

A Modeled Approach for Online Adversarial Test of Operational Vehicle Safety

Linda Capito^{*1}, Bowen Weng^{*1},

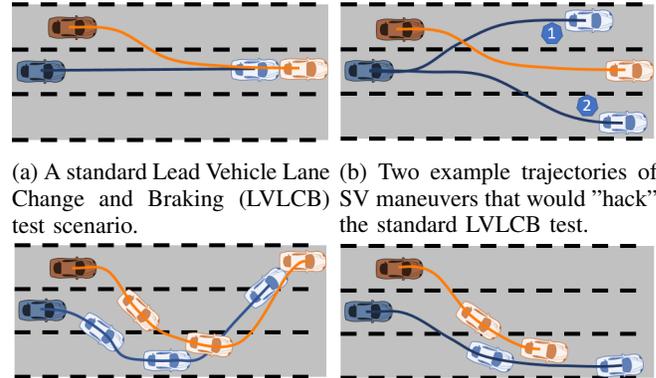
Umit Ozguner¹, *IEEE Life Fellow*, Keith Redmill¹, *IEEE Senior Member and SIAM Member*

Abstract—The scenario-based testing of operational vehicle safety presents a set of principal other vehicle (POV) trajectories that seek to force the subject vehicle (SV) into a certain safety-critical situation. Current scenarios are mostly (i) statistics-driven: inspired by human driver crash data, (ii) deterministic: POV trajectories are pre-determined and are independent of SV responses, and (iii) overly simplified: defined over a finite set of actions performed at the abstracted motion planning level. Such scenario-based testing (i) lacks severity guarantees, (ii) has predefined maneuvers making it easy for an SV with intelligent driving policies to game the test, and (iii) is inefficient in producing safety-critical instances with limited and expensive testing effort. We propose a model-driven online feedback control policy for multiple POVs which propagates efficient adversarial trajectories while respecting traffic rules and other concerns formulated as an admissible state-action space. The approach is formulated in an anchor-template hierarchy structure, with the template model planning inducing a theoretical SV capturing guarantee under standard assumptions. The planned adversarial trajectory is then tracked by a lower-level controller applied to the full-system or the anchor model. The effectiveness of the methodology is illustrated through various simulated examples with the SV controlled by either parameterized self-driving policies or human drivers.

I. INTRODUCTION

A scenario-based evaluation method "attacks" the test subject through enforcing it into a safety-critical situation and observes the subject response and testing outcomes. For operational vehicle safety evaluation, especially concerning Advanced Driver Assistance System (ADAS) [1] and Automated Driving Systems (ADS) [2], the design and execution of test scenarios have raised wide attention from the research community [3], automotive industry, and government sectors for policy making [2], [4], [5]. Typically, one first extracts initialization conditions and vehicle maneuvers inspired by the statistical observation of human driver crash data. The scenario is then specified with a given initialization state for all vehicles and a set of pre-determined principle other vehicle (POV) trajectories. Note that the POV trajectory is generated through a series of high-level planning commands and is often independent of the subject vehicle (SV) response. For example, one can refer to Fig. 1a that shows a Lead Vehicle Lane Change and Braking (LVLCB) scenario [4] for ADAS safety evaluation.

The problems of the aforementioned scenario designs are threefold. First, the current scenarios lack severity guaran-



(a) A standard Lead Vehicle Lane Change and Braking (LVLCB) test scenario. (b) Two example trajectories of SV maneuvers that would "hack" the standard LVLCB test.

(c) An adversarial POV will re-strict the scenario propagation in avoid collisions and force extreme evasive maneuvers that the SV into collisions if it is not beyond the expectation of the responding appropriately. (d) One option for the SV to strict the scenario propagation in avoid collisions is to perform a safety-critical status and force extreme evasive maneuvers that the SV into collisions if it is not beyond the expectation of the responding appropriately.

Fig. 1: Various maneuvers inspired by the standard LVLCB scenario [4]: inspired by a commonly observed cause of rear-end accident from human driver crash data, the POV performs a single lane-change to get in front of the SV, followed by an immediate braking maneuver.

tees. The assumption that a scenario extracted as described above and brought to a test track will generate similar SV responses, similar safety-critical situations, and similar outcomes ignores the significant differences in sensing, control algorithms, control authority and implementation on the different SVs being tested. Second, with pre-determined POV trajectories defined over a simplified action space it is easy for the SV to game the test, especially given that the SV is operating with intelligent decision making both longitudinally and laterally (e.g., Fig. 1b). Finally, the current approach does not create sufficient safety-critical instances within the limited and expensive real-world testing effort. While virtual simulations have provided the opportunity to enhance autonomous vehicle testing efficiency, the sim-to-real gap remains a challenge in general. Furthermore, the simplified action space is incapable of creating agile, dynamic maneuvers, which confines the scenario propagation to relatively simple configurations. This also limits the testing efficiency in both real-world tests and simulations.

Inspired by the problems mentioned above, the adversarial adaptive testing scheme has become a natural alternative. Existing approaches mostly focus on learning inspired methods [6], [7], [8], which typically require the algorithm

^{*}These authors contributed equally.

¹The authors are with the Department of Electrical and Computer Engineering at Ohio State University, OH, USA.

Corresponding author: Keith Redmill, redmill.1@osu.edu

to interact with a fixed SV driving algorithm for many iterations. These are sampling inefficient and would only be applicable in virtual simulations. Some also seek for guided sampling from an offline-built scenario library [3]. While this approach resolves the traditional scenario-based testing approach's problem by making it more difficult for the SV to game the test, the other two problems mentioned above remain valid.

To address the aforementioned problems of both the traditional scenario-based testing methods and the adaptive adversarial testing approaches, the proposed methodology is inspired by investigating the following two questions.

- *Q1: If we can predict the SV motion trajectory with sufficient accuracy, can we design the POV motion to reach a safety-critical scenario?*

In general, if the infinite-time future planning of the SV is perfectly known and the POV is sufficiently capable in terms of its admissible space for control, the POV is guaranteed to at least asymptotically capture the SV, i.e. create conditions that lead to inevitable collisions between the POV and the SV. This is achieved by taking the known SV motion trajectory as a motion tracking reference for the POV, which naturally leads to a minimization problem closely related to the Model Predictive Control (MPC) formulation [9].

- *Q2: If the SV motion trajectory is unknown or difficult to predict, can we design the POV motion to arrive at a safety-critical scenario?*

The reachability analysis from differential games [10] and optimal control has answered the above question analytically. In general, the capture of the SV cannot be guaranteed if the SV's policy is unknown. However, suppose the SV and POV are sufficiently close such that the relative state belongs to the maximal backward reachable set (see equation (6) in [11]). In that case, there exists a feedback control policy for the POV such that for all possible motion trajectories of the SV, the policy renders a motion trajectory that guarantees the capturing of the SV. By assigning such a worst-case policy to the POV, one can ensure the generation of a safety-critical scenario.

In practice, one should also factorize the admissible action constraints, the operable domain allowed by traffic rules, and other concerns to propagate a "reasonable" dangerous scenario. Combining the solutions for the above questions, we propose an online feedback control policy for POVs that allows rendering adversarial paths without breaking traffic rules or the conventions of traffic accident fault or liability. The contribution of this work is threefold:

a) A Modeled Approach: The proposed framework is model-driven and induces theoretical guarantees for a certain safety-critical level of the derived scenario under standard assumptions.

b) Adaptive Scenarios: The derived motion trajectory of each POV varies adaptively based on the SV response throughout the propagation of a testing scenario. This creates more challenging testing instances with the same amount of testing effort as many of the current ADS testing scenario designs [5], [4] where the POV actions are independent

of the SV behavior. The proposed adaptive framework acts on the POV control level, which also creates agile motion trajectories that make it more difficult for the SV to game the scenario.

c) Online Execution: The proposed method formulates an anchor-template hierarchy control structure [12], with the simplified template planning naturally leading to a series of quadratic programming problems with computationally tractable solutions. This is, to the best of our knowledge, the first introduction of an online path planning algorithm for POVs which constructs adversarial testing scenarios.

Furthermore, we extensively evaluate the proposed framework's performance in a multi-lane straight-segment environment with parameterized SV driving policies and human drivers of various levels of aggressiveness. We observe that the adaptive framework enforces the SV into various safety-critical situations. One particularly interesting observation is that the SV typically remains collision-free if (i) the driving policy is conservatively safe, or (ii) the SV is willing to take extreme evasive maneuvers that are beyond the expectation of the modeled formulation (see Fig. 1d).

The remainder of this paper is organized as follows. Section II introduces the background setting of this work, including the problem formulation. Section III presents details on the proposed method. Section IV explains the experimental settings and the results analysis and finally, section V presents conclusions and future work.

II. PRELIMINARIES

We present the problem formulation of operational vehicle safety evaluation through the scenario-based test. We also revisit the basics of the anchor-template hierarchy control.

A. Problem formulation

Consider a heterogeneous multi-vehicle system of one subject vehicle (SV) and k principal other vehicles (POVs). Each vehicle's motion is subject to a certain ordinary differential equation in general. Without loss of generality, for the i -th individual agent ($i \in \{0, \dots, k\}$),

$$\dot{\mathbf{s}}_i = \mathbf{f}_i(\mathbf{s}_i, \mathbf{u}_i, t), \mathbf{s}_i \in \mathcal{S}_i \subset \mathbb{R}^n, \mathbf{u}_i \in \mathcal{U}_i \subset \mathbb{R}^m. \quad (1)$$

For the remainder of this section, the subscript i is omitted for general discussions of vehicle motion. We consider that the map $f : S \times U \rightarrow S$ is uniformly continuous, bounded and Lipschitz continuous in \mathbf{s} for fixed control \mathbf{u} . Hence, given a measurable control action \mathbf{u} , there exists a unique trajectory solving (1) for a given $\mathbf{u} \in \mathcal{U}$.

Let a *snapshot* $\sigma \in \mathcal{S}^{k+1}$ be the combined states of all vehicles at a time instance [13]. For example, the snapshot or state space of a test that involves two vehicles on straight segments would consist of the initial velocity of both vehicles and the relative position between them.

We then have the following definition of a *testing scenario* for $k + 1$ participants:

Definition 1. A scenario is a automaton denoted by a tuple of $\langle \Sigma, \Pi, f, \sigma_0, \Lambda \rangle$, where

- $\Sigma = \mathcal{S}^{k+1}$ is the set of snapshots, with each snapshot containing one SV and k POVs as specified above.
- Π is the "alphabet" of the automaton, in particular, the control action space of all vehicular participants in the scenario, i.e., $\Pi = \mathcal{U}^{k+1}$.
- f is the motion transition function as defined in (1).
- σ_0 specifies the initial condition, i.e. the arbitrary given snapshot.
- Λ is the subset of Σ deemed acceptable.

We define the acceptable set of snapshots Λ , containing all the snapshots that comply with the safety requirement, as

$$\Lambda = \{\sigma \in \Sigma \mid \|\mathbf{p}_0 - \mathbf{p}_j\|_2 \geq c, \forall j \in \{1, \dots, k\}\}, \quad (2)$$

where $\mathbf{p} \in \mathbb{R}^2$ denotes the position states and c is referred to as the *capture diameter*. Note that only the safety of the SV is of interest here. Correspondingly, the unsafe set of states is

$$\Omega = \Sigma \setminus \Lambda. \quad (3)$$

Given a string of control inputs, we designate the *run* of a scenario be a chronological sequence of snapshots $\mathcal{R} = \{\sigma_i\}_{i=0, \dots, T}$ from initial to final time T .

Most of the elements that define a scenario are given, with only the initial condition σ_0 and the automaton alphabet Π subject to various possibilities of configurations. Given that SV control is one of the subjects of the test, only the POV control actions are of interest for scenario designs. Intuitively, one seeks to derive a feedback control policy

$$\pi : \Sigma \times \mathcal{U}^k \rightarrow \Sigma \quad (4)$$

for the group of POVs such that the SV is forced into a sufficiently dangerous situation. Although the POV control is specified through a general class of functions that could be state-dependent and time-variant, it is worth emphasizing that the typical implementation of the POV control in practice remains fairly simple. POV action trajectories are mostly independent of state propagation, as shown in Fig. 1a.

In general, the multi-vehicle system is subject to highly nonlinear motion dynamics and complex state-action constraints. From the model point of view, it remains a challenge to derive agile motion planning algorithms for multiple POVs to cooperate in forcing the SV into dangerous situations. In this paper, we consider an anchor-template framework, as introduced in the following section.

B. Anchor-template framework for vehicle motion planning

When a full system model or anchor model is too complex to permit the design an appropriate controller, it is possible to use a simplification or template model that still captures the anchor model's essential characteristics and properties. Such a framework has been widely adopted in bipedal locomotion [14] and vehicle control [15].

A typical template model captures the motion features of the anchor model with lower dimensional state space and simplified dynamics. In this paper, the complex anchor model state is replaced with the simpler template state of discrete-time locally linearized motion equation formulated as

$$\mathbf{s}(t + \Delta) = \mathbf{A}\mathbf{s}(t) + \mathbf{B}\mathbf{u}(t), \quad (5)$$

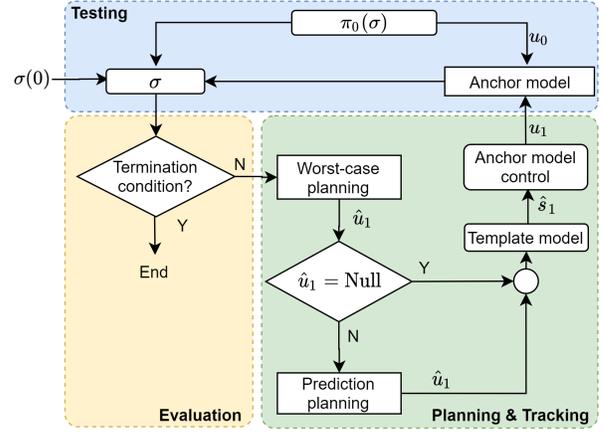


Fig. 2: Proposed method flowchart under the pairwise interactive setting with one POV and one SV.

with time-step Δ and system matrices

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & \Delta & 0 \\ 0 & 1 & 0 & \tilde{v}\Delta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \Delta & 0 \\ 0 & \frac{\Delta}{\tilde{v}} \end{bmatrix}. \quad (6)$$

The state vector $\mathbf{s} \in \mathbb{R}^4$ for each vehicle consists of the Cartesian coordinates of its position $\mathbf{p} = [x, y]$, speed v and heading angle ϕ . For the linearization, a constant speed \tilde{v} is also used. The control action includes longitudinal and lateral accelerations $\mathbf{u} = [a_x, a_y]^T$.

The model above is adapted from [16] and [13] assuming small local course angles (i.e. the angular difference between the combined velocity vector and the vehicle's running direction) and small changes in speed for local linearization. Theoretically, the control input \mathbf{u} is subject to the Kamm's circle [17] induced bounds. The state constraints are typically determined by speed limits, road topology, and other concerns regarding the admissible operable domain of vehicles. In this paper, we approximate the state-action constraints with a linear architecture in the form of $\mathbf{G}_u \mathbf{u} \leq \mathbf{h}_u$ and $\mathbf{G}_s \mathbf{s} \leq \mathbf{h}_s$.

III. METHOD

Fig. 2 presents a general overview of the proposed framework. The adversarial scenario propagation starts with an initial snapshot of the system σ_0 . The test ends if a certain termination condition is satisfied (e.g., encountering a collision); otherwise, one executes the template-based planning stage. For adversarial motion planning, one first seeks to create safety-critical scenarios assuming the SV policy is unknown. If the POV-to-SV distances are not sufficiently small, such a worst-case planning algorithm may not be applicable. One then moves on to the predictive planning stage, where the POV's adversarial policy is derived based on a certain assumed predicted behavior of the SV. The planned reference trajectory is then tracked by a lower-level controller such as MPC.

For the remainder of this section, we introduce the worst-case planning and predictive planning in a pairwise manner

with one POV and one SV. Without loss of generality, let the subscript 1 denote the adversarial POV and 0 the SV. Finally, we will extend the idea to multiple cooperative POVs.

A. Worst-case planning

Intuitively, the worst-case planning stage seeks to derive the POV adversarial trajectory without knowing the SV intentions. By the definition of unsafe snapshots Ω as specified in (3), starting from a safe initial condition $\sigma_0 \in \Sigma$ at time 0, if the pair of control strategies $(\pi_0(\sigma), \pi_1(\sigma))$ renders the trajectory converging to Ω at some finite time, we denote such a capture time as $T(\pi_0(\sigma), \pi_1(\sigma), \sigma_0)$. The pairwise SV and POV form a zero-sum game, wherein the SV seeks to maximize $T(\cdot)$ and the POV seeks to minimize $T(\cdot)$. Correspondingly, we have the minimax optimal feedback strategy $\pi_0^*(\sigma)$ and $\pi_1^*(\sigma)$ satisfying the saddle condition as

$$T(\pi_0(\sigma), \pi_1^*(\sigma), \sigma_0) \leq T^* \leq T(\pi_0^*(\sigma), \pi_1(\sigma), \sigma_0), \quad (7)$$

with $T^* = T(\pi_0^*(\sigma), \pi_1^*(\sigma), \sigma_0)$ referred to as the *minimal capture time*. Theoretically, T^* exists if and only if the initialization condition σ_0 is inside the maximal backward reachable set [11], i.e., there exists a POV control strategy such that regardless of the SV responses, one can ensure the capturing of the SV. This is guaranteed when the POV has moved to the reachable area of the SV, in this case, the inevitable collision zone. In practice, for a general nonlinear system, T^* can be obtained through a discrete approximation of the Hamilton-Jacobian-Bellman partially differential equation (HJB-PDE) [18]. In this paper, we seek to derive the T^* , if applicable, within a local look-ahead time horizon up to \bar{T} seconds. With the simplified template dynamics, one can approximate T^* by iteratively solving a series of minimax quadratic programming problems as discussed in [13]. One can also derive the function of mapping a snapshot to the corresponding T^* offline. Given the minimal capture time T^* , we can then formulate the optimal feedback policy for the POV as

$$\hat{\mathbf{u}}_1^* = \underset{\hat{\mathbf{u}}_1}{\operatorname{argmin}} \underset{\hat{\mathbf{u}}_0}{\operatorname{max}} \|\mathbf{p}_0(T^*) - \mathbf{p}_1(T^*)\|^2 \quad (8a)$$

$$\text{s.t. } \mathbf{s}_i(t + \Delta) = \mathbf{A}\mathbf{s}_i(t) + \mathbf{B}\mathbf{u}_i(t), \forall i \in \{0, 1\}, \quad (8b)$$

$$\mathbf{G}_s \mathbf{s}_i(t) \leq \mathbf{h}_s, \forall t \in \{0, \dots, T^* - \Delta\}, \quad (8c)$$

$$\mathbf{G}_u \mathbf{u}_i(t) \leq \mathbf{h}_u, \forall t \in \{0, \dots, T^* - \Delta\}, \quad (8d)$$

$$\sigma(0) = \sigma_0. \quad (8e)$$

In practice, the above optimization is solved at each planning stage presenting the current snapshot as the initialization condition. The instantaneous reference motion trajectory $\hat{\mathbf{s}}_1$ is then obtained by propagating the template-model motion for T^* seconds with the sequence of optimal actions $\hat{\mathbf{u}}_1^*$.

Note that the capturing guarantee at the template-model level is subject to some assumed state constraints (8c) and action constraints (8d) of the SV. That is, the capturing will undoubtedly fail if the SV is willing to perform extreme evasive maneuvers that are beyond its assumed capability. While such an outcome may be deemed aggressive in terms of passenger comfort and vehicle dynamic stability, it is

still technically a safe choice of action from the operational safety perspective. Furthermore, consider that the definition of safety is induced by the l_2 -norm distance between vehicles as indicated by (2), which is not necessarily equivalent to vehicle-to-vehicle collisions in real-world driving. This leads to a trade-off phenomenon when determining the capture diameter c . A larger choice of c results in a larger maximal BRS and hence increases the likelihood of deriving a minimal capture time T^* . However, this also makes it difficult for a real vehicle-to-vehicle collision to occur, leading to snapshots that are not sufficiently safety-critical. A more detailed comparison regarding this trade-off will be presented in Section IV between Fig. 5 and Fig. 6.

B. Predictive planning

If the worst-case planning is not applicable (i.e., one cannot find a qualified minimal capture time $T^* \leq \bar{T}$), this implies the current snapshot is not severe enough to enable a guaranteed capturing. We propose to consider an alternative based on the assumed predictive motion of the SV.

Historically, vehicle motion prediction has been studied extensively [19], [20]. The various self-driving algorithms that are model-based [21], [22], [23] and learning-inspired [24], [25] are also applicable to serve as the SV motion predictor. In this paper, we predict the SV motion based on the steady-state assumption. The traditional steady-state assumption assumes fixed velocity and heading (i.e. control $\mathbf{u} = \mathbf{0}$). It is also the fundamental assumption for various vehicle safety related methodologies including Time-to-Collision [26] for safety analysis and human driver behavior characterization such as gap acceptance [27] and lead-vehicle following distance [28]. In most of the examples presented in the Section IV, we adopt a modified steady-state assumption which assumes that the SV is maintaining the current control action up to the assigned time horizon. Let such a predictive policy for the SV be $\hat{\pi}_0$, one can then take the predicted motion of the SV as the tracking reference for the POV. With the running cost and termination cost determined by $\mathbf{Q}_r \in \mathbb{R}^{4 \times 4}$ and $\mathbf{Q}_f \in \mathbb{R}^{4 \times 4}$ respectively, and considering $\tilde{\mathbf{s}}(t) = \mathbf{s}_1(t) - \mathbf{s}_0(t)$, we have the following optimization problem for the predictive planning:

$$\hat{\mathbf{u}}_1^* = \underset{\hat{\mathbf{u}}_1}{\operatorname{argmin}} \sum_{t=0}^{\bar{T}-\Delta} (\tilde{\mathbf{s}}(t)^T \mathbf{Q}_r \tilde{\mathbf{s}}(t)) + \tilde{\mathbf{s}}(\bar{T})^T \mathbf{Q}_f \tilde{\mathbf{s}}(\bar{T}) \quad (9a)$$

$$\text{s.t. } \mathbf{s}_1(t + \Delta) = \mathbf{A}\mathbf{s}_1(t) + \mathbf{B}\mathbf{u}_1(t), \quad (9b)$$

$$\mathbf{s}_0(t + \Delta) = \mathbf{A}\mathbf{s}_0(t) + \mathbf{B}\hat{\pi}_0(\mathbf{s}_0(t)), \quad (9c)$$

$$\mathbf{G}_s \mathbf{s}_i(t) \leq \mathbf{h}_s, \forall t \in \{0, \dots, T^* - \Delta\}, \quad (9d)$$

$$\mathbf{G}_u \mathbf{u}_i(t) \leq \mathbf{h}_u, \forall t \in \{0, \dots, T^* - \Delta\}, \quad (9e)$$

$$\sigma(0) = \sigma_0. \quad (9f)$$

The reference trajectory for the POV is generated by a similar procedure as that shown in the worst-case planning section.

Finally, an MPC is used as the navigation controller that allows the POV to follow the reference path obtained from the planning stage. It is worth emphasizing that although both MPC for vehicle motion tracking and the template-model based planning rely on a certain linearized motion

equation, the two formulations are not necessarily the same. In our case, the template-model adopts the formulation in (6) with the control action specified as accelerations, which are easier to identify through perceivable states with quantifiable constraints. On the other hand, the MPC adopts the formulation from [15] with the control action defined as acceleration a and steering wheel angle ω , which is easier to implement for direct vehicle control tasks. The various state-action constraints deployed in the planning stage as specified in (8a) and (9a) are also included in the MPC formulation. Therefore, even if mild discrepancies may occur between the template-model used for planning and the anchor-model adopted for motion tracking, the constraints remain valid.

C. Multi-POVs

With the adversarial pairwise interactive scenario presented above, we are now ready to extend the framework to work with multiple POVs. In practice, it is overly aggressive to assume that all POVs present in the snapshot will cooperatively attack the SV. The absolute cooperative POV planning also poses an extra challenge to the scenario design given none of the POVs are supposed to crash into each other during the scenario propagation. On the other hand, some existing work [13] considers a partially non-cooperative assumption allowing at most one adversarial POV with the rest of the POVs complying with the SV for collision avoidance. Although this assumption simplifies the analysis significantly and enables a pairwise study between the SV and each POV, the partial non-cooperativeness remains conservative.

In this work, we propose a multi-POV scenario generation scheme which allows multiple POVs propagating cooperative safety-critical scenarios in a controlled manner. This is done by assigning each POV a dedicated set of state-action constraints which ensures the non-interactive motion between POVs. Section IV will present more detailed examples for multi-POV adversarial scenarios propagated with the proposed method.

IV. EXPERIMENTAL RESULTS

We illustrate the performance of the proposed approach with two types of simulations. We first demonstrate the adversarial testing scheme against a series of parameterized self-driving policies in a customized simulation environment with the nonholonomic bicycle model of vehicle dynamics. We then expand the test to human-driven SVs performed in the CARLA simulator [29]. Throughout all simulations, we have $\bar{T} = 2$ for the planning and MPC horizon limit. We also assume an identical admissible action space for all POVs and the SV with $a_x^{\max} = 0.67$, $a_x^{\min} = -1.7$, $a_y^{\max} = 1$, $a_y^{\min} = -1$. Considering the straight-segment environment we are studying, the cost matrices in (9a) are defined to emphasize the lateral tracking accuracy satisfying $\mathbf{Q}_r = \mathbf{Q}_f = \text{diag}(1, 100, 0.1, 0.1)$.

A. Testing parameterized self-driving policies

The parameterized self-driving policy is a combination of the Intelligent Driving Model (IDM) and a set of customized lane-change heuristics. The IDM formulation is adapted

from [30]. The lane change heuristics consist of two stages, decision making and lane-change execution. The lane-change decision is determined by the current SV velocity and the lead-lag gap acceptance with parameters adopted from the naturalistic behavioral study of [27]. The lane-change execution is divided into two parts. The longitudinal velocity is adapted from the IDM model. The lateral motion is controlled by a parameterized PD controller subject to yaw-rate constraints. We first present a set of pairwise interactive scenarios with one POV and one SV. We then extend to the multi-POV settings.

1) *Pairwise Interactions*: Throughout the pairwise interactive scenarios presented in this section, we consider the same initialization condition σ_0 inspired by the standard Crash Imminent Brake (CIB) system test as shown in Fig. 3. Both SV and POV are confined to the three-lane operable domain, with extra state constraints for restricting the POV to initiate a rear-end collision against the SV within the same lane. The admissible velocity range is confined to $v \in [5, 45](m/s)$. Vehicles are assumed identical with a length of $5m$ and a width of $2m$. Lane width is set to $3.7m$. The adversarial planning algorithm and all vehicle motion controllers are executed at 10 Hz. A scenario terminates at 50 seconds after the initialization, or earlier if a vehicle-to-vehicle collision is detected. Unless stated otherwise, the mentioned configurations remain the same for all other examples. Simulation results with various self-driving policies of different hyper-parameters are shown in Fig. 4, Fig. 5, and Fig. 6. Here it is worth emphasizing the performance comparison between Fig. 5 and Fig. 6. The SV policy, initialization conditions and state-action constraints are all set identically in both tests with the only difference lying in the choice of the capture diameter c . In Fig. 5 we use a smaller value of $c = 7$. The POV switches from the predictive planning mode to worst-case planning mode about 10 seconds after the test initialization and the collision occurs 2.1 seconds later. On the other hand, in Fig. 6 we take a larger value of $c = 12$, which makes it easier to initiate the worst-case planning mode (2 seconds earlier than the case in Fig. 5). However, the relatively large capture space makes the scenario run less severe than the example in Fig. 5 and fails to force a vehicle-to-vehicle collision. Furthermore, note that the SV performs an abrupt acceleration of $2.6m/s^2$ at $t = 14.2s$ during a single lane-change stage to its left. This behavior is significantly outside the expectation as the worst-case planning algorithm assumes the maximum longitudinal

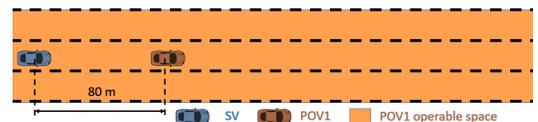


Fig. 3: The standard Crash Imminent Brake (CIB) system testing procedure [5]: the lead-POV brakes and expects the SV to also brake for collision avoidance: both vehicles start within the same lane with an initial velocity of $18m/s$.

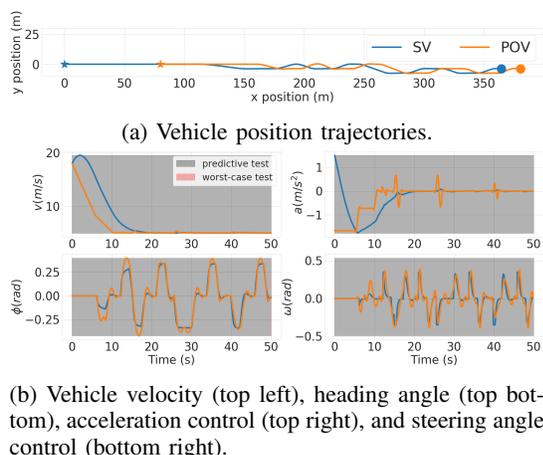


Fig. 4: An adversarial POV persistently blocks an SV follower with a conservative parameterization.

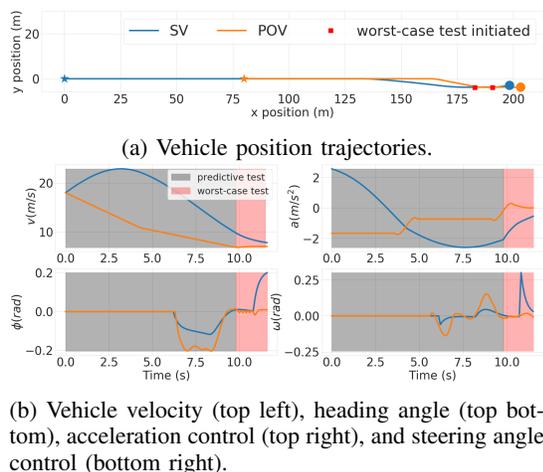


Fig. 5: An adversarial POV crashes the SV follower with the capture diameter set to $c = 7$.

acceleration of the SV is only $0.67m/s^2$. As we have mentioned in Section III, although such an extreme evasive maneuver is deemed aggressive by common sense, it leads the SV to a lower-risk driving status in terms of collision avoidance. Hence, the proposed online adversarial framework succeeds in creating a sufficiently dangerous scenario such that the proper response of the SV can be tested.

2) *Multi-POV scenario*: We further present a multi-POV scenario as shown in Fig. 7. Each POV is assigned a unique operable space that does not overlap with other POVs. The capture diameter is set to $c = 7$. Admissible velocity range is modified to $v \in [12, 45](m/s)$. The assumed admissible action space for the SV and all POVs are identical during the adversarial trajectory planning stage, but in practice, the SV is often more capable than a POV both longitudinally and laterally. This is particularly obvious in Fig. 7 where the worst-case planning transits back to the predictive planning stage 7 times for POV1 before the SV collides with POV2 from the rear side. Some extended examples with the cooperative multi-POV configuration are included in [31].

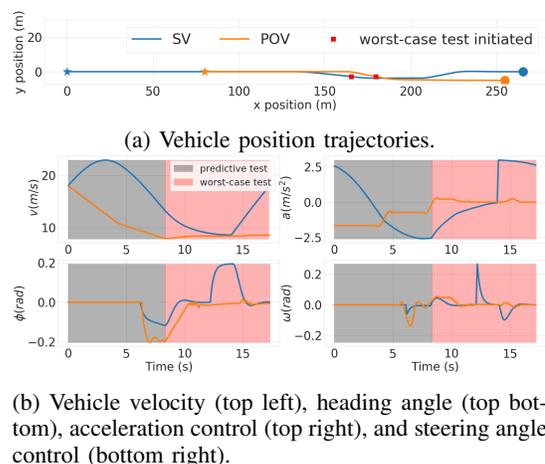


Fig. 6: An adversarial POV fails to force the SV follower into a crash with the capture diameter set to $c = 12$.

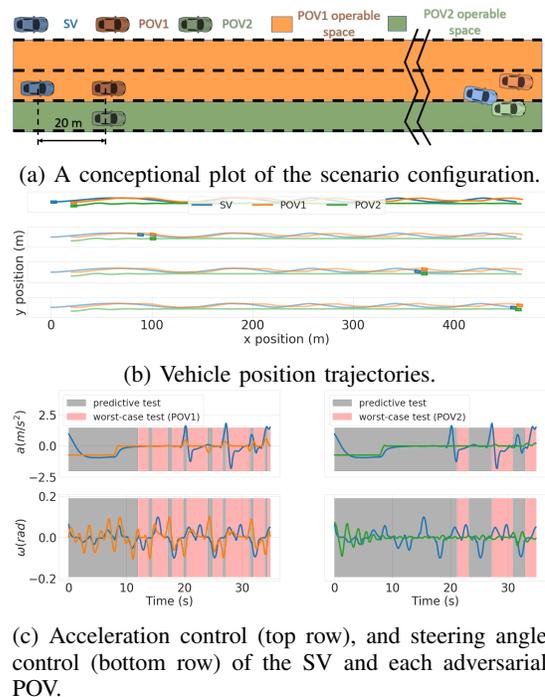


Fig. 7: Two cooperative lead-POVs interact with the SV follower: SV ends up with a rear-end collision against POV2 due to the aggressive choice of lead gap acceptance during the single lane-change stage.

B. Testing human drivers

The proposed method is also implemented in the CARLA simulator. We replicate several of the initial test runs in the previous subsection with similar parameters. Fig. 8 shows a top view in the simulator from the initial condition shown in Fig. 3, where the trajectory followed by both the SV (blue) and the POV (red) is shown on-screen. The waypoints from both vehicles for a time horizon of 2 seconds are also drawn. Additionally, this implementation allows a human driver to take over the SV control while the POV is dynamically reacting to the SV's actions. It is shown that the predictive

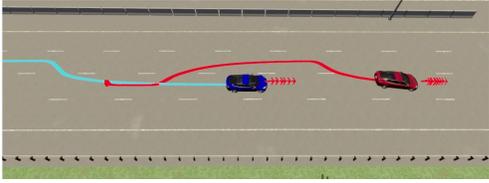


Fig. 8: Top view of a pairwise interactive scenario in CARLA, the POV is red and the SV is blue.

planning has allowed the POV to react appropriately to provoke near-collision situations even when the SV's policy is completely unexpected (human driver).

V. CONCLUSIONS

This paper presents an online adversarial framework for the scenario-based testing of operational vehicle safety. The proposed method generates sufficiently dangerous testing scenarios in an efficient and controlled manner. It is applicable for the safety evaluation of human-driven, as well as ADS and ADAS equipped, vehicles. Various simulated examples are also presented to show the empirical effectiveness. It is of future interest to extend the method to different environment configurations such as intersections, roundabout, and parking-lot. We will also improve the methodology with real-world experiments in proving ground tests.

ACKNOWLEDGMENT

This work was funded by the United States Department of Transportation under award number 69A3551747111 for Mobility21: the National University Transportation Center for Improving Mobility. Any findings, conclusions, or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the United States Department of Transportation, Carnegie Mellon University, or The Ohio State University.

REFERENCES

- [1] A. Paul, R. Chauhan, R. Srivastava, and M. Baruah, "Advanced driver assistance systems," SAE Technical Paper, Tech. Rep., 2016.
- [2] E. Thorn, S. C. Kimmel, M. Chaka, B. A. Hamilton *et al.*, "A framework for automated driving system testable cases and scenarios," United States. Department of Transportation. National Highway Traffic Safety, Tech. Rep., 2018.
- [3] S. Feng, Y. Feng, C. Yu, Y. Zhang, and H. X. Liu, "Testing scenario library generation for connected and automated vehicles, part i: Methodology," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [4] "Traffic jam assist (TJA) system confirmation test," DOT HS 812 759, Tech. Rep., 2019.
- [5] C. A. M. Partnership, "Crash imminent braking (CIB) first annual report," DOT HS 811 340, Tech. Rep., 2010.
- [6] M. Klischat and M. Althoff, "Generating critical test scenarios for automated vehicles with evolutionary algorithms," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 2352–2358.
- [7] B. Chen and L. Li, "Adversarial evaluation of autonomous vehicles in lane-change scenarios," *arXiv preprint arXiv:2004.06531*, 2020.
- [8] W. Ding, M. Xu, and D. Zhao, "Learning to collide: An adaptive safety-critical scenarios generating method," *arXiv preprint arXiv:2003.01197*, 2020.
- [9] E. F. Camacho and C. B. Alba, *Model predictive control*. Springer Science & Business Media, 2013.

- [10] Y. Ho, A. Bryson, and S. Baron, "Differential games and optimal pursuit-evasion strategies," *IEEE Transactions on Automatic Control*, vol. 10, no. 4, pp. 385–389, 1965.
- [11] I. M. Mitchell, "Comparing forward and backward reachability as tools for safety analysis," in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2007, pp. 428–443.
- [12] P. Wensing and S. Revzen, "Template models for control," *Bioinspired Legged Locomotion*. Elsevier, pp. 240–266, 2017.
- [13] B. Weng, S. J. Rao, E. Deosthale, S. Schnelle, and F. Barickman, "Model predictive instantaneous safety metric for evaluation of automated driving systems," in *2020 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2020, pp. 1899–1906.
- [14] P.-B. Wieber, "Trajectory free linear model predictive control for stable walking in the presence of strong perturbations," in *2006 6th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2006, pp. 137–142.
- [15] R. Rajamani, *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- [16] P. Junietz, F. Bonakdar, B. Klamann, and H. Winner, "Criticality metric for the safety validation of automated driving using model predictive trajectory optimization," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 60–65.
- [17] S. M. Erlien, "Shared vehicle control using safe driving envelopes for obstacle avoidance and stability," Ph.D. dissertation, Stanford University, 2015.
- [18] O. Bokanowski, N. Forcadel, and H. Zidani, "Reachability and minimal times for state constrained nonlinear problems without any controllability assumption," *SIAM Journal on Control and Optimization*, vol. 48, no. 7, pp. 4292–4316, 2010.
- [19] A. Houenou, P. Bonnifait, V. Cherfaoui, and W. Yao, "Vehicle trajectory prediction based on motion model and maneuver recognition," in *2013 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2013, pp. 4363–4369.
- [20] M. Koschi and M. Althoff, "SPOT: A tool for set-based prediction of traffic participants," in *2017 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2017, pp. 1686–1693.
- [21] A. Petrovskaya and S. Thrun, "Model based vehicle tracking for autonomous driving in urban environments," *Proceedings of robotics: science and systems IV, Zurich, Switzerland*, vol. 34, 2008.
- [22] L. Capito, K. Redmill, and U. Ozguner, "Optical flow based visual potential field for autonomous driving," *Proceedings of the 31st IEEE Intelligent Vehicles Symposium (IV'20)*, June 2020.
- [23] J. Park, A. Kurt, and U. Ozguner, "A game theoretic approach to controller design for cyber-physical systems: collision avoidance," in *2013 ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)*. IEEE Computer Society, 2013, pp. 254–254.
- [24] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [25] E. Yurtsever, L. Capito, K. Redmill, and U. Ozguner, "Integrating deep reinforcement learning with model-based path planners for automated driving," *Proceedings of the 31st IEEE Intelligent Vehicles Symposium (IV'20)*, June 2020.
- [26] D. N. Lee, "A theory of visual control of braking based on information about time-to-collision," *Perception*, vol. 5, no. 4, pp. 437–459, 1976.
- [27] M. Yang, X. Wang, and M. Qudus, "Examining lane change gap acceptance, duration and impact using naturalistic driving data," *Transportation research part C: emerging technologies*, vol. 104, pp. 317–331, 2019.
- [28] D. Swaroop and K. Rajagopal, "A review of constant time headway policy for automatic vehicle following," in *ITSC 2001. 2001 IEEE Intelligent Transportation Systems. Proceedings (Cat. No. 01TH8585)*. IEEE, 2001, pp. 65–69.
- [29] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," *arXiv preprint arXiv:1711.03938*, 2017.
- [30] M. Liebner, M. Baumann, F. Klanner, and C. Stiller, "Driver intent inference at urban intersections using the intelligent driver model," in *2012 IEEE Intelligent Vehicles Symposium*. IEEE, 2012, pp. 1162–1167.
- [31] L. Capito, B. Weng, U. Ozguner, and K. Redmill, "A modeled approach for online adversarial test of operational vehicle safety," *arXiv preprint arXiv:2009.12222*, 2020.