

MIT Open Access Articles

Probabilistic Risk Metrics for Navigating Occluded Intersections

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: McGill, Stephen G. et al. "Probabilistic Risk Metrics for Navigating Occluded Intersections." IEEE robotics and automation letters, vol. 4, no. 4, 2019, pp. 4322 - 4329 © 2019 The Author(s)

As Published: 10.1109/LRA.2019.2931823

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/126542>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Probabilistic Risk Metrics for Navigating Occluded Intersections

Stephen G. McGill¹, Guy Rosman¹, Teddy Ort², Alyssa Pierson², Igor Gilitschenski²,
Brandon Araki², Luke Fletcher¹, Sertac Karaman², Daniela Rus², John J. Leonard^{1,2}

Abstract—Among traffic accidents in the US, 23% of fatal and 32% of non-fatal incidents occurred at intersections. For driver assistance systems, intersection navigation remains a difficult problem that is critically important to increasing driver safety. In this paper, we examine how to navigate an unsignalized intersection safely under occlusions and faulty perception. We propose a real-time, probabilistic, risk assessment for parallel autonomy control applications for occluded intersection scenarios. The algorithms are implemented on real hardware and are deployed in a variety of turning and merging topologies. We show phenomena that establish go/no-go decisions, augment acceleration through an intersection and encourage nudging behaviors toward intersections.

Index Terms—Intelligent Transportation Systems; Human Factors and Human-in-the-Loop; Autonomous Vehicle Navigation

I. INTRODUCTION

INTERSECTIONS present one of the most challenging driving scenarios because a vehicle must interact with others to navigate safely. In 2016, 23% of fatal and 32% of non-fatal traffic incidents in the U.S. occurred at intersections¹. Recent advances in robotic perception and control routines promise to enhance the safety of passengers on the road through fully autonomous cars or by augmenting the human with advanced driver assistance systems (ADAS). However, reasoning about intersections remains a major challenge.

Perception systems may fail to detect other vehicles as a consequence of occluded views. Causes of these occlusions may stem from cross-traffic behavior, buildings or road geometries that cause poor visibility [2]. As a consequence, tracking and detecting other road objects may fail in unpredictable ways. New models must account for the resulting uncertainty and risk while ensuring robustness to imperfect perception and remaining computationally efficient.

Manuscript received: February 24, 2019; Revised May 27, 2019; Accepted July 10, 2019.

This paper was recommended for publication by Editor Youngjin Choi upon evaluation of the Associate Editor and Reviewers' comments. This work was supported in part by the Office of Naval Research (ONR) Grant N00014-18-1-2830, and by Toyota Research Institute (TRI). This article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity. Their support is gratefully acknowledged.

¹Toyota Research Institute, Cambridge, MA 02139, USA [stephen.mcgill, guy.rosman, luke.fletcher, john.leonard]@tri.global

²Computer Science & Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA [apierson, teddy, igilitschenski, araki, sertac, rus]@mit.edu

Digital Object Identifier (DOI): see top of this page.

¹Incidents totalled 6,462 fatal and 10,119 non-fatal crashes, per the National Highway Traffic Safety Administration [1].




Fig. 1: An unsignalized intersection with an occluded view, shown above, presents a risky scenario for performing unprotected turns across traffic. Below, vertical bars show the risk per discretized road segment, while horizontal bars show the occupancy estimation belief, described in Section II.

With respect to human drivers, collisions at intersections often occur due to inattention or misjudgment of the other cars' dynamics [3]. This remains an open problem for autonomous vehicles, which can struggle to navigate intersections without incident [4] or to interact naturally with cars driven by humans [5]. In this paper, we aim to address this challenge. Our goals include modeling the risk of collision at junctions and deploying both full and shared autonomous systems to enhance vehicle safety. Our driver-assistance framework embodies a parallel autonomy system, where the human maintains primary control of the vehicle, but the autonomous system can intervene for safety.

We propose a risk model that reasons about several crucial aspects of road interactions. It accounts for cross traffic, occlusions, sensor errors and driver attentiveness. This work focuses on intersections that do not contain a stop light to regulate the traffic flow, and where the ego-vehicle must yield to all cross

traffic. Figure 1 provides an example of such an occluded, unsignalized intersection that presents difficulty for human-operated and autonomous systems alike. In this scenario, the ego-vehicle attempts a left turn maneuver into crossing traffic. Houses abutting the intersection obstruct the views of the ego-vehicle and crossing vehicles, creating uncertainty in the estimation of each others' poses.

The proposed algorithm estimates the risk for the ego-vehicle and we demonstrate the use of this estimate in a parallel autonomy framework [6] to increase driver safety. Generalizing beyond the left turn scenario, the approach we present extends to various junction topologies, including roundabouts and merge lanes. The main contributions of the paper are summarized as follows:

- 1) A model and an algorithm to estimate the risk at intersections in the presence of occlusions and uncertain sensor measurements;
- 2) Integration of our estimator into an online, shared control method for negotiating intersections; and
- 3) Demonstration on a physical hardware platform in both fully autonomous and parallel autonomy modes.

Related Work: Human drivers negotiate intersections through arrival times of oncoming traffic [7]. Outlined by [8], usual intersection behavior consists of identifying gaps, assessing risk and mitigating threats. Gap estimation [9] is a crucial aspect in making go/no-go decisions at junctions, providing a common metric for making turns. In assessing risk, the critical times to cross an intersection also play a vital role in decision making [10]. This risk aware decision making process varies based on the attention that human drivers give to the situation [11] and these differences have a major influence on how active safety systems interact with human drivers [8], [12] to reduce threats. Notably, current gap estimation methods lack real time applicability [13].

Successful ADAS deployment requires accurate assessment of the intersection topology and turning prediction. Increased fidelity in lane-level maps can improve turning predictions [14], and studies on human drivers have identified typical distances for turning decisions [15]. Predicting when to warn the driver serves as a key feature of ADAS systems. In [16], the authors combine SVM-based intention prediction with an RRT-based threat assessment to evaluate potential risks and warn the driver. Augmented reality systems may prove useful in warning and assisting the driver [12], [17].

We use techniques from fully-autonomous control to improve our parallel autonomy system. Autonomous vehicle systems often utilize a model-based cost function in reasoning about lane changes and merges [18], [19]. Recent work shows promise in using deep reinforcement learning for intersection navigation [20]. In addition, if cars communicate via a vehicle-to-vehicle (V2V) network, higher-level intersection management techniques can mediate this traffic flow from reservations or auction systems [21]. Our presented work relies on local information and existing topology maps, instead of V2V communication.

Similar to previous work [22], [23], we address the problem of occluded traffic in the environment. Our method accommodates sensor and motion uncertainty, as opposed to over-

approximation of predicted occupancies and possible states in the set based verification method of [22]. Additionally, our model can estimate risk due to unlikely distracted or highly speeding traffic agents. Like [23], we provide a simple control mechanism for accelerating and decelerating in the presence of risk before and after an intersection, while approaches in similar occluded scenarios use only braking [24].

The remainder of the paper is organized as follows: in Section II, we describe our model of intersections and probabilistic approaches to traffic arrival, sensor observations and driver inattention. We quantify risk by using a probabilistic graphical model and considering the likelihoods at all potential cross traffic locations. Section III details our control strategy for leveraging the estimation of risk at an intersection. In Section IV, we describe our experimental design, including algorithm and discretization choices. We show risk estimation results for both simulated and real world scenarios. By testing on one-tenth model scale cars, we show the performance of the model for both fully autonomous and driver assistance (parallel autonomy) operating modes. Finally, we present our conclusions in Section V.

II. RISK ASSESSMENT MODEL

In this section, we describe our model for assessing risk at unsignalized intersections. We define unsignalized intersections as road junctions that do not have lights or signs mediating traffic. This includes an unprotected left turn across traffic, as well as merges and roundabouts. We represent the intersection as a junction node with lanes of traffic entering and exiting the node, and Figure 2 illustrates various topological representations. We compute our risk assessment by incorporating (i) traffic density, (ii) sensor noise and physical occlusions that hinder observations of other vehicles and (iii) attention limitations of other drivers. We use this risk assessment in determining a “go” or “no-go” decision at an intersection.

One common approach to intersection risk assessment is to use gap estimation, where the car uses a pre-selected critical gap size to determine its go and no-go decisions. This gap is defined by the spacing between two vehicles in the desired traffic lane, with the critical gap size being the minimum tolerated spacing for safe maneuvers [25]. When the car sees a critical gap opening, go and no-go decisions are equally likely [26]. Gap estimation algorithms perform well in intersections without occlusions and under the assumption of perfect observability. By incorporating occlusions and noisy observations, in Section IV we show our model outperforms a baseline gap estimation algorithm.

A. Modeling Risk

Consider an ego-car moving across an intersection. During its maneuver, it must move from an origin lane to a new lane, passing through a junction, while avoiding incidents with other vehicles. We define an incident to include collisions, near-miss braking incidents, traffic conflicts [27], small gap spacing [15], or the threat of accident from other vehicles [13]. The risk to the ego-car at time t , denoted r_t , is computed as the expected




Fig. 2: The ego-car traverses from its original lane (blue lines), over crossed lanes (green lines), to its new lane (dashed green and blue). Red circles mark lane origins, with thin red lines showing the ego-car’s path from the original lane to the new lane. Exemplary road distances, marked with d , and an occlusion, Ω_1 , are visualized. For each discretized road segment, shown in (b), we maintain the belief state $P(C_i)$ of other cars and the conditional risk $P(E_i|C_i)$.

number of incidents that will occur if the ego-car enters the intersection at time t .

For an intersection with n_l lanes, we discretize each lane $l \in \{1, \dots, n_l\}$, into n_s segments of length δ meters. For a segment i in lane l , the distance from the segment to the intersection is denoted d_i . We assume that traffic in each lane, l , has some average velocity v_l , such that for sampling intervals Δt , $\delta = v_l \Delta t$. For the ego-car, we denote its velocity as v_e and its distance to the intersection as d_e .

By construction, we compute the overall risk by summing over all segments in all lanes. Thus, we formulate the properties of incident probability, occupancy and observation noise per road segment. This formulation is Eulerian, computing risk properties for traversible space, unlike Lagrangian techniques [28], that reasons about vehicle tracks [29]. As the identity of the individual objects is unimportant, this Eulerian model provides computational efficiency [30].

We denote the event of an incident within segment i of lane l at time t as $E_{i,l}^t \in \{0, 1\}$. We condition $E_{i,l}^t$ on the occupancy of a segment, with the occupancy denoted $C_{i,l}^t \in \{0, 1\}$. We denote $O_{i,l}^t \in \{0, 1\}$ as the observation of another vehicle in segment i of lane l at time t . For clarity, when the target lane and time are fixed, we use the shorthand notation E_i and C_i . Figure 2 illustrates the lane topologies and maintained probabilities within an intersection.

The overall risk r_t is the sum over all segments and lanes,

$$r_t = \sum_{l=1}^{n_l} \sum_{i \in l} \mathbb{E}[E_{i,l}^t | O^{-t}], \quad (1)$$

where O^{-t} denotes the observations seen until time t over all lanes in n_l . Note that in this model, we consider the expected number of incidents, $\mathbb{E}[E_{i,l}^t | O^{-t}]$, in separate lanes and segments. In Section III, we use this risk to determine when it is safe for the vehicle to move through the intersection.




Fig. 3: A graphical model describing the belief about cars per road segment over time. Shaded nodes represent observations of vehicles, and a node represents the event of a vehicle occupying road segment i at time t .

Our formulation of approximate risk serves as a union bound on risky incidents, since the event of at least one risky incident is no more likely than the sum of all risky events. Thus, r_t is an upper bound on expected incidents. It becomes accurate in the case of low risk levels, where correct estimation of the risk is important. For high-risk situations, the ego-car should avoid the intersection, determined as risk r_t greater than a safety threshold, r_{go} .

B. Occupancy Estimation

Here, we discuss how the probability of occupancy depends on the velocity and traffic. We use a dynamic Bayesian network [31], which allows us to reason about probability and model the likelihood of occupied segments. This network is depicted in Figure 3. We initialize the probability of occupancy at the lane origin, $P(C_{0,l}^t)$, with emission rate λ_l since theoretical and real world studies show that the Poisson model captures vehicle arrivals at uncontrolled intersections [32], [33]. With $\lambda_l = 1$, road segments are marked occupied until a free space observation is made.

We define occupancy belief update as $P(C_i^{t+1})$, based on $P(C_i^t | C^t)P(C^t)$, where C^t denotes the occupancy over all surrounding segments, in the following manner. Lane velocities at each road segment are normally distributed, such that $v_l \sim \mathcal{N}(\hat{v}_l, \sigma_l)$. With K_m as a discrete Gaussian convolution kernel with variance σ_l and window size, m , we compute $P(C_i^{t+1} | C^t) = \sum_m P(C_{i-m}^t) \cdot K_m$. Thus, the estimate of occupancy, $P(C_i^{t+1})$, is conditioned on the occupancy belief of surrounding segments at the previous timestep.

C. Driver Attention

When the ego-car enters the intersection, other vehicles may not notice due to limited attention. Failure to react to the ego-car may result in an incident. Here, we present how our model of the conditional risk of an incident, $P(E_i|C_i)$, incorporates driver attention. Consider a vehicle at distance d_i from the intersection. Let t_c denote the time it takes the ego-car to clear




Fig. 4: The risk conditioned on an unaware car as a function of location, with an exact constant velocity (blue line), and with a velocity distribution (red error bars). It is reasonable to assume constant velocity for expected risk.

the intersection, which is a function of the path length through the intersection and velocity v_e .

We define the conditional risk as

$$P(E_i|C_i) = \begin{cases} 0 & d_i > v_l \cdot t_c \\ 1 & d_i < d_s \\ \exp(-\lambda_a(d_i - d_s)) & \text{otherwise} \end{cases}, \quad (2)$$

where $\lambda_a > 0$ is a parameter modeling the attention of the other vehicle and d_s is defined as the comfortable stopping distance. This exponential decay is motivated by [34]. If the other vehicle will take longer than t_c to arrive at the intersection, $P(E_i|C_i) = 0$. Note that, as t_c increases, so does the number of road segments needed to model the lane. Conversely, if the other vehicle is within its comfortable stopping distance d_s , then $P(E_i|C_i) = 1$. This does not imply a collision, but possible cut-offs, tailgating and other unsafe situations included in our definition of incidents. The conditional risk for an ego vehicle crossing a lane of traffic is shown in Figure 4.

We also define d_s^F as the required forward distance between the ego-car and other vehicles in the target lane. For incidents with other vehicles in the new lane, the conditional risk is an indicator function, $P(E_i|C_i) = \mathbf{1}(d_i < d_s^F)$. This implies the ego-car must also maintain a comfortable stopping distance behind the car it is following.

D. Occlusions and Noisy Observations

The occupancy probability depends on both noisy observations as well as line-of-sight occlusions that may prevent the ego vehicle from seeing the road segment. Road segments beyond the vehicle's sensor range are considered occluded. We compute the likelihood of an observation conditioned on the presence of a vehicle at segment i with probability $P(O_{i,l}^t|C_{i,l}^t)$ by incorporating both occlusions and a model of noisy observations. For an environment with n_o occlusions, let Ω_k for $k = \{1, \dots, n_o\}$ represent an occlusion. We define the occlusion as blocking some segments a to b from the ego car's view, $\Omega_k = [d_k^a, d_k^b]$. We refer to the set of all occlusions as Ω , and use the notation $i \in \Omega$ to refer to a segment i that is occluded from view. For occluded segments, $P(O_{i,l}^t|C_{i,l}^t, i \in \Omega) = 0$.

For unoccluded segments, we model standard perception and tracking pipelines with noisy observations. Observed cars are associated with a single segment in the lane, and a belief update is computed for unoccluded $i \notin \Omega$ as

$$P(C_i^t = 1|O_i^t, i \notin \Omega) = \frac{P(O_i^t|C_i^t = 1)P(C_i^t = 1)}{\sum_{j \in [0,1]} P(O_i^t|C_i^t = j)P(C_i^t = j)}, \quad (3)$$

where $P(C_i^t = 1)$ is computed based on the distribution of v_l , and the belief from the previous time step. In our experiments presented in Section IV, we set $P(O_i = 1|C_i^t = 1) = 0.85$ and $P(O_i = 1|C_i^t = 0) = 0.05$.

Using the probability of an event of an incident in each lane, we act based on the expected number of incidents. The expected risk from lane l is the inner product of the conditional risk and vehicle occupancy estimate,

$$\mathbb{E}_l[E^t|O^{-t}] = \sum_i P(E_i^t|C_i^t)P(C_i^t|O^{-t}), \quad (4)$$

where $P(C_i^t|O^{-t})$ denotes the belief of a car in segment i given observations until time t at all segments observed by the ego-car.

E. Combined Risk Algorithm

From (4), we sum the expected number of incidents over the new and crossed lanes to find the overall risk, r_t in (1), $r_t = \sum_{l=1}^{n_l} \sum_{i \in l} \mathbb{E}[E_{i,l}^t|O^{-t}]$. Algorithm 1 outlines our process for updating the conditional risk using driver attention, occupancy, occlusions and noisy observations. Observations O_i^t of a car in segment i are made for a discrete time t . Our risk model is evaluated every Δt seconds, and in both our simulations and experiments, $\Delta t = 0.1$. An example of our risk calculation in simulation is shown in Figure 6.

Algorithm 1 Risk model estimation

Input: Observations at time t

Output: Overall risk estimate at time t

for $t=1..$ **do**

Update car beliefs $P(C_i)$ according to (3)

Compute conditional risks $P(E_i|C_i)$ according to (2)

for all t_c **do**

for all l **do**

Integrate risk, finding $\mathbb{E}_l[E^t|O^{-t}]$, per (4)

end for

Integrate risk to find r_t according to (1)

end for

end for

Key variables in our model include vehicle velocities and environment visibility. Here, we discuss how our model behaves as these variables are taken to their limits.

1) *No Occlusions:* With no occluded regions, $\Omega = \emptyset$, the problem is akin to gap estimation. Here, our model also includes the uncertainty estimates from $P(C_i|O_i)$, $v_l \sim N(v_l^0, \sigma_l)$ and λ_a .

2) *Fully-Occluded View:* For $\Omega = \{(-\infty, 0]\}$, the other vehicles and ego-car cannot observe each other until the ego-car enters the intersection. The model then estimates $P(C_i) = \lambda_l$ for all points i since $P(O_i = 1) = 0$ when i is occluded, and the lane is initialized with likelihood λ_l for a Poisson emission rate. Thus, $r_t = \sum_l \lambda_l \sum_i P(E_i|C_i)$.




Fig. 5: Left: For all times to clear the intersection, t_c , the risk decreases as the ego-car approaches the intersection. Right: Shown are the occluded road segments (blue) and visible road segments (red) for various positions of the ego-car, d_e .

3) *High-Speed Traffic*: When other vehicles travel at high velocity, $v_l \rightarrow \infty$, the model yields the same expected risk as in the no observability case, since the ego car can never clear the intersection in time with $d_s > t_c \cdot v_l$.

4) *High-Speed Ego Car*: As $t_c \rightarrow 0$, the risk integration is performed over a single road segment per lane. Per (1), if the accumulated risk over intersecting lanes l becomes less than r_{go} , the car can occupy that road segment on the lane.

5) *Low-Speed Ego Car*: With $t_c \rightarrow \infty$, the model integrates risk over an infinite number of road segments. Thus, any other cars in view of the ego car, regardless of position, would prevent the ego car from entering the intersection.

III. SHARED CONTROL APPROACHES

We leverage our risk model for control policies before and after entering an intersection. Before the intersection, the ego-car uses the risk to determine if it can proceed, or if it must to nudge toward the intersection to reduce uncertainty. Once the ego-car enters the intersection, we use the risk to adjust its velocity along the path through the intersection. We summarize the overall control algorithm in Algorithm 2.

A. Nudging into Intersections

Figure 5 shows an ego-car approaching an intersection with occlusions on either side of its origin lane. The risk decreases as the car approaches the intersection, due to the increase in visible road segments. The risk is dependent on the time it takes the ego-car to clear the intersection, t_c . Figure 5 plots several values of t_c and the corresponding risks.

In the presence of occlusions or noisy observations, the ego-car should spend time in high visibility positions to reduce uncertainty in occupancy estimation. This desire to reduce uncertainty and to prevent risky go/no-go decisions motivates our nudging control policy,

$$v'_e = \begin{cases} \frac{d_e}{d_{\text{nudge}}} v_e & (r_t > r_{go}) \text{ and } (d_e < d_{\text{nudge}}) \\ v_e & \text{otherwise} \end{cases}, \quad (5)$$

where v'_e is the regulated velocity, bounded by the original desired velocity v_e , and d_{nudge} is the location where the vehicle begins to nudge towards the intersection. The value r_{go} provides a threshold for how much risk is tolerated before going into an intersection. When $d_e > d_{\text{nudge}}$, or the risk is low, the regulated velocity matches the commanded velocity, $v'_e = v_e$. Otherwise, the vehicle slows to a stop before the intersection, while gathering more observations.

B. Clearing the Intersection

Once the risk $r_t < r_{go}$, the ego-car enters the intersection and follows a pre-computed trajectory to clear the intersection. Let L_p represent the length of the remaining path, p , from the original lane to the new lane through the intersection. We assume the ego-car cannot move backwards along its path, thus L_p monotonically decreases over time. The minimum time it takes the ego-car to cross the intersection is given by $t_c^{\text{min}} = \frac{L_p}{v_{\text{max}}}$, for some maximum velocity v_{max} . The maximum time to clear the intersection, while keeping risk under the threshold r_{go} , is t_c^{max} , and depends on other traffic approaching the intersection. We impose an upper bound on t_c^{max} to require the ego-car to complete its turn and not to stop within the intersection. Thus, the regulated velocity after the ego-car enters the intersection is given as

$$v'_e = \max\left(\frac{L_p}{t_c^{\text{max}}}, v_e\right). \quad (6)$$

We combine the nudging policy and clearing policy together in Algorithm 2 for the ego-car's overall control policy through the intersection. In Algorithm 2, we assume the path through the intersection p is pre-defined, and a Pure Pursuit algorithm [35] is used to control the steering commands θ_e .

While the above methods establish low level control mechanisms, opportunities exist to utilize high level planners that utilize risk during planning e.g., to maximize t_c for comfort, as in [23].

IV. EXPERIMENTAL RESULTS AND VALIDATION

To evaluate this model, we perform both simulations and physical experiments on the one-tenth scale MIT racecar platform [36]. We show that full and parallel autonomy control

Algorithm 2 Ego-Car Control

Input: Turning path, p , desired velocity, v_e

Output: Steering, θ_e , and modified velocity, v'_e , commands

while $L_p > 0$ **do**

 Update the risk according to Algorithm 1

if Before intersection ($d_e < 0$) **then**

 Bound velocity v_e to v'_e according to (5)

else

 Bound velocity v_e to v'_e according to (6)

end if

 Find pure pursuit steering, θ_e , based on p

 Command vehicle according to v'_e and θ_e

end while

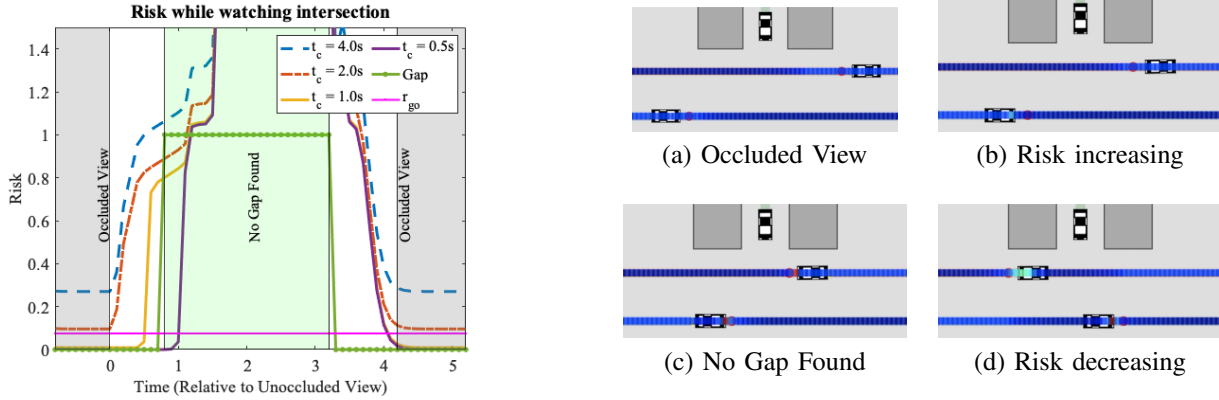


Fig. 6: Until (a), the ego-car sees only an occluded view of the traffic. At point (b), the ego-car observes cars, and the risk begins to increase. At point (c), the baseline system finds no gap. At point (d), the risk decreases, while the baseline system finds a gap between exiting vehicles. Our method assesses risk for multiple values of t_c , and establishes a go/no-go threshold, r_{go} for shared control.

systems, when leveraging our model’s integrated risk, can avoid accidents at occluded junctions.

A. Simulations

We simulate topologically distinct environments to prototype our methods. To verify risk calculations, we vary cross traffic densities (λ_l), velocities (v_l) and environment occlusions (Ω), while generating test trajectories for the ego-cars and other cars. At an unsignalized, occluded left turn junction, the ego car approaches an intersection with two lanes of cross traffic.

Topological adaptations, such as those shown in Figure 7, demonstrate the flexibility of our theory and how the risk model extends to other motions and occlusions. Mimicking blind merges exiting tunnels or on-ramps, the merging scenario requires the ego-car to enter a lane of traffic with occlusions blocking the view until near the merge. For the roundabout, the ego-car must enter the roundabout, but features within the roundabout occlude the car from knowing the positions of the other cars. We refer the reader to our video supplement for animations of risk, akin to Figure 6, for these topologies.

B. Experiments

To benchmark our system, we consider five separate scenarios, including two full autonomy scenarios, two parallel autonomy scenarios and a fully manual scenario. In the parallel autonomy and manual control cases, the expert human operator utilizes a first-person view from cameras mounted on the vehicle. Each autonomy scenario is evaluated with our risk-based controller and a baseline gap-acceptance controller. All scenarios include a pre-planned path, with the computer and human supplying velocity commands in time.

- **Human Expert:** An expert fully controls the vehicle velocity.
- **Full Autonomy (Ours):** The car uses Algorithm 2, setting the desired velocity v_e to v_l of the new lane.

- **Full Autonomy (Baseline):** The car uses the baseline gap-acceptance algorithm, similar to the warning system of [37], where other vehicles need to stop only within a time gap. In place of a warning, the vehicle stops. The time gap is chosen based on lane speed and car length.
- **Parallel Autonomy (Ours):** An expert sets the desired velocity, v_e , which the risk assessment system may override, per Algorithm 2.
- **Parallel Autonomy (Baseline):** An expert controls velocity, but the gap-acceptance system may override to stop the vehicle.

Tables I and II summarize the run time and success rates, where a success indicates that a vehicle completes a turn without incurring contact with another vehicle.

C. Full Autonomy

To validate our model on a fully autonomous system, we first implement a motion controller that executes turns at junctions. Given the task of moving across traffic into a new lane, the control policy executes a pure pursuit controller [35] to follow a path, with no ability to stop for cross traffic. Pure pursuit sets the steering angle, θ_e , but not velocity v_e , which is set to v_l in the case of full autonomy. It is set by the human in the case of parallel autonomy and manual control.

We then run two turning aids for making a go/no-go decision: a gap estimator and our risk model. We set $d_s = 0.2m$ as the gap size in the baseline estimator and set $d_s^F = 0.6m$ to accommodate the wheel-base length between vehicle pose and its front bumper. To model arrivals and attention, we set $\lambda_l = 0.05$ and $\lambda_a = 0.5$. We evaluated six t_c values between 2 and 4.5 seconds, with a risk tolerance of $r_{go} = 0.02$, such that the ego-vehicle will not wait indefinitely at the intersection. By applying (5) and (6), the controller will slow down or speed up the ego-car along the trajectory to minimize risk. For the gap estimator, if a gap in traffic exists, then the car will attempt to turn; otherwise, the ego velocity is set to zero.




Fig. 7: Merge (a) and roundabout (b) scenarios of Figure 2 have been tested in addition to left turns (c).

We tested making left turns into a crowded two lane roadway with an unsignalized intersection, shown in Figure 2b, and applied our control approach. To simulate a stream of traffic, we used six other vehicles driving in a loop at $v_l = 0.75$ m/s under ground truth information from a motion tracking system, shown in Figure 1. Virtual occlusions were added to preclude observations. A computer program chose the start time for the ego-car to approach the intersection uniformly at random, with no advance information about the positions of the looping cars. For every example, the ego car started approximately 75 cm before the junction, with its view occluded. The baseline gap estimation system completed 70% of left turns, while the risk based model successfully completed 90% of trials.

Slowing v_l to 0.5 m/s, the risk based model performance dropped to 70%. However, the gap estimation system could not complete a single trial without collision when other cars were in proximity. This is because the gap estimation system attempts a turn as soon as it sees a gap without regard for the risk of occluded vehicles. To test this, we increased d_s to 0.6m, which effectively causes the gap estimation system to require a larger gap in traffic. This resulted in the gap estimation system scoring 80%, while our risk estimation system maintained a 70% success rate.

TABLE I: Results Across Control Modes

	Trials	Success	Avg Time	Std Dev
Human Expert	20	90%	9.96s	5.25s
Parallel Autonomy (Ours)	20	85%	9.68s	3.36s
Full Autonomy (Ours)	30	77%	5.89s	1.29s

TABLE II: Comparing the Risk Model and the Gap Baseline

	Autonomy	Trials	Success	Avg Time	Std Dev
Ours	Full	30	77%	5.89s	1.29s
Baseline	Full	20	75%	3.79s	1.03s
Ours	Parallel	20	85%	9.68s	3.36s
Baseline	Parallel	20	65%	8.93s	3.43s
Ours	Full & Parallel	50	80%	7.41s	2.12s
Baseline	Full & Parallel	40	70%	6.36s	2.23s

D. Parallel Autonomy

To validate our model on a parallel autonomy system, we leverage the same motion controller as in the fully autonomous case. However, a human controls velocity v_e with a remote



Fig. 8: The ego-vehicle viewpoint from three cameras corresponds to the configuration shown in Figure 1. Note there is a vehicle to the left completely occluded by the blue house.

controller. For visual feedback, three camera streams are sent from the car to remote displays to provide the human with situational awareness, as shown in Figure 8. The human operator cannot see the physical cars, which ensures they share the occluded viewpoint from the vehicle’s sensors. On average, over 60 trials, the human took 9.96 seconds to execute a turn (including collisions). In contrast, for over 50 autonomous trials, the autonomous system took only 4.49 seconds to execute a turn. Human operators took far longer because they often missed opportunities that the autonomous controllers deemed safe.

As an additional test, the human was required to steer the vehicle in addition to choosing its speed, v_e . 80% of turns were completed safely without any aid. However, when the gap estimation system was active in the parallel autonomy mode, the human operator achieved only 60% success. A similar set of parallel autonomy trials with our risk based model allowed the human to safely complete 70% of turns. This experiment demonstrates that a parallel autonomy system that intervenes only on speed, with the human operator responsible for steering, does not perform well, since the vehicle controls are not coordinated. Parallel autonomy performed best when the human and system collaborated on the same control inputs.

We set v_l to 0.75 m/s for the parallel autonomy tests. In these tests, the gap estimation system enabled the human to complete 80% of trials without a collision, while our risk model based controller enabled the human to complete 100% of trials. The unaided human operator, controlling only the speed, performed almost as well with a success rate of 90% but took an average 22% longer to complete the maneuver.

V. CONCLUSIONS

In this paper, we have presented a probabilistic model for assessing risk at occluded intersections with uncertain perception data. Whereas previous approaches typically adopt a Lagrangian perspective, based on the trajectories of perceived vehicles, our Eulerian approach reasons about the risk for road segments. This enables us to handle occluded junctions efficiently with limited perceptual data. We have deployed the risk assessment algorithm on scale model cars, testing for both autonomous control and remote human interaction (parallel autonomy). When a vehicle enters an occluded intersection, the risk assessment algorithm augments the commanded inputs to improve the safety of maneuvers, while also reducing the amount of time the vehicle waits before crossing traffic.

For future work, we will explore integration of our risk metrics into complete planning approaches, including more

diverse vehicle platforms. Additionally, we plan to explore ways to adapt our algorithm with respect to additional road topologies and naturalistic data, as well as learning optimal model parameters.

REFERENCES

- [1] N. H. T. S. A. U.S. Department of Transportation, “Fatality Analysis Reporting System (FARS) Encyclopedia,” accessed 2018-08-28. [Online]. Available: <https://www-fars.nhtsa.dot.gov>
- [2] B. D. Greenshields, D. Schapiro, and E. L. Ericksen, “Traffic performance at urban street intersections,” Yale Bureau Highway Traffic, Tech. Rep., 1946.
- [3] E.-H. Choi, “Crash factors in intersection-related crashes: An on-scene perspective,” National Highway Traffic Safety Administration, Tech. Rep., 2010.
- [4] L. Fletcher, S. Teller, E. Olson, D. Moore, Y. Kuwata, J. How, J. Leonard, I. Miller, M. Campbell, D. Huttenlocher, *et al.*, “The MIT–Cornell collision and why it happened,” *Journal of Field Robotics*, vol. 25, no. 10, pp. 775–807, 2008.
- [5] A. Efrati, “Waymo’s Foes: Left Turns and the Mean Streets of Phoenix,” *The Information*, 2017. [Online]. Available: <https://www.theinformation.com/articles/waymos-foes-left-turns-and-the-mean-streets-of-phoenix>
- [6] W. Schwarting, J. Alonso-Mora, L. Paull, S. Karaman, and D. Rus, “Parallel autonomy in automated vehicles: Safe motion generation with minimal intervention,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1928–1935.
- [7] J. Caird and P. Hancock, “The perception of arrival time for different oncoming vehicles at an intersection,” *Ecological Psychology*, vol. 6, no. 2, pp. 83–109, 1994.
- [8] D. R. Ragland, S. Arroyo, S. E. Shladover, J. A. Misener, and C.-Y. Chan, “Gap acceptance for vehicles turning left across on-coming traffic: implications for intersection decision support design,” *UC Berkeley: Safe Transportation Research & Education Center*, 2006.
- [9] J. Caird and P. Hancock, “Left-turn and gap acceptance crashes,” *Human factors in traffic safety*, pp. 613–652, 2002.
- [10] G. Marti, A. H. Morice, and G. Montagne, “Drivers’ decision-making when attempting to cross an intersection results from choice between affordances,” *Frontiers in human neuroscience*, vol. 8, p. 1026, 2015.
- [11] P. J. Cooper and Y. Zheng, “Turning gap acceptance decision-making: the impact of driver distraction,” *Journal of safety research*, vol. 33, no. 3, pp. 321–335, 2002.
- [12] M. L. Rusch, M. C. Schall Jr, J. D. Lee, J. D. Dawson, and M. Rizzo, “Augmented reality cues to assist older drivers with gap estimation for left-turns,” *Accident Analysis & Prevention*, vol. 71, pp. 210–221, 2014.
- [13] M. S. Shirazi and B. T. Morris, “Looking at intersections: a survey of intersection monitoring, behavior and safety analysis of recent studies,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 1, pp. 4–24, 2017.
- [14] B. Tang, S. Khokhar, and R. Gupta, “Turn prediction at generalized intersections,” in *Intelligent Vehicles Symposium (IV), 2015 IEEE*. IEEE, 2015, pp. 1399–1404.
- [15] K. Smith, A. Thome, C. Blåberg, and J. Bårgman, “An invariant may drive the decision to encroach at unsignalized intersections,” in *Proceedings of the 5th International Driving Symposium on Human Factors in Driver Assessment and Design*. University of Iowa, 2009.
- [16] G. S. Aoude, B. D. Luders, K. K. Lee, D. S. Levine, and J. P. How, “Threat assessment design for driver assistance system at intersections,” in *Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on*. IEEE, 2010, pp. 1855–1862.
- [17] C. Tran, K. Bark, and V. Ng-Thow-Hing, “A left-turn driving aid using projected oncoming vehicle paths with augmented reality,” in *Proceedings of the 5th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, 2013, pp. 300–307.
- [18] J. Levinson, J. Askeland, J. Becker, J. Dolson, D. Held, S. Kammel, J. Z. Kolter, D. Langer, O. Pink, V. Pratt, *et al.*, “Towards fully autonomous driving: Systems and algorithms,” in *Intelligent Vehicles Symposium (IV), 2011 IEEE*. IEEE, 2011, pp. 163–168.
- [19] A. G. Cunningham, E. Galceran, R. M. Eustice, and E. Olson, “MPDM: Multipolicy decision-making in dynamic, uncertain environments for autonomous driving,” in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 1670–1677.
- [20] D. Isele, R. Rahimi, A. Cosgun, K. Subramanian, and K. Fujimura, “Navigating occluded intersections with autonomous vehicles using deep reinforcement learning,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2034–2039.
- [21] D. Carlino, S. D. Boyles, and P. Stone, “Auction-based autonomous intersection management,” in *Intelligent Transportation Systems-ITSC), 2013 16th International IEEE Conference on*. IEEE, 2013, pp. 529–534.
- [22] P. F. Orzechowski, A. Meyer, and M. Lauer, “Tackling occlusions & limited sensor range with set-based safety verification,” in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 1729–1736.
- [23] M.-Y. Yu, R. Vasudevan, and M. Johnson-Roberson, “Occlusion-aware risk assessment for autonomous driving in urban environments,” *IEEE Robotics and Automation Letters*, 2019.
- [24] M. Schratte, M. Bouton, M. J. Kochenderfer, and D. Watenig, “Pedestrian collision avoidance system for scenarios with occlusions,” *arXiv preprint arXiv:1904.11566*, 2019.
- [25] W. Brilon, R. Koenig, and R. J. Troutbeck, “Useful estimation procedures for critical gaps,” *Transportation Research Part A: Policy and Practice*, vol. 33, no. 3–4, pp. 161–186, 1999.
- [26] R. Troutbeck, “Estimating the mean critical gap,” *Transportation Research Record: Journal of the Transportation Research Board*, no. 2461, pp. 76–84, 2014.
- [27] S. Older, T. J Shippey, and R. R. Laboratory, “Traffic conflict studies in the United Kingdom,” in *Proceedings: First Workshop on Traffic Conflicts*, C. H. Finn Harald Admundsen, Ed., vol. 1, November 1977, pp. 1–16.
- [28] Z. Sun, W.-L. Jin, and S. G. Ritchie, “Simultaneous estimation of states and parameters in newells simplified kinematic wave model with Eulerian and Lagrangian traffic data,” *Transportation Research Part B: Methodological*, vol. 104, pp. 106–122, 2017.
- [29] A. Petrovskaya and S. Thrun, “Model based vehicle detection and tracking for autonomous urban driving,” *Autonomous Robots*, vol. 26, no. 2–3, pp. 123–139, 2009.
- [30] E. A. Lee and M. Sirjani, “What good are models?” in *International Conference on Formal Aspects of Component Software*. Springer, 2018, pp. 3–31.
- [31] T. Dean and K. Kanazawa, “A model for reasoning about persistence and causation,” *Computational intelligence*, vol. 5, no. 2, pp. 142–150, 1989.
- [32] V. Rengaraju and V. T. Rao, “Vehicle-arrival characteristics at urban uncontrolled intersections,” *Journal of transportation engineering*, vol. 121, no. 4, pp. 317–323, 1995.
- [33] R. J. Troutbeck and W. Brilon, “Unsignalized intersection theory,” *Revised Monograph on Traffic Flow Theory*, 1997.
- [34] J. S. McCarley, C. D. Wickens, J. Goh, and W. J. Horrey, “A computational model of attention/situation awareness,” in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 46, no. 17. SAGE Publications Sage CA: Los Angeles, CA, 2002, pp. 1669–1673.
- [35] R. C. Coulter, “Implementation of the pure pursuit path tracking algorithm,” Carnegie-Mellon UNIV Pittsburgh PA Robotics INST, Tech. Rep., 1992.
- [36] S. Karaman, A. Anders, M. Boulet, J. Connor, K. Gregson, W. Guerra, O. Guldner, M. Mohamoud, B. Plancher, R. Shin, and J. Vivilechia, “Project-based, collaborative, algorithmic robotics for high school students: Programming self-driving race cars at MIT,” in *2017 IEEE Integrated STEM Education Conference (ISEC)*, March 2017, pp. 195–203.
- [37] S. Lefèvre, C. Laugier, and J. Ibañez-Guzmán, “Evaluating risk at road intersections by detecting conflicting intentions,” in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 4841–4846.