

Safe Occlusion-aware Autonomous Driving via Game-Theoretic Active Perception

Zixu Zhang and Jaime F. Fisac
Department of Electrical and Computer Engineering
Princeton University, Princeton, NJ 08544
Email: zixuz@princeton.edu; jfisac@princeton.edu

Abstract—Autonomous vehicles interacting with other traffic participants heavily rely on the perception and prediction of other agents’ behaviors to plan safe trajectories. However, as occlusions limit the vehicle’s perception ability, reasoning about potential hazards beyond the field of view is one of the most challenging issues in developing autonomous driving systems. This paper introduces a novel analytical approach that poses safe trajectory planning under occlusions as a hybrid zero-sum dynamic game between the autonomous vehicle (evader) and an initially hidden traffic participant (pursuer). Due to occlusions, the pursuer’s state is initially unknown to the evader and may later be discovered by the vehicle’s sensors. The analysis yields optimal strategies for both players as well as the set of initial conditions from which the autonomous vehicle is guaranteed to avoid collisions. We leverage this theoretical result to develop a novel trajectory planning framework for autonomous driving that provides worst-case safety guarantees while minimizing conservativeness by accounting for the vehicle’s ability to actively avoid other road users as soon as they are detected in future observations. Our framework is agnostic to the driving environment and suitable for various motion planners. We demonstrate our algorithm on challenging urban and highway driving scenarios using the open-source CARLA simulator.

I. INTRODUCTION

Line-of-sight sensors, such as cameras, radars, and LiDARs, are widely used on autonomous vehicles and other mobile robots to perceive surroundings and make real-time decisions. The robustness of trajectory planning in dynamic environments, such as highway driving (Fig. 1), has primarily depended on predicting and analyzing other agents’ future actions (we refer the reader to recent thorough surveys [1, 2] on these topics). In everyday driving scenarios, occlusions are ubiquitous and may be encountered in various forms (common types are shown in Fig. 2). They limit the sensors’ perception abilities and introduce uncertainty in motion predictions. For example, some occlusions are due to the road’s geometric nature, where landscapes and buildings block the view of potential road hazards ahead. Human drivers usually slow down in these situations, to earn more reaction time and allow an eventual stop if needed. Other occlusion scenarios also involve dynamic agents, where more complicated interactions than solely slowing down are necessary to maintain safety and achieve trajectory efficiency.

For autonomous vehicles to operate safely, risks outside their field of view must be evaluated while planning for future actions. Previous works [3–6] utilize forward reachability analysis to over-approximate potential occluded objects’




Fig. 1: The autonomous vehicle (evader) verifies the safety of two potential trajectories \tilde{x}_e and \hat{x}_e against collisions with a hidden agent (pursuer) outside the current field of view Z^{t_0} . Our framework utilizes a game-theoretic approach to determine the existence of a closed-loop evasive strategy if an agent is observed at t_n , by checking the intersection between the danger zone $\mathcal{D}(\tilde{x}_e(t_n))$ of the proposed way-point and the forward hidden set $\bar{\mathcal{H}}(t_n; \mathcal{H}_o^{t_0}, \tilde{x}_e)$. Using this approach, we can determine that the trajectory \hat{x}_e shown in b) is safe, while \tilde{x}_e shown in c) is not safe due to the intersection between two sets. Supplementary materials can be found in https://saferobotics.princeton.edu/research/occlusion_hybrid_game.

future occupancy based on the current sensor observations and generate collision-free trajectories by avoiding the entire forward reachable set (FRS) over the planning horizon. However, this approach ignores the autonomous agent’s future ability to respond to objects, which are currently occluded but may later become visible through new sensor observations. Furthermore, especially when considering potential dynamic agents outside of the field of view, their FRS grows rapidly as the planning time horizon increases. Thus, while FRS-based method provides a sufficient condition to achieve safe behaviors in the presence of occlusions, this may cause the ego agent to choose suboptimal and unnecessarily conservative actions, which can negatively impact driving performance, as well as trust and perceived safety by the public [7].

A reasonable strategy to overtake a slow-moving truck on



(a) A blind intersection



(b) A blind summit in San Francisco



(c) Overtaking a truck

Fig. 2: Common occlusion scenarios that create challenges and safety risks for autonomous vehicles

an undivided highway, as shown in Fig. 1 and 2c, is first to gain the visibility of the incoming traffic by slowing down and cautiously moving towards the oncoming lane. During this process, the ego vehicle must reason about whether it can safely move back to the original lane if any objects show up from the occluded regions. While slowly edging sideways does not maximize the instantaneous progress, it helps to gather additional information regarding the occluded regions, opening up the possibility of overtaking, and potentially making more progress over a longer time horizon. Inspired by the above observations, we set out to formally analyze the decision-making problem faced by an autonomous system navigating a partially occluded environment, considering the worst-case existence and behavior of potential unobserved objects defined within the Operational Design Domain (ODD) [8, Ch. 2] and, critically, accounting for the ability to detect and actively avoid them once they are observed.

For each potential object, we find that this decision-making problem takes the form of a two-mode zero-sum hybrid dynamic game between an evader, the autonomous vehicle, with limited information, and an initially hidden pursuer with perfect state information. The terminal mode of the game is a closed-loop feedback game in which both players are able to observe each other’s state at each time instant and adapt their actions strategically. The first mode of the game, in which the pursuer is invisible to the evader, can be posed as an open-loop “trajectory game”, in which the pursuer may choose its worst-case trajectory in response to *each* possible trajectory chosen by the evader. While the evader is at a strong informational disadvantage in the first mode of the game, it can plan strategically to force a transition to the second mode by having the invisible pursuer fall into its field of view. The pursuer can only win the overall game (force the evader to collide or violate constraints) if it enters the *capture basin* of the second mode (from which the evader cannot escape capture) without first being detected.

Based on this theoretical analysis we propose a novel occlusion-aware trajectory planning framework for autonomous vehicles that leverages this two-mode dynamic game structure. The closed-loop planner solves a reach-avoid game [9–11], computing an evasive feedback policy towards an invariant safe set while avoiding collision with both observed and hidden objects. The open-loop game solution is im-

plemented by a receding horizon planner, which generates optimal trajectories according to some cost function (e.g., progress, comfort, and efficiency), based on (possibly conservative) predictions of objects currently inside the vehicle’s field of view and under the assumption (for optimization purposes) that no hidden objects will appear.¹ To provide a safety guarantee, we introduce a constraint on the future field of view, obtained from the game-theoretic analysis, to ensure that the closed-loop planner will find a collision-free evasive strategy under the worst-case scenario from all waypoints along the planned trajectory.

The remainder of the paper is structured as follows: In Section II, we place our contribution in the context of related efforts in the existing literature. In Section III, we introduce some important theoretical building blocks. In Section IV, we formalize the problem of safe planning under occlusions as a two-mode hybrid dynamic game, and characterize its solution in terms of optimal strategies and winning conditions. In Section V, we introduce the algorithmic trajectory planning framework based on the game-theoretic analysis. In Section VI, we evaluate our proposed framework in a diverse set of representative driving scenarios using the open-source CARLA simulator [12] and demonstrate the benefits of our hybrid open/closed-loop game-theoretic solution with respect to recently proposed, more conservative methods. Finally, in Section VII, we provide a closing discussion and briefly consider promising directions for future work.

II. RELATED WORK

Perception-aware motion planning has been studied for several decades. In particular, the risk of potential objects beyond the sensor’s field of view has received attentions from the robotics and autonomous vehicle communities recently. Early works [13, 14] plan trajectories for mobile robots in the free space represented by a shrunk field of view considering objects that may emerge from occlusions. To formally analyze the collision avoidance behavior with occluded objects, Bouraine et al. [15] propose the braking Inevitable Collision States to achieve a passive-Safety guarantee [16]. Nager et al.

¹This assumption is made for simplicity of the performance optimization and may be replaced by an alternative criterion, e.g. a probabilistic model of hidden object existence. We note that this only affects performance, and has no bearing on safety guarantees, which are our focus in this work.

[17] employs the above ideas in the autonomous driving scenarios by predicting different traffic participants' future occupancy in the current field of view and avoid their braking Inevitable Collision States. However, the passive safety is a weak guarantee since other traffic participants may inevitably collide with the stopped autonomous vehicle.

Recent works [3–6, 18] tackle the risk assessment problem under occlusion by introducing “phantom” objects beyond the field of view and predicting their future states through the forward reachability analysis. Yu et al. [5, 6] apply a sampling method to represent potentially occluded objects' states as particles. Bidirectional reachability analysis is used to determine the distribution of dangerous “phantom” objects. While their methods efficiently represent the distribution of hidden obstacles, they do not provide guarantees to avoid collisions.

To achieve safety guarantee, Orzechowski et al. [4] over-approximate the hidden objects' forward reachable set (FRS) by assuming an object may emerging from occlusion with the maximum possible velocity at any moment. A trajectory is then planned to avoid the intersection with the FRS. To improve FRS approximation, Neel and Saripalli [19] use sequential observations to update the velocity bound that an object can emerge from the occlusion. Koschi and Althoff [3] also propose a set-based approach to approximate the FRS of occluded objects by considering road structure and traffic rules. While the added structure helps reducing conservativeness, these approaches still ignore the autonomous vehicle's future ability to respond to the other road participants' presence and actions once new sensor observations are obtained. Therefore, they require the trajectory to avoid the computed FRS based on current information as a sufficient condition for safety, and the resulting behaviors are often overly conservative.

Considering autonomous vehicle's future ability of sensing and reaction, previous works [20–22] formulate the occlusion-aware trajectory planning as a Partially Observable Markov Decision Process (POMDP) and carefully design the observation model in the belief state to take occlusion into account. While POMDP solutions can generate less conservative strategies for autonomous driving under occlusion, safety guarantees are not provided, and a reduction in the state-space is required to make the algorithm computationally tractable.

The idea of active perception [23–25] attempts to close the loop between robot perception and motion planning by generating trajectories that will explore the environment and gain information to reduce uncertainty and improve future planning. Some recent works [26, 27] also tackle the occlusion-aware planning problem through active perception by incorporating the prediction of the autonomous vehicle's future field-of-view into motion planning. Andersen et al. [26] solve the truck-overtaking problem through an MPC trajectory planner guided by a finite state-machine to determine overtaking behavior and maximize visibility. Wang et al. [27] propose a visibility risk metric based on predicted field-of-view to penalize the trajectories with low visibility and high velocity when approaching

an occluded road section. While their results show improved the trajectory efficiency through exploration behaviors that increase visibility, their safety guarantees still rely on avoiding the over entire FRS of hidden objects.

III. PRELIMINARIES

A. System Dynamics and Reachable Sets

We consider an ego agent e and other dynamic objects o_1, o_2, \dots in a three-dimensional environment $\mathcal{W} \subseteq \mathbb{R}^3$. The evolution of each agent's state $x_i^t \in \mathbb{R}^{n_i}$ is governed by continuous-time dynamics

$$\dot{x}_i^t = f_i(x_i^t, u_i^t), \quad i \in \{e, o_1, o_2, \dots\}, \quad (1)$$

where $u_i^t \in \mathcal{U}_i \subseteq \mathbb{R}^{m_i}$ is the agent's control input. The agent's allowable workspace \mathcal{W} and input set \mathcal{U}_i are defined by the Operational Design Domain.

The input sets \mathcal{U}_i are assumed compact, and the dynamics $f_i : \mathbb{R}^{n_i} \times \mathcal{U}_i \rightarrow \mathbb{R}^{n_i}$ are assumed bounded, continuous in u_i and Lipschitz-continuous in x_i . Under these assumptions, for any initial agent configuration $x_i^0 \in \mathbb{R}^{n_i}$ and measurable control signal $\mathbf{u}_i : [0, \infty) \rightarrow \mathcal{U}_i$, there exists a unique, continuous state trajectory $\mathbf{x}_i(\cdot; 0, x_i^0, \mathbf{u}_i)$ satisfying (cf. [28, Ch. 2]):

$$\mathbf{x}_i(0; 0, x_i^0, \mathbf{u}_i) = x_i^0 \quad (2a)$$

$$\dot{\mathbf{x}}_i(t; 0, x_i^0, \mathbf{u}_i) = f_i(\mathbf{x}_i(t; 0, x_i^0, \mathbf{u}_i), \mathbf{u}_i(t)), \text{ a.e. } t \geq 0. \quad (2b)$$

To formally analyze the safe interaction between agents, we will often need to reason about *constrained* reachable sets, in which we require that agents avoid entering certain (possibly time-varying) forbidden regions of their state space, which may be encoded through a set-valued map $\mathcal{F}_i : [0, \infty) \rightrightarrows \mathbb{R}^{n_i}$, with $\mathcal{F}_i(\tau) \subseteq \mathbb{R}^{n_i}$ for each τ .

Definition 1 (Forward reach-avoid set). *An agent i 's forward reach-avoid set at time $t \geq t_0$ from an initial set $\mathcal{C}_i^{t_0}$ under failure set $\mathcal{F}_i(\cdot)$ is the set of all states that agent i can reach at time t starting from some state $x_i^{t_0} \in \mathcal{C}_i^{t_0}$ at time t_0 while avoiding the set $\mathcal{F}_i(\tau)$ at all intermediate times τ :*

$$\begin{aligned} \overrightarrow{\mathcal{R}}_{\mathcal{A}_i}(t; t_0, \mathcal{C}_i^{t_0}, \mathcal{F}_i) = \{ & x_i^t \mid \exists x_i^{t_0} \in \mathcal{C}_i^{t_0}, \exists \mathbf{u}_i, \mathbf{x}(t; t_0, x_i^{t_0}, \mathbf{u}_i) = x_i^t \\ & \wedge \mathbf{x}(\tau; t_0, x_i^{t_0}, \mathbf{u}_i) \notin \mathcal{F}_i(\tau) \forall \tau \in [t_0, t] \} \end{aligned}$$

B. Observations and Field of View

The set-valued *footprint* $\phi_i : \mathbb{R}^{n_i} \rightrightarrows \mathcal{W}$ maps agent i 's state x_i to its occupied workspace $\phi_i(x_i) \subseteq \mathcal{W}$. We will also denote by ϕ_i^{-1} the inverse image induced by the footprint map: $\phi_i^{-1}(M) := \{x_i \mid \phi_i(x_i) \cap M \neq \emptyset\}$. The occupied workspace by all non-ego objects at time t is then given by

$$\mathcal{O}(t) := \bigcup_i \phi_i(s_i^t) \subseteq \mathcal{W}, \quad i \in \{o_1, o_2, \dots\} \quad (3)$$

For compactness, we will often write this set as \mathcal{O}^t .

We assume that the ego agent is equipped with a sensor (or suite of sensors) that enables it to accurately determine the presence or absence of other agents within certain well-observed regions of the workspace. For the purposes of our analysis, we refer to this overall region as the ego agent's *field*

of view (FOV), denoted $\mathcal{Z}(x_e^t, \mathcal{O}^t) \subseteq \mathcal{W}$. The variation of the FOV over time as a function of the ego agent's configuration x_e^t and the footprints of dynamic objects in the environment \mathcal{O}^t will be central to our analysis.

For example, an autonomous vehicle's perception system, in the context of this analysis, may be abstractly modeled through a line-of-sight sensor with range r , centered around the position $p_e^t(x_e^t)$ of the vehicle. This sensor's field of view $\mathcal{Z}(x_e^t, \mathcal{O}^t)$ is then a set of points $y' \in \mathcal{W}$, such that $\|p_e^t - y'\| \leq r$ and the line segment connecting p_e^t and y' does not intersect with \mathcal{O}^t (excepting end points).

Moreover, at the moment when the ego agent starts operating, it estimates a set of hypothetical hidden agents' states:

$$\mathcal{H}_o^0 := \phi_o^{-1}(\mathcal{Z}(x_o^0, \mathcal{O}^0) \cup \mathcal{O}^0)^c \quad (4)$$

In the later section, we will discuss how the hidden set \mathcal{H}_o^0 is updated through observations and reachability analysis.

IV. THE HIDDEN PURSUER HYBRID GAME

The theoretical contribution in this work is characterizing the least-restrictive guarantees for safe trajectory planning under nondeterministic (worst-case) uncertainty and dynamic environment occlusions through the solution to a hybrid differential game combining open-loop and closed-loop information patterns.

We focus on the problem of strategically detecting and avoiding a single (potentially) adversarial hidden object, in the midst of an arbitrary number of static and dynamic obstacles with bounded motion predictions, and note that the theoretical analysis may be extended to multiple simultaneous hidden adversaries, at the cost of increased complexity.

Assumption 1 (Separation of hidden adversaries.). *There exists at most one unobserved pursuer for the duration of the differential game.*

We clarify that this assumption does not render our results inapplicable to scenarios with multiple hidden objects, which are common in driving settings. Rather, Assumption 1 restricts the ego vehicle to *strategically* plan to detect and avoid one hidden pursuer at any given planning cycle. Any additional potentially-existing objects can be represented as dynamic obstacles whose entire forward-reachable sets must be avoided, following the current state-of-the-art approach. Our framework therefore enables reducing conservativeness around strategically important occlusions, while maintaining safety guarantees with respect to any number of objects.

The problem at hand is then a two-player zero-sum differential game of pursuit [29, 30], in which the pursuer \bar{o} seeks to capture (result in collision with) the evader e , who in turn wishes to avoid capture. The hidden pursuer's initial state x_o^0 may be anywhere within an initial existence set \mathcal{H}_o^0 . The evader and pursuer are each required to avoid a time-varying obstacle set $\mathcal{O}_e : [0, T] \rightrightarrows \mathcal{W}$, $\mathcal{O}_{\bar{o}} : [0, T] \rightrightarrows \mathcal{W}$,

which may be different². The outcome of the game is binary, with the pursuer winning if it achieves capture, by driving the joint state $(x_e, x_{\bar{o}})$ into a capture set $\text{Cap} \subset \mathbb{R}^{n_e} \times \mathbb{R}^{n_{\bar{o}}}$, or if the evader violates its obstacle-induced constraint; the evader wins the game if the pursuer violates its obstacle constraint or if capture does not take place within the horizon T of the game (in practice we will often consider the infinite-horizon case $T \rightarrow \infty$).

While the general solution of this game when both players have perfect state information is well studied (cf. [31, 32]), our problem of interest introduces an important new consideration: the information structure of the game depends on the state history. In particular:

- The pursuer \bar{o} has perfect state feedback information, i.e. can observe the full state $(x_e, x_{\bar{o}})$ at each time.
- The evader e can only observe its own state x_e^t until the pursuer enters its field of view, i.e. $\phi_{\bar{o}}(x_{\bar{o}}^t) \in \mathcal{Z}(x_e^t, \mathcal{O}^t)$; after such an event, the evader also has perfect state feedback information.

Because of this history-dependent information structure, it is not possible to formulate this problem as either a closed-loop feedback game or an open-loop trajectory game [9, Ch. 8]. However, we can instead decompose its solution through a two-mode hybrid differential game [33] with an initial open-loop mode and a terminal closed-loop mode. As a result of the dynamic programming principle, the solution to the initial game depends on that of the terminal game; we will therefore consider them in backward order.

A. Closed-Loop Game: Perfect-Information Pursuit-Evasion

In the terminal game, the pursuer has been detected by the evader, who can now adapt its evasive strategy to the observed pursuer motion over time. This is a well-studied problem, corresponding to the classical *games of pursuit* first studied by Isaacs in the 1950s [29, 30]. The general setting in which players are required to avoid dynamic obstacles was studied in [31, 32].

The winning set for the pursuer, or *capture basin*, at time t is the set of joint conditions from which, regardless of the evader's strategy, either the pursuer can enforce capture or the evader otherwise collides, without the pursuer violating its obstacle constraints. Formally:

$$\mathcal{B}(t) = \left\{ (x_e^t, x_{\bar{o}}^t) \mid \exists \mathbf{v}_{\bar{o}}, \forall \mathbf{u}_e, \exists \tau \geq t \right. \\ \left. \left(\phi_e(\mathbf{x}_e(\tau; t, x_e^t, \mathbf{u}_e)) \cap \phi_{\bar{o}}(\mathbf{x}_{\bar{o}}(\tau; t, x_{\bar{o}}^t, \mathbf{v}_{\bar{o}}[\mathbf{u}_e])) \neq \emptyset \right. \right. \\ \left. \left. \vee \phi_e(\mathbf{x}_e(\tau; t, x_e^t, \mathbf{u}_e)) \cap \mathcal{O}_e(\tau) \neq \emptyset \right) \right. \\ \left. \wedge \forall s \in [t, \tau], \phi_{\bar{o}}(\mathbf{x}_{\bar{o}}(s; t, x_{\bar{o}}^t, \mathbf{v}_{\bar{o}}[\mathbf{u}_e])) \cap \mathcal{O}_{\bar{o}}(s) = \emptyset \right\} \quad (5)$$

where $\mathbf{v}_{\bar{o}}$ is a non-anticipative strategy for the pursuer [34]; this is equivalent to both players performing state feedback³.

²For example, we may require our ego vehicle, but not the hidden pursuer, to avoid other potentially existing objects.

³The equivalence is due to *Isaacs' condition*; we direct interested readers to [9, Ch. 8].

The capture basin $\mathcal{B} : [0, T] \rightrightarrows \mathbb{R}^{n_e} \times \mathbb{R}^{n_{\bar{o}}}$ and players' feedback strategies $\pi_e^t, \pi_{\bar{o}}^t$ are encoded by the solution to a Hamilton-Jacobi-Isaacs variational inequality [31]. In practice, the numerical solution (which provides a necessary and sufficient condition for safety) may be computationally prohibitive for real-time planning in changing environments; however, as we will see in Section V, conservative approximations can be efficiently computed, yielding a sufficient safety condition for the evader.

A critical observation is that in this closed-loop feedback game the evader need not have any single available trajectory that will avoid *all* possible motions of the pursuer (i.e. its forward-reachable set) in order to avoid capture; instead, it can *adapt* its course of action to the unfolding pursuer trajectory, since different pursuer futures are mutually exclusive. This is a key distinction between our analysis and existing methods based on avoidance of the entire forward-reachable set.

B. Open-Loop Game: Preventing an Invisible Pursuer from Becoming Unavoidable

Having established the key properties of the closed-loop, perfect information game solution, we can now reason about players' optimal strategies for the previous, open-loop stage. Indeed, the winning conditions of this first game are directly related to the solution of the second.

In order to make the autonomous vehicle lose the overall game, the pursuer needs to drive the joint state into the capture basin of the closed-loop game without being previously detected: at this point even if the evader becomes aware of its state, the pursuer can enforce capture. Conversely, if the evader can detect the pursuer before the state enters the capture basin, it can immediately apply its full-state feedback policy π_e^t to avoid capture.

In this first stage, however, the evader has no knowledge of the pursuer's whereabouts, nor can it obtain any information about the pursuer until the pursuer's footprint enters its field of view. As a result, the evader needs to blindly commit to an open-loop trajectory, which (in the worst case) will be met with an optimally adversarial initialization and execution of the pursuer trajectory. We have the following results (proven in Appendix A):

Proposition 1 (Hybrid game solution). *Let the closed-loop capture basin $\mathcal{B} : [0, T] \rightrightarrows \mathbb{R}^{n_e} \times \mathbb{R}^{n_{\bar{o}}}$ be as given in (5). Then, the evader starting from state x_e^0 can maintain safety and avoid capture by a hidden pursuer starting anywhere in $\mathcal{H}_{\bar{o}}^0$ if and only if there exists an open-loop control signal \mathbf{u}_e for the evader such that $\forall t \geq 0, \phi_e(\mathbf{x}_e(t; 0, x_e^0, \mathbf{u}_e)) \cap \mathcal{O}_e(t) = \emptyset$ and, additionally, for all initial pursuer states $x_{\bar{o}}^0 \in \mathcal{H}_{\bar{o}}^0$ and control signals $\mathbf{u}_{\bar{o}}$:*

$$\begin{aligned} & \left[\exists t \geq 0, (\mathbf{x}_e(t; 0, x_e^0, \mathbf{u}_e), \mathbf{x}_{\bar{o}}(t; 0, x_{\bar{o}}^0, \mathbf{u}_{\bar{o}})) \in \mathcal{B}(t) \right] \implies \\ & \left[\exists \tau \in [0, t) : \phi_{\bar{o}}(\mathbf{x}_{\bar{o}}(\tau; 0, x_{\bar{o}}^0, \mathbf{u}_{\bar{o}})) \in \mathcal{Z}(\mathbf{x}_e(\tau; 0, x_e^0, \mathbf{u}_e), \mathcal{O}_e(\tau)) \right. \\ & \quad \left. \vee \phi_{\bar{o}}(\mathbf{x}_{\bar{o}}(\tau; 0, x_{\bar{o}}^0, \mathbf{u}_{\bar{o}})) \in \mathcal{O}_{\bar{o}}(\tau) \right] \end{aligned} \quad (6)$$

Proposition 1 states that, in order for the evader to avoid capture, it must be that any pursuer trajectory that eventually drives the state to the capture basin previously enters the evader's field of view or violates the pursuer's obstacle constraints.

The following result will serve as a bridge between the theoretical solution of the hybrid game and a safety-preserving onboard decision-making scheme for autonomous driving systems.

Definition 2 (Forward Hidden Set). *The pursuer's forward hidden set at time t for a given evader trajectory \mathbf{x}_e is the set of states that the pursuer could reach at time t without violating its obstacle constraints or being discovered by the evader.*

$$\overrightarrow{\mathcal{H}}_{\bar{o}}(t; \mathcal{H}_{\bar{o}}^0, \mathbf{x}_e) = \overrightarrow{\mathcal{R}\mathcal{A}}_{\bar{o}}(t; 0, \mathcal{H}_{\bar{o}}^0, \phi_{\bar{o}}^{-1}(\mathcal{Z}(\mathbf{x}_e(\cdot), \mathcal{O}(\cdot)) \cup \mathcal{O}_{\bar{o}}(\cdot)))$$

Definition 3 (Danger Zone). *The danger zone for a given evader state x_e at time t is the set of pursuer states from which the evader cannot safely avoid capture in the closed-loop game.*

$$\mathcal{D}(t; x_e^t) = \{x_{\bar{o}} \mid (x_e^t, x_{\bar{o}}^t) \in \mathcal{B}(t)\}$$

Theorem 1 (Forward-Backward Reach-Avoid Set Characterization). *The evader starting from state x_e^0 can maintain safety and avoid capture by a hidden pursuer starting anywhere in $\mathcal{H}_{\bar{o}}^0$ if and only if there exists \mathbf{u}_e such that $\forall t \geq 0, \phi_e(\mathbf{x}_e(t; 0, x_e^0, \mathbf{u}_e)) \notin \mathcal{O}_e(t)$ and, additionally, for all $t \geq 0$*

$$\overrightarrow{\mathcal{H}}_{\bar{o}}(t; \mathcal{H}_{\bar{o}}^0, \mathbf{x}_e(\cdot; 0, x_e^0, \mathbf{u}_e)) \cap \mathcal{D}(t; \mathbf{x}_e(t; 0, x_e^0, \mathbf{u}_e)) = \emptyset$$

The above characterization provides a direct safety checking criterion for onboard decision-making. If we can ensure that at all states in the autonomous vehicle's trajectory plan *the forward hidden set and the danger zone are disjoint*, then the autonomous vehicle will be able to safely respond (through π_e^t) to maintain safety in the event that a hidden object appears.

V. TECHNICAL FRAMEWORK

Our theoretical results concern a two-mode hybrid dynamic game in continuous time. However, in real-world operation, autonomous vehicles receive sensor observations, conduct safety verification, and make decisions at discrete moments in time. Moreover, computing the winning condition for an infinite-horizon pursuit-evasion game is intractable for general cases. This section introduces an algorithmic trajectory planning framework based on the theoretical analysis in Section IV to overcome these challenges.

A. Finite Horizon Approximation of the Closed-loop Game

Robots, including autonomous vehicles, utilize receding horizon control [35] and partial motion planning [36] to compute partial trajectories efficiently and frequently replan to deal with uncertainty evolution over a long time. However, collision avoidance over a short planning horizon is not able to assure persistent feasibility and safety. Instead, previous works provide infinite horizon safety guarantees by planning partial

trajectories within a *safe set* [37, 38], or to avoid the *inevitable collision set (ICS)* [15, 39, 40].

Assumption 2 (Invariant Safe Set). *There exists a (possibly time-varying) invariant safe set $\Omega : [0, \infty) \Rightarrow \mathbb{R}^{n_e}$ of collision-free states of the ego agent, from which, for any possible actions from other agents, the ego agent has a **known** strategy $\pi_e^\Omega : [0, \infty) \times \mathbb{R}^{n_e} \rightarrow \mathcal{U}_e$ that enables it to remain inside this set for an infinite time horizon:*

$$\forall x_e^t \in \Omega(t), \phi(x_e^t) \cap \mathcal{O}_e(t) = \emptyset, \quad \forall t \geq 0 \quad (7a)$$

$$\forall x_e^t \in \Omega(t), \forall x_o^t \in \mathcal{H}_o^t, (x_e^t, x_o^t) \notin \mathcal{B}(t) \quad (7b)$$

$$x_e^t \in \Omega(t) \implies \forall \tau \geq t, \mathbf{x}_e(\tau; t, x_e^t, \pi_e^\Omega(\cdot)) \in \Omega(\tau) \quad (7c)$$

The control barrier function [41] is introduced to verify and enforce such set. It can be approximated through optimization [42], learned through demonstrations [43], and analytically expressed for specific dynamics [44].

To retain an infinite horizon safety guarantee, we conservatively approximate the pursuit-evasion game in Section IV-A as a finite horizon reach-avoid game, whose winning set for the evader is:

$$\begin{aligned} \tilde{\mathcal{B}}^c(t) = \{ & (x_e^t, x_o^t) \mid \forall \mathbf{v}_o, \exists \mathbf{u}_e, \exists \tau \in [t, T], \\ & \mathbf{x}_e(\tau; t, x_e^t, \mathbf{u}_e) \in \Omega, \wedge \forall s \in [t, \tau], \\ & \phi_e(\mathbf{x}_e(s; t, x_e^t, \mathbf{u}_e)) \cap \phi_o(\mathbf{x}_o(s; t, x_o^t, \mathbf{v}_o[\mathbf{u}_e])) = \emptyset \\ & \wedge \phi_e(\mathbf{x}_e(s; t, x_e^t, \mathbf{u}_e)) \cap \mathcal{O}_e(s) = \emptyset \} \end{aligned} \quad (8)$$

The ego agent must be able to reach a state in the invariant safe set before T while avoiding capture by the pursuer and collisions with other objects \mathcal{O}_e . After reaching a state in Ω , the ego agent can apply a known feedback policy to achieve an infinite-horizon safety guarantee. It can be seen by inspection of (5), (7), (8) that the winning condition for the ego agent is an under-approximation of the complement of the original *capture basin* defined in (5), i.e. $\tilde{\mathcal{B}}^c(t) \subseteq \mathcal{B}^c(t)$.

Therefore, we are able to over-approximate the danger zone in Definition 3 using finite-time reach-avoid game, and maintain safety over infinite horizon. In general, the winning condition for the reach-avoid problem can be obtained through level-set methods [45], using Hamilton-Jacobi (HJ) reachability analysis [32], and applying sum-of-square optimization [46].

B. Open-loop Game with Discrete Sensor Observation

Consider the ego agent receives new real-time sensor observation every Δt seconds. At $t_n = n\Delta t$, the ego agent tries to determine whether it can safely move towards $x_e^{t_{n+1}}$. As shown in Theorem 1, the safety criterion is checking whether the forward hidden set $\mathcal{H}_o(t_{n+1}; \mathcal{H}_o^0, \mathbf{x}_e)$ and the danger zone $\mathcal{D}(x_e^{t_{n+1}})$ are disjoint. As we do not have the access to future sensor observation, the forward hidden set can be over-approximated as (9) by the forward-reach avoid set of hidden set that avoiding being observed up to t_n .

$$\begin{aligned} \tilde{\mathcal{H}}_o(t_{n+1}; \mathcal{H}_o^0, \mathbf{x}_e) = \overrightarrow{\mathcal{R}\mathcal{A}}_o(t_{n+1}; 0, \mathcal{H}_o^0, \\ \phi_o^{-1}(\mathcal{Z}(\mathbf{x}_e(0 : t_n; \cdot), \mathcal{O}(0 : t_n; \cdot)) \cup \mathcal{O}_o(\cdot))) \end{aligned} \quad (9)$$

In order to achieve the long term goal and provides an efficient motion strategy that minimizes any given cost function, the open-loop game in our framework plans trajectories in a receding horizon fashion. Without loss of generality, we consider the ego agent at state $x_e^{t_0}$ plans an open-loop trajectory with $N\Delta t$ seconds time-horizon. At each intermediate step, the planner tries to determine if a proposed sub-trajectory, which connects the waypoints $\tilde{x}^{t_{n-1}}$ to \tilde{x}^{t_n} , can be expanded to $\tilde{x}^{t_{n+1}}$. By assuming no new agent will emerge from the occlusion at t_n , it first predicts the states of all observed dynamic agents in the environment, and then estimates the field of view to predict the forward hidden set $\tilde{\mathcal{H}}_o(t_{n+1}; \mathcal{H}_o^0, \mathbf{x}_e)$. If proposed waypoint $\tilde{x}^{t_{n+1}}$ satisfies the safety constraint defined in Theorem 1, the sub-trajectory is successfully expanded.

With a long planning horizon, the open-loop game inexplicitly encourages the ego agent to choose safe trajectories that may lead to information gain and potentially open up more free space for planning by predicting the future information state optimistically. This allows the ego agent to plan trajectories, which may originally be rejected by previous methods that avoid the entire FRS. On the other hand, the constraint introduced in Theorem 1 also guarantees the existence of evasive trajectories when the actual sensor observation disagrees with the prediction or a new object emerged from the occlusion. Therefore, the proposed two-mode hybrid game framework provides sufficient conditions for safe trajectory planning in a partially occluded environment and can generate less conservative behaviors than previous methods.

C. Framework of the Two-Mode Hybrid Game

This subsection summarizes the proposed two-mode hybrid game in the Algorithm 1 for its application in trajectory planning. Our framework consists following functions:

- **Sense**: detects obstacles, conducts state estimations and identify the field of view at a fixed rate.
- **UpdateHidden**: updates the forward hidden set described in Definition 2.
- **PlayOpenLoop**: generates a receding-horizon open-loop trajectory described in Section V-B.
- **PlayClosedLoop**: plans the evasive action to win the reach-avoid game introduced in the Section V-A.
- **PlanSafe**: checks the safety criterion in Theorem 1.
- **InvSafeState**: determines whether the ego agent wins the closed loop game by reaching Ω .

Our framework operates under a finite state machine with two modes - `OpenLoop` and `ClosedLoop`. By assuming the vehicle starts inside Ω , the system is initialized in the `OpenLoop` mode and frequently replan to improve efficiency. At each time step in the `OpenLoop` mode, the planned waypoint is verified based on the current observation using Theorem 1. If the safety verification failed, the system immediately switches to the `ClosedLoop` mode and plays the reach-avoid to avoid a collision and shield Ω . Once the ego agent reaches a safe state, the system switches back to the `OpenLoop` mode and replan for future actions.

Algorithm 1 Occlusion-aware Trajectory Planning

```
1: Mode  $\leftarrow$  OpenLoop
2:  $t \leftarrow 0$ 
3:  $t_{\text{replan}} \leftarrow 0$ 
4:  $\mathcal{I} \leftarrow \mathcal{C}_{\bar{o}}$ 
5: while running at the timestep  $t$  do
6:    $x_e, \mathcal{Z}, \mathcal{O}_e \leftarrow \text{Sense}()$ 
7:    $\mathcal{H}_{\bar{o}} \leftarrow \text{UpdateHidden}(\mathcal{H}, x_e, \mathcal{Z})$ 
8:   if Mode is OpenLoop then
9:     if  $t \geq t_{\text{replan}}$  then
10:       $\mathbf{x} \leftarrow \text{PlayOpenLoop}(x, \mathcal{O}_e, \mathcal{H}_{\bar{o}})$ 
11:       $t_{\text{replan}} \leftarrow t + T_{\text{replan}}$ 
12:    end if
13:     $x_{\text{plan}} \leftarrow \mathbf{x}[t + 1]$ 
14:    if not PlanSafe( $x, x_{\text{plan}}, \mathcal{O}_e, \mathcal{H}_{\bar{o}}$ ) then
15:       $x_{\text{plan}} \leftarrow \text{PlayClosedLoop}(x, \mathcal{O}_e, \mathcal{H}_{\bar{o}})$ 
16:      Mode  $\leftarrow$  ClosedLoop
17:    end if
18:  else
19:     $x_{\text{plan}} \leftarrow \text{PlayClosedLoop}(x, \mathcal{O}_e, \mathcal{H}_{\bar{o}})$ 
20:    if InvSafeState( $x_{\text{plan}}, \mathcal{O}_e, \mathcal{H}_{\bar{o}}$ ) then
21:      Mode  $\leftarrow$  OpenLoop
22:       $t_{\text{replan}} \leftarrow t + 1$ 
23:    end if
24:  end if
25:  DriveTo( $x_{\text{plan}}$ )
26:   $t \leftarrow t + 1$ 
27: end while
```

VI. SIMULATION RESULTS AND EVALUATION

This section analyzes the benefit of modeling occlusion-aware trajectory planning problems as a two-mode hybrid game by comparing our approach with the previous work that over-approximates the hidden agent’s occupancy through forward-reachability analysis. We describe our implementation details regarding our prototype planner and the resultant behaviors in three typical driving scenarios involving occlusion.

A. Dynamics and Trajectory Planner

Our experiment uses the road centerline as the reference path and represents all agents’ states in the Frenet frame, where p is the agent’s progress, and l is the lateral position relative to the road centerline. Their longitudinal and lateral velocities are denoted as v_{lon} and v_{lat} . Assuming the agent’s heading angle with respect to the reference path is small, we can represent the system’s dynamics with a double integrator shown in the (10).

$$\dot{p}(t) = v_{lon}(t), \quad \dot{v}_{lon}(t) = a_{lon}(t) \quad (10a)$$

$$\dot{l}(t) = v_{lat}(t), \quad \dot{v}_{lat}(t) = a_{lat}(t) \quad (10b)$$

The autonomous vehicle is assumed to have direct control over its longitudinal and lateral accelerations denoted by a_{lon} and a_{lat} . The accelerations of ego and hidden agents are

subjected to constraints listed in the (11).

$$a_{lon} \in [-8, 4] \text{ m s}^{-2}, \quad \sqrt{a_{lon}^2 + a_{lat}^2} \leq 8 \text{ m s}^{-2} \quad (11)$$

We construct a spatio-temporal lattice [47, 48] over a pre-defined route with 0.25 m spatial resolution in both longitudinal and lateral dimension. The open-loop game planner uses an A* search in a receding horizon fashion and generates a trajectory satisfying winning conditions stated in Theorem 1 and maximizing the progress at the end of the time horizon.

At each step, we conduct a ray-casting computation to measure the current field of view and use ground truth data from CARLA for observed objects’ state. During the planning of the open-loop game, we predict the field of view from a pre-constructed Octomap [49] and estimate the occlusion due to dynamic obstacles from their predicted occupancy. Following the method proposed in [19], the forward hidden set $\vec{\mathcal{H}}_{\bar{o}}^t$ is approximated by the velocity upper bound of hidden agents based on the occluded region’s length and observation history. Although obtaining the solution for the closed-loop reach-avoid game is computationally challenging in general case, our prototype planner takes the advantage of the simplified dynamics and approximates the closed-form solution through a finite set of motion primitives, as detailed in Appendix B.

While our prototype planner utilizes an Octomap to predict future visibility, a precise representation of the environment is not required for safety guarantees. Our framework always verifies the safety of the next step in a planned trajectory before executing it, by checking the Theorem 1 condition for the current sensor observation and updated hidden set. In the event that the verification is negative, a closed-loop strategy still exists to avoid the collision under our safety criterion. However, accurate prediction of future visibility can improve planning efficiency and avoid frequent evasive behaviors.

B. Invariant Safe Set

In our experiments, we assume the ego agent is safe when its entire footprint occupies a single lane with correct traffic direction, and the vehicle is either stopped or keeps a safe distance with agents ahead. The safe distance can be obtained by determining the minimum required distance to avoid collisions, when both the ego agent and other traffic participants apply the maximum deceleration for emergency brake. However, stopping in opposite lanes or in the intersection is not considered safe.

C. Scenario 1: Blind Summit

In the first experiment shown in Fig 3, we consider the ego vehicle approaches an uphill slope in the CARLA map Town05. The initial velocity of the ego agent is 30 m s^{-1} . The hill has a 7.2° slope with 10 m vertical difference. The planning horizon of the open-loop game is 10 steps, 0.5 s apart. The open-loop trajectory is replanned every time step. By considering the worst-case scenario of a potential obstacle stopped in the occluded region, the closed-loop game reduces to determining whether the ego agent can stop before colliding with the hidden obstacle. As shown in Fig. 3, our




Fig. 3: The road geometry and ego vehicle’s progress versus its longitudinal velocity in the blind summit example

proposed method, which considers the vehicle’s future sensing and decision-making ability, allows the ego agent to plan a trajectory with higher velocity and efficiency when compared with the baseline method.

D. Scenario 2: Intersection With Partial Occlusion

In the second experiment as shown in Fig. 4, the ego agent drives in an urban environment of CARLA map Town01 and approaches a T intersection 60 m ahead, where a part of oncoming traffic is occluded by parked trucks. We set the speed limits as 15 m s^{-1} for all agents in the scene. The planning horizon of the open-loop game is 6 steps, 0.5 s apart. The open-loop game is replanned every time step to improve planning efficiency.




Fig. 4: The ego agent approaches an open T intersection, with the oncoming lane partially occluded by parked trucks.




Fig. 5: Ego agent’s progress versus its longitudinal velocity in the blind T intersection example

Assuming that the ego vehicle remains in its lane, the closed-loop game solution in this example is analytically obtainable, with the optimal evasive policy always consisting in one of two “bang-bang” actions. The vehicle can apply the maximum deceleration to either enter the intersection as late as possible or stop before the junction. Alternatively, the vehicle can apply the maximum acceleration to pass the intersection as early as possible.

We test two intersection examples with hidden objects entering the field of view at different times. The longitudinal velocity profile of the proposed algorithm under each scenario and the velocity profile of the baseline method are shown in Fig. 5. In the first example, the ego agent initially slows down and then maintains its speed as it approaches the intersection. At $p = 52 \text{ m}$, as shown by the same-lane danger zone in Fig. 6a, the ego vehicle would be able to stop just before entering the junction. However, as the danger zone does not intersect with the recently detected agent, the ego vehicle can safely traverse the intersection by applying full acceleration.

In the second example, the hidden agent is set to enter the field of view earlier than the previous case. As shown in Fig. 6b, the ego agent detects the previously hidden agent at $p = 45 \text{ m}$ and the safety condition indicates the (would-be) pursuer will enter the danger zone if the ego agent attempts to accelerate. Therefore, the ego agent takes an evasive trajectory and stops before the junction. Once the agent is about to exit the intersection, the ego vehicle slowly begins to enter it, and then proceeds by acceleration when the traffic is clear. The planner benefits from our two-mode game framework since the analysis predicts that the ego vehicle will gain sufficient visibility to safely pass the intersection. As shown in Fig. 6c, at $p = 60 \text{ m}$ when the ego vehicle enters the intersection from static, the danger zone is barely disjoint with the hidden set. In the baseline result, the ego vehicle always freezes before the intersection due to the rapidly growing forward-reachable set of hidden states over the planning horizon (uninhibited by future measurements).

E. Scenario 3: Truck Overtaking

Finally, we set up the overtaking example in a long undivided highway of CARLA map Town01 as shown in Fig. 7. The ego agent is initially 10 m behind the truck, and both vehicles start with the velocity of 12 m s^{-1} . The speed limit in both traffic directions is 20 m s^{-1} . The open-loop game’s planning horizon is 15 time steps, 0.5 s apart. Trajectories are replanned every time step to improve planning efficiency.

We formulate the closed-loop strategy by first finding two shortest evasive trajectories reaching invariant safe states behind and ahead of the truck. Since our vehicle dynamics are represented as a double integrator, both trajectories and danger zone can be analytically expressed.

In Fig. 7 and 8, we show the ego agent’s overtaking trajectory and corresponding velocity profile using the proposed algorithm. Initially, the autonomous vehicle slows down to increase the distance behind the truck and gain visibility of oncoming traffic. After the lane change, the vehicle briefly accelerates and then maintains a constant velocity to observe the driving lane ahead of the truck. By doing so, the ego agent maintains the option to return to a safe state behind the truck, in the event of an oncoming vehicle being observed later. Once it determines that it is safe to merge in the front of the truck, the ego agent accelerates immediately, completes the overtaking, and then returns back to its original lane. During this process, even though an oncoming agent enters the field

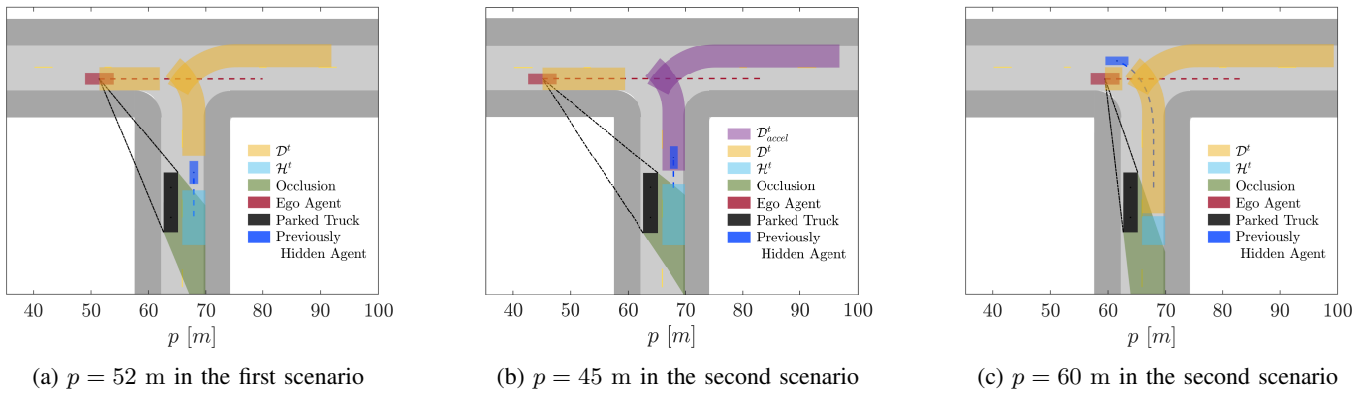


Fig. 6: Three distinctive interaction scenarios in our intersection example. a) the ego agent approaches the intersection with the pursuer outside its danger zone. b) the ego agent approaches the intersection and starts slowing down as the pursuer enters the intersection. c) the ego agent continues to proceed after the pursuer exited the intersection




Fig. 7: The ego agent safely overtakes the truck, merges back to the driving lane, and avoids collisions with oncoming agent




Fig. 9: The ego agent attempts to overtake the truck and safely aborted the overtaking after detecting the oncoming agent




Fig. 8: Ego agent's progress versus its longitudinal velocity during the overtaking attempts

of view, the ego agent still successfully finishes the overtaking since the agent never enters the danger zone. For comparison, we test the baseline algorithm in the same scenario: velocity profiles are shown in Fig. 8. The vehicle using the baseline method-keeps a constant velocity behind the truck, because the forward reachable set of hypothetical oncoming traffic limits its ability to change lanes.

In another overtaking example shown in Fig. 9, we set the oncoming traffic to enter the ego agent's field of view earlier than the previous scenario. As shown in the velocity profile of Fig. 8, the ego agent follows the same strategy initially and moves towards the opposite lane. After the oncoming traffic is detected, the ego agent determines it is unsafe to overtake

the truck. It slows down immediately, merges back into the driving lane, and avoids colliding with the oncoming vehicle.

VII. CONCLUSION AND FUTURE WORK

In this paper, we introduce a novel game-theoretic analysis for autonomous vehicles' motion planning under densely-occluded environments, and derive a decision-making framework that is compatible with a wide class of motion planning algorithms and provides worst-case safety guarantees. We demonstrate our approach in the CARLA simulator.

This work opens the door to multiple promising research directions. In real driving settings, the autonomous vehicle often encounters situations that require interaction and coordination with multiple agents to navigate efficiently and safely. Although obtaining a closed-loop solution for dynamic games involving multiple agents is computationally intractable in real time, the proposed framework may become a building block to efficiently approximate solutions for such multi-sided encounters by combining open- and closed-loop strategies from multiple sub-games.

While our implementation assumes perfect perception and state estimation, the game-theoretic analysis can be extended to scenarios in which sensing limitations are induced not

only by obstacles but also by the physical properties and eventual failures of the sensor itself. When incorporating a more sophisticated sensor model, our hybrid game setup may help prevent unsafe behavior due to perception uncertainty and sensor malfunctions.

Finally, while our game-theoretic analysis computes safety measures under worst-case scenarios, probabilistic modeling of hidden agents can be built upon the current framework to generate more efficient vehicle motion, by considering the observed traffic density and other prior or acquired knowledge about traffic participants' behavior.

REFERENCES

- [1] B. Paden, M. p, S. Z. Yong, D. Yershov, and E. Frazzoli. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles*, 1(1):33–55, 2016. doi: 10.1109/TIV.2016.2578706.
- [2] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Darius M Gavrila, and Kai O Arras. Human motion trajectory prediction: a survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020. doi: 10.1177/0278364920917446. URL <https://doi.org/10.1177/0278364920917446>.
- [3] Markus Koschi and Matthias Althoff. Set-based prediction of traffic participants considering occlusions and traffic rules. *IEEE Transactions on Intelligent Vehicles*, 2020.
- [4] 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)*, pages 1729–1736, 2018.
- [5] M. Y. Yu, R. Vasudevan, and M. Johnson-Roberson. Risk assessment and planning with bidirectional reachability for autonomous driving. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5363–5369, 2020.
- [6] Ming-Yuan Yu, Ram Vasudevan, and Matthew Johnson-Roberson. Occlusion-aware risk assessment for autonomous driving in urban environments. *IEEE Robotics and Automation Letters*, 4(2):2235–2241, 2019.
- [7] Kanwaldeep Kaur and Giselle Rampersad. Trust in driverless cars: Investigating key factors influencing the adoption of driverless cars. 48:87–96. ISSN 09234748. doi: 10.1016/j.jengtecman.2018.04.006. URL <https://linkinghub.elsevier.com/retrieve/pii/S0923474817304253>.
- [8] Eric Thorn, Shawn C Kimmel, Michelle Chaka, Booz Allen Hamilton, et al. A framework for automated driving system testable cases and scenarios. Technical report, United States. Department of Transportation. National Highway Traffic Safety , 2018.
- [9] Tamer Başar and Geert Jan Olsder. *Dynamic noncooperative game theory*. SIAM, 1998.
- [10] Y. Ho, A. Bryson, and S. Baron. Differential games and optimal pursuit-evasion strategies. *IEEE Transactions on Automatic Control*, 10(4):385–389, 1965. doi: 10.1109/TAC.1965.1098197.
- [11] Rufus Isaacs. *Differential games: a mathematical theory with applications to warfare and pursuit, control and optimization*. Courier Corporation, 1999.
- [12] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pages 1–16, 2017.
- [13] Woojin Chung, Seokgyu Kim, Minki Choi, Jaesik Choi, Hoyeon Kim, Chang-bae Moon, and Jae-Bok Song. Safe navigation of a mobile robot considering visibility of environment. *IEEE Transactions on Industrial Electronics*, 56(10):3941–3950, 2009.
- [14] Sara Bouraine, Thierry Fraichard, and Hassen Salhi. Provably safe navigation for mobile robots with limited field-of-views in dynamic environments. *Autonomous Robots*, 32(3):267–283, 2012.
- [15] Sara Bouraine, Th Fraichard, Ouahiba Azouaoui, and Hassen Salhi. Passively safe partial motion planning for mobile robots with limited field-of-views in unknown dynamic environments. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3576–3582. IEEE, 2014.
- [16] Kristijan Macek, Dizan Alejandro Vasquez Govea, Thierry Fraichard, and Roland Siegwart. Towards safe vehicle navigation in dynamic urban scenarios. *Automatika—Journal for Control, Measurement, Electronics, Computing and Communications*, 2009.
- [17] Y. Nager, A. Censi, and E. Frazzoli. What lies in the shadows? safe and computation-aware motion planning for autonomous vehicles using intent-aware dynamic shadow regions. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5800–5806, 2019. doi: 10.1109/ICRA.2019.8793557.
- [18] Minchul Lee, Kichun Jo, and Myoung-ho Sunwoo. Collision risk assessment for possible collision vehicle in occluded area based on precise map. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6. IEEE, 2017.
- [19] Garrison Neel and Srikanth Saripalli. Improving bounds on occluded vehicle states for use in safe motion planning. In *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 268–275. IEEE, 2020.
- [20] Sebastian Brechtel, Tobias Gindele, and Rüdiger Dillmann. Probabilistic decision-making under uncertainty for autonomous driving using continuous pomdps. In *17th international IEEE conference on intelligent transportation systems (ITSC)*, pages 392–399. IEEE, 2014.
- [21] Maxime Bouton, Alireza Nakhaei, Kikuo Fujimura, and Mykel J Kochenderfer. Scalable decision making with sensor occlusions for autonomous driving. In *2018 IEEE International Conference on Robotics and Automation*

- (ICRA), pages 2076–2081. IEEE, 2018.
- [22] C. Hubmann, N. Quetschlich, J. Schulz, J. Bernhard, D. Althoff, and C. Stiller. A pomdp maneuver planner for occlusions in urban scenarios. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 2172–2179, 2019. doi: 10.1109/IVS.2019.8814179.
- [23] Ruzena Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):966–1005, 1988.
- [24] Ruzena Bajcsy, Yiannis Aloimonos, and John K Tsotsos. Revisiting active perception. *Autonomous Robots*, 42(2): 177–196, 2018.
- [25] A. Unterholzner, M. Himmelsbach, and H. Wuensche. Active perception for autonomous vehicles. In *2012 IEEE International Conference on Robotics and Automation*, pages 1620–1627, 2012. doi: 10.1109/ICRA.2012.6224879.
- [26] Hans Andersen, Wilko Schwarting, Felix Naser, You Hong Eng, Marcelo H Ang, Daniela Rus, and Javier Alonso-Mora. Trajectory optimization for autonomous overtaking with visibility maximization. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8. IEEE, 2017.
- [27] Lingguang Wang, Carlos Fetnandez Lopez, and Christoph Stiller. Generating efficient behaviour with predictive visibility risk for scenarios with occlusions. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7. IEEE, 2020.
- [28] Earl A Coddington and Norman Levinson. *Theory of ordinary differential equations*. Tata McGraw-Hill, 1955.
- [29] Rufus Isaacs. Games of pursuit. Technical report, RAND Corporation, Santa Monica, California, 1951. URL <http://www.rand.org/pubs/papers/P257.html>.
- [30] Rufus Isaacs. Differential games. Technical report, RAND Corporation, Santa Monica, California, 1954. URL https://www.rand.org/pubs/research_memoranda/RM1391.html.
- [31] J. F. Fisac and S. S. Sastry. The pursuit-evasion-defense differential game in dynamic constrained environments. In *IEEE Conference on Decision and Control (CDC)*, pages 4549–4556. doi: 10.1109/CDC.2015.7402930.
- [32] Jaime F. Fisac, Mo Chen, Claire J. Tomlin, and S. Shankar Sastry. Reach-avoid problems with time-varying dynamics, targets and constraints. In *International Conference on Hybrid Systems: Computation and Control (HSCC)*, HSCC '15, page 1120, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450334334. doi: 10.1145/2728606.2728612. URL <https://doi.org/10.1145/2728606.2728612>.
- [33] C.J. Tomlin, J. Lygeros, and S. Shankar Sastry. A game theoretic approach to controller design for hybrid systems. 88(7):949–970. doi: 10.1109/5.871303. URL <http://ieeexplore.ieee.org/document/871303/>.
- [34] L. C. Evans and P. E. Souganidis. Differential games and representation formulas for solutions of Hamilton-Jacobi-Isaacs equations. 33(5):773–797. ISSN 0022-2518.
- [35] J. Bellingham, A. Richards, and J. P. How. Receding horizon control of autonomous aerial vehicles. In *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*, volume 5, pages 3741–3746 vol.5, 2002. doi: 10.1109/ACC.2002.1024509.
- [36] Stephane Petti and Thierry Fraichard. Partial motion planning framework for reactive planning within dynamic environments. In *Proc. of the IFAC/AAAI Int. Conf. on Informatics in Control, Automation and Robotics*, 2005.
- [37] David Q Mayne. Model predictive control: Recent developments and future promise. *Automatica*, 50(12): 2967–2986, 2014.
- [38] Tom Schouwenaars. *Safe trajectory planning of autonomous vehicles*. PhD thesis, Massachusetts Institute of Technology, 2006.
- [39] Stéphane Petti and Thierry Fraichard. Safe motion planning in dynamic environments. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2210–2215. IEEE, 2005.
- [40] Albert Wu and Jonathan P How. Guaranteed infinite horizon avoidance of unpredictable, dynamically constrained obstacles. *Autonomous robots*, 32(3):227–242, 2012.
- [41] Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Genaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *2019 18th European Control Conference (ECC)*, pages 3420–3431. IEEE, 2019.
- [42] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2016.
- [43] Alexander Robey, Haimin Hu, Lars Lindemann, Hanwen Zhang, Dimos V Dimarogonas, Stephen Tu, and Nikolai Matni. Learning control barrier functions from expert demonstrations. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 3717–3724. IEEE, 2020.
- [44] Christian Pek and Matthias Althoff. Efficient computation of invariably safe states for motion planning of self-driving vehicles. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3523–3530. IEEE, 2018.
- [45] Ian M. Mitchell. The Flexible, Extensible and Efficient Toolbox of Level Set Methods. *Journal of Scientific Computing*, 35(2):300–329, 2008. ISSN 1573-7691. doi: 10.1007/s10915-007-9174-4.
- [46] Benoit Landry, Mo Chen, Scott Hemley, and Marco Pavone. Reach-avoid problems via sum-of-squares optimization and dynamic programming. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4325–4332. IEEE, 2018.
- [47] Matthew McNaughton, Chris Urmson, John M Dolan, and Jin-Woo Lee. Motion planning for autonomous driving with a conformal spatiotemporal lattice. In *2011 IEEE International Conference on Robotics and Automation*, pages 4889–4895. IEEE, 2011.
- [48] J. Ziegler and C. Stiller. Spatiotemporal state lattices

for fast trajectory planning in dynamic on-road driving scenarios. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1879–1884, 2009. doi: 10.1109/IROS.2009.5354448.

- [49] Armin Hornung, Kai M. Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*, 2013. doi: 10.1007/s10514-012-9321-0. URL <http://octomap.github.com>. Software available at <http://octomap.github.com>.

A. Proofs of Theoretical Results

Proof of Proposition 1: The proof of this proposition is constructive, by enumerating the winning and losing conditions for both players. The evader's winning conditions in the hybrid game are those from which it can indefinitely avoid capture by the pursuer while remaining collision-free. Since the evader will not obtain new information about the pursuer's state until detection, it must (at least initially) commit to an open-loop strategy \mathbf{u}_e . Because we are concerned with worst-case analysis, the pursuer can choose its initial state x_o^0 and open-loop strategy \mathbf{u}_o after the evader has declared \mathbf{u}_e .

P1: evader collision. An evader control signal \mathbf{u}_e can only be a winning strategy if the resulting trajectory avoids collisions with environment obstacles $\mathcal{O}_e(t)$ for all future time. Failure of this condition at any future time means that the evader cannot maintain safety and therefore loses the game.

P2: capture basin reached. If at some time $t \geq 0$ the state enters $\mathcal{B}(t)$, however, then by definition the pursuer can subsequently achieve capture in the closed-loop game; this further implies that the pursuer can achieve capture *regardless* of whether or not the evader eventually detects it. Indeed, at time t , the pursuer has an available non-anticipative strategy \mathbf{v}_o for which capture ensues for *all* possible evader control signals, including the currently declared \mathbf{u}_e . (Conversely if the joint open-loop trajectory never enters the capture basin of the closed-loop game, then the evader's open-loop strategy is guaranteed to avoid capture regardless of whether or not the pursuer is detected.)

The only way for the evader to win the game if the joint open-loop trajectory enters the capture basin at time t is for its winning conditions to take place at some $\tau \in [0, t)$.

E1: pursuer collision. If the pursuer has violated its obstacle constraints $\mathcal{O}_o(\tau)$, it has lost the game and its subsequent open-loop trajectory is irrelevant (and physically meaningless).

E2: pursuer detected. If the pursuer has been detected by entering the evader's field of view while the joint state is outside of the capture basin $\mathcal{B}(\tau)$, the remainder of the open-loop trajectory is not realized, and the game instead transitions to the closed-loop mode, in which the evader is guaranteed to avoid capture and remain collision-free by applying a feedback strategy.

Since the above classification of game executions is complete (every possible game outcome belongs to one of the above categories), we have that the evader wins the game if, and only if, $\neg \mathbf{P1} \wedge (\mathbf{P2} \implies (\mathbf{E1} \vee \mathbf{E2}))$. This concludes the proof. ■

Proof of Theorem 1: We need to show that, for any given \mathbf{u}_e , characterization (1) imposed for all $t \geq 0$ is equivalent to condition (6). Indeed, let us require that (1) holds $\forall t \geq 0$, that is, the forward hidden set $\vec{\mathcal{H}}_o(t; \mathcal{H}_o^0, \mathbf{x}_e)$ and danger zone $\mathcal{D}(t; x_e^t)$ are disjoint for all t . By Definition 2, this condition is equivalent to requiring that any states that the pursuer can reach at time t without first being detected or violating the obstacle constraints must be outside of the danger zone, that is,

by Definition 3, the joint state $(x_e^t, x_o^t) \notin \mathcal{B}(t)$. This requires that there can be no undetected, collision-free trajectory of the pursuer for which the joint state eventually enters the closed-loop capture basin, which is exactly the contrapositive of (6). This concludes the proof. ■

B. Approximate reach-avoid solution

To further simplify the closed-loop game, we can under-approximate the winning condition of the reach-avoid problem defined in (8) by searching among a finite set of open-loop strategies $\mathbf{U} = \{\mathbf{u}_1, \mathbf{u}_2, \dots\}$.

We first define a set of reaching trajectory $\tilde{\mathbf{x}}_i$, which drives the ego agent to a state in the target invariant safe set Ω before T and avoid colliding with other objects using one of the open-loop strategy \mathbf{u}_i .

$$\begin{aligned} \mathcal{X}_e(x_e^t, T, \mathbf{U}) = \{ & \tilde{\mathbf{x}}_i(\cdot; t, x_e^t, \mathbf{u}_i) \mid \exists \mathbf{u}_i \in \mathbf{U}, \exists \tau \in [t, T], \\ & \tilde{\mathbf{x}}_i(\tau; t, x_e^t, \mathbf{u}_i) \in \Omega \wedge \forall s \in [t, \tau], \\ & \phi_e(\tilde{\mathbf{x}}_i(\tau; t, x_e^t, \mathbf{u}_i)) \cap \mathcal{O}_e(s) = \emptyset \} \end{aligned} \quad (12)$$

Assume the ego agent is committing to one of the reaching trajectories $\tilde{\mathbf{x}}_i$, the corresponding trajectory-specific danger zone $\bar{\mathcal{D}}(\tilde{\mathbf{x}}_i)$ is defined as the set of current pursuer states from which the pursuer can achieve capture against this ego trajectory within the time horizon T .

$$\begin{aligned} \bar{\mathcal{D}}(\tilde{\mathbf{x}}_i) = \{ & x_o^t \mid \exists \mathbf{v}_o, \exists \tau \in [t, T], \\ & \phi_e(\tilde{\mathbf{x}}_i(\tau; t, x_e^t, \mathbf{u}_i)) \cap \phi_o(\tilde{\mathbf{x}}_o(\tau; t, x_o^t, \mathbf{v}_o)) \neq \emptyset \} \end{aligned} \quad (13)$$

Therefore, we can over-approximate the danger zone $\mathcal{D}(x_e^t)$ of this reach-avoid game as

$$\mathcal{D}(x_e^t) = \bigcap_{\tilde{\mathbf{x}}_i \in \mathcal{X}_e(x_e^t, T, \mathbf{U})} \bar{\mathcal{D}}(\tilde{\mathbf{x}}_i). \quad (14)$$