An occlusion- and interaction-aware safe control strategy for autonomous vehicles

Siddharth Gangadhar[1] * **Zhuoyuan Wang**[1] * **Kofi Poku** *
Naoya Yamada † **Kohei Honda** † **Yorie Nakahira** *
Hiroyuki Okuda † **Tatsuya Suzuki** †

* *Electrical and Computer Engineering Department, Carnegie Mellon University, PA 15213, USA (e-mail: {sgangadh, zhuoyuaw, kpoku, ynakahir}@andrew.cmu.edu).*
† *Mechanical Systems Engineering Department, Nagoya University, Japan, (e-mail: {yamada.naoya.n9, honda.kohei.b0}@s.mail.nagoya-u.ac.jp, {h_okuda, t_suzuki}@nuem.nagoya-u.ac.jp)*

Abstract: We consider the problem of safe autonomous driving in the presence of occlusions. Dealing with latent risks arising from occlusions is challenging because there does not exist direct mapping from sensor input to visible threats; attempts to ensure safety for all worst-case latent threats can be infeasible or overly conservative, and accounting for a multitude of latent risks for sufficient future horizon may require prohibitive computation in real-time. To address these issues, in this paper, we propose to use a probability-based predictive controller to make safe decisions for autonomous vehicles. We prove that the proposed safety controller can generate vehicle control profiles that yield the desired safety probability. Numerical and onboard experiments on a visual occluded pedestrian crossing scenario verifies the efficacy of the proposed method in real-time. The merits of the proposed control strategy include being able to guarantee long-term safety under occlusions without being over-conservative, handling latent risks caused by on-road interactions in real-time, and ease of design with transparency to the exposed risks.

Keywords: Autonomous vehicles, Human and vehicle interaction, Occlusion-aware control, Safe motion control, Adaptive and robust control of automotive systems.

1. INTRODUCTION

Visual occlusions impose huge challenges to autonomous driving because most sensors can not see through opaque objects, leading to large unobserved regions and potentially unsafe behaviors of the ego vehicle (Yu et al., 2019; Poncelet et al., 2020). Besides, the stochastic nature of all road users (other vehicles, pedestrians, *etc.*) introduces uncertainties into the system, which further increases the difficulty of dealing with occlusions (Zhang and Fisac, 2021). Given such uncertainties caused by occlusions and complex interactions between road users, designing a safe controller for the ego vehicle is very difficult (Koç et al., 2021). In this study, we focus on the problem of safe autonomous driving under scenarios with visual occlusions. The challenge of this problem includes

(1) It is hard to control the level of safety of autonomous vehicles when there are visual occlusions and potential interactions with other road users, because the latent risks are difficult to measure (Kahn et al., 2022).
(2) Considering all aspects of the uncertainties in the system with long time horizons could be computationally intractable, which imposes difficulties for applications in real-world scenarios (Lyu et al., 2021).
(3) Approximately accounting for the latent risks will lead to over-conservative behaviors which compromise performance to certain extents (Brüdigam et al., 2021).

To resolve the abovementioned issues, we propose a model-based probabilistic safe control strategy to regulate the vehicle's speed and steering profiles under visual occlusions. Inspired by (Gangadhar et al., 2022), the proposed method encodes future safety information into a probability value and imposes a linear constraint on the control input to guarantee the long-term safety of the system. The resulting optimization-based controller can be solved efficiently via quadratic programs (QPs) while meeting other goals and constraints. The technical merits of the proposed method are

(1) Guaranteed long-term safety accounting for latent risks that are invisible and occluded (see Theorem 1).
(2) Balancing competing for safety and performance objectives, ensuring robustness to large uncertainty

[1] These authors contributed equally.

[2] This project is funded in part by Carnegie Mellon University's Mobility21 National University Transportation Center, which is sponsored by the US Department of Transportation, in part by Manufacturing Futures Institute, in part by Oracle Cloud credits and related resources provided by the Oracle for Research Program, and in part by JST, PRESTO Grant Number JPMJPR2136, Japan.

without being over-conservative (see Fig. 7 and Fig. 8).
(3) Ease of design and transparency to the exposed risks (see the proposed optimization-based controller (17), remark 1 and remark 2).
(4) Fast real-time response that ensures long-term safety using onboard resources (see section 5.2 for real-time hardware implementation).

The rest of the paper is organized as follows. In section 2, we discuss the related works, in section 3, we formulate the safe control problem of interest, in section 4, we introduce the proposed occlusion-aware control framework, and in section 5 we present the numerical and on-board simulations to validate the proposed method, and at the end, we conclude the paper in section 6.

2. RELATED WORKS

Previous studies for the safe control of vehicles in the presence of occlusions can be categorized into the following approaches.

(1) External perception from infrastructures. With external perceptions from infrastructures, the autonomous vehicle directly gets information from other vehicles behind the occlusion, thus can execute safe control (Müller et al., 2022). However, such pervasive perception requires expensive infrastructure systems, but autonomous vehicles may also need to operate in regions when such infrastructure is unavailable. The proposed method compasses this issue by leveraging the system model and data to acquire the latent risk of any given road situation with only onboard sensors.
(2) Learning-based control. The learning-based method leverages the data from expert drivers and builds a mapping from the vehicle sensory input to the desired control of the vehicle (Isele et al., 2018; Sama et al., 2020). Even if the amount of data is enormous, generalization to the different scenarios is not guaranteed. The relationship between occlusion and vehicle motion is obscure due to the black-box nature of neural networks. In comparison, the proposed method characterizes the exact risk of different scenarios and can produce desired safety probability as specified.
(3) Partially observed Markov decision process (POMDP). POMDP uses a belief state to represent the latent dynamics of occluded vehicles and solve the optimal control (Hubmann et al., 2019). This framework considers the uncertainty of perception, but the computation is expensive for continuous control and often cannot be implemented in real-time (Zhang and Fisac, 2021). On the other hand, the proposed method only requires solving a quadratic program to get the vehicle control, which has efficient online implementation.

The features of the existing methods and the proposed method is summarized in table 1.

3. PROBLEM FORMULATION

In this section, we introduce the general problem statement, which includes the vehicle dynamics in section 3.1, interaction model in section 3.2, the occlusion model in 3.3, and the safety specifications in section 3.4.

Table 1. Features of the existing methods and the proposed method.

Method	On-board sensing	Transparency in design	Real-time computation
Infrastructures		✓	✓
Learning-based	✓		✓
POMDP	✓	✓	
Proposed	✓	✓	✓

3.1 Vehicle Dynamics and Nominal Control

We consider a general discrete-time control-affine dynamical model for the vehicle as follows:

$$x_{k+1} = f(x_k) + g(x_k)u_k \tag{1}$$

where $x \in \mathbb{R}^n$ is the vehicle's state, $u \in \mathbb{R}^m$ is the control input, and $f : \mathbb{R}^n \to \mathbb{R}^n$ and $g : \mathbb{R}^n \to \mathbb{R}^{n\times m}$ encompass the system dynamics, and k is the time step. We consider discrete-time dynamics throughout the paper because all real-world vehicle controls are achieved with digital systems. One can discretize any continuous dynamics into the form of (1) as in (Ogata, 1995). The choice of model can range from simple double-integrators (Liang and Peng, 2000) to complete 6 DoF models (Kiencke and Nielsen, 2000). The control action u is determined by a predetermined control law $N : \mathbb{R}^n \to \mathbb{R}^m$:

$$u = N(x) \tag{2}$$

This nominal controller will satisfy desired performance specifications, such as ensuring that the vehicle follows a planned trajectory and can be obtained via MPC, backstepping, machine learning, or other techniques but may not guarantee safety. The closed-loop vehicle dynamics with the nominal controller will be:

$$x_{k+1} = f(x_k) + g(x_k)N(x_k) \tag{3}$$

The specific realizations of the vehicle model and the nominal controller for experiments are introduced in section 5.

3.2 Interaction Model

We model road users' behavior as a combination of decision-making and motion dynamics. The decision-making model characterizes the high-level decisions of the agent based on the surrounding situation and the context. Let $X \in \mathbb{R}^z$ be the joint state of all agents involved in the interaction, $\mathcal{Z}$ be the external factors that affect the decision-making, such as physical context, traffic characteristic and social contexts, as described in (Rasouli and Tsotsos, 2019), and D be the decision-making function that outputs a distribution of the intentions of the agent (e.g., go/wait, lane-keep/lane-change) with respect to the joint state X conditioned on the context $\mathcal{Z}$. The general decision-making process can be formulated as follows:

$$d_k \sim D(X_k|\mathcal{Z}_k) \tag{4}$$

where d is the road users' decision, and we assume that d can take on a finite number of decision values. In practice, this decision-making process can be modeled as a finite state machine (Kielar et al., 2014), a POMDP (Hubmann et al., 2018), an interactive multiple model (IMM) filter (Burger et al., 2020), or a neural network (Rasouli et al., 2017).

The motion model characterizes the agents' behavior given the intention d. Specifically, the motion model is written as:

$$X_{k+1} \sim f_z(X_k \mid d_k) \tag{5}$$

where f_z is the distribution of the state update function of the agents. Previous studies have used social force model (Helbing and Molnár, 1995) and recurrent neural network (Camara et al., 2020) to represent the motion model f_z of the road users. For a given decision d, many models assume that f_z for each state X follows a Gaussian distribution, as different sources of noise and uncertainties will add up to a Gaussian due to the Central Limit Theorem (Johnson, 2004). The specific formulation of the interaction model used for the case study is described in section 5.

3.3 Occlusion Model

Occlusion is defined by the space where the ego vehicle is not visible. Visibility, here, broadly includes images and videos, radars, sonars, and other sensing devices. Occlusion $\mathcal{H}_k$ is defined in map space $\mathcal{M}$, where $\mathcal{O}_t$ is the occupied space by objects and $\mathcal{V}(x(k), \mathcal{O}_k)$ is the visible space in the field of view (FOV) of the ego vehicle at time k. Then the occlusion $\mathcal{H}_k$ is defined as follow:

$$\mathcal{H}_k = (\bar{\mathcal{O}}_k \cap \bar{\mathcal{V}}(x(k), \mathcal{O}_k)) \in \mathcal{M} \tag{6}$$

where $\bar{\mathcal{O}}$ and $\bar{\mathcal{V}}$ are the exclusive space of $\mathcal{O}$ and $\mathcal{V}$ from the map space $\mathcal{M}$, and the occlusion $\mathcal{H}_k$ is the space excluding obstacle and visible spaces in the map. In Eq. (6), the method for identifying $\mathcal{O}_k$ and $\mathcal{V}(x(k), \mathcal{O}_k)$ is highly dependent on the configuration of the sensor. With the use of LiDAR, it is typical that $\mathcal{O}_k$ is detected using neural networks (Lang et al., 2019), and $\mathcal{V}(x(k), \mathcal{O}_k)$ is calculated virtually by ray casting algorithm (Zhang et al., 2019). Although it is expected that the infrastructure-to-vehicle (I2V) or vehicle-to-vehicle(V2V) communication systems can compensate a part of the occlusion (Müller et al., 2022), not all occlusions can be covered in various driving scenes. Occlusion detection is out of scope in this paper. However, the proposed method can incorporate the size and shape of the detected occlusion as parameters.

3.4 Safety Specification

Our goal is to ensure the long-term safety of all road users. We assume that there are B safety specifications for the overall interaction system, indexed by $j \in \{1, 2, \cdots, B\}$, and each specification is represented as follows: specification j is defined by the event

$$\mathcal{C}_j = \{X \in \mathbb{R}^z : \phi_j(X) \geq 0\}, \tag{7}$$

where $\phi_j(X) : \mathbb{R}^z \to \mathbb{R}$ is a continuous mapping. The definition can capture various safety requirements in autonomous driving, e.g., all road users do not collide with each other, and the vehicle's speed should be less than a certain value when it is close to other vehicles. Let

$$S = \{X_\tau \in \mathcal{C}_j, \forall \tau \in \{k, k+1, \cdots, k+T\}, \forall j\}, \tag{8}$$

where T is the outlook time horizon. The long-term safety we aim to ensure is defined as

$$\mathbb{P}(S) \geq 1 - \epsilon, \ \forall k \geq 0. \tag{9}$$

We will present the specific choice of safe event used in the experiments in section 5.

4. PROPOSED METHOD

In this section, we present the safe condition to ensure long-term safety in section 4.1 and its realization as the safe occlusion-aware control algorithm in section 4.2.

4.1 Condition for Assuring Safety

In this subsection, we present a sufficient condition for the long-term safety specifications (9). Let

$$\mathbf{\Psi}(I) := \mathbb{P}(S|I) \in \mathbb{R} \tag{10}$$

be the sequence of probability of event S conditioned on the information I. We define a new notion of conditional discrete-time generator as below.

Definition 1. (Conditional discrete-time generator). The conditional discrete-time generator A of a discrete-time stochastic process $\{x_k\}_{k\in\mathbb{Z}_+}$ conditioned on another process $\{y_k\}_{k\in\mathbb{Z}_+}$ with sampling interval Δt evaluated at time k is given by

$$A\phi(x_k|y_k) = \frac{\mathbb{E}[\phi(x_{k+1})|y_k] - \mathbb{E}[\phi(x_k)|y_k]}{\Delta t} \tag{11}$$

whose domain is the set of all functions $\phi : \mathbb{R}^n \to \mathbb{R}$ of the stochastic process.

When $x_k = y_k$, this generator becomes the discrete-time counterpart of the continuous-time infinitesimal generator. We add the conditioning of y_k to capture the ego vehicle's limited information due to occlusions. Although the value of $A\phi(y_k)$ depends on both x_k and y_k, with a slight abuse of notation, for the rest of the paper, we will use $A\phi(y_k)$ where the discrete-time stochastic process x_k in Definition 1 is the full state of the interaction system, *i.e.,* X_k in (5).

Let Q_k be the information that the ego vehicle can acquire at time k. This information Q_k will be $Q_k = [x_k, x_k^o]$ with x_k^o being the observed state of all other road users by the ego vehicle at time k. Note that $Q_k = x_k$ if no other road users appear from the occlusions.

We consider the following condition at all time k:

$$A\mathbf{\Psi}(Q_k) \geq -\gamma(\mathbf{\Psi}(Q_k) - (1-\epsilon)), \ \forall k \geq 0. \tag{12}$$

Here, $\gamma : \mathbb{R} \to \mathbb{R}$ is a function that satisfies the following 2 design requirements:

Requirement 1: $\gamma(h)$ is linear and increasing in h.
Requirement 2: $\gamma(h) \leq h$ for any $h \in \mathbb{R}$.

The probability measure of $\mathbb{P}(S|I)$ is taken over X, the global state, conditioned on Q, the information that can be accessed by the ego vehicle. Therefore, the values on both sides of (12) can be computed using Q.

Theorem 1. Consider systems (1) and (5). We assume the initial condition $x_0 = x$ satisfies $\mathbb{P}(S|x_0 = x) \geq 1 - \epsilon$. If at each time k, the ego vehicle generates a control policy that satisfies (12), then the following condition holds:

$$\mathbb{P}(S) = \mathbb{E}[\mathbb{P}(S|x_k)] \geq 1 - \epsilon, \ \forall k \geq 0. \tag{13}$$

See (Jing and Nakahira, 2022; Wang et al., 2021) for the proof. Theorem 1 says the long-term safety of the system is guaranteed by the proposed safe condition (12) for all time with desired probability.

4.2 Proposed Safe Occlusion-Aware Control

In this section, we propose a control strategy that imposes (12) to ensure long-term safety of the system. We start by approximating $A\boldsymbol{\Psi}(Q)$, the infinitesimal generator of long-term safety. Since only the ego's vehicle's state can be controlled, with a slight abuse of notation, we use $\boldsymbol{\Psi}(x)$ to represent $\boldsymbol{\Psi}(Q)$ in the control design phase for the rest of the paper. This $\boldsymbol{\Psi}(x)$ will refer to different $\boldsymbol{\Psi}(Q)$ under specific situations (*e.g.,* there are no other road users in sight, or pedestrians are currently crossing the street). Let $\mathcal{D}(x) \in \mathbb{R}^n$ denote the finite-difference approximation of the gradient $\nabla_x \boldsymbol{\Psi}(x)$, *i.e.,*

$$\mathcal{D}_j(x) = \frac{\boldsymbol{\Psi}(x + \Delta e_j) - \boldsymbol{\Psi}(x - \Delta e_j)}{2\Delta} \tag{14}$$

where $\mathcal{D}_j$ is the j^{th} element of $\mathcal{D}$, Δ is the step-size, and e_j denotes a vector that takes a scalar value of 1 in the j^{th} entry and 0 otherwise.

Lemma 1. If $A\boldsymbol{\Psi}(x)$ exists, then:

$$A\boldsymbol{\Psi}(x) = \lim_{\Delta \to 0} \mathcal{D}_j^T(x) \left(f(x) + g(x)u\right). \tag{15}$$

Lemma 1 is a result of the chain rule, with the left-hand side of (15) being the time derivative of $\boldsymbol{\Psi}(x)$, and the right-hand side being the state derivative of $\boldsymbol{\Psi}(x)$ multiplies dx/dt which is the dynamics (1). With this, we obtain the inequality constraint on the control-action u:

$$-\mathcal{D}_j^T(x) g(x) u \leq \mathcal{D}_j^T(x) f(x) + \gamma\left(\boldsymbol{\Psi}(x) - (1 - \epsilon)\right) \tag{16}$$

Utilizing the safety condition (16) and the nominal controller N from section 3.1, we can formulate the safe controller $K : \mathbb{R}^n \to \mathbb{R}^m$ as a constrained quadratic optimization problem:

$$\begin{aligned} K(x) \coloneqq \; & \underset{u \,\in\, \mathbb{R}^m}{\arg\min} \quad \|u - N(x)\|_2 \\ & \text{s.t.} \qquad\quad (16) \end{aligned} \tag{17}$$

The optimization problem penalizes deviation from the nominal control action (minimally invasive) while ensuring the specified constraints are satisfied, complying with requisite safety specifications.

Remark 1. The proposed optimization-based safe control (17) is easy to design and implement because it only contains function γ and the desired risk tolerance ϵ as tunable parameters and only imposes linear constraints on control, which forms an efficient quadratic program (QP).

Remark 2. The variable $\boldsymbol{\Psi}(x)$ has the physical meaning of the safety probability of the system in the long term. Its value at x indicates how risky the system will be in the future, evolving from state x. This property of $\boldsymbol{\Psi}(x)$ can also guide the control design when necessary (*e.g.,* one can directly specify the expected future state and control based on $\boldsymbol{\Psi}(x)$ when the control constraint (16) is numerically infeasible).

4.3 Algorithm Description

We present the overall safe control strategy in Algorithm 1. The safety probability, in procedure $\boldsymbol{\Psi}(x)$, is numerically estimated using Monte-Carlo simulations; we loop over the number of specified MC-episodes (N_E), and at the k^{th} iteration, we do the following. In line 3, we initialize the safety check p_k that switches to 0 when a violation is detected. In line 4, we initialize the state of vehicle dynamics initial value problem (IVP) with the current state estimate of the actual vehicle. Next, in line 5, we jointly solve the time-invariant closed-loop vehicle dynamics IVP (3) with the initial condition from line 4 and the interaction motion model (5) over the specified time interval $\mathcal{T} = \{k, k+1, \cdots k+T\}$, giving us a forward rollout. Since both the vehicle and the interaction models are time-invariant, the start and end times of the interval $\mathcal{T}$ are irrelevant. In line 7, we check for a safety violation in the forward rollout. Finally, at the end of the procedure, we compute and return the mean of p_k over all the episodes. Since $\boldsymbol{\Psi}(x)$ gives the safety probability of the system over the time horizon $\mathcal{T}$, it encodes information of prediction on the future as well as the levels of uncertainty.

Algorithm 1 Occlusion and interaction-aware safe controller

Input: x ▷ Vehicle State
Output: u ▷ Safe Control Action
Parameters: T, N_E ▷ Preview Horizon, # Episodes

1: **procedure** $\boldsymbol{\Psi}(x)$
2: **for** $k \in \{1, 2, \ldots, N_E\}$ **do** ▷ MC Episodes
3: $p_k \leftarrow 1$ ▷ Initialize safety check
4: $x_0 \leftarrow x$ ▷ Initialize state
5: Solve (3) and (5) in $\mathcal{T}$ ▷ Forward Rollout
6: **if** not S **then**
7: $p_k \leftarrow 0$ ▷ Safety Violation
8: **end if**
9: **end for**
10: **return** $\frac{1}{N_E} \sum_{k=1}^{N_E} p_k$ ▷ Safe probability
11: **end procedure**
12:
13: **procedure** $K(x)$
14: $u_N \leftarrow N(x)$ ▷ Compute nominal control action
15: Obtain $\mathcal{D}(x)$ using (14) ▷ Gradient of $\boldsymbol{\Psi}$
16: Obtain u by solving QP (17) with constraint (16)
17: **return** u
18: **end procedure**

Procedure $K(x)$ encompasses the constrained optimization controller outlined in (17), which involves evaluating the nominal control action u_N in line 14, computing the finite difference approximation of the gradient of safe probability $\mathcal{D}(x)$, and finally obtaining the safe control action by solving the QP (17). This control action ensures that condition (9) is met.

Therefore, Algorithm 1 can account for long-term safety and guarantees to steer system trajectories toward the direction of non-decreasing long-term safe probability in the presence of latent risks, eliminating potential myopic decision-making, typically seen in traditional safe control techniques (Ames et al., 2019), that may result in unsafe behaviors in the future.

5. SIMULATIONS AND EXPERIMENTS

For the remainder of this paper, we focus on a case study of the safe control strategy on a pedestrian-vehicle interaction scenario at occluded crossing intersections. We introduce the design of the simulation in section 5.1, the experiments on a 1/10th scale autonomous vehicle in section 5.2, and present the results in section 5.3. We point

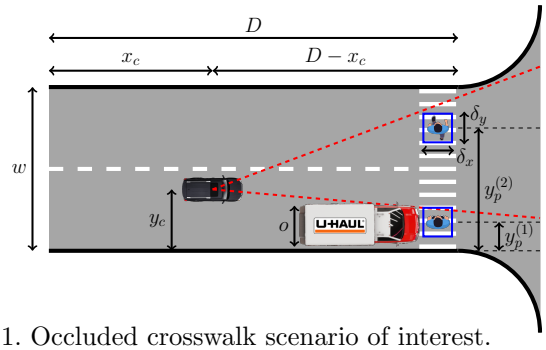

Fig. 1. Occluded crosswalk scenario of interest.

out that even though we have chosen to examine the safety of pedestrian-vehicle interactions at crossing intersections as the case study, the proposed control strategy can be applied to different road scenarios without redesign.

5.1 Case Study Scenario

We consider a case of an intersection with an unsignalized marked crosswalk (zebra crossing) as seen in Fig. 1. We use local coordinates whose origin is D away from the edge of the crosswalk. Here, the vehicle's position is (x_c, y_c) , the distance from the car to the crosswalk is $D - x_c$, and the i^{th} pedestrian's position along the crosswalk is denoted by $y_{\text{p}}^{(i)}$. We assign a bounding box to each pedestrian with dimensions (δ_x, δ_y), centered around the pedestrian's position, as seen in Fig. 1. The obstruction, which is a parked truck, has size o and occludes pedestrians' presence from the field of view of the vehicle until they have moved past it.

1) Vehicle Model: For the simulation, we start by setting up the dynamics of the vehicle based upon (Liang and Peng, 2000). Assuming the vehicle is a point-mass moving along a straight line, we get the following dynamics

$$m\frac{dv}{dt} = F_t - F_r, \tag{18}$$

where m is the mass of the vehicle, v is the longitudinal velocity of the vehicle, F_t is the tire force generated by the engine/motor, and F_r is the net drag force due to the tire's rolling resistance and aerodynamic drag. Here, we assume F_t to be our control input. With this, we can write down the governing equations of motion as an ODE. Let $x = [x_1 \ x_2]^T$, where x_1 is the vehicle's position, and $\dot{x}_1 = x_2 = v$ the longitudinal velocity. With this, assuming a time-step of Δt, we get the following discrete-time system dynamics using the forward-Euler method

$$x_{k+1} = \underbrace{x_k + \Delta t \begin{bmatrix} x_{2k} \\ -\frac{1}{m}F_r \end{bmatrix}}_{f(x_k)} + \underbrace{\Delta t \begin{bmatrix} 0 \\ \frac{1}{m} \end{bmatrix}}_{g(x_k)} u_k. \tag{19}$$

2) Nominal Cruise Controller: The cruise controller maintains a set cruising speed while ensuring comfortable acceleration/deceleration. When a pedestrian is visible in the vehicle's field of view, the cruise controller attempts to reduce the vehicle's speed based on the calculated time-to-collision T_{TTC} to that pedestrian

$$T_{\text{TTC}} = \frac{r}{\max\left(-\dot{r}, 0^+\right)}, \tag{20}$$

where 0^+ is a small positive constant and r is the estimated range to the pedestrian, note that the vehicle's auxiliary automatic emergency braking system will take precedence over the nominal cruise controller in a catastrophic situation.

3) Pedestrian Model: In our study, we model the behaviors of pedestrians as the combination of decision-making and motion dynamics described in section 3.2. Specifically, in the scenario shown in Fig. 1, we model that the pedestrians keep crossing without detecting the ego vehicles with a certain probability. This pedestrian model allows us to account for the worst-case scenarios when evaluating safety (*e.g.,* distracted pedestrian keeps crossing even as a car approaches them).

The decision process for pedestrian i, assumes that after the pedestrian has crossed the occlusion, they will recognize the oncoming vehicle and stop with a probability α:

$$d(k) \sim D(y_p^{(i)} \mid \mathcal{Z}_k) = \begin{cases} \text{Bernoulli}(\alpha) & \text{if} \quad y_p^i \geq o \\ 0 \quad \text{if} \quad y_p^i < o \end{cases} \tag{21}$$

where $\mathcal{Z}_k$ is the event of a vehicle currently approaching the crosswalk at time step k and $d(k) \in \{0, 1\}$ is a Bernoulli stochastic process, indicating whether or not the pedestrian recognizes the vehicle and stops.

For the pedestrian motion model, we have that when pedestrian i begins crossing, they travel at a fixed speed sampled from a normal distribution $v_{\text{ped}}^{(i)} \sim \mathcal{N}(\bar{v}_{\text{ped}}, \sigma_{\text{ped}}^2)$ (Onelcin and Alver, 2017) till they've reached the end of the crosswalk or stop after recognizing the oncoming car:

$$\begin{aligned} y_{p_{(k+1)}}^{(i)} &\sim f_z(y_{p_{(k)}}^{(i)} \mid d(k)) \\ &= \begin{cases} 0 \quad \text{if} \quad d(k) = 1 \\ v_{\text{ped}}^{(i)} \Delta t + y_{p_{(k)}}^{(i)} \quad \text{if} \quad d(k) = 0 \end{cases} \end{aligned} \tag{22}$$

Further, we assume that pedestrians arrive at the crossing independently of each other and at random with a mean interarrival time T_a (Cox, 2020; Lartey et al., 2014). That is to say, within an infinitesimally small time interval dt, the probability of a pedestrian arriving at the crossing is dt/T_a. As a consequence, the time interval ΔT between two successive pedestrians arriving at the crosswalk will follow an exponential distribution, and the number of pedestrians N_p that arrive within a time interval Δt will follow a Poisson distribution:

$$\Delta T \sim \text{Exponential}\left(1/T_a\right) \tag{23}$$

$$N_{\text{ped}} \sim \text{Poisson}\left(\Delta t/T_a\right) \tag{24}$$

5) Safety Specification: For this case study, we define the safety criteria in terms of the set $\mathcal{B}^i \subset \mathbb{R}^2$, which is set of the i^{th} pedestrian's bounding box as seen in Fig. 1:

$$\mathcal{B}^i = \left[x_p^{(i)} - \delta_x, x_p^{(i)} + \delta_x\right] \times \left[y_p^{(i)} - \delta_y, y_p^{(i)} + \delta_y\right] \tag{25}$$

with this, we specify the safe event $\mathcal{C}_t^i$ at time t as the set of outcomes wherein the vehicle is not present in $\mathcal{B}^i$, i.e., $(x_c, y_c) \notin \mathcal{B}^i$.

5.2 Hardware Experimental Setup

We evaluate the feasibility and efficacy of the proposed safe control algorithm in a real-world scaled setting of our case-study environment. In particular, we implement the proposed method on a 1/10th scale autonomous vehicle (AV) (see (O'Kelly et al., 2020) for specific details on the

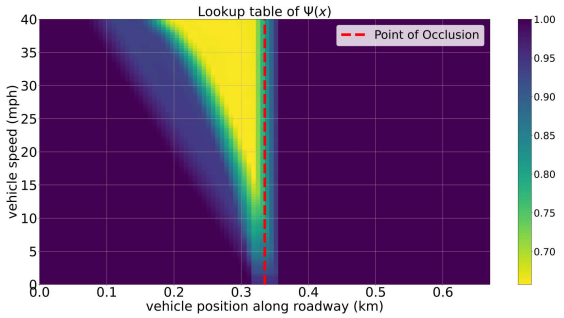

Fig. 2. Precomputed lookup table of $\mathbf{\Psi}(x)$.

Fig. 3. Case-study environment for hardware experiments.

hardware platform). We used an NVIDIA Jetson Xavier NX with a CPU SPEC2006 score of 18.9 GFLOPS as our mobile-embedded computing platform and ROS-2 in C++ for software. The computing capability of our platform is relatively limited. One major challenge of the hardware experiments includes added stochasticity due to imperfect state estimation, which in our case, comes from a particle filter. Another challenge is that most autonomous driving stacks rely on kinematic control inputs, i.e., velocity and not acceleration. To account for this, we use an online double integrator model of the vehicle dynamics (19) to translate our proposed controller's acceleration outputs into velocity commands, allowing seamless integration into existing state-of-the-art platforms.

1) Hardware Nominal Controller: For path tracking, we use a kinematic model-based linear time-varying MPC (Kong et al., 2015). We use the cruise controller outlined in section 5.1 for speed control.

2) Proposed Controller Implementation: The limited compute capability of the embedded platform poses a challenge to real-time implementation. To address this issue, we propose modifying procedure $\mathbf{\Psi}(x)$ in Algorithm 1 as follows. First, we offline precompute the values of $\mathbf{\Psi}(x)$ with a mesh grid of x to form a lookup table as shown in Fig. 2. We then store this lookup table and use it to build a sinc interpolator. With the interpolator, it is possible to achieve real-time performance.

3) Experiment Setup: Fig. 3 shows the experimental setup for the hardware evaluation. The goal of the nominal controller, in this case, would be to drive the vehicle along the corridor and through the occluded intersection (blue bin) as seen in Fig. 3.

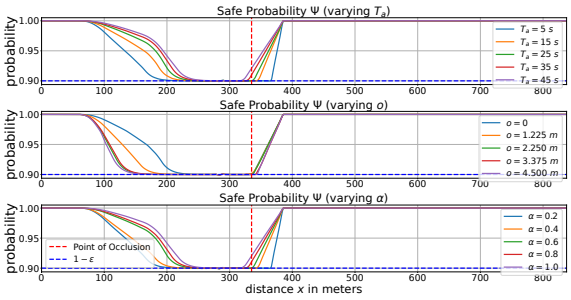

Fig. 4. Long-term safe probability of proposed method for three parameter cases. The top subplot shows the effect of varying the mean interarrival time T_a, the middle varies in occlusion size o, and the bottom-most varies in pedestrian awareness probability α.

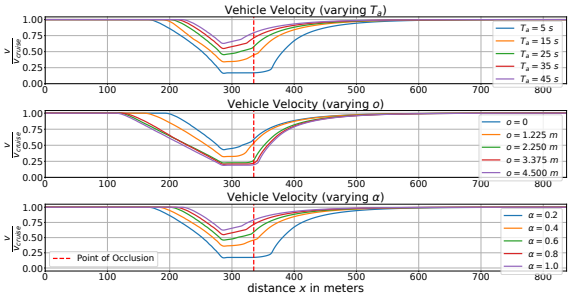

Fig. 5. Safe vehicle velocity generated for three parameter cases. The top subplot shows the effect of varying the mean interarrival time T_a, the middle varies in occlusion size o, and the bottom-most varies in pedestrian awareness probability α.

5.3 Results and Analyses

In this section, we present empirical results of proposed method's performance with numerical simulations and hardware experiments. These serve to quantify and corroborate the technical merits of the proposed method, specifically the guaranteed long-term safety, balancing opposing safety and performance objectives and robustness to large uncertainties. For all experiments, we pick an outlook horizon of $T = 10$ s and a baseline cruising velocity of $v_{\text{cruise}} = 35$ mph.

1) Long-term safety guarantee: We choose a risk tolerance of $\epsilon = 0.10$. We run the proposed controller with different pedestrian interarrival time T_a, occlusion size o, and pedestrian awareness α. We choose $T_a = 25$ s, $o = 2.25$ m, and $\alpha = 0.4$ as the baseline parameter values and vary one of the parameter values for ablation experiments. Fig. 4 shows the proposed controller's ability to guarantee long-term safety under all tested circumstances. In the first subplot of Fig. 4, the safety probability drops earlier in response to shorter interarrival times, larger occlusion sizes, or lower pedestrian awareness, thereby demonstrating the algorithm's predictive capabilities.

Fig. 5 shows the vehicle velocities generated with the proposed safe controller on different parameter settings. We see that the proposed method modulates the vehicle's velocity tantamount to perceived latent risks. For cases with higher perceived latent risk, we see that the speed reduces well ahead of the occlusion's position and main-

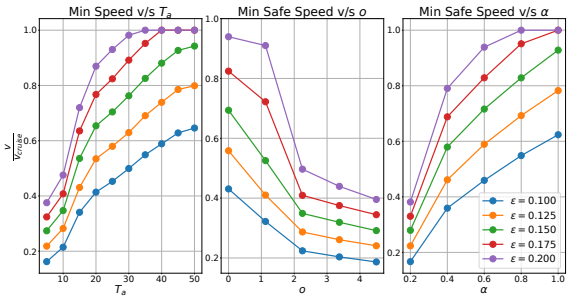

Fig. 6. Minimum safe speeds achieved with the proposed controller over all three parameters and for different risk tolerances ϵ.

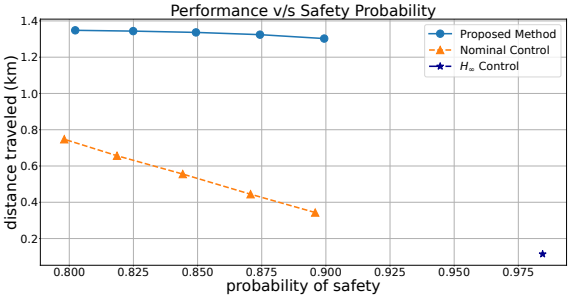

Fig. 7. Performance v/s safety tradeoff.

tains that speed until past it, showcasing the proposed controller's ability to look into the future and impose a more effective safe control on the system once the safety probability tends to drop. Further, when perceived risks are lower, the controller does not slow down more than necessary, not compromising the desired performance. This phenomenon is made more evident in Fig. 6, where we compare the minimum safe speeds of the vehicle with varying risk tolerance values ϵ.

2) Safety v/s Performance Trade-off: We begin by quantifying safety, performance, and uncertainty in the context of our case study. We use the distance traveled by the vehicle over 2 minutes as the performance metric; for safety, we use the minimum safe probability achieved over the run, and for uncertainty in the pedestrian arrival process, we evaluate the Shannon Entropy of (24). To obtain the trade-off, in the case of the nominal controller, we vary the desired cruising speed, and for the proposed controller, we range over the risk tolerance ϵ with baseline parameters. In this case, we have chosen [6.433, 8.467, 10.411, 12.223, 14.258] mph as the desired speeds for the nominal controller to match the risk levels achieved by the proposed method closely. For the proposed controller, we choose [0.1, 0.125, 0.150, 0.175, 0.200] as values for the risk tolerance parameter ϵ. And finally, we also consider $\mathcal{H}_\infty$ control to demonstrate the effects of using a deterministic worst-case safe controller. Fig. 7 exhibits our proposed method's ability to fulfill the desired safety requirements while not being excessively conservative, and Fig. 8 shows that the proposed method's performance degrades gracefully with increasing uncertainty. Furthermore, we see that accounting for all possible worst-cases without considering causality, as in the case of the $\mathcal{H}_\infty$, produces overly conservative behaviors that compromises performance or in some cases induces infeasibility.

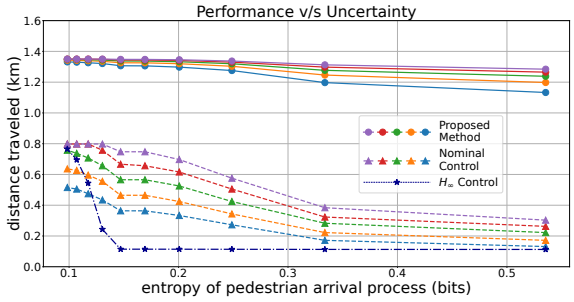

Fig. 8. Performance v/s uncertainty in pedestrian arrival process. Matching colors for plots of the nominal and proposed controller correspond to equivalent safety probabilities.

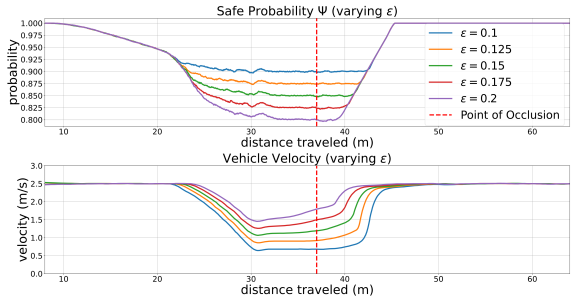

Fig. 9. Experimental results of long-term safe probability (top) and safe vehicle velocities (bottom) for five cases of ϵ.

3) Hardware Experiment Results

A video of the experiments can be found at https://youtu-.be/bII1aARKA-o. We run repeated identical experiments on the 1/10th scale AV over five values of the risk tolerance parameter ϵ and record estimated state and sensor data. Fig. 9 demonstrates that the performance of the proposed controller on hardware is consistent with our simulation results. Further, benchmarking the proposed algorithm on the embedded platform yields the following statistics for computational throughput: average rate of 67.924 Hz, maximum rate of 166.67 Hz, minimum rate of 52.631 Hz, with a standard deviation of 0.02661 Hz, over 61130 samples.

6. CONCLUSION

This paper proposes an occlusion- and interaction-aware safe control strategy that ensures long-term safety in the presence of latent risks without overly compromising performance. We demonstrate its reliability and computational efficiency via numerical simulations and hardware experiments. Finally, we show that the proposed controller is modular and can seamlessly integrate into existing control frameworks, vastly improving its applicability. Future work includes conducting real-world experiments with the proposed method, and comparing the results with human driving behaviors.

REFERENCES

Ames, A.D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., and Tabuada, P. (2019). Control barrier functions: Theory and applications. In *European control conference*, 3420–3431. IEEE.

Brüdigam, T., Olbrich, M., Wollherr, D., and Leibold, M. (2021). Stochastic model predictive control with a safety guarantee for automated driving. *Transactions on Intelligent Vehicles*.

Burger, C., Schneider, T., and Lauer, M. (2020). Interaction aware cooperative trajectory planning for lane change maneuvers in dense traffic. In *International Conference on Intelligent Transportation Systems*, 1–8. IEEE.

Camara, F., Bellotto, N., Cosar, S., Weber, F., Nathanael, D., Althoff, M., Wu, J., Ruenz, J., Dietrich, A., Markkula, G., et al. (2020). Pedestrian models for autonomous driving part ii: high-level models of human behavior. *Transactions on Intelligent Transportation Systems*, 22(9), 5453–5472.

Cox, D.R. (2020). *Queues*. Chapman and Hall/CRC.

Gangadhar, S., Wang, Z., Jing, H., and Nakahira, Y. (2022). Adaptive safe control for driving in uncertain environments. In *Intelligent Vehicles Symposium*, 1662–1668. IEEE.

Helbing, D. and Molnár, P. (1995). Social force model for pedestrian dynamics. *Physical Review E*, 51(5), 4282–4286.

Hubmann, C., Quetschlich, N., Schulz, J., Bernhard, J., Althoff, D., and Stiller, C. (2019). A pomdp maneuver planner for occlusions in urban scenarios. In *Intelligent Vehicles Symposium*, 2172–2179. IEEE.

Hubmann, C., Schulz, J., Xu, G., Althoff, D., and Stiller, C. (2018). A belief state planner for interactive merge maneuvers in congested traffic. In *International Conference on Intelligent Transportation Systems*, 1617–1624. IEEE.

Isele, D., Rahimi, R., Cosgun, A., Subramanian, K., and Fujimura, K. (2018). Navigating occluded intersections with autonomous vehicles using deep reinforcement learning. In *International Conference on Robotics and Automation*, 2034–2039. IEEE.

Jing, H. and Nakahira, Y. (2022). Probabilistic safety certificate for multi-agent systems. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, 5343–5350. IEEE.

Johnson, O. (2004). *Information theory and the central limit theorem*. World Scientific.

Kahn, M., Sarkar, A., and Czarnecki, K. (2022). I know you can't see me: Dynamic occlusion-aware safety validation of strategic planners for autonomous vehicles using hypergames. In *International Conference on Robotics and Automation*, 11202–11208. IEEE.

Kielar, M, P., Handel, O., Biedermann, H, D., and Borrmann, A. (2014). Concurrent hierarchical finite state machines for modeling pedestrian behavioral tendencies. *Transportation Research Procedia*, 2, 576–584.

Kiencke, U. and Nielsen, L. (2000). Automotive control systems: for engine, driveline, and vehicle.

Koç, M., Yurtsever, E., Redmill, K., and Özgüner, Ü. (2021). Pedestrian emergence estimation and occlusion-aware risk assessment for urban autonomous driving. In *International Intelligent Transportation Systems Conference*, 292–297. IEEE.

Kong, J., Pfeiffer, M., Schildbach, G., and Borrelli, F. (2015). Kinematic and dynamic vehicle models for autonomous driving control design. In *Intelligent Vehicles Symposium*, 1094–1099.

Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., and Beijbom, O. (2019). Pointpillars: Fast encoders for object detection from point clouds. In *Conference on Computer Vision and Pattern Recognition*, 12697–12705.

Lartey, J.D. et al. (2014). Predicting traffic congestion: A queuing perspective. *Open Journal of Modelling and Simulation*, 2(02), 57.

Liang, C.Y. and Peng, H. (2000). String stability analysis of adaptive cruise controlled vehicles. *International Journal Series C Mechanical Systems, Machine Elements and Manufacturing*, 43(3), 671–677.

Lyu, Y., Luo, W., and Dolan, J.M. (2021). Probabilistic safety-assured adaptive merging control for autonomous vehicles. In *International Conference on Robotics and Automation*, 10764–10770. IEEE.

Müller, J., Strohbeck, J., Herrmann, M., and Buchholz, M. (2022). Motion planning for connected automated vehicles at occluded intersections with infrastructure sensors. *Transactions on Intelligent Transportation Systems*.

Ogata, K. (1995). *Discrete-time control systems*. Prentice-Hall, Inc.

Onelcin, P. and Alver, Y. (2017). The crossing speed and safety margin of pedestrians at signalized intersections. *Transportation Research Procedia*, 22, 3–12.

O'Kelly, M., Zheng, H., Jain, A., Auckley, J., Luong, K., and Mangharam, R. (2020). Tunercar: A superoptimization toolchain for autonomous racing. In *International Conference on Robotics and Automation*, 5356–5362.

Poncelet, R., Verroust-Blondet, A., and Nashashibi, F. (2020). Safe geometric speed planning approach for autonomous driving through occluded intersections. In *International Conference on Control, Automation, Robotics and Vision*, 393–399. IEEE.

Rasouli, A., Kotserub, I., and Tsotsos, K, J. (2017). Are they going to cross? a benchmark dataset and baseline for pedestrian crosswalk behavior. In *International Conference on Computer Vision Workshops*. IEEE.

Rasouli, A. and Tsotsos, J.K. (2019). Autonomous vehicles that interact with pedestrians: A survey of theory and practice. *Transactions on Intelligent Transportation Systems*, 21(3), 900–918.

Sama, K., Morales, Y., Liu, H., Akai, N., Carballo, A., Takeuchi, E., and Takeda, K. (2020). Extracting human-like driving behaviors from expert driver data using deep learning. *Transactions on Vehicular Technology*, 69(9), 9315–9329.

Wang, Z., Jing, H., Kurniawan, C., Chern, A., and Nakahira, Y. (2021). Myopically verifiable probabilistic certificate for long-term safety. *arXiv preprint arXiv:2110.13380*.

Yu, M.Y., Vasudevan, R., and Johnson-Roberson, M. (2019). Occlusion-aware risk assessment for autonomous driving in urban environments. *Robotics and Automation Letters*, 4(2), 2235–2241.

Zhang, X., Fu, H., and Dai, B. (2019). Lidar-based object classification with explicit occlusion modeling. In *International Conference on Intelligent Human-Machine Systems and Cybernetics*. IEEE.

Zhang, Z. and Fisac, F, J. (2021). Safe occlusion-aware autonomous driving via game-theoretic active perception. In *Robotics: Science and Systems*. RSS.