
A Scenario-Based Platform for Testing Autonomous Vehicle Behavior Prediction Models in Simulation

Francis Indaheng¹
findaheng@berkeley.edu

Edward Kim¹
ek65@eecs.berkeley.edu

Kesav Viswanadha¹
kesav@berkeley.edu

Jay Shenoy¹
jayshenoy@berkeley.edu

Jinkyu Kim²
jinkyukim@korea.ac.kr

Daniel J. Fremont³
dfremont@ucsc.edu

Sanjit A. Seshia¹
sseshia@eecs.berkeley.edu

¹University of California, Berkeley

²Korea University

³University of California, Santa Cruz

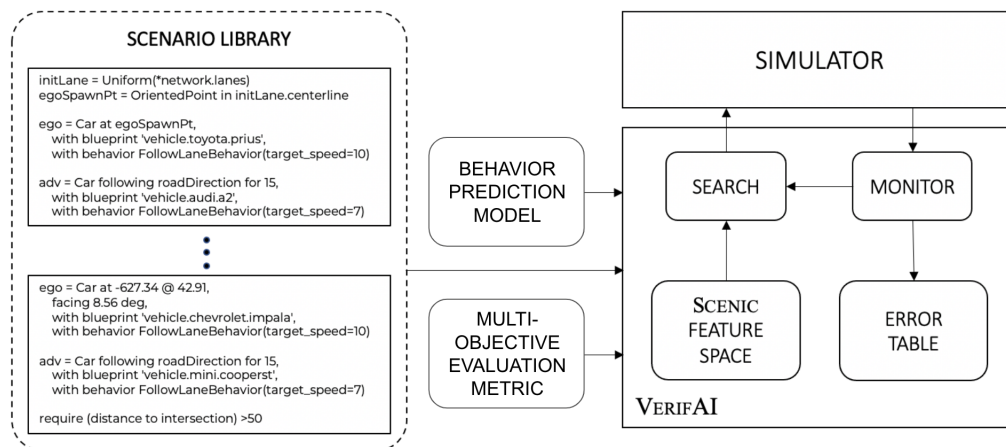


Figure 1. An architecture diagram demonstrating the flow of components in the testing platform. A behavior prediction model, library of SCENIC test scenarios, and multi-objective evaluation metric serve as inputs to VERIFAI, which monitors and searches for failure cases.

Abstract

Behavior prediction remains one of the most challenging tasks in the autonomous vehicle (AV) software stack. Forecasting the future trajectories of nearby agents plays a critical role in ensuring road safety, as it equips AVs with the necessary information to plan safe routes of travel. However, these prediction models are data-driven and trained on data collected in real life that may not represent the full range of scenarios an AV can encounter. Hence, it is important that these prediction models are extensively tested in various test scenarios involving interactive behaviors prior to deployment. To support this need, we present a simulation-based testing platform that supports (1) intuitive scenario modeling with a probabilistic programming language called SCENIC, (2) specifying a multi-objective evaluation metric with a partial priority ordering, (3) falsification of the provided metric, and (4) parallelization of simulations for scalable testing. As a part of the platform, we

provide a library of 25 SCENIC programs that model challenging test scenarios involving interactive traffic participant behaviors. We demonstrate the effectiveness and the scalability of our platform by testing a trained behavior prediction model and searching for failure scenarios.

1 Introduction

Behavior (or trajectory) prediction, which is a component of an autopilot, plays a critical role in ensuring road safety, as an autonomous vehicle (AV) that can accurately forecast the motions of surrounding vehicles is better equipped to plan safe routes of travel. This prediction task is challenging and is currently an active field of study. The difficulty of this task stems from the multifaceted nature of dynamic traffic scenarios: an effective model must appropriately account for both the AV’s surrounding environment and the social interactions between various agents over a duration of time.

Models for such tasks are commonly based on neural network (NN) architectures, such as Recurrent Neural Networks (RNNs) [13] and Long Short-Term Memory (LSTM) networks [12, 16]. These are data-driven models trained on data collected on roads. Capturing data related to all possible scenarios on roads is difficult, let alone characterizing all possible scenarios itself. There are edge-case (i.e. rare) scenarios that, by definition, are rarely encountered on roads and, consequently, are only sparsely reflected in the training data, if at all. Hence, testing behavior prediction models in, at least, *known* edge case scenarios is crucial. Tending to this issue in the real world would involve reconstructing such scenarios in a controlled testing facility, which can be expensive and labor-intensive, resulting in a lack of scalability of testing. On the contrary, simulation offers an inexpensive mechanism for reconstructing scenarios of interest, making for a more scalable approach. Yet, simulation-based testing of behavior prediction models requires support for (1) modeling interactive scenarios of interest in simulation, (2) generating necessary input synthetic sensor data from those simulations, and (3) evaluating behavior prediction models using multi-objective evaluation metrics. We aim to address this need in this paper.

Contributions We provide a platform to test behavior prediction models that is efficient, scalable, and scenario-based. The main contributions of this work include the following:

- An open-source platform¹ for researchers and developers to evaluate behavior prediction models in the CARLA simulator [4] under scenarios modeled in the SCENIC language [9, 10]. The platform allows users to specify custom multi-objective evaluation metrics and falsify their models with respect to these metrics, parallelizing simulations for scalable testing;
- An open-source library of 25 SCENIC programs to evaluate behavior prediction models in various dynamic scenarios;²
- An experimental evaluation of a behavior prediction model against five traffic scenarios to demonstrate the efficacy and scalability of our platform.

2 Background

2.1 SCENIC: Scenario Modeling Language

SCENIC [9, 10] is a probabilistic programming language whose syntax and semantics are designed to intuitively *model* and *generate* scenarios involving interactive agents. A SCENIC program represents an *abstract* scenario which defines distributions over the initial states and dynamic behaviors of agents in the scenario. Sampling from these distributions yields a *concrete* scenario that assigns specific values to these parameters. The SCENIC tool can generate concrete scenarios from a SCENIC program and execute them in a simulator to generate data and test a system. SCENIC is simulator-agnostic and has interfaces to a variety of simulators for driving, aviation, robotics, and other domains [10].

¹<https://github.com/BehaviorPredictionTestingPlatform/TestingPipeline>

²<https://github.com/BehaviorPredictionTestingPlatform/ScenarioLibrary>

Scenario #	minADE	minFDE	MR	CR	SD
1	0.037	0.15	0.017	0.034	0.58
2	0.044	0.19	0.016	0.025	0.54
3	0.081	0.39	0.034	0.24	0.45
4	0.13	0.74	0.31	0.36	0.52
5	0.19	1.1	0.49	0.51	0.56

Table 1. Results of evaluating the LaneGCN model against five SCENIC programs. The minADE, minFDE, and MR metrics indicate the model’s error of predictions with respect to ground-truth trajectories. The CR and SD metrics reflect the ability of the VERIFAI sampler to identify a diverse number of counterexample scenarios.

2.2 VERIFAI: Formal Analysis Software Toolkit

VERIFAI [5, 20] is a toolkit for formal design and analysis of AI/ML-based systems. Its architecture is shown in Figure 1. It takes three inputs: (1) the system of interest (e.g. behavior prediction model), (2) an environment model (e.g. SCENIC program), and (3) a multi-objective evaluation metric.

VERIFAI performs falsification by sampling concrete scenarios using a variety of techniques, simulating each one and monitoring the system’s performance with respect to the given evaluation metric. The sampled test parameters and corresponding system performance are logged for further analysis, and also provided to the sampling algorithms to more intelligently select the next scenario. To support multi-objective evaluation metrics, VERIFAI provides a multi-armed bandit sampler, which formulates the search process as a multi-armed bandit problem with a reward function favoring concrete scenarios that pessimize multiple metrics (taking priority into account) [20].

3 Overview

The overall architecture of our platform is visualized in Figure 1. The user needs to provide three inputs: (1) a trained behavior prediction model, (2) a multi-objective evaluation metric, and (3) a library of SCENIC programs modeling a set of abstract scenarios (refer to Section 2.1).

The evaluation pipeline for behavior prediction is shown in Figure 2 in the Appendix. As stated in Section 2.1, a SCENIC program in conjunction with the CARLA simulator is used to generate various concrete scenarios. We implement support for collecting various types of data from each simulation to enable testing different behavior prediction models, which can take various types of inputs. Users can collect trajectories (i.e. timestamped position and headings of all traffic participants in the scenario) and a stream of sensor data from RGB camera(s), depth camera(s), and/or LiDAR sensor(s). We can also collect corresponding ground-truth labels such as segmentation labels for RGB images and LiDAR, as well as 3D bounding boxes.

From these collected data, we extract two different types of data: (1) historical data, which is the segment of the data which is input to the behavior prediction model, and (2) ground-truth future trajectory data, which is used to evaluate the model’s predicted trajectories using the provided multi-objective evaluation metric, which outputs $\vec{\rho}$, a tuple containing a *score* for each metric specified. VERIFAI’s multi-armed bandit sampler is used to falsify the given multi-objective evaluation metric. Please refer to Section 2.2 for more detail on the sampler.

4 Experimental Evaluation

We used our platform to evaluate the LaneGCN model [15]; for details of our setup, see Appendix D.1.

In Table 1, we enumerate the results from our testing platform after evaluating LaneGCN against the five scenarios in Appendix D.2. The first three metrics (minADE, minFDE, MR) quantify the behavior prediction model’s performance, averaged over all samples generated by VERIFAI, whereas the latter two metrics, counterexample rate (CR) and scenario diversity (SD), provide insight about the effectiveness of different sampling methods. The CR expresses the number of counterexamples found in VERIFAI with respect to the total number of samples run. The SD relates the observed variance to the defined range of values for a set of feature distributions. For further details on these metrics, refer to Appendix B.

While SD remained relatively stable around 0.53 ± 0.08 , we observed some variance in the other metrics. LaneGCN made noticeably more accurate predictions for Scenarios 1, 2, and 3 than for Scenarios 4 and 5, reflected by their substantially lower minADE, minFDE, and MR values. We speculate this may be attributed to the differences in complexity of the vehicle’s behavior in each of these scenarios. In Scenarios 1, 2, and 3, the oncoming vehicles had a relatively straight trajectory, either continuing a straight path or making a single lane change. Meanwhile, in Scenarios 4 and 5, the oncoming vehicles makes an unprotected left turn across an intersection. Scenario 5 is even more complex than Scenario 4, with the ego vehicle making a right turn at the intersection instead of continuing straight. Recall that LaneGCN takes historical trajectories for *all* traffic agents as input, hence why the ego vehicle’s own behavior could influence the accuracy of its predictions. Overall, these results suggest that our scenario-based testing platform is effective for evaluating behavior prediction model performance and identifying difficult driving scenarios for the model.

To demonstrate the efficiency of our approach, we benchmarked the runtime of the testing platform for both sequential and parallel executions. We conducted these experiments by evaluating LaneGCN against Scenario 1 for N iterations, where $N \in \{25, 50, 75, 100\}$. We ran the pipeline sequentially, during which there was no overlap between sampling and simulation, and then we ran the pipeline in parallel, first with 2 workers, then with 5 workers.

The results are plotted in Figure 4 in the Appendix. On average, parallel execution yielded a 1.51x speedup with 2 workers and a 3.09x speedup with 5 workers. There are a number of reasons the speedup did not quite achieve a factor consistent with the number of parallel workers. One definite cause is that VERIFAI only parallelizes simulation and evaluation [20]. Thus, while simulation of the concrete scenarios can occur simultaneously on several instances of the CARLA simulator, the sampling process remains sequential.

5 Related Work

Although there have been increasing efforts to develop behavior prediction models for AVs, there has been little work on formally testing them in simulation. Some work has evaluated AV planners using simulated perception and prediction inputs [21] from a set of scenarios, but the focus was to demonstrate applying different noise models to improve the realism of simulated data. Our efforts are directed towards creating an efficient and scalable testing platform that provides customizable metrics and data generation capabilities, with the capacity for users to model their own scenarios.

VERIFAI has been used in a number of other case studies on AI systems [5], including: identifying relevant tests for AV track testing [11]; falsifying, debugging, and retraining an autonomous aircraft taxiing system [8]; and evaluating a full autopilot stack in the IEEE AV Test Challenge [19]. Our work is the first to use VERIFAI for testing behavior prediction models.

6 Conclusion

In this paper, we presented a platform for testing behavior prediction models with SCENIC and VERIFAI. As part of the platform, we provided a library of interactive scenarios encoded as SCENIC programs. We demonstrated our platform by evaluating the LaneGCN model against well-known motion forecasting metrics. Lastly, we showcased the scalability of running the pipeline in a parallelized fashion with a runtime benchmark comparison to a sequential execution.

Acknowledgments and Disclosure of Funding

The authors would like to express thanks to Hazem Torfah, Sebastian Junges, Marcell Vazquez-Chanlatte, Yash Pant, Justin Wong, and Ameesh Shah for their helpful discussions and feedback. This work is supported in part by NSF grants 1545126 (VeHICaL) and 1837132, by the DARPA contracts FA8750-18-C-0101 (AA) and FA8750-20-C-0156 (SDCPS), by Berkeley Deep Drive, by the Toyota Research Institute, and by Toyota under the iCyPhy center.

References

- [1] Argo AI. Argoverse motion forecasting competition. <https://eval.ai/web/challenges/challenge-page/454/overview>, 2019. Published: 2019-09-27.
- [2] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [3] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff G. Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. *2019 International Conference on Robotics and Automation (ICRA)*, pages 2090–2096, 2019.
- [4] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pages 1–16, 2017.
- [5] Tommaso Dreossi, Daniel J. Fremont, Shromona Ghosh, Edward Kim, Hadi Ravanbakhsh, Marcell Vazquez-Chanlatte, and Sanjit A. Seshia. VerifAI: A toolkit for the formal design and analysis of artificial intelligence-based systems. In *31st International Conference on Computer Aided Verification (CAV)*, July 2019.
- [6] Marius Dupuis and Hans Grezlikowski. OpenDRIVE - An open standard for the description of roads in driving simulations. In *Proceedings of the Driving Simulation Conference Europe*, pages 25–35, 2006.
- [7] Marius Dupuis, Martin Strobl, and Hans Grezlikowski. OpenDRIVE 2010 and beyond - Status and future of the de facto standard for the description of road networks. In *Proceedings of the Driving Simulation Conference Europe*, pages 231–242, 2010.
- [8] Daniel J. Fremont, Johnathan Chiu, Dragos D. Margineantu, Denis Osipychev, and Sanjit A. Seshia. Formal analysis and redesign of a neural network-based aircraft taxiing system with verifai. In *Computer Aided Verification - 32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21-24, 2020, Proceedings, Part I*. Springer, 2020.
- [9] Daniel J. Fremont, Tommaso Dreossi, Shromona Ghosh, Xiangyu Yue, Alberto L. Sangiovanni-Vincentelli, and Sanjit A. Seshia. Scenic: A language for scenario specification and scene generation. *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, June 2019.
- [10] Daniel J. Fremont, Edward Kim, Tommaso Dreossi, Shromona Ghosh, Xiangyu Yue, Alberto L. Sangiovanni-Vincentelli, and Sanjit A. Seshia. Scenic: A language for scenario specification and data generation. *arXiv:2010.06580 [cs.PL]*, 2020. To appear in *Machine Learning*.
- [11] Daniel J. Fremont, Edward Kim, Yash Vardhan Pant, Sanjit A. Seshia, Atul Acharya, Xantha Bruso, Paul Wells, Steve Lemke, Qiang Lu, and Shalin Mehta. Formal scenario-based testing of autonomous vehicles: From simulation to the real world. *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8, 2020.
- [12] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, November 1997.
- [13] Namdar Homayounfar, Wei-Chiu Ma, Shrinidhi Kowshika Lakshmikanth, and Raquel Urtasun. Hierarchical recurrent attention networks for structured online maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [14] Siddhesh Khandelwal, William Qi, Jagjeet Singh, Andrew Hartnett, and Deva Ramanan. What-if motion prediction for autonomous driving. *arXiv preprint arXiv:2008.10587*, 2020.
- [15] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *ECCV*, 2020.

- [16] Jean Mercat, Thomas Gilles, Nicole El Zoghby, Guillaume Sandou, Dominique Beauvois, and Guillermo Pita Gil. Multi-modal simultaneous forecasting of vehicle position sequences using social attention. *CoRR*, abs/1910.03650, 2019.
- [17] Jean Pierre Mercat, Thomas Gilles, Nicole El Zoghby, Guillaume Sandou, Dominique Beauvois, and Guillermo Pita Gil. Multi-head attention for multi-modal joint vehicle motion forecasting. *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9638–9644, 2020.
- [18] Elizabeth D. Swanson, Frank Foderaro, Mikio Yanagisawa, Wassim G. Najm, and Philip Azeredo. Statistics of light-vehicle pre-crash scenarios based on 2011-2015 national crash data. *National Highway Traffic Safety Administration (Report No. DOT HS 812 745)*, August 2019.
- [19] Kesav Viswanadha, Francis Indaheng, Justin Wong, Edward Kim, Ellen Kalvan, Yash Vardhan Pant, Daniel J. Fremont, and Sanjit A. Seshia. Addressing the IEEE AV test challenge with Scenic and VerifAI. *2021 IEEE International Conference on Artificial Intelligence Testing (AITest)*, pages 136–142, 2021.
- [20] Kesav Viswanadha, Edward Kim, Francis Indaheng, Daniel J. Fremont, and Sanjit A. Seshia. Parallel and multi-objective falsification with Scenic and VerifAI. In Lu Feng and Dana Fisman, editors, *Runtime Verification*, pages 265–276, Cham, 2021. Springer International Publishing.
- [21] K. Wong, Qiang Zhang, Ming Liang, Binh Yang, Renjie Liao, Abbas Sadat, and Raquel Urtasun. Testing the safety of self-driving vehicles by simulating perception and prediction. In *ECCV*, 2020.
- [22] Wei Zhan, Liting Sun, Di Wang, Haojie Shi, Aubrey Clausse, Maximilian Naumann, Julius Kümmerle, Hendrik Königshof, Christoph Stiller, Arnaud de La Fortelle, and Masayoshi Tomizuka. INTERACTION Dataset: An INTERnational, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps. *arXiv:1910.03088 [cs, eess]*, September 2019.

A Description of the Library of SCENIC Programs

As a part of our behavior prediction evaluation platform, we provide a library of 25 SCENIC programs as test scenario models. We referred to the INTERACTION dataset [22], the Argoverse Motion Forecasting dataset [2], and a publication by the National Highway Traffic Safety Administration (NHTSA) of pre-crash scenarios based on 2011-2015 crash data recorded by the U.S. Department of Transportation [18] and encoded scenarios as SCENIC programs. These datasets were chosen based on our efforts to capture a diverse selection of *interactive* scenarios with *critical* trajectories (e.g., near-collision situations). These scenarios covered a wide variety of road structures—such as highways, intersections, and roundabouts—as well as agent interactions—such as bypassing, merging, and unprotected left turns.

B Metrics

B.1 Multi-Objective Evaluation Metrics

In VERIFAI, a multi-objective function defines a set of metrics used to determine the efficacy of the system being evaluated. As stated in Section 3, the output of the function is $\vec{\rho}$, a tuple containing a *score* for each metric specified, the calculation of which is defined in the function. By convention, a negative score for any element in $\vec{\rho}$ indicates a failure, in which case the counterexample is stored in the error table for offline analysis; accordingly, a non-negative score for all elements in $\vec{\rho}$ indicates the model succeeded in meeting the evaluation criteria.

For the purpose of testing behavior prediction models, we apply three commonly-used motion forecasting metrics. *Average Displacement Error* (ADE) is the Euclidean distance between the predicted trajectories, (\hat{x}_t, \hat{y}_t) , and the ground-truth trajectories, (x_t, y_t) , averaged over all n timesteps:

$$ADE(x, y) = \frac{1}{n} \sum_{t=1}^n \sqrt{(x_t - \hat{x}_t)^2 + (y_t - \hat{y}_t)^2} \quad (1)$$

Final Displacement Error (FDE) is the Euclidean distance between the predicted trajectory, (\hat{x}_n, \hat{y}_n) , and the ground-truth trajectory, (x_n, y_n) , at the last timestep:

$$FDE(x, y) = \sqrt{(x_n - \hat{x}_n)^2 + (y_n - \hat{y}_n)^2} \quad (2)$$

For the i^{th} sample in a dataset, we take the minimum ADE (minADE) and minimum FDE (minFDE) across the k most probable predictions, P_k^i , generated by the model, as done in similar approaches accommodating multi-modal trajectories [3, 14, 15, 16]. Doing so avoids penalizing the model for producing plausible predictions that don't align with the ground truth.

$$\min ADE(P_k^i) = \min_{(x,y) \in P_k^i} ADE(x, y) \quad (3)$$

$$\min FDE(P_k^i) = \min_{(x,y) \in P_k^i} FDE(x, y) \quad (4)$$

The third metric, *Miss Rate* (MR), is the ratio of missed predictions over the entire test set of N data points, P_k . In this context, a missed prediction occurs when all k proposed trajectories in prediction P_k^i differ in final position from the ground truth by more than some distance d threshold. This metric indicates the trade-off between accuracy and diversity [17].

$$MR(P_k, d) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\min FDE(P_k^i) > d\} \quad (5)$$

In our platform, we enabled a command-line argument for the user to specify the threshold values for minADE, minFDE, and MR. Recall that $\vec{\rho}$ is a tuple of scores, and if any score in $\vec{\rho}$ is negative, then the scenario is deemed a failing counterexample. To conform to this convention, we store the difference between the thresholds and their corresponding metrics as the scores in $\vec{\rho}$ (e.g., we store $(2 - \min ADE)$ to enforce a minADE threshold of 2).

B.2 Platform Metrics

In section B.1, we discussed the metrics that would be used to evaluate the behavior prediction model (Eq. 3, 4, 5). Now, we describe two metrics for evaluating the effectiveness of our testing platform: *Counterexample Rate* (CR) and *Scenario Diversity* (SD). The CR expresses the number of counterexamples found in VERIFAI with respect to the total number of samples run. Thus, the CR reflects the ability of the multi-armed bandit sampler to search the feature space and identify failing scenarios in which the behavior prediction model’s predictions did not meet the criteria of the multi-objective function:

$$CR = \frac{\text{Number of counterexamples}}{\text{Number of all examples}} \quad (6)$$

We can also formalize the overall SD produced by SCENIC, which relates the observed variance to the defined range of values for a set of feature distributions. For the i^{th} feature, f_i , let $\sigma(f_i)$ denote the standard deviation of values observed and $L(f_i)$ denote the interval length, where the interval length is the size of the range of values specified in the feature’s distribution. The SD for a set of N features is as follows:

$$SD = \frac{2 \cdot \sum_{i=1}^N \sigma(f_i)}{\sum_{i=1}^N L(f_i)} \quad (7)$$

C LaneGCN Behavior Prediction Model

We evaluate LaneGCN behavior prediction model by Uber ATG [15]. This behavior prediction model ranked 1st in a motion forecasting competition [1] held by Argo AI, based on the Argoverse dataset [2]. Furthermore, it was one of few working open-sourced behavior prediction models. We evaluate the pre-trained LaneGCN model, which was trained on over 300,000 real-life scenarios collected for the aforementioned Argoverse dataset.

To test the model in CARLA, we implemented a support to make it operational with the OpenDRIVE format [6, 7]. We applied the necessary computations to bridge the gap between map representations in SCENIC and the model. Due to the modular nature of the model’s architecture, this functional exchange could be completed with minimal effect to any dependencies.

The model expects a CSV file containing 20 timesteps of historical trajectories for all traffic agents as input. It produces 6 predictions, each containing 15 timesteps of predicted trajectories, for a designated single traffic agent. All trajectories use a sample rate of 10 Hz.

D Experimental Details

D.1 Experimental Settings

For these experiments, we utilized the multi-armed bandit sampler in VERIFAI for searching the feature space. As demonstrated in a recent case study [20], the multi-armed bandit sampler strikes a strong balance between falsification capabilities (i.e., the number of counterexamples found) and feature space diversity (i.e., the concrete values sampled for each feature). Please refer to Section 2 for more detail.

For each scenario, we sample and generate 120 concrete scenarios, divided into 4 batches. The difference between batches was the initial point in time we wished to start the prediction, referred to as the parameter *timepoint*. For example, setting `timepoint=40` indicates historical trajectories encompass timesteps in the range [20, 40) and predicted trajectories encompass timesteps in the range [40, 55). Note that since historical trajectories include 20 timesteps, there is a minimum constraint of `timepoint=20`. This parameter allows users to test a model on different parts of a scenario, which may involve greater or lesser complexity of agent-to-agent interactions. Thus, we generated 30 samples (i.e., a batch) for each of `timepoint ∈ {20, 40, 60, 80}` to evaluate the behavior prediction model’s capability to predict in different segments of a scenario. We set the minADE threshold to 0.1 meters, the minFDE threshold to 1.0 meter, and the MR distance threshold to 1.0 meter.

These experiments were conducted on a machine running Ubuntu 18.04.5, equipped with an Intel Core i9-9900X @ 3.50GHz central processing unit (CPU), as well as an Nvidia GeForce GTX and two Nvidia Titan RTX graphics processing units (GPUs).

D.2 Scenario Descriptions

We illustrate use of our testing platform by evaluating the LaneGCN model against five scenarios of varying complexity from our library of SCENIC programs, described as the following:

1. The ego vehicle waits at a 3-way intersection for another vehicle from the lateral lane to pass before making a right turn.
2. The ego vehicle makes a left turn at a 3-way intersection and must maneuver to avoid collision when another vehicle from the lateral lane continues straight.
3. A trailing vehicle performs a lane change to bypass the leading ego vehicle before returning to its original lane.
4. The ego vehicle drives straight through a 4-way intersection and must suddenly stop to avoid collision when another vehicle from the oncoming lane makes an unprotected left turn.
5. The ego vehicle makes a right turn at a 4-way intersection while another vehicle from the oncoming lane makes an unprotected left turn, such that both vehicles are turning to the same outgoing leg of an intersection.

Here, *ego vehicle* refers to the reference point for the behavior prediction (i.e., the vehicle in the scenario that would be administering the behavior prediction model). The SCENIC source code of these scenarios is included in the appendix as Figures 5-9.

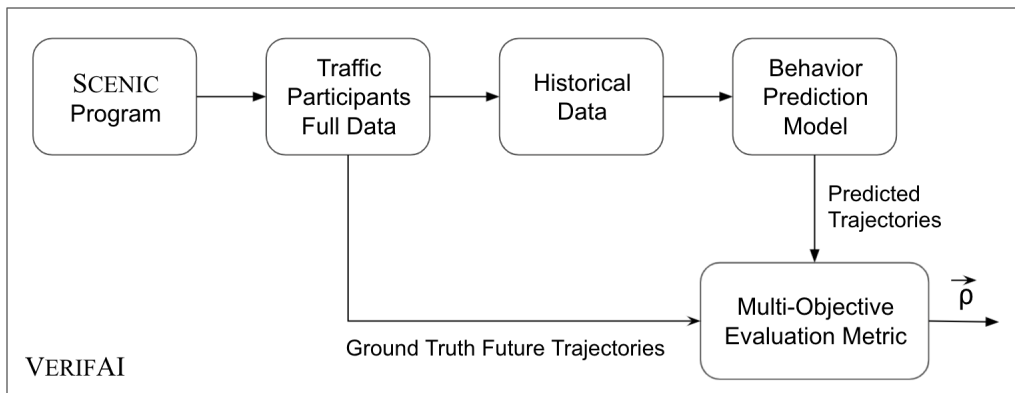


Figure 2. An overview of the testing pipeline in VERIFAI. Historical data is extracted and serves as input to the behavior prediction model. The predicted and ground-truth trajectories are compared against the multi-objective evaluation metric to derive a tuple of scores, referred to as \vec{p} .

```

1 behavior PullIntoRoad(laneToFollow, target_speed):
2     while (distance from self to ego) > CUT_IN_TRIGGER_DISTANCE:
3         wait
4         do FollowLaneBehavior(laneToFollow=ego.lane, target_speed)
5
6 behavior EgoBehavior(target_speed):
7     try:
8         do FollowLaneBehavior(target_speed)
9         interrupt when (distance to intersection) < 5: # meters
10        do TurnRight()
11        interrupt when withinDistanceToAnyObjs(self, SAFETY_DISTANCE):
12        do Uniform(CollisionAvoidance(), LaneChangeBehavior())
13
14 target_speed = 5 # meters/second
15 ego = Car with behavior EgoBehavior(target_speed)
16 spot = OrientedPoint on visible curb
17 badAngle = Uniform(-1,1) * Range(10, 20) deg
18 parkedCar = Car left of spot by 0.5, # meter
19             facing badAngle relative to roadDirection,
20             with behavior PullIntoRoad(ego.lane, target_speed)
21 require (distance from ego to parkedCar) < 20
22 require eventually (ego.lane is parkedCar.lane)

```

Figure 3. An example SCENIC program in which a badly parked car cuts into the ego vehicle’s lane.

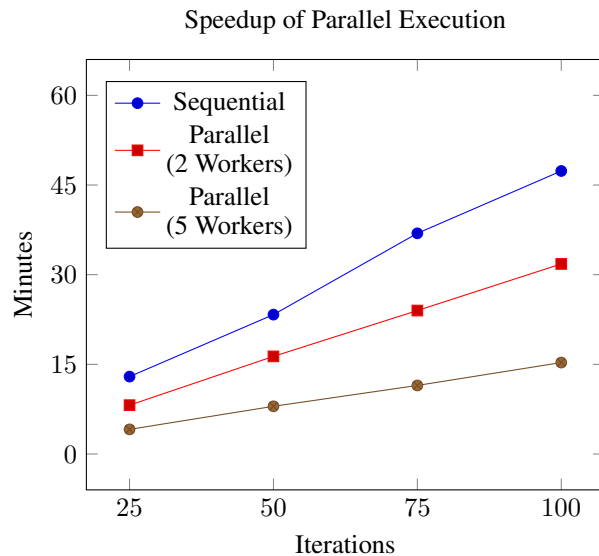


Figure 4. Runtime benchmark of sequential and parallel executions in the testing pipeline. Iterations is the number of simulations ran.

```

#####
# MAP AND MODEL #
#####

param map = localPath('.././.././../tests/formats/opendrive/maps/CARLA/Town05.xodr')
param carla_map = 'Town05'
model scenic.simulators.carla.model

#####
# CONSTANTS #
#####

MODEL = 'vehicle.lincoln.mkz2017'

EGO_INIT_DIST = [20, 25]
param EGO_SPEED = VerifaiRange(7, 10)
EGO_BRAKE = 1.0

ADV_INIT_DIST = [10, 15]
param ADV_SPEED = VerifaiRange(7, 10)

param SAFETY_DIST = VerifaiRange(10, 20)
CRASH_DIST = 5
TERM_DIST = 70

#####
# AGENT BEHAVIORS #
#####

behavior EgoBehavior(traj):
  try:
    do FollowTrajectoryBehavior(
      target_speed=globalParameters.EGO_SPEED, trajectory=traj)
  interrupt when withinDistanceToAnyObjs(self, globalParameters.SAFETY_DIST):
    take SetBrakeAction(EGO_BRAKE)
  interrupt when withinDistanceToAnyObjs(self, CRASH_DIST):
    terminate

#####
# SPATIAL RELATIONS #
#####

intersection = Uniform(*filter(lambda i: i.is3Way, network.intersections))

egoManeuver = Uniform(*filter(lambda m:
  m.type is ManeuverType.RIGHT_TURN, intersection.maneuvers))
egoLane = egoManeuver.startLane
egoTraj = [egoLane, egoManeuver.connectingLane, egoManeuver.endLane]
egoSpawnPt = OrientedPoint in egoLane.centerline

advManeuver = Uniform(*filter(lambda m:
  m.type is ManeuverType.STRAIGHT, egoManeuver.conflictingManeuvers))
advLane = advManeuver.startLane
advTraj = [advLane, advManeuver.connectingLane, advManeuver.endLane]
advSpawnPt = OrientedPoint in advLane.centerline

#####
# SCENARIO SPECIFICATION #
#####

ego = Car at egoSpawnPt,
  with blueprint MODEL,
  with behavior EgoBehavior(egoTraj)

adv = Car at advSpawnPt,
  with blueprint MODEL,
  with behavior FollowTrajectoryBehavior(
    target_speed=globalParameters.ADV_SPEED, trajectory=advTraj)

require EGO_INIT_DIST[0] <= (distance to intersection) <= EGO_INIT_DIST[1]
require ADV_INIT_DIST[0] <= (distance from adv to intersection) <= ADV_INIT_DIST[1]
terminate when (distance to egoSpawnPt) > TERM_DIST

```

Figure 5. SCENIC program used in Scenario 1

```

#####
# MAP AND MODEL #
#####

param map = localPath('.././.././../tests/formats/opendrive/maps/CARLA/Town05.xodr')
param carla_map = 'Town05'
model scenic.simulators.carla.model

#####
# CONSTANTS #
#####

MODEL = 'vehicle.lincoln.mkz2017'

EGO_INIT_DIST = [20, 40]
param EGO_SPEED = VerifaiRange(7, 10)
param EGO_BRAKE = VerifaiRange(0.5, 1.0)

ADV_INIT_DIST = [0, 20]
param ADV_SPEED = VerifaiRange(7, 10)

param SAFETY_DIST = VerifaiRange(10, 20)
CRASH_DIST = 5
TERM_DIST = 70

#####
# AGENT BEHAVIORS #
#####

behavior EgoBehavior(traj):
  try:
    do FollowTrajectoryBehavior(
      target_speed=globalParameters.EGO_SPEED, trajectory=traj)
  interrupt when withinDistanceToAnyObjs(self, globalParameters.SAFETY_DIST):
    take SetBrakeAction(globalParameters.EGO_BRAKE)
  interrupt when withinDistanceToAnyObjs(self, CRASH_DIST):
    terminate

#####
# SPATIAL RELATIONS #
#####

intersection = Uniform(*filter(lambda i: i.is3Way, network.intersections))

egoLane = Uniform(*intersection.incomingLanes)
egoManeuver = Uniform(*filter(lambda m:
  m.type is ManeuverType.LEFT_TURN, egoLane.maneuvers))
egoTraj = [egoLane, egoManeuver.connectingLane, egoManeuver.endLane]
egoSpawnPt = OrientedPoint in egoLane.centerline

advManeuver = Uniform(*filter(lambda m:
  m.type is ManeuverType.STRAIGHT, egoManeuver.conflictingManeuvers))
advLane = advManeuver.startLane
advTraj = [advLane, advManeuver.connectingLane, advManeuver.endLane]
advSpawnPt = OrientedPoint in advLane.centerline

#####
# SCENARIO SPECIFICATION #
#####

ego = Car at egoSpawnPt,
  with blueprint MODEL,
  with behavior EgoBehavior(egoTraj)

adv = Car at advSpawnPt,
  with blueprint MODEL,
  with behavior FollowTrajectoryBehavior(
    target_speed=globalParameters.ADV_SPEED, trajectory=advTraj)

require EGO_INIT_DIST[0] <= (distance to intersection) <= EGO_INIT_DIST[1]
require ADV_INIT_DIST[0] <= (distance from adv to intersection) <= ADV_INIT_DIST[1]
terminate when (distance to egoSpawnPt) > TERM_DIST

```

Figure 6. SCENIC program used in Scenario 2

```

#####
# MAP AND MODEL #
#####

param map = localPath('.././.././../tests/formats/opendrive/maps/CARLA/Town03.xodr')
param carla_map = 'Town03'
model scenic.simulators.carla.model

#####
# CONSTANTS #
#####

MODEL = 'vehicle.lincoln.mkz2017'

param EGO_SPEED = VerifaiRange(2, 4)

param ADV_DIST = VerifaiRange(-25, -10)
param ADV_SPEED = VerifaiRange(7, 10)

BYPASS_DIST = [15, 10]
INIT_DIST = 50
TERM_TIME = 5

#####
# AGENT BEHAVIORS #
#####

behavior AdversaryBehavior():
  try:
    do FollowLaneBehavior(target_speed=globalParameters.ADV_SPEED)
  interrupt when withinDistanceToAnyObjs(self, BYPASS_DIST[0]):
    fasterLaneSec = self.laneSection.fasterLane
    do LaneChangeBehavior(
      laneSectionToSwitch=fasterLaneSec,
      target_speed=globalParameters.ADV_SPEED)
    do FollowLaneBehavior(
      target_speed=globalParameters.ADV_SPEED,
      laneToFollow=fasterLaneSec.lane) \
      until (distance to adv) > BYPASS_DIST[1]
    slowerLaneSec = self.laneSection.slowerLane
    do LaneChangeBehavior(
      laneSectionToSwitch=slowerLaneSec,
      target_speed=globalParameters.ADV_SPEED)
    do FollowLaneBehavior(target_speed=globalParameters.ADV_SPEED) \
      for TERM_TIME seconds
    terminate

#####
# SPATIAL RELATIONS #
#####

initLane = Uniform(*network.lanes)
egoSpawnPt = OrientedPoint in initLane.centerline

#####
# SCENARIO SPECIFICATION #
#####

ego = Car at egoSpawnPt,
  with blueprint MODEL,
  with behavior FollowLaneBehavior(target_speed=globalParameters.EGO_SPEED)

adv = Car following roadDirection for globalParameters.ADV_DIST,
  with blueprint MODEL,
  with behavior AdversaryBehavior()

require (distance to intersection) > INIT_DIST
require (distance from adv to intersection) > INIT_DIST
require always (ego.laneSection._fasterLane is not None)

```

Figure 7. SCENIC program used in Scenario 3

```

#####
# MAP AND MODEL #
#####

param map = localPath('.././.././../tests/formats/opendrive/maps/CARLA/Town05.xodr')
param carla_map = 'Town05'
model scenic.simulators.carla.model

#####
# CONSTANTS #
#####

MODEL = 'vehicle.lincoln.mkz2017'

EGO_INIT_DIST = [20, 25]
param EGO_SPEED = VerifaiRange(7, 10)
param EGO_BRAKE = VerifaiRange(0.5, 1.0)

ADV_INIT_DIST = [15, 20]
param ADV_SPEED = VerifaiRange(7, 10)

param SAFETY_DIST = VerifaiRange(10, 20)
CRASH_DIST = 5
TERM_DIST = 70

#####
# AGENT BEHAVIORS #
#####

behavior EgoBehavior(traj):
  try:
    do FollowTrajectoryBehavior(
      target_speed=globalParameters.EGO_SPEED, trajectory=traj)
  interrupt when withinDistanceToAnyObjs(self, globalParameters.SAFETY_DIST):
    take SetBrakeAction(globalParameters.EGO_BRAKE)
  interrupt when withinDistanceToAnyObjs(self, CRASH_DIST):
    terminate

#####
# SPATIAL RELATIONS #
#####

intersection = Uniform(*filter(lambda i: i.is4Way, network.intersections))

egoLane = Uniform(*intersection.incomingLanes)
egoManeuver = Uniform(*filter(lambda m:
  m.type is ManeuverType.STRAIGHT, egoLane.maneuvers))
egoTraj = [egoLane, egoManeuver.connectingLane, egoManeuver.endLane]
egoSpawnPt = OrientedPoint in egoLane.centerline

advLane = Uniform(*filter(lambda m:
  m.type is ManeuverType.STRAIGHT,
  egoManeuver.reverseManeuvers)
).startLane
advManeuver = Uniform(*filter(lambda m:
  m.type is ManeuverType.LEFT_TURN, advLane.maneuvers))
advTraj = [advLane, advManeuver.connectingLane, advManeuver.endLane]
advSpawnPt = OrientedPoint in advLane.centerline

#####
# SCENARIO SPECIFICATION #
#####

ego = Car at egoSpawnPt,
  with blueprint MODEL,
  with behavior EgoBehavior(egoTraj)

adv = Car at advSpawnPt,
  with blueprint MODEL,
  with behavior FollowTrajectoryBehavior(
    target_speed=globalParameters.ADV_SPEED, trajectory=advTraj)

require EGO_INIT_DIST[0] <= (distance to intersection) <= EGO_INIT_DIST[1]
require ADV_INIT_DIST[0] <= (distance from adv to intersection) <= ADV_INIT_DIST[1]
terminate when (distance to egoSpawnPt) > TERM_DIST

```

Figure 8. SCENIC program used in Scenario 4

```

#####
# MAP AND MODEL #
#####

param map = localPath('.././.././../tests/formats/opendrive/maps/CARLA/Town05.xodr')
param carla_map = 'Town05'
model scenic.simulators.carla.model

#####
# CONSTANTS #
#####

MODEL = 'vehicle.lincoln.mkz2017'

EGO_INIT_DIST = [20, 25]
param EGO_SPEED = VerifaiRange(7, 10)
param EGO_BRAKE = VerifaiRange(0.5, 1.0)

ADV_INIT_DIST = [5, 15]
param ADV_SPEED = VerifaiRange(7, 10)

param SAFETY_DIST = VerifaiRange(10, 20)
CRASH_DIST = 5
TERM_DIST = 70

#####
# AGENT BEHAVIORS #
#####

behavior EgoBehavior(traj):
  try:
    do FollowTrajectoryBehavior(
      target_speed=globalParameters.EGO_SPEED, trajectory=traj)
  interrupt when withinDistanceToAnyObjs(self, globalParameters.SAFETY_DIST):
    take SetBrakeAction(globalParameters.EGO_BRAKE)
  interrupt when withinDistanceToAnyObjs(self, CRASH_DIST):
    terminate

#####
# SPATIAL RELATIONS #
#####

intersection = Uniform(*filter(lambda i: i.is4Way, network.intersections))

egoLane = Uniform(*intersection.incomingLanes)
egoManeuver = Uniform(*filter(lambda m:
  m.type is ManeuverType.RIGHT_TURN, egoLane.maneuvers))
egoTraj = [egoLane, egoManeuver.connectingLane, egoManeuver.endLane]
egoSpawnPt = OrientedPoint in egoLane.centerline

advLane = Uniform(*filter(lambda m:
  m.type is ManeuverType.STRAIGHT,
  Uniform(*filter(lambda m:
    m.type is ManeuverType.STRAIGHT,
    egoLane.maneuvers)
  ).reverseManeuvers)
).startLane
advManeuver = Uniform(*filter(lambda m:
  m.type is ManeuverType.LEFT_TURN, advLane.maneuvers))
advTraj = [advLane, advManeuver.connectingLane, advManeuver.endLane]
advSpawnPt = OrientedPoint in advLane.centerline

#####
# SCENARIO SPECIFICATION #
#####

ego = Car at egoSpawnPt,
  with blueprint MODEL,
  with behavior EgoBehavior(egoTraj)

adv = Car at advSpawnPt,
  with blueprint MODEL,
  with behavior FollowTrajectoryBehavior(
    target_speed=globalParameters.ADV_SPEED, trajectory=advTraj)

require EGO_INIT_DIST[0] <= (distance to intersection) <= EGO_INIT_DIST[1]
require ADV_INIT_DIST[0] <= (distance from adv to intersection) <= ADV_INIT_DIST[1]
terminate when (distance to egoSpawnPt) > TERM_DIST

```

Figure 9. SCENIC program used in Scenario 5