the pedestrian as understanding this as indicating that the vehicle intends to let the pedestrian cross first. Additionally, to have the vehicle's policy be deliberate in choosing to communicate intent in establishing ξ_t, we penalize using the lights.

The result, in Fig. 7, shows that the vehicle opts to flash its lights (quite frequently) despite the negative reward incurred. Moreover, as we compare the third graph in Figs. 2b and 7, it is clear that knowledge of the pedestrian's characteristic is recognized faster with explicit communication. The fourth graph in both figures shows that ξ_t stabilizes sooner too. And, as the ambiguity is resolved, the result is that both the vehicle and pedestrian cross the crosswalk more efficiently.

VIII. Conclusion

Social interaction is valuable in resolving ambiguity in traffic; it is necessary for autonomous vehicles if they

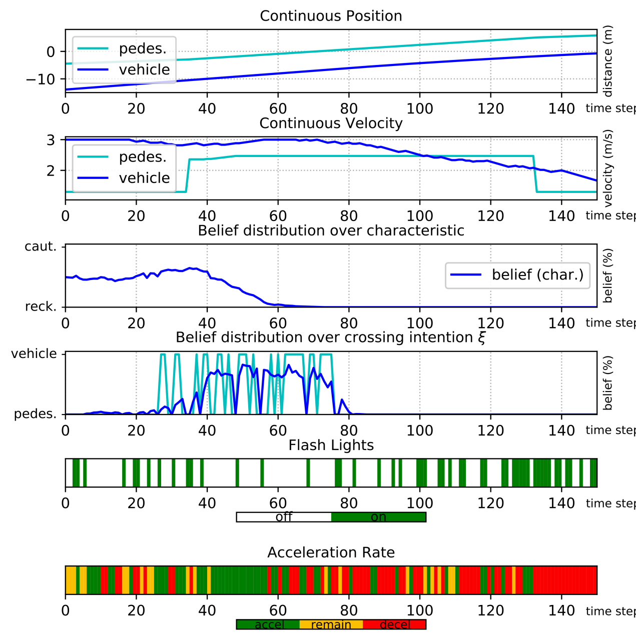

Fig. 7: A reckless pedestrian interacts with a vehicle equipped to flash its lights, communicating explicitly.

are to operate harmoniously within the existing systems, infrastructure, and norms. This paper has examined how an uncertainty-aware planner can help an autonomous vehicle interact competently. We have examined a scenario involving ambiguity, formulating a decision-theoretic model that captures elements related to a pedestrian's intent as a form of uncertainty, which is ultimately expressed as non-determinism and partial observability. A state-of-the-art solver is then used to produce a plan. Our implementation (planning and executing in both simulation and on a Ford Lincoln MKZ) produces results indicating that the vehicle's belief of the pedestrian's intention converges correctly.

Further work might consider multiple pedestrians and different sensing assumptions, and conduct more extensive real-world experiments. Determining whether this simple 'lumped parameter' model will suffice for more complex scenarios, and whether other indicators might help identify intentions, are interesting next steps. Much research remains to done to better realize social and understandable behavior in vehicles among multiple other road users.

IX. Disclaimer & Acknowledgment

This document reflects the authors' views. It is disseminated in the interest of information exchange. Support for this research was provided in part by a grant from the U.S. Department of Transportation, University Transportation Centers Program to the Safety through Disruption University Transportation Center (451453-19C36). However, the U.S. Government assumes no liability for the contents or use thereof.

References

[1] M. Richtel and C. Dougherty, "Google's driverless cars run into problem: Cars with drivers," https://www.nytimes.com/2015/09/02/technology/personaltech/google-says-its-not-the-driverless-cars-fault-its-other-drivers.html?_r=2, 2015, accessed: 2020-02-27.

[2] S. E. Anthony, "The trollable self-driving car," https://slate.com/technology/2016/03/google-self-driving-cars-lack-a-humans-intuition-for-what-other-drivers-will-do.html, 2016, accessed: 2020-02-27.

[3] N. Guéguen, S. Meineri, and C. Eyssartier, "A pedestrians stare and drivers' stopping behavior: A field experiment at the pedestrian crossing," *Safety science*, vol. 75, pp. 87–89, 2015.

[4] A. Varhelyi, "Drivers' speed behaviour at a zebra crossing: a case study," *Accident Anal. & Prev.*, vol. 30, no. 6, pp. 731–743, 1998.

[5] McGee, Hugh W., "Unsignalized intersection improvement guide," https://toolkits.ite.org/uiig/types.asp, 2015, accessed: 2020-02-25.

[6] D. Rothenbücher, J. Li, D. Sirkin, B. Mok, and W. Ju, "Ghost driver: A field study investigating the interaction between pedestrians and driverless vehicles," in *Int. Sym. Robot and Human Interactive Comm.*, 2016, pp. 795–802.

[7] 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, Ann Arbor, Trans. Research Institute, Tech. Rep., 2014.

[8] A. Rasouli and J. K. Tsotsos, "Autonomous vehicles that interact with pedestrians: A survey of theory and practice," *Transactions on Intel. Trans. Sys.*, vol. 21, no. 3, pp. 900–918, 2019.

[9] R. L. Moore, "Pedestrian choice and judgment," *J. of the Operational Research Society*, vol. 4, no. 1, pp. 3–10, 1953.

[10] W. Reilly, *Highway capacity manual 2000*, Trans. Research Board, 1997.

[11] J. Cohen, E. Dearnaley, and C. Hansel, "The risk taken in crossing a road," *J. of the Operational Research Society*, vol. 6, no. 3, pp. 120–128, 1955.

[12] M. A. Brewer, K. Fitzpatrick, J. A. Whitacre, and D. Lord, "Exploration of pedestrian gap-acceptance behavior at selected locations," *Trans. Research Record*, vol. 1982, no. 1, pp. 132–140, 2006.

[13] W. A. Harrell, "Factors influencing pedestrian cautiousness in crossing streets," *The J. of Social Psy.*, vol. 131, no. 3, pp. 367–372, 1991.

[14] F. M. Khan, M. Jawaid, H. Chotani, and S. Luby, "Pedestrian environment and behavior in Karachi, Pakistan," *Accident Anal. & Prev.*, vol. 31, no. 4, pp. 335–339, 1999.

[15] G. Tiwari, S. Bangdiwala, A. Saraswat, and S. Gaurav, "Survival analysis: Pedestrian risk exposure at signalized intersections," *Trans. Research Part F: Traffic Psy. and Behav.*, vol. 10, no. 2, pp. 77–89, 2007.

[16] J. A. Oxley, 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 Anal. & Prev.*, vol. 37, no. 5, pp. 962–971, 2005.

[17] J. Oxley, B. Fildes, E. Ihsen, J. Charlton, and R. Day, "Differences in traffic judgements between young and old adult pedestrians," *Accident Anal. & Prev.*, vol. 29, no. 6, pp. 839–847, 1997.

[18] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann *et al.*, "Stanley: The robot that won the DARPA Grand Challenge," *J. of Field Robo.*, vol. 23, no. 9, pp. 661–692, 2006.

[19] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer *et al.*, "Autonomous driving in urban environments: Boss and the urban challenge," *J. of Field Robo.*, vol. 25, no. 8, pp. 425–466, 2008.

[20] A. Gorrini, G. Vizzari, and S. Bandini, "Towards modelling pedestrian-vehicle interactions: Empirical study on urban unsignalized intersection," *arXiv preprint arXiv:1610.07892*, 2016.

[21] F. Schneemann and I. Gohl, "Analyzing driver-pedestrian interaction at crosswalks: A contribution to autonomous driving in urban environments," in *Intel. Vehicles Sym.*, 2016, pp. 38–43.

[22] F. Camara, O. Giles, R. Madigan, M. Rothmüller, P. H. Rasmussen, S. Vendelbo-Larsen, G. Markkula, Y. Lee, L. Garach, N. Merat *et al.*, "Predicting pedestrian road-crossing assertiveness for autonomous vehicle control," in *Int. Conf. on Intel. Trans. Sys.*, 2018, pp. 2098–2103.

[23] S. Gupta, M. Vasardani, and S. Winter, "Negotiation between vehicles and pedestrians for the right of way at intersections," *Transactions on Intel. Trans. Sys.*, vol. 20, no. 3, pp. 888–899, 2018.

[24] C. Ackermann, M. Beggiato, L.-F. Bluhm, A. Löw, and J. F. Krems, "Deceleration parameters and their applicability as informal communication signal between pedestrians and automated vehicles," *Trans. Research Part F: Traffic Psy. and Behav.*, vol. 62, pp. 757–768, 2019.

[25] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Int. Sym. on Robo. Research*, 2011, pp. 3–19.

[26] D. Vasquez, T. Fraichard, and C. Laugier, "Growing hidden markov models: An incremental tool for learning and predicting human and vehicle motion," *The Int. J. of Robo. Research*, vol. 28, no. 11-12, pp. 1486–1506, 2009.

[27] G. Ferrer and A. Sanfeliu, "Bayesian human motion intentionality prediction in urban environments," *Pattern Recognition Letters*, vol. 44, pp. 134–140, 2014.

[28] H. Bai, S. Cai, N. Ye, D. Hsu, and W. S. Lee, "Intention-aware online POMDP planning for autonomous driving in a crowd," in *Int. Conf. on Robo. and Auto.*, 2015, pp. 454–460.

[29] T. Bandyopadhyay, C. Z. Jie, D. Hsu, M. H. Ang, D. Rus, and E. Frazzoli, "Intention-aware pedestrian avoidance," in *Experimental Robo.* Springer, 2013, pp. 963–977.

[30] M. Clamann, M. Aubert, and M. L. Cummings, "Evaluation of vehicle-to-pedestrian communication displays for autonomous vehicles," Trans. Research Board, WA, DC, USA, Tech. Rep., 2017.

[31] T. Lagstrom and V. M. Lundgren, "Avip-autonomous vehicles interaction with pedestrians," Master's thesis, Chalmers Uni. of Tech., 2015.

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

[33] C. P. Urmson, I. J. Mahon, D. A. Dolgov, and J. Zhu, "Pedestrian Notifications," https://patents.google.com/patent/US8954252B1/en, accessed: 2020-02-26.

[34] Mercede-Benz, "Mercedes-Benz F 015 Luxury in Motion," https://www.mercedes-benz.com/en/mercedes-benz/innovation/research-vehicle-f-015-luxury-in-motion/, accessed: 2020-02-24.

[35] Nissan Motor Corporation, "Nissan IDS Concept: Nissan's vision for the future of EVs and autonomous driving," https://www.alliance-2022.com/blog/nissan-ids-concept-nissans-vision-future-evs-and-autonomous-driving/, accessed: 2020-02-25.

[36] R. D. Smallwood and E. J. Sondik, "The optimal control of partially observable Markov processes over a finite horizon," *Operations Research*, vol. 21, no. 5, pp. 1071–1088, 1973.

[37] Wikipedia, "Preferred walking speed," https://en.wikipedia.org/wiki/Preferred_walking_speed, accessed: 2020-02-24.

[38] A. Somani, 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.

[39] N. Ye, A. Somani, D. Hsu, and W. S. Lee, "DESPOT: Online POMDP planning with regularization," *J. of Artificial Intelligence Research*, vol. 58, pp. 231–266, 2017.

[40] M. Quigley, J. Faust, T. Foote, and J. Leibs, "ROS: an open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, vol. 3, no. 3.2, 2009, p. 5.

[41] Y.-C. Hsu, S. Gopalswamy, S. Saripalli, and D. A. Shell, "An MDP model of vehicle-pedestrian interaction at an unsignalized intersection," *Vehicular Tech. Conf. (VTC-Fall)*, pp. 1–6, 2018.

[42] S. Gopalswamy and S. Rathinam, "Infrastructure enabled autonomy: A distributed intelligence architecture for autonomous vehicles," in *Intel. Vehicles Sym.*, 2018, pp. 986–992.