

Worst-Case Hidden-Vehicle Trajectory Search in Spatiotemporal Occlusion Regions

Ruichen Tan¹, Zengxiang Lei¹, and Satish Ukkusuri^{1*}

Lyles School of Civil and Construction Engineering, Purdue University,
West Lafayette, IN 47907, USA
{tan479, lei67, sukkusur}@purdue.edu

Abstract. Occlusion creates fundamental uncertainty in autonomous driving. Existing methods often propagate frame-wise hypotheses or optimize ego behavior against prescribed hidden-agent predictions, leaving the worst history-consistent interaction unexplored. We introduce History-Conditioned Minimax Trajectory Search (HC-MTS), which combines temporal occlusion reasoning with response-aware search. First, HC-MTS constructs finite hidden-state modes, each certified by a backward witness satisfying multi-frame visibility, occupancy, semantic-map support, and class-specific kinematic constraints. It then solves a bilevel minimax problem: an inner finite oracle maximizes the ego driving score over destination attainment and ride comfort, while the outer search selects the legal hidden-vehicle trajectory that minimizes this best-response value. Across eight Waymo Open Motion Dataset scenarios, increasing the visibility-memory horizon from $K = 1$ to $K = 20$ reduces the mean per-scenario vehicle, pedestrian, and total retained hidden-seed counts by 18.12%, 21.67%, and 18.45%, respectively. HC-MTS identifies six avoidable counterexamples, while no legal collision-producing attacker is found in the remaining two scenes within the finite search budget.

Keywords: Occlusion reasoning · Hidden traffic participants · Critical trajectory generation · Autonomous-driving safety

1 Introduction

Occlusion introduces a fundamental source of uncertainty in autonomous driving. Vehicles and pedestrians concealed by traffic, roadside structures, or road geometry may be absent from the current perception output while remaining capable of entering the ego vehicle’s path within the planning horizon. Treating occluded space as free is therefore unsafe, whereas assuming that every blind location is occupied leads to unnecessarily conservative behavior. Effective occlusion reasoning must distinguish physically plausible hidden participants from hypotheses that are incompatible with the ego vehicle’s observation history. In particular, a state that is geometrically feasible within the current blind area should be

* Corresponding author.

discarded when no admissible trajectory could have reached that state without passing through previously visible and unoccupied space [9, 16, 19, 21, 22].

A common class of occlusion-aware methods initializes hypothetical participants from the current field of view and propagates their possible states for risk estimation, safety verification, or motion planning [9, 13, 14, 16, 22]. Probabilistic and learning-based approaches estimate likely hidden occupancy or motion, while POMDP, game-theoretic, and contingency-based methods optimize ego behavior under partial observability [1, 6, 17, 23]. Although these approaches address hidden-state estimation and decision-making under uncertainty, safety-critical evaluation requires jointly reasoning about both components. A hidden trajectory that collides with the nominal ego plan may impose little actual risk if the ego vehicle can avoid it without substantially sacrificing progress or comfort. Conversely, a less obvious trajectory may remain highly disruptive even after the ego vehicle selects its strongest available response. Identifying the most consequential hidden trajectory therefore requires evaluating each candidate against the ego vehicle’s optimized response rather than against a fixed nominal plan.

We introduce *History-Conditioned Minimax Trajectory Search* (HC-MTS), a two-stage framework combining temporal occlusion reasoning with response-aware worst-case trajectory discovery. In Stage I, HC-MTS samples class-specific current speeds and backward control hypotheses for candidate states within the current blind area. A candidate is retained as a finite hidden-state mode only when it admits at least one backward witness satisfying multi-frame visibility, observed-occupancy, semantic-map, and class-specific kinematic constraints over the selected memory window. All memory horizons are evaluated from the same sampled modes using a prefix-validity test. Consequently, extending the observation history can eliminate unsupported hypotheses but cannot introduce new ones.

In Stage II, the certified modes initialize a collision-conditioned bilevel minimax search. Candidate hidden trajectories must satisfy forward dynamics, geometry, semantic-map, occupancy, and legality constraints. To concentrate the finite search on safety-relevant interactions, each proposal must also collide with the nominal ego plan. This collision condition serves only as a proposal mechanism; the final severity of a hidden trajectory is determined after the ego vehicle is allowed to respond. Let $\mathcal{T}_K$ denote the set of legal hidden-vehicle trajectories initialized from modes certified over a K -frame history, and let $\mathcal{R}_e(\tau_h)$ denote the finite set of ego responses considered for a hidden trajectory τ_h . HC-MTS approximates

$$\tau_h^* \in \arg \min_{\tau_h \in \mathcal{T}_K} \max_{\tau_e \in \mathcal{R}_e(\tau_h)} J(\tau_e, \tau_h), \quad (1)$$

where J measures task-level driving performance through destination attainment and ride comfort. For each candidate hidden trajectory, the inner response oracle identifies the highest-scoring ego maneuver available within its finite response set. The outer search then seeks a legal hidden trajectory that minimizes this best-response value. The refinement process consequently prioritizes trajectories that remain damaging after accounting for the strongest ego response found by the oracle. The implemented method is a budgeted finite-oracle approximation

using collision-conditioned proposals and a Gaussian elite refinement; it is not presented as a converged solution to a continuous bilevel optimization problem.

We evaluate HC-MTS on eight vectorized scenarios from the Waymo Open Motion Dataset (WOMD) [3]. The experiments examine how multi-frame visibility memory contracts the set of admissible hidden states and how the minimax search differentiates nominal-plan collisions that can be resolved by an ego response from counterexamples that remain unresolved under the current finite oracle. Together, these evaluations demonstrate the importance of combining temporal evidence with response-aware adversarial search when assessing risks created by occluded traffic participants. Our contributions are threefold:

1. We formulate hidden-state inference over multiple frames as a finite history-certification problem combining visibility, observed occupancy, semantic-map, and class-specific kinematic constraints. Every retained mode carries an explicit backward witness, and the retained set contracts monotonically with the memory horizon.
2. We formulate hidden-trajectory discovery as a minimax problem: the ego response oracle maximizes task-level driving performance, while the adversarial search minimizes the resulting best-response score over legal history-certified trajectories.
3. We provide an exact interface between inference and search: counterexample generation starts from the certified current modes without resampling their states or using hidden-agent future labels. Experiments on WOMD quantify both the contraction from temporal evidence and the counterexamples exposed by minimax search.

2 Related Work

2.1 History-Aware Reasoning about Occluded Participants

Occlusion-aware methods represent unseen participants using sampled hypotheses, reachable occupancy sets, phantom agents, or learned probability distributions. Geometry- and set-based approaches initialize possible agents from the current field of view and propagate their states for risk estimation, safety verification, or motion planning [9, 13, 14, 16, 22]. Learning-based methods instead infer likely hidden occupancy and motion from traffic data [1, 10, 11]. Although these methods capture plausible future hazards, estimates derived primarily from the current blind area may retain states that could not have remained hidden throughout earlier observations. Temporal approaches address this limitation by tracking occluded regions or sequentially removing hidden states that conflict with newly available visibility evidence [12, 15, 19, 21]. HC-MTS builds on this temporal principle through an explicit history certificate: every retained current mode must admit a backward witness satisfying multi-frame visibility, observed-occupancy, semantic-map, and class-specific kinematic constraints. The certified modes then serve directly as initial conditions for worst-case trajectory search rather than only as occupancy bounds or inputs to an ego planner.

2.2 Response-Aware Planning under Occlusion

Partially observable, information-aware, and game-theoretic planners account for uncertainty when selecting ego behavior. Hubmann et al. incorporate potentially occluded vehicles and future visibility into a POMDP maneuver planner [6], while Gilhuly et al. optimize trajectories for both safety and information acquisition [4]. Zhang and Fisac formulate occlusion-aware driving as a hybrid dynamic game that accounts for adversarial hidden behavior and the ego vehicle’s future response [23], and Qiu and Fridovich-Keil infer occluded-agent behavior within a receding-horizon contingency-game framework [17]. These methods primarily seek safe or informative ego policies under uncertainty. HC-MTS instead treats the hidden trajectory as the outer decision variable and evaluates its severity only after optimizing the ego response. Specifically, the inner finite oracle maximizes a task-level driving score based on destination attainment and comfort, while the outer search identifies a legal, history-certified hidden trajectory that minimizes this best-response value. The resulting output is therefore an interpretable worst-case hidden trajectory together with the strongest ego response found against it, rather than only a belief-conditioned ego action or binary safety guarantee.

2.3 Safety-Critical Scenario and Trajectory Generation

Safety-critical scenario-generation methods search for rare interactions that expose failures more efficiently than naturalistic sampling. Adaptive Stress Testing optimizes stochastic environment disturbances [8]; AdvSim perturbs actor trajectories while preserving physical plausibility [20]; STRIVE searches the latent space of a learned traffic model for realistic collision-inducing scenes [18]; and KING differentiates through a kinematic proxy to modify surrounding traffic adversarially [5]. Recent work also combines partial-observability risk estimation with generative models for adversarial scenario synthesis [7]. These approaches generally perturb visible actors, learned scene representations, or simulator parameters. HC-MTS addresses a different source of risk: an actor that may never have been observed. Its current state must first be supported by a feasible hidden history, and its future trajectory must satisfy dynamic, geometric, occupancy, map, and legality constraints. Furthermore, collision with the nominal ego plan is used only to generate critical candidates; final severity is determined by the optimized ego-response score. HC-MTS therefore extends safety-critical trajectory generation to history-conditioned hidden participants and evaluates adversariality through a response-aware minimax objective.

3 Methodology

3.1 HC-MTS Overview

History-Conditioned Minimax Trajectory Search (HC-MTS) identifies a history-consistent hidden trajectory that remains most detrimental after the ego vehicle




Fig. 1: Overview of *History-Conditioned Minimax Trajectory Search* (HC-MTS). Stage I certifies current hidden-agent modes using backward witnesses that satisfy multi-frame visibility, observed-occupancy, semantic-map, and class-specific kinematic constraints. Stage II initializes legal forward trajectories from the certified modes and performs a finite response-aware minimax search: an inner oracle maximizes the ego driving score, while the outer search seeks the hidden trajectory that minimizes this best-response value. Collision with the current ego plan is used to generate critical proposals, not to define their final severity.

is allowed to respond. At the current time t_0 , HC-MTS receives the ego state, a vectorized semantic map, and observed road-user tracks from the current frame and the preceding K frames. All scene elements are expressed in an ego-fixed bird's-eye-view frame anchored at t_0 , with the ego vehicle at the origin, the $+x$ axis pointing forward, and the $+y$ axis pointing left.

HC-MTS consists of two stages. Stage I performs *history-conditioned hidden-mode inference*. It samples candidate current states in the blind area, augments them with class-specific speed and backward-control hypotheses, and retains only modes supported by at least one feasible backward witness through the observation history. Stage II performs *response-aware minimax trajectory search*. Each certified mode initializes legal forward hidden trajectories. For every candidate hidden trajectory, a finite ego-response oracle searches for the response that maximizes task-level driving performance. The outer search then favors hidden trajectories that reduce this optimized ego score. Figure 1 separates feasibility from adversarial search: Stage I identifies history-consistent hidden modes, and Stage II optimizes their future trajectories for worst-case interaction with the ego vehicle.

3.2 Problem Definition

For each frame t , let Ω_t denote the field of regard, let $\mathcal{B}_t^{\text{obs}}$ be the union of observed road-user footprints, and let $\mathcal{F}_t \subseteq \Omega_t$ denote the *visible-free region*: space in which a road user would be expected to have been observed under the adopted geometric visibility model and that is not occupied by a detected participant. The unresolved non-visible region is

$$\mathcal{O}_t = \Omega_t \setminus (\mathcal{F}_t \cup \mathcal{B}_t^{\text{obs}}). \quad (2)$$

Thus, $\mathcal{F}_t$, $\mathcal{B}_t^{\text{obs}}$, and $\mathcal{O}_t$ distinguish observed free space, observed occupied space, and space whose occupancy remains unresolved. A hidden participant cannot intersect either of the first two sets, but it may occupy $\mathcal{O}_t$.

Given the observation history from $t_0 - K$ to t_0 , Stage I of HC-MTS returns a finite set of history-certified hidden-state modes

$$\mathcal{M}_K = \{m \mid \mathcal{C}_K(m) = 1\}, \quad (3)$$

where $\mathcal{C}_K$ is the K -frame consistency predicate defined in Section 3.4. Each $m \in \mathcal{M}_K$ contains an exact current hidden state and a backward witness demonstrating that the participant could have reached that state without violating the available visibility, occupancy, map, or kinematic evidence.

Let $\mathcal{T}(m)$ denote the set of forward hidden trajectories initialized from mode m that satisfy the class-specific dynamics, footprint, semantic-map, observed-occupancy, and traffic-legality constraints. The history-conditioned admissible trajectory set is

$$\mathcal{T}_K = \bigcup_{m \in \mathcal{M}_K} \mathcal{T}(m). \quad (4)$$

We make the ego-response objective explicit in a parameterized form. Let τ_{ref} denote the logged ego future, which is used only as a reference trajectory. The nominal ego cost is defined as

$$\begin{aligned} C_{\text{nom}}(\tau_E; \theta_E) = & w_g e_{\text{goal}}^2 + w_p e_{\text{prog}}^2 + w_t e_{\text{time}} + w_r e_{\text{route}} \\ & + w_a \overline{a^2} + w_\kappa \overline{\kappa^2} + w_s \overline{(\Delta a)^2}, \end{aligned} \quad (5)$$

where $\theta_E = \{w_g, w_p, w_t, w_r, w_a, w_\kappa, w_s\}$ contains nonnegative weighting parameters. The first four terms measure goal-reaching, progress, time-aligned reference deviation, and route deviation, respectively, while the remaining terms regularize control effort and smoothness.

The ego utility is then defined as

$$U_E(\tau_E, \tau_A) = \begin{cases} -\infty, & \text{if contact occurs,} \\ U_{\text{max}} - s_U^{-1} [C_{\text{nom}}(\tau_E; \theta_E) + \psi(d_{\text{min}})], & \text{otherwise.} \end{cases} \quad (6)$$

where d_{min} is the minimum time-aligned oriented-box gap, $w_d > 0$ controls the clearance penalty, $\sigma_d > 0$ determines its spatial decay, $s_U > 0$ is a normalization factor, and U_{max} denotes the maximum collision-free utility. This

formulation prioritizes collision avoidance while favoring efficient, smooth, and well-separated ego responses.

For a fixed hidden trajectory $\tau^A \in \mathcal{T}_K$, let $\mathcal{R}_E(\tau^A)$ denote the ego responses available to the response oracle. We set $J(\tau_E, \tau_A) \equiv U_E(\tau_E, \tau_A)$. The best-response value of a hidden trajectory is

$$V_E(\tau^A) = \max_{\tau^E \in \mathcal{R}_E(\tau^A)} J(\tau^E, \tau^A). \quad (7)$$

HC-MTS targets the hidden trajectory that minimizes the ego vehicle's best achievable score:

$$\tau^{A*} \in \arg \min_{\tau^A \in \mathcal{T}_K} \max_{\tau^E \in \mathcal{R}_E(\tau^A)} J(\tau^E, \tau^A) = \arg \min_{\tau^A \in \mathcal{T}_K} V_E(\tau^A). \quad (8)$$

3.3 Ego-Centric Occlusion Representation

HC-MTS constructs $\mathcal{F}_t$ and $\mathcal{O}_t$ with a geometry-based visibility procedure. Line-of-sight rays are cast within the field of regard, and the oriented footprints of observed road users are treated as occluders. The resulting partition is used consistently across all frames in the selected history window.

Current hidden-agent proposals are restricted to non-visible, semantically valid support. A spatial seed is

$$s = (\mathbf{p}_0, \psi_0, c), \quad (9)$$

Here, $\mathbf{p}_0 \in \mathbb{R}^2$ is the current position and ψ_0 is the heading. The class variable c denotes either a vehicle or a pedestrian. Vehicle centers are sampled from lane or drivable corridors, with headings induced by the local lane direction. Pedestrian centers are sampled from crosswalk, walkable, and road-edge support. A seed is retained only if its class-specific footprint is hidden under the sampled-footprint visibility test and does not overlap an observed road user. This produces

$$\mathcal{S}_{\text{base}} = \{s = (\mathbf{p}_0, \psi_0, c)\}, \quad (10)$$

which describes geometrically plausible current positions and headings but does not yet certify temporal feasibility.

3.4 HC-MTS Stage I: History-Conditioned Hidden-Mode Inference

Stage I of HC-MTS converts the spatial seeds in $\mathcal{S}_{\text{base}}$ into finite hidden-state modes with explicit historical witnesses.

Finite mode construction. A current spatial seed may admit multiple kinematic histories, including straight motion, acceleration, braking, turning, and piecewise control. HC-MTS therefore represents sampled state-history combinations explicitly rather than assigning a continuous speed interval to each seed. A mode is

$$m = (z_0, \mathbf{u}^-, \tau^-), \quad z_0 = (s, v_0), \quad (11)$$

where v_0 is the sampled speed at t_0 , $\mathbf{u}^-$ is a sampled historical control sequence, and $\boldsymbol{\tau}^-$ is the corresponding backward state trace. This explicit representation is necessary because feasibility of one speed-control combination does not imply feasibility of nearby combinations under the visibility, occupancy, map, and dynamics constraints.

For vehicles, the backward-control sequence is

$$\mathbf{u}^- = \{(a_{t_0-j}, \kappa_{t_0-j})\}_{j=1}^{K_{\max}}, \quad (12)$$

where a_t and κ_t denote longitudinal acceleration and path curvature. Constant and piecewise-constant templates represent straight, accelerating, braking, and turning histories. Pedestrian modes use class-specific speed and direction hypotheses without the vehicle curvature model.

Backward witness propagation. For a vehicle state $(\mathbf{p}_t, \psi_t, v_t)$ and control (a_t, κ_t) , HC-MTS applies the sampled backward update

$$v_{t-1} = v_t - a_t \Delta t, \quad \bar{v}_t = \frac{1}{2}(v_t + v_{t-1}), \quad (13)$$

$$\psi_{t-1} = \psi_t - \kappa_t \bar{v}_t \Delta t, \quad \bar{\psi}_t = \frac{1}{2}(\psi_t + \psi_{t-1}), \quad (14)$$

$$\mathbf{p}_{t-1} = \mathbf{p}_t - \bar{v}_t \begin{bmatrix} \cos \bar{\psi}_t \\ \sin \bar{\psi}_t \end{bmatrix} \Delta t. \quad (15)$$

Backward propagation terminates at the first contradiction with visibility, observed occupancy, semantic-map support, or class-specific dynamics. The longest valid prefix is denoted by $K^*(m)$, and the corresponding trace is retained as the backward witness for mode m .

History-consistency constraints. At every frame required by the selected memory horizon, a mode must satisfy

$$Q(B_t(m)) \cap \mathcal{F}_t = \emptyset, \quad \text{visibility}, \quad (16)$$

$$\text{Footprint}(m_t) \cap \mathcal{B}_t^{\text{obs}} = \emptyset, \quad \text{observed occupancy}, \quad (17)$$

$$\text{Footprint}(m_t) \subseteq \mathcal{D}_c, \quad \text{semantic map}, \quad (18)$$

where $B_t(m)$ is the oriented participant footprint, $Q(B_t(m))$ is its finite visibility-sampling set, $\mathcal{B}_t^{\text{obs}}$ is the union of observed road-user footprints, and $\mathcal{D}_c$ is the semantic support for class c . Vehicle modes additionally satisfy

$$v_t \in [0, v_{\max}^{\text{veh}}], \quad |\kappa_t v_t| \leq \omega_{\max}^{\text{veh}}, \quad |\kappa_t v_t^2| \leq a_{\text{lat}, \max}^{\text{veh}}, \quad (19)$$

while pedestrian modes satisfy $v_t \in [0, v_{\max}^{\text{ped}}]$ and $\mathbf{p}_t \in \mathcal{D}_{\text{ped}}^{\text{hist}}$.

The combined K -frame certificate is

$$\mathcal{C}_K(m) = C_{\text{cur}}^{t_0}(m) \prod_{j=1}^K C_{\text{vis}}^{t_0-j}(m) C_{\text{occ}}^{t_0-j}(m) C_{\text{map}}^{t_0-j}(m) C_{\text{dyn}}^{t_0-j}(m). \quad (20)$$

Because every requested memory horizon is evaluated from the same sampled modes and the same prefix-validity record,

$$\mathcal{M}_{K_2} \subseteq \mathcal{M}_{K_1}, \quad K_2 > K_1. \quad (21)$$

Longer visibility memory can therefore preserve or eliminate an existing sampled hypothesis, but it cannot introduce a new one.

3.5 HC-MTS Stage II: Iterative Mini-max Trajectory Search

Stage II of HC-MTS (Algorithm 1) is implemented as a sequence of finite mini-max rounds. It is therefore not a single evaluation over a fixed attacker set. Each round starts from the ego trajectory accepted in the preceding round, searches for hidden trajectories that challenge that trajectory, optimizes the ego response to each candidate, and continues only when the selected attacker further reduces the ego vehicle's optimized driving score.

Let τ_i^E denote the ego trajectory at the beginning of round i , with τ_0^E equal to the nominal ego plan. Let $\mathcal{P}$ be the finite library of candidate attacker motion patterns. For every certified mode $m \in \mathcal{M}_K$ and pattern $p \in \mathcal{P}$, HC-MTS searches a bounded set of pattern parameters to instantiate a forward hidden trajectory. Every generated trajectory must satisfy the class-specific dynamics, footprint, semantic-map, observed-occupancy, known-agent, and traffic-legality constraints. The attacker search is collision-seeking: whenever feasible within the finite search budget, the pattern parameters are chosen so that the resulting hidden trajectory collides with the current ego trajectory τ_i^E . The legal collision-producing candidates found in round i form

$$\hat{\mathcal{A}}_i = \left\{ \tau^A \mid \begin{array}{l} \tau^A \text{ is generated from some } m \in \mathcal{M}_K \text{ and } p \in \mathcal{P}, \\ \mathcal{C}_{\text{fwd}}(\tau^A) = 1, \quad \text{Coll}(\tau_i^E, \tau^A) = 1 \end{array} \right\}, \quad (22)$$

where $\mathcal{C}_{\text{fwd}}$ collects the forward feasibility and legality constraints. Because $\hat{\mathcal{A}}_i$ is produced by a bounded search over finite patterns, an empty set means that HC-MTS did not discover a legal collision-producing attacker in that round; it does not prove that none exists in the continuous trajectory space.

Inner ego-response optimization. For every candidate $\tau^A \in \hat{\mathcal{A}}_i$, HC-MTS independently optimizes an ego response. Let $\hat{\mathcal{R}}_{E,i}(\tau^A)$ be the finite response set searched against that candidate. The optimized ego value is

$$V_i(\tau^A) = \max_{\tau^E \in \hat{\mathcal{R}}_{E,i}(\tau^A)} J(\tau^E, \tau^A), \quad (23)$$

with corresponding best response

$$\tau_i^{E*}(\tau^A) \in \arg \max_{\tau^E \in \hat{\mathcal{R}}_{E,i}(\tau^A)} J(\tau^E, \tau^A). \quad (24)$$

Here, larger J indicates better ego performance. Collision receives strict priority, while collision-free responses are ranked by the destination-attainment and ride-comfort terms defined in the driving objective.

Within-round minimization. After solving the inner response problem for every candidate, HC-MTS selects the attacker that produces the lowest optimized ego score:

$$\tau_i^{A*} \in \arg \min_{\tau^A \in \hat{\mathcal{A}}_i} V_i(\tau^A), \quad (25)$$

$$S_i = V_i(\tau_i^{A*}) = J(\tau_i^{E*}(\tau_i^{A*}), \tau_i^{A*}). \quad (26)$$

Equations (23)–(26) define the mini-max problem solved within round i : the inner maximization finds the strongest ego response to each fixed attacker, and the outer minimization selects the attacker that leaves the ego with the lowest best-response score.

Algorithm 1 HC-MTS Stage II: Iterative mini-max trajectory search

Require: Certified modes $\mathcal{M}_K$, attacker-pattern library $\mathcal{P}$, nominal ego plan τ_0^E , score tolerance ϵ , and round budget L

Ensure: Last accepted attacker–response pair and its optimized ego score

- 1: $S_{\text{prev}} \leftarrow +\infty$; accepted pair $\leftarrow \emptyset$
- 2: **for** $i = 0, \dots, L - 1$ **do**
- 3: Generate a finite set $\hat{\mathcal{A}}_i$ of legal, collision-seeking attacker trajectories from $\mathcal{M}_K$ and $\mathcal{P}$ against τ_i^E
- 4: **if** $\hat{\mathcal{A}}_i = \emptyset$ **then**
- 5: **break**
- 6: **end if**
- 7: **for each** $\tau^A \in \hat{\mathcal{A}}_i$ **do**
- 8: Optimize the ego response $\tau_i^{E*}(\tau^A)$
- 9: $V_i(\tau^A) \leftarrow J(\tau_i^{E*}(\tau^A), \tau^A)$
- 10: **end for**
- 11: Optionally refine low-value attacker patterns; revalidate each refined trajectory and re-optimize its ego response
- 12: $\tau_i^{A*} \leftarrow \arg \min_{\tau^A \in \hat{\mathcal{A}}_i} V_i(\tau^A)$
- 13: $\tau_i^{E*} \leftarrow \tau_i^{E*}(\tau_i^{A*})$; $S_i \leftarrow V_i(\tau_i^{A*})$
- 14: **if** $i > 0$ **and** $S_i \geq S_{\text{prev}} - \epsilon$ **then**
- 15: **break** ▷ the optimized ego score no longer decreases
- 16: **end if**
- 17: Accept $(\tau_i^{A*}, \tau_i^{E*}, S_i)$; $S_{\text{prev}} \leftarrow S_i$
- 18: **if** τ_i^{E*} is not collision-free **then**
- 19: Mark the accepted pair as unresolved and **break**
- 20: **end if**
- 21: $\tau_{i+1}^E \leftarrow \tau_i^{E*}$
- 22: **end for**
- 23: **return** the last accepted pair and S_{prev}

Cross-round descent and termination. The selected pair is accepted only when it produces a sufficient decrease relative to the preceding accepted round. For

$i \geq 1$, the acceptance condition is

$$S_i < S_{i-1} - \epsilon, \quad (27)$$

where $\epsilon \geq 0$ is the score-decrease tolerance. If Eq. (27) holds, HC-MTS accepts $(\tau_i^{A*}, \tau_i^{E*}(\tau_i^{A*}))$ and uses the optimized ego response as the plan for the next round:

$$\tau_{i+1}^E = \tau_i^{E*}(\tau_i^{A*}). \quad (28)$$

The next attacker search then attempts to reduce the ego score again by finding a new trajectory against this updated plan. If $S_i \geq S_{i-1} - \epsilon$, the current round does not improve the adversarial objective, the pair is not accepted, and HC-MTS terminates with the last accepted pair.

HC-MTS also terminates when no legal collision-producing attacker is discovered, when the round budget is exhausted, or when the selected attacker admits no collision-free response. The last case is reported as an unresolved counterexample because no collision-free ego trajectory is available to initialize another round.

Finite candidate search and refinement. Within each round, HC-MTS evaluates only a finite number of attacker trajectories generated from the candidate patterns. A cluster-balanced shortlist may be used to distribute this budget across distinct current hidden modes or spatial seeds. Candidate patterns associated with low best-response scores can be refined by one round of clipped Gaussian elite sampling inspired by the cross-entropy method [2]. Every refined attacker is revalidated and, crucially, the inner ego-response optimization is rerun for that refined trajectory before its score is compared with other candidates. The refinement, therefore, seeks a lower *optimized* ego score rather than exploiting a response computed for a different attacker.

4 Numerical Results

4.1 Experiment Setups

Research questions. The experiments are designed to answer two research questions: **RQ1:** How does the length of the visibility history affect the set of hidden states that remain consistent with the available observations? **RQ2:** Can HC-MTS discover legal, history-consistent hidden trajectories that reduce the ego vehicle’s best achievable driving score after response optimization? The first question evaluates the history-conditioned inference stage, whereas the second evaluates the iterative response-aware minimax search.

Dataset and scenario selection. We evaluate HC-MTS on eight vectorized scenarios from the WOMD [3]. The scenarios are randomly sampled from the subset for which the geometric visibility procedure produces a nonempty current blind region relevant to the ego vehicle’s planning horizon. The scenarios are selected

once before inspecting the outputs of HC-MTS and are retained regardless of the resulting hidden-area reduction or counterexample-search outcome. The same eight scenarios are used in all experiments. For each scene, the method uses the vectorized semantic map, the ego state, and the observed road-user tracks up to the current time t_0 . No future state of a hypothetical hidden participant is used for either history certification or adversarial trajectory search. All scene elements are represented in the ego-fixed coordinate system defined in Section 3.3.

4.2 Effect of Visibility Memory

Table 1 reports the number of history-consistent hidden seeds retained under different visibility-memory lengths. Increasing K consistently reduces the retained hypothesis set. From $K = 1$ to $K = 20$, the mean per-scenario vehicle, pedestrian, and total seed counts decrease by 18.12%, 21.67%, and 18.45%, respectively. The larger reduction for pedestrians is consistent with their smaller backward-reachable displacement, which makes their historical trajectories more likely to conflict with previously visible-free regions. In contrast, vehicle hypotheses can remain feasible by originating farther along continuously occluded road corridors.

Table 1: Average number of retained history-consistent hidden seeds over eight WOMD scenarios. Reduction is the mean per-scene percentage decrease relative to $K = 1$. Total denotes all retained vehicle and pedestrian seeds.

K	Vehicle		Pedestrian		Total	
	Count	Reduction	Count	Reduction	Count	Reduction
1	3392.12	0.00%	342.62	0.00%	3734.75	0.00%
5	3155.12	6.99%	311.12	9.19%	3466.25	7.19%
10	2992.62	11.78%	289.25	15.58%	3281.88	12.13%
20	2777.38	18.12%	268.38	21.67%	3045.75	18.45%

Fig. 2 provides the corresponding qualitative comparison. Longer histories remove spatially coherent seed clusters rather than uniformly reducing the entire current occlusion mask. Scenarios 1 and 3 show substantial contraction, indicating that ego motion reveals large portions of the candidate hidden corridors. Scenarios 5 and 8 change less because their occlusions remain persistent over the evaluated history. Accordingly, the $K = 1$ to $K = 20$ reduction varies from 4.57% to 35.44% across the eight scenarios, showing that the value of temporal visibility depends strongly on ego motion, occluder geometry, and road topology.

In Fig. 2, darker markers represent higher mean history-feasible current speeds. The color is used only for visualization: each seed may retain multiple feasible motion modes, and the complete mode set, rather than the displayed

mean speed, is passed to the subsequent trajectory optimization. The persistence of several high-speed seeds in continuously occluded corridors also shows that the method does not remove candidates based on danger alone. Seeds are pruned only when their historical existence is contradicted by visibility, dynamics, or physical occupancy constraints.




Fig. 2: Comparison of hidden-seed distributions using different temporal history lengths K across eight WOMB scenarios. Marker color indicates the mean history-feasible current speed, with darker markers representing higher speeds.

4.3 Response-Aware Critical-Trajectory Search

In Fig. 3, the eight scenes exhibit two of the three possible outcomes. In Scenes 1–5 and 7, HC-MTS finds a legal, history-consistent attacker that collides with the current ego trajectory, while the finite ego-response oracle still returns a collision-free response. The corresponding final ego scores are 86.39, 91.44, 95.20, 84.71, 93.44, and 96.26, respectively, with a mean of 91.24. Scene 4 shows the largest degradation, with a score of 84.71. These results show that collision with the reference ego trajectory alone can overstate attacker severity: all six attackers remain avoidable, although replanning incurs losses in destination attainment, or ride comfort.

In Scenes 6 and 8, no legal collision-producing attacker is found within the finite search budget; consequently, the ego score remains at the nominal value of 100, and the outer game terminates without an update. No scene yields a score of $-\infty$, meaning that the finite oracle finds at least one collision-free response for every attacker identified in this evaluation. Current outcomes are budget-dependent and do not constitute formal guarantees over the continuous hidden-state and control spaces. Overall, the results highlight the value of response-

aware evaluation for distinguishing avoidable reference-trajectory collisions from threats that remain unresolved after ego replanning.




Fig. 3: Response-aware hidden-attacker search across eight WOMD scenes. Each panel shows the final result for one scene and reports the ego score and the number of outer best-response iterations. The dashed blue trajectory is the nominal ego trajectory, the solid blue trajectory is the optimized ego response, and the red trajectory is the selected hidden attacker. Vehicle footprints are displayed at the corresponding evaluation time. A finite ego score below 100 indicates that a legal collision-producing attacker was found, but the ego response oracle could still avoid the collision at the cost of reduced progress, and route tracking. A score of 100 indicates that no legal collision-producing attacker was identified under the current search budget. A score of $-\infty$ indicates that no collision-free ego response could be found against the selected attacker (not found within the finite search budget).

5 Conclusion

This work introduced *History-Conditioned Minimax Trajectory Search* (HC-MTS), a unified framework for identifying safety-critical trajectories of traffic participants concealed by occlusion. HC-MTS couples history-consistent hidden-state inference with response-aware adversarial search, enabling candidate trajectories to be evaluated not only for physical and observational feasibility, but also for their effect on the ego vehicle’s best achievable task-level performance.

The experimental results demonstrate two key advantages of the proposed formulation. First, conditioning on multi-frame visibility and occupancy evidence eliminates hidden-state hypotheses that are compatible with the current occluded region but inconsistent with the observation history. Second, optimizing the ego response before assessing trajectory criticality yields a more principled evaluation than measuring collision risk against a fixed nominal plan. This response-aware criterion distinguishes hazards that can be mitigated through replanning from those that remain safety-critical under the available ego strategy, thereby providing a stronger and more informative basis for evaluating autonomous-driving planners.

The current limitation of HC-MTS is its reliance on a heuristic solution algorithm, which does not guarantee global optimality and may become computationally demanding as the scenario complexity and search space increase. Future work will therefore focus on two directions. First, we will evaluate HC-MTS on larger-scale datasets and a broader range of occlusion-critical driving scenarios to assess its robustness and generalizability. Second, we will refine the solution algorithm to improve computational efficiency, optimization stability, and solution quality, and to reduce its dependence on heuristic search.

References

1. Christianos, F., Karkus, P., Ivanovic, B., Albrecht, S.V., Pavone, M.: Planning with occluded traffic agents using bi-level variational occlusion models. arXiv preprint arXiv:2210.14584 (2022)
2. De Boer, P.T., Kroese, D.P., Mannor, S., Rubinstein, R.Y.: A tutorial on the cross-entropy method. *Annals of operations research* **134**(1), 19–67 (2005)
3. Ettinger, S., Cheng, S., Caine, B., Liu, C., Zhao, H., Pradhan, S., Chai, Y., Sapp, B., Qi, C.R., Zhou, Y., et al.: Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9710–9719 (2021)
4. Gilhuly, B., Sadeghi, A., Yedemellat, P., Rezaee, K., Smith, S.L.: Looking for trouble: Informative planning for safe trajectories with occlusions. In: *2022 International Conference on Robotics and Automation (ICRA)*. pp. 8985–8991. IEEE (2022)
5. Hanselmann, N., Renz, K., Chitta, K., Bhattacharyya, A., Geiger, A.: King: Generating safety-critical driving scenarios for robust imitation via kinematics gradients. In: *European Conference on Computer Vision*. pp. 335–352. Springer (2022)
6. Hubmann, C., Quetschlich, N., Schulz, J., Bernhard, J., Althoff, D., Stiller, C.: A pomdp maneuver planner for occlusions in urban scenarios. In: *2019 IEEE Intelligent Vehicles Symposium (IV)*. pp. 2172–2179. IEEE (2019)
7. Jia, J., Su, Y., Bao, Z., Hong, Y., Gao, B., Gan, Z., Ding, W.: Learning a unified risk map for autonomous driving in partially observable environments. *IEEE Robotics and Automation Letters* (2026)
8. Koren, M., Alsaif, S., Lee, R., Kochenderfer, M.J.: Adaptive stress testing for autonomous vehicles. In: *2018 IEEE Intelligent Vehicles Symposium (IV)*. pp. 1–7. IEEE (2018)
9. Koschi, M., Althoff, M.: Set-based prediction of traffic participants considering occlusions and traffic rules. *IEEE Transactions on Intelligent Vehicles* **6**(2), 249–265 (2020)
10. Lange, B., Li, J., Kochenderfer, M.J.: Scene informer: Anchor-based occlusion inference and trajectory prediction in partially observable environments. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 14138–14145. IEEE (2024)
11. Mahjourian, R., Kim, J., Chai, Y., Tan, M., Sapp, B., Anguelov, D.: Occupancy flow fields for motion forecasting in autonomous driving. *IEEE Robotics and Automation Letters* **7**(2), 5639–5646 (2022)
12. Moller, K., Schwarzmeier, L., Betz, J.: From shadows to safety: Occlusion tracking and risk mitigation for urban autonomous driving. In: *2025 IEEE Intelligent Vehicles Symposium (IV)*. pp. 1883–1890. IEEE (2025)

13. Moller, K., Trauth, R., Betz, J.: Overcoming blind spots: Occlusion considerations for improved autonomous driving safety. In: 2024 IEEE Intelligent Vehicles Symposium (IV). pp. 819–826. IEEE (2024)
14. Nager, Y., Censi, A., Frazzoli, E.: 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). pp. 5800–5806. IEEE (2019)
15. Nyberg, T., Sánchez, J.M.G., Pek, C., Tumova, J., Törngren, M.: Evaluating sequential reasoning about hidden objects in traffic. In: 2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPS). pp. 306–307. IEEE (2022)
16. Orzechowski, P.F., Meyer, A., Lauer, M.: Tackling occlusions & limited sensor range with set-based safety verification. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC). pp. 1729–1736. IEEE (2018)
17. Qiu, T., Fridovich-Keil, D.: Inferring occluded agent behavior in dynamic games from noise corrupted observations. *IEEE Robotics and Automation Letters* **9**(12), 11489–11496 (2024)
18. Rempe, D., Phillion, J., Guibas, L.J., Fidler, S., Litany, O.: Generating useful accident-prone driving scenarios via a learned traffic prior. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 17305–17315 (2022)
19. Sánchez, J.M.G., Nyberg, T., Pek, C., Tumova, J., Törngren, M.: Foresee the unseen: Sequential reasoning about hidden obstacles for safe driving. In: 2022 IEEE Intelligent Vehicles Symposium (IV). pp. 255–264. IEEE (2022)
20. Wang, J., Pun, A., Tu, J., Manivasagam, S., Sadat, A., Casas, S., Ren, M., Urtasun, R.: Advsim: Generating safety-critical scenarios for self-driving vehicles. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9909–9918 (2021)
21. Wang, L., Burger, C., Stiller, C.: Reasoning about potential hidden traffic participants by tracking occluded areas. In: 2021 IEEE International Intelligent Transportation Systems Conference (ITSC). pp. 157–163. IEEE (2021)
22. Yu, M.Y., Vasudevan, R., Johnson-Roberson, M.: Occlusion-aware risk assessment for autonomous driving in urban environments. *IEEE Robotics and Automation Letters* **4**(2), 2235–2241 (2019)
23. Zhang, Z., Fisac, J.F.: Safe occlusion-aware autonomous driving via game-theoretic active perception. *arXiv preprint arXiv:2105.08169* (2021)