

Rapidly-exploring Random Trees for Testing Automated Driving Systems

Cumhur Erkan Tuncali and Georgios Fainekos

Abstract—Autonomous vehicles are in an intensive research and development stage, and the organizations developing these systems are targeting to deploy them on public roads in a very near future. One of the expectations from fully-automated vehicles is never to cause an accident. However, an automated vehicle may not be able to avoid all collisions, *e.g.*, the collisions caused by other road occupants. Hence, it is important for the system designers to understand the boundary case scenarios where an autonomous vehicle can no longer avoid a collision. In this paper, an automated test generation approach that utilizes Rapidly-exploring Random Trees is presented. A comparison of the proposed approach with an optimization-guided falsification approach from the literature is provided. Furthermore, a cost function that guides the test generation toward almost-avoidable collisions or near-misses is proposed.

I. INTRODUCTION

Autonomous vehicles are safety-critical systems, and they should be tested thoroughly before they are deployed on public roads. Although the testing in real traffic environments is not fully replaceable, simulation-based testing provides many advantages such as fully controllable environments, ground-truth information, ability to try a massive number of scenarios, and creating risky scenarios without risking human life or the vehicle under test.

Optimization-guided falsification techniques utilize optimization engines to generate challenging scenarios for an Autonomous Vehicle (AV) under test with the ultimate goal of finding a scenario in which the Vehicle Under Test (VUT), also called as *Ego vehicle*, fails to satisfy its requirements [1], [2], [3], [4]. In our previous work [2], we have used an optimization-guided falsification tool, S-TALIRO [1], to search for maneuvers of agent actors (other road occupants) with the aim of identifying the boundary between the safe and unsafe operation of VUT.

In this paper, we propose an alternative approach to the falsification-based approach that is described in [2]. We consider the test generation problem as a path planning problem for agent actors with the aim of creating interesting collisions with vehicles under test by utilizing Rapidly-exploring Random Trees (RRTs). Because of the nature of the driving environment, an automated vehicle cannot avoid all collisions. For instance, a collision with a vehicle that loses control and drives into the automated vehicle may not be avoidable. Figure 1 illustrates an unavoidable collision example. However, identifying the boundaries between the collisions that are barely avoided or the collisions that could have been avoided with minor changes in the control or environment would be valuable for the engineering teams to improve the safety-related capabilities of the system. We focus on finding test cases that lead to behaviors at the proximity of the boundary between collisions and near-collisions. For this purpose, we propose a cost function that can guide the test generation toward that boundary.

This work has been partially supported by award NSF CNS 1350420
C.E. Tuncali, and G. Fainekos are with School of Computing, Informatics and Decision Systems, Arizona State University, Tempe, AZ 85281, USA
(e-mail:{etuncali, fainekos}@asu.edu)




Fig. 1. An unavoidable collision example. Red car suddenly moves into the Ego vehicle's lane at a very short distance.

RRTs have been first developed for robot path/motion planning problems [5], [6], [7]. However, thanks to their ability to efficiently search over high-dimensional spaces, RRT-based approaches also deliver promising results in the test generation domain [8], [9], [10], [11], [12], [13].

Since their first introduction, many variants of RRTs have been proposed. In [6], a method called Transition-based RRT (T-RRT) was introduced. Transition-based RRT method extends the classical RRT by incorporating additional cost criteria to the explored paths rather than only aiming to reach a target configuration. T-RRT borrows the notion of transitions tests from stochastic optimization approaches. Hence, it can be considered as a method that is merging RRTs with stochastic optimization. Furthermore, T-RRT controls exploration versus refinement using a method called *minimal expansion control* which helps to promote the expansion of a tree to the unexplored areas of the search space. Efficiency of T-RRT on continuous cost spaces is studied in [6]. An optimal RRT approach, RRT*, was proposed, and the optimality was analyzed by Karaman and Frazzoli in [7].

II. OVERVIEW OF THE APPROACH

In our approach, we utilize notions from RRT* [7] and T-RRT [6] with a custom cost function that we propose to find boundary-case collisions. We implement our version of *minimal expansion control* using the notion of *sparseness* from evolutionary algorithms that perform novelty search [14], [15].

With this approach, which we will refer to as “RRT-based approach”, we address the following limitations of our falsification-based approach [2]:

1. In [2], we use a limited number of control points over the longitudinal position axis as the specific points where the lateral axis of the vehicle trajectories are sampled. As the number of control points decreases, possible variations in shapes of the trajectories are limited. On the other hand, increasing the number of control points also increases the dimension of the search space, which makes the problem more challenging.
2. In [2], the duration of the simulations is fixed. With the RRT-based approach, although there is an inherent limit on the maximum simulation duration that is dictated by the maximum number of RRT nodes, there is flexibility in the duration of the simulations. So, non-promising simulations are stopped earlier while more promising ones can be executed for longer times.
3. With the RRT-based approach, we can minimize the need for hand-designing a test scenario in detail and allow more freedom in the exploration of the space compared to the falsification-based approach.

4. RRT-based approach is promising to avoid local minima that can be challenging to the falsification-based approach proposed in [2].

The flowchart of our RRT-based approach is shown in Fig. 2. In the rest of this section, we will describe the key components of our approach.

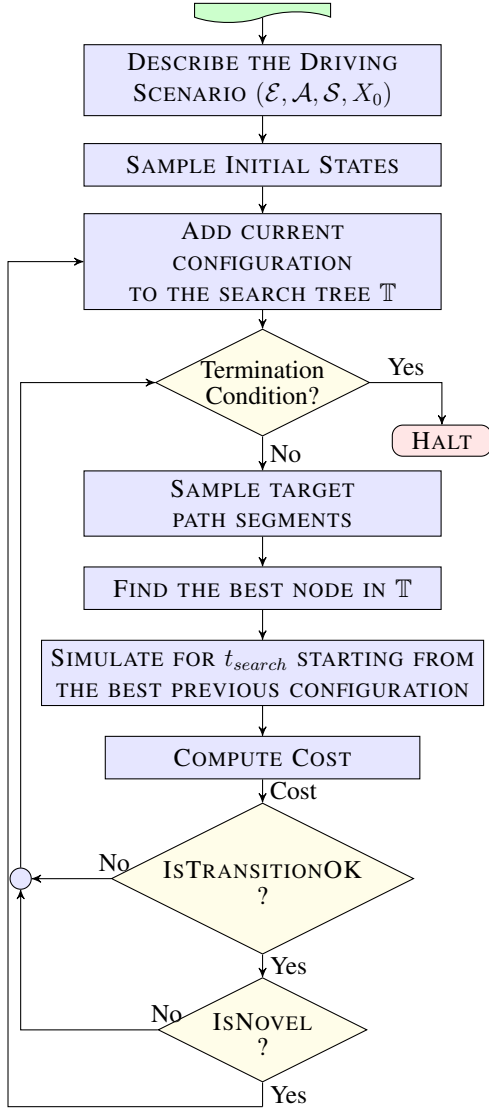


Fig. 2. Flowchart illustrating the RRT-based approach.

A. Initializing the Search

A scenario, \mathcal{T} , is described by the sets of Ego vehicles \mathcal{E} , agent actors \mathcal{A} and environment \mathcal{S} as $\mathcal{T} = \mathcal{E} \cup \mathcal{A} \cup \mathcal{S}$. The model of the driving scenario is given as $\mathcal{M}_{\mathcal{T}} = (X_{\mathcal{T}}, U_{\mathcal{T}}, P_{\mathcal{T}}, sim_{\mathcal{T}})$ where $X_{\mathcal{T}}$ is the set of system states, $U_{\mathcal{T}}$ is the set of system inputs, $P_{\mathcal{T}}$ is the set of system parameters, and $sim_{\mathcal{T}}$ is the system simulation function, respectively, for the overall scenario. Given $x \in \mathcal{X}$, $\hat{x} = sim(x, u, p, t)$ is the state reached starting from the system state x after time $t \in T$ under the input vector $u \in \mathcal{U}$ and the parameter vector $p \in \mathcal{P}$. Hence, the system simulation function is a mapping $sim : \mathcal{X} \times \mathcal{U} \times \mathcal{P} \times T \rightarrow \mathcal{X}$ where T is a discrete set of time samples t_0, t_1, \dots, t_N , with $t_i < t_{i+1}$.

In this approach, the set of inputs, $U_{\mathcal{T}}$, can simply be a set of target paths for the agent actors, as well as inputs to the other entities in the simulation environment such as models

of road and weather conditions. After general outlines of the driving scenario are described, the sampling space for the initial states, i.e., $\mathcal{X}_{0,\mathcal{T}}$, is used to sample initial configurations of the simulation entities, including Ego vehicles and agent actors.

B. Information Stored on RRT Tree Nodes

A tree grows while seeking to discover interesting behaviors. While growing the tree, instead of executing simulation traces starting from the initial configuration, only a partial simulation is executed starting from an existing node in the tree. For that purpose, the state of the system, the state of the controllers, and the simulation time are stored on the tree nodes. Table I provides more details on the data stored on the tree nodes after each partial simulation. The history-related fields will be empty for the root nodes of trees.

C. Sampling a Target Path Segment

A sample target path segment is simply a set of waypoints which is used as an immediate target for the corresponding vehicle. A waypoint is denoted as $w = (x_w, y_w, \theta_w, v_w) \in \mathcal{P}_w$ where x_w, y_w, θ_w, v_w are the x -coordinate, y -coordinate, target driving direction and the target speed at the waypoint. The sampling space for the waypoints is defined by a corresponding parameter space $\mathcal{P}_w = \mathcal{P}_{w,x} \times \mathcal{P}_{w,y} \times \mathcal{P}_{w,\theta} \times \mathcal{P}_{w,v}$ that describe the limits on the $x - y$ coordinates, driving direction, and target speed where $\mathcal{P}_w \subseteq \mathcal{P}_{\mathcal{T}}$. An example waypoint sampled on a straight road is shown in Fig. 3. A coordinate transformation can be applied for sampling from curved roads. Although the example waypoint in Fig. 3 is sampled from a road, the sample space of the waypoint doesn't have to be the same as the area of a road in the simulation. It may be defined to go beyond the road limits, it may be limited to only a part of a road, or it may be completely irrelevant to a road in the simulation environment.




Fig. 3. Sampling a waypoint.

Once a waypoint is sampled, the next step in sampling a target path segment is to add an endpoint at a predefined distance d_{leg} from the waypoint, along the direction of the waypoint. Figure 4 shows a target path segment formed using this approach. The sampled target path segment for this example can be denoted by $p = ((x_w, y_w, \theta_w, v_w), (x_{w2}, y_{w2}, \theta_w, v_w))$.




Fig. 4. Sampling a target path segment.

If the endpoint of a target path segment is outside the sampling space of the waypoint, we simply break the segment

Data	Description
State	The state of the system at the end of the corresponding part of the simulation is stored. It serves as the initial state for a new partial simulation that is starting from the configuration represented in the current node.
State History	(optional) The state history over the corresponding part of the simulation is used to (i) reproduce agent behaviors after the search is over (ii) to simulate any sensor delays for the next simulation step.
Input History	(optional) The history of inputs to the system is used as the past input data for the next simulation step, which may be useful for the systems that need to remember past inputs. This data is also valuable for analysis purposes when the search is over.
Controller State	The final state of controllers, if available, are stored so that the next step of the simulation can be started from the current node without needing to run the whole simulation from the starting node to the current node.
Simulation Time	The time at the end of the corresponding part of the simulation is stored.

TABLE I
DATA STORED ON THE RRT NODES.

at the boundary of the sampling space and add a second leg along the boundary in the direction closest to the waypoint direction. Figure 5 shows an example target path section for a longer $d_{leg} = d_{leg1} + d_{leg2}$ compared to the one in Fig. 4. The sampled target path segment for this example can be denoted by $\mathbf{p} = ((x_w, y_w, \theta_w, v_w), (x_{w2}, y_{w2}, \theta_{w2}, v_w), (x_{w3}, y_{w3}, \theta_{w2}, v_w))$.




Fig. 5. Sampling a target path segment with space constraints.

After the shape of the target segment is decided, we also sample a target speed, v_w , for the target path segment.

This approach is applied to each agent vehicle for sampling their target path segments. Note that the framework we propose allows using a different algorithm for selecting a target path and/or any other type of input to the agent vehicles.

D. Selecting the Best Node from the Search Tree

Once target path segments for the agent vehicles are sampled, we pick one node from the existing tree, as the initial configuration for the simulation that will be executed with the sampled target path segments. There is no single correct approach to decide which node of the tree would be the best choice.

In our study, the notion of selecting the optimal node is adopted from the RRT* method [7]. In [7], when adding a new node to the tree, existing nodes in a neighborhood of the new node are all checked and the one which minimizes the cost is selected as the previous node. Our approach has similarities to the RRT* one. We compute the sum of distances from each vehicle to the starting point of their target path segment with the constraint that the configurations of all vehicles are behind the start position of the initial target waypoint with respect to the driving direction of that waypoint. Then, we execute partial simulations from the best n candidate previous nodes ($n = 5$ for our case studies) and pick the one which gives the minimum cost. We believe that this approach is promising to create relatively natural-looking vehicle trajectories while still allowing enough randomness in the maneuvers.

In [7], after adding the new node to the tree, there is a rewiring step which modifies the connections to the other existing nodes in the neighborhood of the new node. The rewiring step checks whether reaching to an existing node in that neighborhood from the newly added node would result in a reduction in the cost without violating any constraints such as obstacles. If so, it replaces the existing edge incoming to

that node with an edge originating from the newly added node. Application of the rewiring step is straightforward for path planning problems on Euclidean spaces where the tree nodes represent planned waypoints on a path instead of the simulated vehicle configurations. In our approach, the tree nodes are the resulting configurations reached by the simulated vehicles with a controller and an input target path. Hence, the rewiring step requires the execution of partial simulations starting from the newly added node to the other nodes in the neighborhood of the newly added node. Since the resulting configuration will most likely be different at the end of such a partial simulation, the configurations on the target nodes will change. This creates the necessity to execute simulations from the updated nodes to all of the remaining tree nodes that can be reached from the modified node. Hence, the rewiring step can be computationally costly in our approach, and so, we do not apply the rewiring step and leave it as a future work for which the applicability should be analyzed.

We would like to emphasize that the function used for selecting the best previous node is user-configurable in our framework and depending on how much randomness is plausible in the generated driving paths, a different algorithm, *e.g.*, simply selecting the closest node, can be utilized.

E. Simulating the System

After obtaining a set of target path segments for agent vehicles and deciding the initial configuration for the simulation, we create the simulation scene in the simulation environment using the data stored in the selected node of the search tree. That is, we set the initial states of the simulation entities and initialize the Ego vehicle controllers with the previous inputs and the saved controller states. We also pass the sampled target path segments to the agent vehicles as inputs. Finally, we simulate the system for t_{search} time and collect state and input histories at each time step of the simulation. For our setup, we use MATLAB simulations, however, this is not mandatory and another simulator can be used. Note that if the simulator and the Ego vehicle controllers allow saving the state and continuing simulation from a saved state, *which is the case in our setup*, the time spent in the simulations can be radically reduced because it would be enough to simulate only the new part of the simulation. Otherwise, the simulation should always start from the root node of the tree and run until the current target time.

F. Cost Function

After a simulation is executed, a cost function is used to compute how close the simulation trace is to an interesting behavior. The approach we describe here can be utilized to discover other types of interesting/failing behavior; however, our target in this work is to explore the behaviors that are on the boundary between safe and unsafe operation. Hence, an interesting behavior for our purposes would be (i) a collision

between an Ego vehicle and an agent that could have been avoided with a minor change in the control applied or agent trajectories, (ii) an almost-collision (near-collision) which could have ended with a collision with a minor change in the control applied or agent trajectories.

The properties of a good cost function that would guide the search toward an interesting behavior for our purposes can be listed as follows:

- Among two similar collisions between an Ego vehicle and an agent vehicle, the one which has the smaller magnitude in the relative speed between the vehicles should have a smaller cost, as a smaller change in the speed of the Ego vehicle would be enough to avoid the collision.
- Among two similar collisions, the one which has the smaller impact area, *i.e.*, the area of the collision surface, should have a smaller cost, as a smaller change in the steering maneuver of the Ego vehicle would be enough to avoid the collision.
- For vehicle paths without a collision, a smaller time-to-collision at any point of the path, and a smaller corresponding collision speed and a smaller area for that collision-to-be should lead to a smaller cost.

We propose the following cost function:

$$\mathcal{R}(y) = (1 + s_{coll,y})(v_{coll,y}^2 + tt_{min,y}^2) \quad (1)$$

where $s_{coll,y} \in [0, 1]$ is the ratio of the collision surface to the overall surface on the collision side of the vehicle, $v_{coll,y}$ is the relative speed of the vehicles at the moment of collision, and $tt_{min,y}$ is the minimum time-to-collision encountered during the simulation output trace y . For the simulations with a collision, $tt_{min,y}$ is 0. For the simulations without a collision, $s_{coll,y}$ and $v_{coll,y}$ are computed at the instance of smallest time-to-collision with the assumption that the vehicles continue their motion without changing their speeds and orientations. When the simulation output trace y contains collision(s) with Ego vehicle, we only consider the first collision of an Ego vehicle with any object for computing Eq. (1). Figure 6 shows the function with respect to the minimum time-to-collision and collision speed variables for a fixed collision surface. The effect of the collision surface to the cost is linear.




Fig. 6. Cost function to guide the search toward the boundary between collisions and near collisions.

G. Transition Check Function

The function we use for accepting a new configuration based on the cost is similar to the one proposed in [6]. Algorithm 1 repeats it for convenience. We denote the newly simulated configuration as the *candidate* and the initial configuration selected from the search tree as *previous* configuration. In the algorithm, K is a constant parameter normalizing the change

in the cost, T is the temperature parameter that is governing the likelihood of acceptance. The temperature parameter T is adaptively tuned by multiplying with or dividing by α with respect to the ratio of rejections to acceptances.

Algorithm 1 Algorithm used to check acceptance of a new configuration based on the change in the cost.

```

1: function ISTRANSITIONOK( $c_{prev}, c_{cand}$ ) ▷
    $c_{prev}$  and  $c_{cand}$  are the costs associated with the previous
   and candidate configurations, respectively.  $\alpha$  and  $K$  are
   constant parameters, and  $T$  is a persistent parameter.
2:   if  $c_{cand} < c_{prev}$  then
3:     return True
4:   end if
5:   if  $rand(0, 1) < e^{(c_{prev} - c_{cand}) / (K * T)}$  then
6:      $T = T / \alpha$ 
7:      $numberOfFails = 0$ 
8:     return True
9:   else
10:    if  $numberOfFails \geq maxNumberOfFails$  then
11:       $T = T * \alpha$ 
12:       $numberOfFails = 0$ 
13:    else
14:       $numberOfFails = numberOfFails + 1$ 
15:    end if
16:    return False
17:   end if
18: end function

```

H. Novelty Function

To get better coverage of the state space and to avoid local minima, we reward novelty in our search. For an Ego vehicle $e_i \in \mathcal{E}$ and an agent $a_j \in \mathcal{A}$, we define $\mathbf{x}_{e_i, a_j}[k]$ as the vector of relative states and change in relative states at discrete simulation time $k \in [0, n]$, *i.e.*, $\mathbf{x}_{e_i, a_j}[k] = (\mathbf{x}_{e_i}[k] - \mathbf{x}_{a_j}[k], (\mathbf{x}_{e_i}[k] - \mathbf{x}_{a_j}[k]) - (\mathbf{x}_{e_i}[k-1] - \mathbf{x}_{a_j}[k-1]))$.

We compute the novelty of $\mathbf{x}_{e_i, a_j}[k]$ as follows:

$$\mathcal{N} = \sum_{l=0}^m dist(\mathbf{x}_{e_i, a_j}[k], \mu_l) \quad (2)$$

where $\mu_l \in X_{rel, k-1}$ is the l^{th} nearest neighbor of $\mathbf{x}_{e_i, a_j}[k]$ in the set $X_{rel, k-1}$ which contains \mathbf{x}_{e_i, a_j} vectors for all ego-agent pairs for all times before k . The function $dist$ computes the Mahalanobis distance between $\mathbf{x}_{e_i, a_j}[k]$ and the elements of its m -nearest neighbors set. We choose to use the Mahalanobis distance because of its ability to provide a dissimilarity measure between two observations by utilizing the sample covariance matrix [16].

As each new configuration has a corresponding partial simulation of length t_{search} starting from a previous configuration, we compute the novelty for the trace of that partial simulation using Algorithm 2.

I. Termination Condition

Our algorithm checks a set of termination conditions to stop the search and returns the configuration which has the minimum cost associated with it. One of the termination conditions we use is a threshold for the minimum interesting cost. Another termination condition is a preset maximum overall time spent. Alternative termination conditions can be used, *e.g.*, a maximum number of nodes in the search tree.

Algorithm 2 Algorithm used to check the novelty of a new configuration.

```

1: function ISNOVEL( $\zeta, \mathcal{E}, \mathcal{A}, k_{start}, k_{end}, c_{prev}, c_{cand}$ )
2:    $N_{last}$  is persistent and keeps the last computed 10
   novelty values. numR is persistent and keeps the number
   of rejections. maxReject is the maximum number of
   consecutive rejections.  $X_{rel,k}$  is persistent and keeps the
   set of all past relative state computations.  $c_{prev}$  and  $c_{cand}$ 
   are the costs associated with the previous and candidate
   configurations, respectively.
3:
4:   Initialize  $N$  as an empty set
5:   for each  $e_i \in \mathcal{E}$  do
6:     for each  $a_j \in \mathcal{A}$  do
7:       for  $k = k_{start}$  to  $k_{end}$  do
8:         Compute  $x_{e_i, a_j}[k]$  from  $\zeta$ 
9:         Compute  $m$ -nearest neighbors of  $x_{e_i, a_j}[k]$ 
   in  $X_{rel, k-1}$ 
10:        Compute  $\mathcal{N}$  (novelty) with Eq. (2)
11:        Add  $\mathcal{N}$  to  $N$ 
12:        Add  $x_{e_i, a_j}[k]$  to  $X_{rel, k}$ 
13:      end for
14:    end for
15:  end for
16:   $\mathcal{N} = \max(N)$ 
17:  Update  $N_{last}$  to keep the last computed 10 novelty
  values
18:  if  $c_{cand} < 0.9c_{prev}$  or  $\text{numR} \geq \text{maxReject}$  or
   $|N_{last}| < 10$  or  $\mathcal{N} > \text{mean}(N_{last})$  then
19:     $\text{numR} = 0$ 
20:    return True
21:  else
22:     $\text{numR} = \text{numR} + 1$ 
23:    return False
24:  end if
25: end function

```

III. CASE STUDIES

Here, we present 2 case studies and compare our RRT-based approach with our falsification-based approach [2].

A. Case Study 1

1) *Scenario Setup*: In this case study, we have 2 agent vehicles and 1 Ego vehicle on a 3-lane straight road, i.e., $\mathcal{A} = \{a_1, a_2\}$ and $\mathcal{E} = \{e\}$. Figure 7 gives an high-level overview of our simulation setup. The initial position of agent vehicle 1 on the x axis is randomly sampled between 0 m and 25 m, the initial x position of agent vehicle 2 is randomly sampled between 10 m and 20 m, and the initial x position of Ego vehicle is randomly sampled between 30 m and 50 m. The initial positions of agent vehicles on the y axis are sampled between the centers of lane 1 and lane 3, i.e., between -3.5 m and 3.5 m, and the initial y position of Ego vehicle is randomly sampled between -1.75 m and 1.75 m, that is the lane markings separating Lane 3 and Lane 1 from Lane 2, respectively. The initial orientation of the Ego vehicle with respect to the x axis is randomly sampled between $-\pi/8$ rad and $\pi/8$ rad. The initial speed of Ego vehicle is sampled between 10 m/s and 15 m/s while the target speed is fixed to 15 m/s. The initial speeds of the agent vehicles are

sampled between 0 m/s and 15 m/s, and their target speeds at each waypoint are sampled between 0 m/s and 30 m/s.




Fig. 7. Vehicle initial states in the simulation setup for Case Study 1.

Ego vehicle has 5 sensors. Figure 8 visualizes the sensor placement and ranges of the sensors. A long-range sensor with a 45° field of view and 60 m range is placed at the front of the vehicle. Two 10 m-range sensors with 90° field-of-view are placed on the sides, facing left and right. Two 10 m-range sensors with 90° field-of-view are placed at the rear-left and rear-right corners with an angle to scan the area behind the rear corners of the vehicle.




Fig. 8. Ego vehicle sensor setup for Case Study 1.

Agent vehicles are controlled by the *move-to-pose* controller described in [17]. Ego vehicle controller has multiple modes based on the collision risk. When there is no estimated collision risk, a proportional control is applied to track target driving speed. If a collision is estimated in front, emergency brakes are applied and at the same time, depending on the occupancy of rear-left and rear-right areas, left or right steering is applied, respectively. If a collision is estimated in front-left (right), emergency brakes are applied and if rear-right (left) area is empty, also steering is applied to the right (left). If a collision is estimated in rear-left (right), if front and front-right (left) areas are empty, vehicle is accelerated with a right (left) steering, if only front area is empty and the vehicle on the front-right (left) is not imposing a risk, vehicle is accelerated without steering, otherwise emergency brakes are applied and right (left) steering is applied if the rear-right (left) area is empty. The collision avoidance algorithm presented here is very simple and it is not comparable to a controller that could be found in a real ADS. However, since our target in this work to study test generation approaches, rather than proposing a controller, we argue that this naive control approach is satisfactory for the purpose of this work. For the lateral control, we use the Stanford Racing Team's controller that was used in the DARPA Grand challenge [18].

2) *Experiment Results*: We have run 200 experiments with the falsification approach and with the RRT-based approach. The minimum, mean and maximum costs achieved by the




Fig. 9. The minimum cost result returned by the RRT-based approach in Case Study 1.




Fig. 10. The minimum cost result returned by the falsification-based approach in Case Study 1.

falsification approach were 0.0001, 12.4794, and 100.6082, respectively. The minimum, mean and maximum costs achieved by the RRT-based approach were 3.9124, 17.7190, and 88.9793, respectively. Figure 9 and Fig. 10 visualize the minimum-cost trajectories returned by the RRT-based and falsification-based approach, respectively. Histories of the vehicles are numbered to show their evolution over time. Figure 11 provides box and whisker diagrams of the minimum costs achieved by the two approaches among the 100 experiments we have carried. The black diamonds plotted on top of the box plots show the mean values for the returned minimum costs. In this case study, the falsification-based approach achieved smaller mean cost values, as well as the smaller minimum cost compared to the RRT-based approach. One reason for this is that, since the space between agent vehicles and Ego vehicle is open, there are not many local minimums that would make the exploration capabilities to achieve better than falsification-based approach. As it is easy to find a trajectory that is in the neighborhood

of an interesting case, falsification approach can focus on that neighborhood and minimize the cost as much as possible while RRT-based approach keeps looking for novel trajectories.

B. Case Study 2

1) *Scenario Setup*: In this case study, we have 4 agent vehicles and 1 Ego vehicle on a multiple-lane straight road, *i.e.*, $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$ and $\mathcal{E} = \{e\}$. Figure 12 gives an high-level overview of our simulation setup. The initial position of agent a_1 on the y axis is randomly sampled between 1.25 m and 6 m, the initial y positions of agent vehicles a_2, a_3, a_4 are randomly sampled between -2.25 m and -1.75 m. The initial speed of a_1 is randomly sampled between 5 m/s and 15 m/s, and its target speed at each waypoint is sampled between 0 m/s and 30 m/s. The initial and target speeds of all other vehicles are set to 15 m/s. All other initial states of the vehicles are fixed.

Ego vehicle has 5 sensors. Figure 13 visualizes the sensor placement and ranges of the sensors. A long-range sensor with a 22.5° field of view and 50 m range is placed at the front of the vehicle. Two 5 m-range sensors with 90° field-of-view are placed on the sides, facing left and right. Two 7 m-range sensors with 90° field-of-view are placed at the rear-left and rear-right corners with an angle to scan the area behind the rear corners of the vehicle.

Agent vehicle a_1 is controlled by the *move-to-pose* controller described in [17]. Agent vehicles $a_2, a_3,$ and a_4 are driven with a constant speed on a straight line. Ego vehicle controller is the same as the one described in Case Study 1.




Fig. 11. Comparison of the minimum cost achieved by the falsification approach and the RRT-based approach in Case Study 1.




Fig. 12. Initial states of the vehicles in the simulation setup for Case Study 2.




Fig. 13. Ego vehicle sensor setup for Case Study 2.

2) *Experiment Results:* For this case study, we only search for an optimal trajectory for a_1 with the target of minimizing the cost function described in Eq. (1). The existence of agent vehicles a_2 , a_3 , and a_4 between Ego vehicle e and agent vehicle a_1 creates many local minima for the selection of a_1 trajectories.

In this case study, we have executed 100 experiments with both the RRT-based approach and the falsification-based approach in MATLAB. Each run of both approaches had a time-out duration of 30 min. Out of 100 runs, only 5 of the minimum-cost trajectories returned by the falsification-based approach was able to make a_1 to move into the lane of Ego vehicle, and only 2 trajectories were able to cause a collision (1 high-speed collision and 1 boundary-case collision) and, other than the collision cases, only 1 trajectory was able to challenge Ego vehicle by activating its collision avoidance system. All other trajectories returned by the falsification approach were stuck in local-minima where a_1 tries to get closer to Ego vehicle and ends up colliding with one of the other agent vehicles. On the other hand in 28 of the minimum-cost trajectories returned by the RRT-based approach, agent a_1 was able to get into the lane of Ego vehicle and it was able to cause Ego vehicle to collide in 11 of those cases.

Figure 14 shows one of the interesting collision cases discovered by the RRT-based approach. Agent a_1 first forces Ego to move to the right to avoid a collision and then to the left where it ends up colliding with Agent a_3 . Histories of a_1 (red) and Ego (yellow) vehicles are numbered to show their evolution over time. Figure 15 shows the only small-speed collision case discovered by the falsification-based approach. Agent a_1 moves into the Ego vehicle's lane, accelerates and rear-ends with Ego vehicle even though Ego vehicle tries to avoid the collision by accelerating and steering away. Figure 16 shows a typical trajectory that is stuck in a local minimum. Agent a_1 tries to move closer to Ego vehicle and reduces the time-to-collision but collides with one of the other agents, which is a_2 in this figure. Although both approaches can get stuck in a local minimum, this case is significantly more common for the falsification-based approach as discussed above.

As a numerical comparison of the minimum costs discovered by the two approaches, the minimum, mean and maximum costs achieved by the falsification approach were 0.0043, 13.7134, and 50.6017, respectively. The minimum, mean and maximum costs achieved by the RRT-based approach were 4.8955, 10.2571, 15.0856. Figure 17 provides box and whisker diagrams of the minimum costs achieved by the two approaches among the 100 experiments we have carried. The black diamonds plotted on top of the box plots show the mean values for the returned minimum costs. While 44 out of 100 RRT-based approach experiments were able to avoid local minima occurring around the cost 10, only 2 of the falsification approach experiments were able to achieve this.

An observation we would like to share is that in the single case where the falsification approach was able to cause a small speed collision, the achieved minimum cost was significantly smaller than any of the 11 collision cases discovered by the RRT-based approach. This observation supports our conclusion in Case Study 1, *i.e.*, when the falsification-based approach can get into the neighborhood of a local or global minimum,

it is more able to get closer to the minimum point than the RRT-based approach. The advantage of the RRT-based approach in avoiding local minima and the advantage of the falsification-based in reaching minima suggests that an approach combining these two has the potential to improve the overall test generation performance. In such an approach, firstly, RRT-based approach would discover neighborhoods of minima and then the falsification-based approach would guide the test toward the minima starting from the results of the RRT-based approach.

IV. CONCLUSION AND FUTURE WORK

We proposed an approach that explores maneuvers for road occupants using rapidly-exploring random trees with the target of minimizing safety/performance cost functions defined for the automated driving system under test. In our approach, we have adopted notions from transition-based RRT [6] and RRT* [7] methods. Our RRT-based approach delivered more promising results compared to the stochastic optimization-based falsification approach which is described in [2] for the problems which contain many local minima that should be avoided for reaching the global minimum.

When formulating the agent trajectory generation problem as an optimization problem, the trajectories are presented with a finite number of parameters, *i.e.*, waypoint parameters for a fixed number of waypoints, which are provided by the test designer. One of the advantages of utilizing the RRT-based exploration is the ability to abandon the finite parameterization of the trajectories. This creates the opportunity to more freely explore the space, and to minimize the need for manually designing the general structure of the trajectory shapes. In an optimization-based approach, choosing a small number of parameters would limit the flexibility in generating critical trajectories, while choosing a large number would increase the dimensionality of the search space. For instance, in our Case Study 2, the number of parameters that the optimization method would need to create the minimum-cost trajectories that are discovered by the RRT-based approach varies from 8 to 68 based on the number of parameters at each waypoint and the number of waypoints that are used to create these trajectories.

Future work will include:

- Exploring new methods for computing the novelty of a trajectory instead of a single state
- Computing a cost on the shape of the trajectory instead of computing the minimum of the instantaneous costs of the points on the trajectory. This may especially be useful for rewarding some types of vehicle paths such as the ones which are closer to real-world driving behaviors.
- Exploring new methods for using multi-objective optimization can be studied for different objectives like time-to-collision, collision speed, collision impact area etc.
- Combining RRT-based approach with the falsification-based approach such that the exploration starts with RRT-based approach and falsification-based approach further minimizes the cost starting from the best cases discovered by the RRT-based approach.

REFERENCES

- [1] G. Fainekos, S. Sankaranarayanan, K. Ueda, and H. Yazarel, "Verification of automotive control applications using S-TaLiRo," in *Proceedings of the American Control Conference*, 2012.
- [2] C. E. Tuncali, T. P. Pavlic, and G. Fainekos, "Utilizing S-TaLiRo as an automatic test generation framework for autonomous vehicles," in *IEEE Intelligent Transportation Systems Conference*, 2016.
- [3] C. E. Tuncali, G. Fainekos, H. Ito, and J. Kapinski, "Simulation-based adversarial test generation for autonomous vehicles with machine learning components," in *IEEE Intelligent Vehicles Symposium (IV)*, 2018.




Fig. 14. One of the collision cases discovered by the RRT-based approach.




Fig. 15. The small-speed collision discovered by the falsification-based approach.




Fig. 16. One of the cases where falsification-based approach was stuck in a local minimum.




Fig. 17. Comparison of the minimum cost achieved by the falsification approach and the RRT-based approach in Case Study 2.

- [4] T. Dreossi, S. Ghosh, A. Sangiovanni-Vincentelli, and S. A. Seshia, "Systematic testing of convolutional neural networks for autonomous driving," in *Reliable Machine Learning in the Wild (RMLW)*, 2017.
- [5] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," Tech. Rep., 1998.
- [6] L. Jaillet, J. Cortés, and T. Siméon, "Transition-based RRT for path planning in continuous cost spaces," in *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*. IEEE, 2008, pp. 2145–2150.
- [7] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The international journal of robotics research*, vol. 30, no. 7, pp. 846–894, 2011.
- [8] J. M. Esposito, J. Kim, and V. Kumar, "Adaptive RRTs for validating hybrid robotic control systems," in *Algorithmic Foundations of Robotics VI*. Springer, 2004, pp. 107–121.
- [9] J. Kim, J. M. Esposito, and V. Kumar, "An RRT-based algorithm for testing and validating multi-robot controllers," Moore School of Electrical Engineering PA, Grasp Lab, Tech. Rep., 2005.
- [10] M. S. Branicky, M. M. Curtiss, J. Levine, and S. Morgan, "Sampling-based planning, control and verification of hybrid systems," *IEE Proceedings-Control Theory and Applications*, vol. 153, no. 5, pp. 575–590, 2006.
- [11] T. Dang, A. Donzé, O. Maler, and N. Shalev, "Sensitive state-space exploration," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*. IEEE, 2008, pp. 4049–4054.
- [12] E. Plaku, L. E. Kavraki, and M. Y. Vardi, "Hybrid systems: from verification to falsification by combining motion planning and discrete search," *Formal Methods in System Design*, vol. 34, no. 2, pp. 157–182, 2009.
- [13] T. Dreossi, T. Dang, A. Donzé, J. Kapinski, X. Jin, and J. V. Deshmukh, "Efficient guiding strategies for testing of temporal properties of hybrid systems," in *7th International Symposium NASA Formal Methods (NFM)*, ser. LNCS, vol. 9058. Springer, 2015, pp. 127–142.
- [14] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [15] J. Lehman and K. O. Stanley, "Exploiting open-endedness to solve problems through the search for novelty," in *ALIFE*, 2008.
- [16] R. De Maesschalck, D. Jouan-Rimbaud, and D. L. Massart, "The mahalanobis distance," *Chemometrics and intelligent laboratory systems*, vol. 50, no. 1, pp. 1–18, 2000.
- [17] P. Corke, *Robotics, Vision and Control: Fundamental Algorithms In MATLAB®, 2nd, Completely Revised*. Springer, 2017, vol. 118.
- [18] G. M. Hoffmann, C. J. Tomlin, M. Montemerlo, and S. Thrun, "Autonomous automobile trajectory tracking for off-road driving: Controller design, experimental validation and racing," in *American Control Conference*, 2007, pp. 2296–2301.