

BARK: Open Behavior Benchmarking in Multi-Agent Environments

Julian Bernhard^{1*}, Klemens Esterle^{1*}, Patrick Hart^{1*}, Tobias Kessler^{1*}

Abstract—Predicting and planning interactive behaviors in complex traffic situations presents a challenging task. Especially in scenarios involving multiple traffic participants that interact densely, autonomous vehicles still struggle to interpret situations and to eventually achieve their own mission goal. As driving tests are costly and challenging scenarios are hard to find and reproduce, simulation is widely used to develop, test, and benchmark behavior models. However, most simulations rely on datasets and simplistic behavior models for traffic participants and do not cover the full variety of real-world, interactive human behaviors. In this work, we introduce BARK, an open-source behavior benchmarking environment designed to mitigate the shortcomings stated above. In BARK, behavior models are (re-)used for planning, prediction, and simulation. A range of models is currently available, such as Monte-Carlo Tree Search and Reinforcement Learning-based behavior models. We use a public dataset and sampling-based scenario generation to show the inter-exchangeability of behavior models in BARK. We evaluate how well the models used cope with interactions and how robust they are towards exchanging behavior models. Our evaluation shows that BARK provides a suitable framework for a systematic development of behavior models.

I. INTRODUCTION

Among other technical and legal challenges, a major problem in the development of autonomous vehicles is solving complex tasks in close interaction with humans. Human drivers have developed strategies to adapt to the behavior of other humans, implicitly embedding potential uncertainty about the behavior of surrounding participants into their own plan.

In our work, we see the behavior B_t^i of an agent i at time t as *its desired future* sequence of physical states encoding the agent’s strategy to reach a short-term goal, e.g. changing lane. A behavior may deviate from the executed motion in the environment due to errors in trajectory tracking or environmental influences. Further details are given in Section III.

A longstanding challenge in the development of decision-making for autonomous vehicles is to mimic interactive behavior when interacting with human drivers. Since most behavior generation algorithms do not consider interactions, the number of disengagements in interactive situations, such as urban areas, is especially large [1]. We refer to planners addressing this challenge as *interactive behavior generation* algorithms. To model interactivity, planners must employ some kind of prediction model of other agents. Though a variety of modeling methods exist [2], it remains unclear how

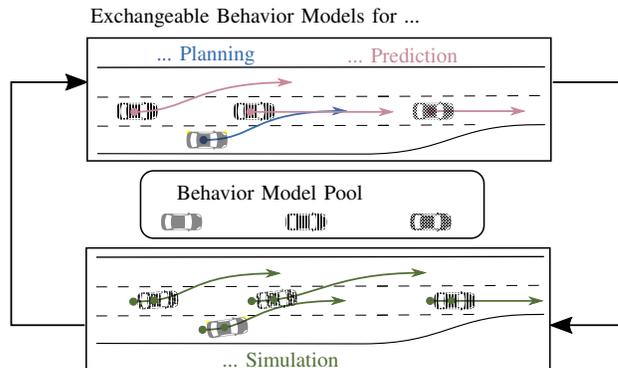


Fig. 1: The core concept of BARK as a development platform for behavior models: Reuse a behavior model interchangeably for planning, simulation, and prediction.

to guarantee their feasibility in domains like autonomous driving. Errors in the prediction due to incomplete and faulty models can cause huge problems in real-world driving situations.

This outlines the importance of and need for considering interactions in behavior generation. It also raises the need for algorithms – either conventional or learning-based – to be benchmarked on how well they generalize and cope with inaccurate prediction models. To tackle these problems, research must be conducted systematically in deterministic environments. It is generally accepted that vehicle dynamics simulation is a suitable tool for speeding up the development and verification of behavior planning algorithms. However, existing frameworks for simulating and benchmarking behavior models rarely provide sophisticated behavior models for other agents. Microscopic traffic simulators can model traffic flow to simulate groups of vehicles but neglect the interactions between them. To simulate interactive behavior, ingredients of both modeling methodologies – vehicle- and traffic-centered – are required. We will elaborate on this in Section II.

We claim that a white-box approach for the simulation of agents is the only way to compare multiple approaches on a scientific level. To systematically improve applied models in these domains, it shall be possible to use a behavior model from a modeling point of view to

- 1) plan the ego motion of the autonomous vehicle,
- 2) predict the motion of other, potentially human-driven, vehicles,
- 3) simulate an agent in a purely virtual environment.

Furthermore, a platform for behavior model development must support full reproducibility of experiments and pro-

¹fortiss GmbH, Research Institute of the Free State of Bavaria, Munich, Germany

*These authors contributed equally to this work.

vide a fast simulation loop. This is especially beneficial for machine-learning-based models that require thousands of training episodes. Fig. 1 sketches the idea of using behavior models interchangeably for planning, prediction, and simulation.

In this work, we introduce the simulation platform *BARK*, an acronym for *Behavior BenchmARK*, which addresses all the needs to develop interactive behavior planners and behavior simulation and prediction models as described in Section III. Besides the tool itself, which is available as open-source software¹, we contribute the methodology to develop and benchmark new behavior models. We also provide and evaluate reference behavior model implementations based on Reinforcement Learning and Monte-Carlo Tree Search, alongside other simplistic traffic model implementations from the literature. Details on this can be found in Section IV.

Besides the contribution to the overall goal of shifting the development efforts from field testing to simulation and modeling, we evaluate the benefits and highlight some of the core features of *BARK* in Section V.

II. RELATED WORK

We will now discuss relevant simulation and benchmarking environments. Most tools can, to some extent, simulate simple behavior models alongside an accurate motion model of the simulated agents.

CommonRoad² was established by Althoff *et al.* [3] as a composable benchmarking tool for motion planning on roads. It provides a collection of scenarios and benchmarks and aims to evaluate and compare different motion planning algorithms. With the focus on assessing the performance of motion planners and an easy-to-use interface for controlling one agent in the statically predefined scene, the tool aims to foster comparability between different motion planners. CommonRoad is fully open-source. However, CommonRoad uses only pre-recorded data for the other agents – only enabling non-interactive behavior planning. In *BARK*, the other agents actively interact with the ego vehicle.

The open-source simulator Carla³ has attracted attention recently. It fills the gap of the research community’s need for a high-fidelity simulator with modern software design and state-of-the-art visualization. Especially for sensor simulations and visualization of different agent types, the tool provides powerful resources. Palanisamy [4] has built a gym environment based on Carla. However, with Carla being based on the Unreal Game Engine, problems like non-determinism and timing issues are introduced, that we consider undesirable when developing and comparing behavior models. Contrary to the 1 Hz simulation frequency in Carla, our step environment can cover any type of frequency. The recently introduced Deepdrive⁴ by Voyage has similar drawbacks as Carla.

¹<https://github.com/bark-simulator/bark/>

²<https://commonroad.in.tum.de/>

³<http://carla.org/>

⁴<https://deepdrive.voyage.auto/>

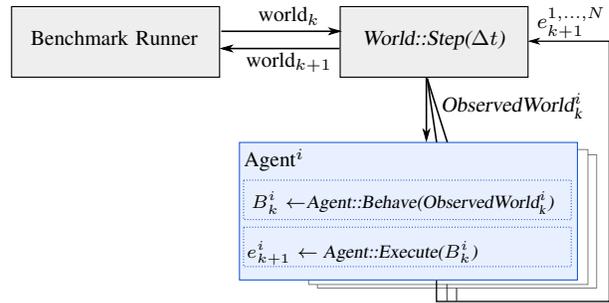


Fig. 2: *BARK*’s simulation loop is handled by the benchmark runner holding the current world state at discrete world time k . In each iteration, the benchmark runner calls $World::Step(\Delta t)$. This function generates an $ObservedWorld$ for each agent, and passes it to the agent’s internal $BehaviorModel$ which generates a behavior B_k^i . The behavior is passed to the agent’s $ExecutionModel$ calculating the next executed agent state e_{k+1}^i . The next world state at time $t_{k+1} = t_k + \Delta t$ integrates the updated agent states for all agents N and is returned to the *Benchmark Runner*.

A variety of commercial software suites are available. For an overview of existing automotive simulators, we refer to Kang *et al.* [5]. Most tools accompany specific industry needs and are highly configurable for specific use cases. However, they are not available as such to the research community and most require significant engineering effort to be adapted to these specific needs.

Microscopic traffic simulators such as SUMO⁵ by Lopez *et al.* [6] allow each agent to be simulated individually, and provide a realistic simulation of the traffic flow in a potentially large traffic network with numerous agents. These tools do not aim for benchmark capacities or tracking the accurate motion of each agent.

III. BARK IN A NUTSHELL

BARK focuses on providing a software framework for the systematic evaluation and improvement of behavior models. The same implementations can be used to either plan the ego motion of a vehicle, predict other vehicles’ motions, and forward simulate a driving scenario. For example, a traffic model, such as the Intelligent Driver Model [7] can on the one hand be used to populate a simulation with agents (cf. Fig. 1) but also as a generative model to predict other agents’ motion from the viewpoint of the ego vehicle. Experiments in *BARK* shall be fully reproducible, independent of the frequency at which the simulation runs. To ensure this, *BARK* models the world as a multi-agent system with agents performing *simultaneous movements* in the simulated world. At fixed, discrete world time-steps, each agent plans its behavior using an agent-specific behavior model, which only has access to the agent’s observed world, but not the simulator’s simulated world. The concept of simultaneous movement ensures that a behavior model can plan based on reproducible input information. It avoids timing artifacts that may occur in message-passing, middleware-based simulation architectures. Fig. 2 visualizes the core concept of *BARK*’s simulation model.

⁵<https://sumo.dlr.de/docs/>

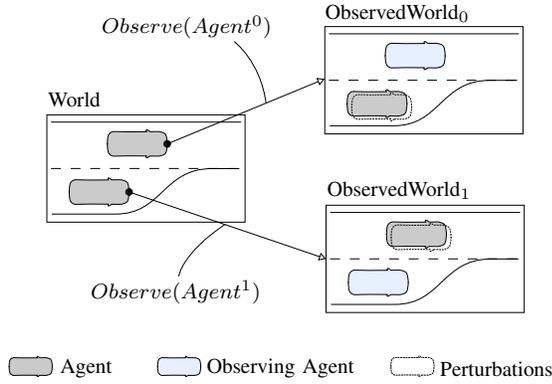


Fig. 3: Each agent in BARK uses an observed world for defining its behavior. In each observed world there is one observing agent (depicted in blue) from whose perspective the observation is being made. Perturbations can be introduced by e.g. exchanging the other agents' behavior models and model parameters in the observed world.

BARK uses behavior models not only for behavior planning, but also for predicting other agents in the world. For example, the observed world of each planner derives from the actual world. All agents in this observed world behave according to their prediction configuration. Fig. 3 visualizes BARK's observed world model. With the concept of deriving an observed world with various prediction configurations from the actual world definition, potential errors of the behavior planner caused by inaccurate internal prediction of other traffic participants can be systematically examined.

BARK provides several state-of-the-art behavior models ranging from conventional planning to machine learning approaches. BARK's current behavior model implementation will be discussed in Section IV. In the following, we take a more detailed look at other BARK components.

a) *World and ObservedWorld Model*: The BARK *World* model contains the map, all objects and agents. Static and dynamic objects are represented in the form of object lists. We will refer to dynamic objects as agents in the following.

The *ObservedWorld* model, on the other hand, reflects the world that is perceived by an agent i . BARK's *ObservedWorld* model accounts for the fact that the observing agent has no access to the true (world) behavior model of other agents. BARK can model different degrees of observability, by either completely restricting access to the world behavior model or only by perturbing the parameters of the world behavior model. Additionally, occlusions and sensor noise can be introduced in the *ObservedWorld* model.

b) *Agent Models*: As shown in Fig. 2, an agent in BARK provides two main interfaces:

- $B_k^i \leftarrow \text{Agent}::\text{Behave}(\text{ObservedWorld}_k^i)$: calls the agent-specific behavior model, which generates a behavior trajectory $B_k^i = (b_1^i, b_2^i, \dots, b_L^i)$, being a sequence of desired future physical agent states $b_l^i = (t, x, y, \theta, v)$ between current simulation world time t_k and at least the end time of the simultaneous movement $b_L^i(t) \geq$

$t_k + \Delta t$. The time discretization of the behavior trajectory can be arbitrary.

- $e_{k+1}^i \leftarrow \text{Agent}::\text{Execute}(B_k^i)$: calls the agent-specific execution model determining the agent's next state e_{k+1}^i in the world based on the generated behavior trajectory B_k^i . An additional execution model allows to systematically examine the robustness of behavior planners against execution errors. Since this is not the current focus of our work, we use an interpolation-based execution model. Models that account for controller inaccuracies will form the subject of future work.

To implement these two main interfaces, an agent holds the following additional agent-specific information:

- *GoalDefinition*: BARK considers agents to be goal-driven. It provides an abstract goal specification with various inherited types of agent goals, e.g. geometric goal regions or lane-based goals. Each agent contains a single goal definition instance.
- *RoadCorridor*: When an agent is initialized, it computes the set of roads and corresponding lanes useful to reach its goal. The topology information how roads are connected is extracted from the map. Also, the geometric information of the map such as lane boundaries are being discretized. This precomputation avoids computational overhead during simulation.
- *Polygon*: A 2D polygon defines the shape of the agent.

c) *Scenario*: A BARK scenario contains a list of agents with their initial states, behavior and execution models as well as a goal definition for each agent. Further, it contains a map file in the OpenDrive⁶ format. To support behavior benchmarking, each scenario specifies which agent is considered the 'controlled' agent during the simulation. A BARK scenario does not explicitly specify how agents will behave over time, e.g. using predefined maneuvers or trajectories. This concept allows us to simulate interactive scenarios in which other agents react to the controlled agent's behavior as well as to data-driven scenarios with fixed agent trajectories. Both interactive and data-driven scenarios can be combined, as we show in our evaluation.

d) *Scenario Generation*: BARK provides a scenario generation module for configuring sets of scenarios. The *ConfigurableScenario* generation used in this work contains source-sink pairs with additional parameter sets. Using the source and sink, a routing determines a lane on which the agents are placed. The parameter sets for a source-sink pair specify the distribution of agent states, their behavior and execution models as well as goal definitions. Using a configuration reader as well as a scenario generation module, the scenario is then built and ready to be executed. This modularized concept allows us to model a variety of scenario types and enables an easy extension with novel scenario types.

e) *Benchmarking*: For the systematic development of behavior models, BARK supports large scale evaluation of behavior models over a collection of scenario sets contained

⁶<http://www.opendrive.org/>

in a *benchmarking database*. Binary serialization of the database before experiments are started ensures that scenarios remain reproducible across different systems⁷.

BARK provides a benchmark runner to evaluate specific behavior models with different parameter configurations over the entire benchmarking database. The evaluation is based on an abstract evaluator interface calculating a Boolean, integer or real-valued metric based on the current simulation world state. The current evaluators available in BARK are:

- **StepCount**: returns the step count the scenario is at.
- **GoalReached**: checks if a controlled agent’s *GoalDefinition* is satisfied.
- **DrivableArea**: checks whether the agent is inside its *RoadCorridor*.
- **Collision(ControlledAgent)**: checks whether any agent or only the currently controlled agent collided.
- **GoalDistance**: calculates an euclidean distance to the *GoalDefinition*.

Evaluators can be used not only for benchmarking, but can also internally by the behavior models, e.g. for the reward calculation in search or reinforcement learning-based planners.

The *BenchmarkRunner* runs each scenario of the database evaluating world states of the simulation. It terminates the scenario run based on criteria defined with respect to the evaluators. The results are dumped into a PANDAS data frame⁸. BARK also provides a *distributed benchmark runner* based on RAY⁹ to perform evaluations in parallel on multiple cores and clusters. This is especially beneficial for computationally expensive planning algorithms.

f) *Software Design*: BARK has a monolithic, single-threaded core, written in C++. The core of BARK is wrapped in Python using PYBIND¹⁰ including pickling support for all C++ types. Scenario generation and benchmarking, as well as service methods for parameter handling and visualization, are implemented in Python.

A simulator planning cycle is entirely deterministic, which enables the simulation and experiments to be reproducible. We argue this to be essential to conduct systematic research in the field of behavior modeling.

The BARK ecosystem is split over multiple Github repositories¹¹. We use Bazel¹² as the build system. Its sandboxed build environment simplifies the reproducibility of experiments since dependency versions can be tracked easily over multiple repositories.

IV. BEHAVIOR MODELS

BARK contains a set of reference implementations for behavior models from the literature. All behavior models

⁷Though BARK supports adjusting random seeds in the scenario generation, differing implementations and versions of pseudorandom number generators across systems may yield differing scenarios using the same random seed.

⁸<https://pandas.pydata.org>

⁹<https://ray.readthedocs.io>

¹⁰<https://github.com/pybind>

¹¹<https://github.com/bark-simulator/>

¹²<https://bazel.build/>

inherit from the abstract BehaviorModel class:

```
class BehaviorModel {
public:
    ...
    virtual Trajectory Plan(float delta_time,
        const ObservedWorld& world) = 0;
    ...
};
```

Behavior models overload the ‘Plan’ function and return a time-dependent state trajectory. Since all behavior models share the same interface, this makes them easily exchangeable.

By accessing the ObservedWorld, a behavior model can query egocentric semantic information, e.g. *GetLeftLane()*→*GetCenterLine()*, as well as semantic relations to other agents, e.g. *GetAgentInFront()*. Together with the *Evaluator* concept and implemented standard behavior models, this allows for rapid prototyping and development of new behavior models in BARK. We list and highlight some of the implemented behavior models in this section.

a) *Intelligent Driver Model*: The very popular Intelligent Driver Model (IDM) by Treiber *et al.* [7] is an easy-to-implement vehicle-following model often used in microscopic traffic simulation. To model maintaining a gap from the vehicle in front, it leverages the spacing between the agents, speed asymmetries, and realistic braking profiles. The model cannot handle lane changes, intersections, and unstructured scenarios. Several extensions of the model exist [8, 9] that consider limitations of human drivers, as well as adaptive cruise control behavior. To support larger world step times, our IDM implementation assumes that the preceding vehicle maintains a constant velocity throughout a world simulation step.

b) *MOBIL Model*: For multi-lane scenarios, we implemented the MOBIL model, a decision-making entity to trigger lane changes to improve traffic flow [10]. It maintains a politeness factor, which captures how much other agents are slowed down by a potential lane change. In case of a lane change decision, we start tracking the center line of the respective left or right driving corridor. The longitudinal motion is controlled using the IDM.

c) *Reinforcement Learning Model*: Reinforcement Learning (RL) learns behaviors implicitly by collecting experiences and updating its policy to maximize the long-term cumulative expected reward. RL can be divided into three main categories: Actor-only, Critic-only, and Actor-Critic (AC) methods [11]. In particular AC-RL, such as the Proximal-Policy-Optimization (PPO) and the Soft-Actor-Critic (SAC) algorithm, have been shown empirically to work well [12, 13] for continuous control problems. In this work, we use the SAC algorithm that uses a stochastic policy $\pi(\underline{s}) = \underline{a}$ with the state \underline{s} and action \underline{a} . The SAC agent acts and collects experiences in BARK and improves its policy iteratively. The other agents may have arbitrary behavior models, such as the IDM model or also use an RL model. We use ‘Observers’ that transform the semantic BARK environment into a representation suitable for RL. The reward r is

calculated using ‘Evaluators’. These modules are available in our Machine Learning module¹³. As it integrates the standard OpenAI Gym-interface¹⁴, various popular RL libraries, such as TF-Agents¹⁵ can be easily integrated used with BARK.

d) Multi-Agent Monte Carlo Tree Search Model: Monte Carlo Tree Search (MCTS) is a well-known tree search method for finding approximate solutions to sequential decision problems via sampling [14]. It has been adapted to interactive driving by using information sets assuming simultaneous, multi-agent movements of traffic participants [15, 16]. They apply it to the context of cooperative planning, meaning that they introduce a cooperative cost function, which minimizes the costs for all agents. Similar to [15], our action set consists of lane changing and lane keeping (CV, CA, CD¹⁶). The MCTS algorithm iteratively performs selection, expansion, rollout, and backpropagation in each search iteration. As a selection strategy, we are currently using Upper Confidence Trees (UCT). After expanding a new node, we obtain an instance of ObservedWorld. During rollout, we randomly select actions for all interacting agents.

e) Single-Agent Monte Carlo Tree Search Model: Similar to the Multi-Agent MCTS implementation, we implement a Single-Agent MCTS, in which the other agents are assumed to follow some kind of modeled behavior. During expansion and rollout, we forward-simulate the observed world using a behavior model for each agent. These are set via the prediction setup of the observed world before the MCTS algorithm is started. Any behavior model from BARK with varying parameter specifications can be used. During expansion, the tree policy is used. During rollout, we randomly select actions for the ego agent.

f) Dataset Tracking Model: To get real-world data into the simulation, BARK contains an agent model that tracks recorded trajectories as close as possible. We can start the simulation at an arbitrary timestamp and include all or only a subset of recorded agents. We evaluate this using the INTERACTION dataset [17].

V. EVALUATION

We demonstrate the capabilities of BARK in two use cases:

- 1) Employing BARK **behavior models for behavior prediction** using the sampling-based scenario generation: We benchmark the effect of an inaccurate prediction model on the performance of an MCTS and RL-based planner.
- 2) Employing BARK **behavior planners for behavior simulation** using dataset-based scenario generation: We benchmark how planners perform when replacing human drivers in recorded traffic scenarios.

Our evaluation demonstrates the advantages of BARK for development, systematic improvement and large scale bench-

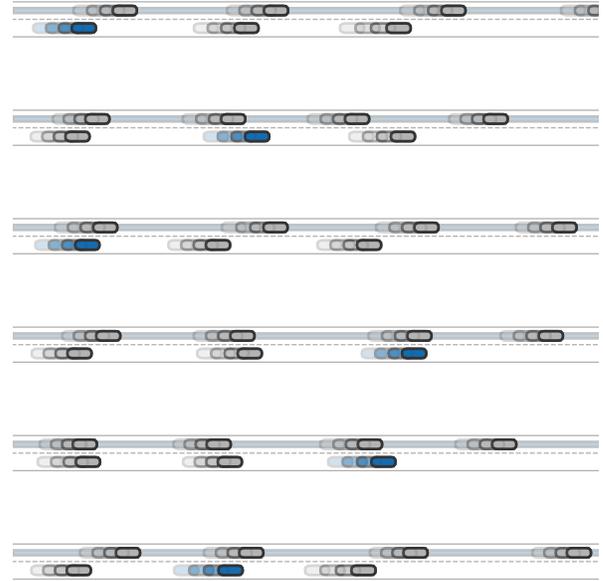


Fig. 4: The world states of the lane change task after the first four time steps ($t = 0.8\text{ s}$, $\Delta t = 0.2\text{ s}$). Scenarios are generated with BARK’s sampling-based scenario generation. The benchmarked agent (blue) must change to the left lane (goal definition is depicted in light blue). Past agent states are shown with increasing transparency. The driving direction is from the left to the right.

marking of behavior models, simulation, and prediction algorithms.

A. Behavior Models for Behavior Prediction

The performance of online planning algorithms, such as the MCTS, depends greatly on the accuracy of the prediction models used internally. RL can be considered as an offline planning algorithm – not relying on a prediction model but requiring a training environment to learn an optimal policy beforehand. The inaccuracy of prediction relates to the amount of behavior model inaccuracy between training and evaluation. In the following, we show to evaluate both sources of errors - inaccurate prediction models and training data - under a common framework.

In our experiment, we consider lane changing in crowded urban roads as depicted in Fig. 4. We evaluate the performance of MCTS and RL planners statistically, over a large set of scenarios. The scenarios are generated as outlined in Section III. We employ uniform sampling of longitudinal distances from the range [20 m, 30 m] and velocities from the range [40 km/h, 60 km/h]. All vehicles except the ego vehicle are controlled using the IDM. In the benchmark database, the following configurations are stored:

- 0% variation: Vehicles in this scenario use the same IDM model parameters as employed by the MCTS for behavior prediction and RL for training. The IDM parameter time headway, accounting for a velocity-dependent safety distance, is set to $T_{\text{Head}} = 3\text{ s}$.

¹³<https://github.com/bark-simulator/bark-ml>

¹⁴<https://github.com/openai/gym>

¹⁵<https://github.com/tensorflow/agents>

¹⁶Constant Velocity, Constant Acceleration, Constant Deceleration

- 20%, 40%, and 80% variation: In three additional scenario sets, we decrease T_{Head} by the given percentages, respectively, e.g. $T_{\text{Head}}^{40\%} = 1.8s$.

In the experiment, BARK’s parallelized benchmark runner runs 600 scenarios for each scenario set. The following planning configurations are applied to the ego vehicle:

- Single-Agent MCTS for 2k, 4k, and 8k search iterations
- SAC policy trained with RL over scenario set 0% variation until results were stable (120k episodes, 6h)

Fig. 5 gives the percentages of each scenario run with the controlled agent reaching the goal, colliding, or exceeding the maximum allowed simulation time steps (>30), as well as the average number of simulation steps in the event that the goal is reached. It is clear that the collision rate rises as the prediction error increases. For 2k iterations, the MCTS fails to explore worst-case outcomes explaining the high success but also high collision rate. For 4k iterations, the worst-case outcomes are explored, yet no goal-reaching plan is found, yielding low collision but a large maximum number of steps. With a higher exploration at 8k iterations, the success rate increases greatly, with a slight increase in collisions. The RL, SAC model outperforms the MCTS model for 0% variation in terms of the success rate. However, when the other agent’s behavior is different from that used in training, the collision rate rises more quickly.

Our experiment showed that a *systematic* evaluation of behavior planners, modeling limited knowledge about other agents’ behavior is necessary to uncover and understand the weaknesses of both conventional and learning-based approaches in such settings. BARK covers the whole development life cycle to tackle this open problem.

B. Behavior Planners for Behavior Simulation

Planning models should be evaluated using realistic real-world data. However, when a planner is inserted into recorded scenarios, others will keep the behavior as specified in the dataset, yielding an open-loop simulation. Since this restricts the evaluation of interactive scenarios, we require

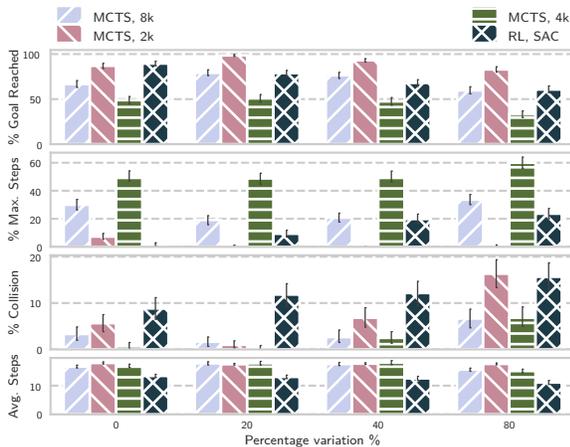


Fig. 5: Statistical benchmark of the MCTS and RL planners for increasing percentages of prediction error.

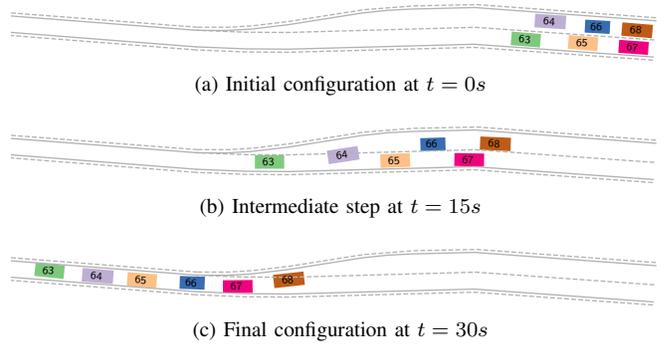


Fig. 6: Replay of the INTERACTION dataset at a simulation step size $\Delta t = 0.1s$.

behavior models to adapt to other dataset-based agent models. We evaluate the suitability of BARK’s existing behavior models for this use case by inserting them into a recorded traffic scenario without fine-tuning the model parameters in any way.

Datasets in the prediction and planning community usually consist of fused object lists. In contrast to current datasets [18, 19], the INTERACTION dataset [17] also provides maps, which are essential for most on-road planning approaches. Here, we analyze the two-way merge scenario DR_DEU_Merging_MT¹⁷. Fig. 6 shows the original scenario at the initial, the intermediate, and the final time-step. The scenario originates in Germany, where vehicles must follow the zip-merge principle. The scenario is particularly challenging as the agents not only have to master a merge, but also maintain the order given by the zip-merge rule. Otherwise collisions may occur.

In Table I, we evaluate whether a collision occurred when an agent from the dataset was replaced with various agent behavior models. We denote the set of replaced agents by \mathcal{A}_{\square} . In the evaluation, we do not account for all possible combinations of agents for replacement, but focus on six challenging ones. We establish that, with the proposed tooling, critical scenarios can be analyzed easily. Algorithm fine-tuning is considered a subsequent step. We compare the performance of the IDM, MOBIL, and Single Agent MCTS described in Section IV. All three models use IDM parameters, as the longitudinal behavior of the MOBIL model is controlled using IDM, while Monte Carlo Tree Search uses an IDM as a prediction model of the other agents. We thus define four different parameter sets \mathcal{P}_{1-4} for the IDM, see Table II. As we aim to provide a methodology for evaluation, we only vary the minimum distance, without fine-tuning of other model parameters for this single scenario. The IDM is a pure lane-following model and thus not capable to perform lane changes. We thus do not replace agent sets $\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_6$ with the IDM.

We observe that, without fine-tuning model parameters, most behavior models fail to coexist next to replayed agents. The IDM can partially solve for \mathcal{A}_4 , as it needs to retain

¹⁷The scenario starts at 232.000s

sufficient gap distance from the preceding vehicle to let merging vehicles in. This is crucial, as the order must be retained to prevent crashes further on in the scenario. MOBIL can partially solve \mathcal{A}_0 , where the agent has to change to the left lane. If the gap distance parameter is too small, the agent will try to accelerate and change lanes before agent 65. However, even if the lane change is executed successfully, a collision occurs at the end of the scenario, due to the incorrect order of the vehicles. This shows the necessity to develop behavior models for simulation that are capable of handling more complex constraints, such as maintaining a specific order.

Although some MOBIL configurations can master \mathcal{A}_4 without collision, the overall scenario is changed, as sometimes lane changes are triggered to go from left to right. Collision metrics often do not account for the changed scenario. A wide range of research has been carried out in the field of prediction [20] that quantifies the measure of similarities between recorded and modeled behavior. We will address this in future work. All prediction parameter sets used in the MCTS successfully master the lane change of \mathcal{A}_1 . Note that we do not constrain the final order, therefore replacing agents in the left lane (65, 67) fails.

We conclude that current rule-based models (IDM, MOBIL) perform poorly in highly dense scenarios, as they do not model obstacle avoidance based on prediction or future interaction. More elaborate models should be investigated, as these allow constraints such as ordering and traffic rules to be modeled. MCTS can be used, but without an accurate model of the prediction, it also leads to crashes, similar to the evaluations in Section V-A.

VI. NEXT STEPS

Besides further development of the platform and fostering a community for the systematic improvement of behavior models, we consider the following topics to be the most relevant.

A. Behavior Modeling

Integrating a more comprehensive set of agent models, e.g. based on the formalization of traffic rules [21] as planner and evaluator would be a big benefit for the platform. Modern optimization-based methods, e.g. mixed-integer programming [22], have proven to be computationally fast for motion and strategic planning, and can act as a dual approach

TABLE I: Agents from the dataset are replaced by a behavior model. \diamond indicates no collision occurred. \blacklozenge indicates a collision occurred. \mathcal{A} denotes the set of agents we replace. \mathcal{P} refers to the respective IDM parameter set being used.

\mathcal{A}	Replaced IDs	IDM				MOBIL				MCTS			
		\mathcal{P}_1	\mathcal{P}_2	\mathcal{P}_3	\mathcal{P}_4	\mathcal{P}_1	\mathcal{P}_2	\mathcal{P}_3	\mathcal{P}_4	\mathcal{P}_1	\mathcal{P}_2	\mathcal{P}_3	\mathcal{P}_4
\mathcal{A}_0	66	-	-	-	-	\blacklozenge							
\mathcal{A}_1	68	-	-	-	-	\blacklozenge	\blacklozenge	\blacklozenge	\blacklozenge	\diamond	\diamond	\diamond	\diamond
\mathcal{A}_2	66, 68	-	-	-	-	\blacklozenge							
\mathcal{A}_3	65	\blacklozenge											
\mathcal{A}_4	67	\blacklozenge	\blacklozenge	\diamond	\diamond	\blacklozenge							
\mathcal{A}_5	65, 67	\blacklozenge											
\mathcal{A}_6	66, 68, 65, 67	-	-	-	-	\blacklozenge							

TABLE II: Parameters of the IDM

Parameter		\mathcal{P}_1	\mathcal{P}_2	\mathcal{P}_3	\mathcal{P}_4
Desired velocity	v_0 [m/s]	5	5	5	5
Maximum accel.	a_{max} [m/s ²]	1.7	1.7	1.7	1.7
Time gap	τ [s]	1	1	1	1
Comfortable decel.	b [m/s ²]	1.7	1.7	1.7	1.7
Minimum distance	s [m]	2	3	4	5

compared to learning-based approaches. A combination of classical and learning-based methods [23] is computationally fast and achieves safe and comfortable motions.

Recently imitation learning has shown great benefits for agent modeling [24, 25]. Such behavior models promise to have great benefits, especially when used in the simulation. Robustness against inaccuracies in prediction should be further evaluated, e.g. by advancing risk-aware planning [26].

B. Agent Types

As of now, only vehicle models are implemented in BARK. However, other agent types can easily be integrated into BARK due to its highly modular and abstract architecture. Having other types of agents, such as cyclists or pedestrians would provide additional benefits. For pedestrians, a variety of datasets and prediction algorithms already exist [27].

C. Scenario Languages

We did not choose to base the scenario definition on a standard like OpenSCENARIO¹⁸, as the introduced complexity is not relevant for our use case. Alternative approaches, such as the work by Damm *et al.* [28] modeling scenarios as sequence charts, do not account for behavior model aspects. However, having interchangeable scenario definitions between different simulation tools would be highly beneficial.

D. Co-Simulation with Other Environments

BARK is not designed to replace any existing simulation tool, but its features can enhance other environments and vice versa. We aim to publish an interface with Carla to benefit from its high-fidelity visualization and contribute sophisticated agent behavior. Also supporting state-of-the-art robotics and automotive middlewares would be beneficial to ease the exchange of components.

VII. CONCLUSION AND FUTURE WORK

In this work, we introduced the behavior benchmarking and simulation tool BARK, an environment tailored for developing interactive behavior models. Behavior models in BARK are exchangeable, which allows them to be used for planning, prediction, and simulation.

In our evaluation, we analyzed how robust the Monte-Carlo Tree Search and Reinforcement Learning behavior models are against prediction inaccuracy. With BARK's data-set tracking model, we evaluated whether sophisticated behavior planners, as well as simplistic traffic models, can drive accurately together with humans in recorded scenarios.

¹⁸<http://www.openscenario.org/>

Both experiments demonstrated the usefulness of having interchangeable behavior models in BARK.

However, the evaluation revealed that more systematic efforts must be made to understand the circumstances under which state-of-the-art behavior models can robustly accomplish prediction, simulation, and planning tasks. BARK is a platform that enables these efforts.

BARK is open-source under the MIT license and can be applied in a wide range of applications and scenarios. We welcome contributions from other researchers in this field.

ACKNOWLEDGMENT

This research was funded by the Bavarian Ministry of Economic Affairs, Regional Development and Energy, project Dependable AI and supported by the Autonomous Intelligent Driving GmbH.

REFERENCES

- [1] F. Favarò, S. Eurich, and N. Nader, "Autonomous vehicles' disengagements: Trends, triggers, and regulatory limitations," *Accident; analysis and prevention*, 2017.
- [2] S. Albrecht and P. Stone, "Autonomous agents modelling other agents: A comprehensive survey and open problems," *Artificial Intelligence*, vol. 258, Sep. 2017.
- [3] M. Althoff, M. Koschi, and S. Manzing, "CommonRoad: Composable benchmarks for motion planning on roads," *IEEE Intelligent Vehicles Symposium, Proceedings, ITSC*, pp. 719–726, 2017.
- [4] P. Palanisamy, "Multi-Agent Connected Autonomous Driving using Deep Reinforcement Learning," 2019.
- [5] Y. Kang, H. Yin, and C. Berger, "Test Your Self-Driving Algorithm: An Overview of Publicly Available Driving Datasets and Virtual Testing Environments," *IEEE Transactions on Intelligent Vehicles*, vol. 4, no. 2, pp. 171–185, 2019.
- [6] P. A. Lopez, M. Behrisch, L. Bieker-Walz, *et al.*, "Microscopic Traffic Simulation using SUMO," *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, vol. 2018-Novem, pp. 2575–2582, 2018.
- [7] M. Treiber, A. Hennecke, and D. Helbing, "Congested Traffic States in Empirical Observations and Microscopic Simulations," *Phys. Rev. E*, vol. 62, no. 2, pp. 1805–1824, 2000.
- [8] M. Treiber, A. Kesting, and D. Helbing, "Delays, inaccuracies and anticipation in microscopic traffic models," *Physica A: Statistical Mechanics and its Applications*, vol. 360, no. 1, pp. 71–88, 2006.
- [9] A. Kesting, M. Treiber, and D. Helbing, "Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4585–4605, 2010.
- [10] A. Kesting, M. Treiber, and D. Helbing, "General lane-changing model MOBIL for car-following models," *Journal of the Transportation Research Board*, 2007.
- [11] J. N. Konda, Vijay R and Tsitsiklis, "On Actor-critic algorithms," *SIAM journal on Control and Optimization*, vol. 42, pp. 1143–1166,
- [12] J. Schulman, F. Wolski, P. Dhariwal, *et al.*, "Proximal Policy Optimization Algorithms,"
- [13] T. Haamoja, A. Zhou, P. Abbeel, *et al.*, "Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor," in *Proceedings of the 35th International Conference on Machine Learning (ICML), Stockholm, Sweden, 2018*, vol. 80, PMLR, 2018, pp. 1856–1865.
- [14] C. B. Browne, E. Powley, D. Whitehouse, *et al.*, "A Survey of Monte Carlo Tree Search Methods," *IEEE Transactions on Computational Intelligence and AI in Games*, 2012, 01024.
- [15] D. Lenz, T. Kessler, and A. Knoll, "Tactical Cooperative Planning for Autonomous Vehicles using MCTS," in *IEEE Intelligent Vehicles Symposium (IV)*, IEEE, Jun. 2016, pp. 447–453.
- [16] C. Paxton, V. Raman, G. D. Hager, *et al.*, "Combining neural networks and tree search for task and motion planning in challenging environments," in *International Conference on Intelligent Robots and Systems*, 00011, IEEE, 2017.
- [17] W. Zhan, L. Sun, D. Wang, *et al.*, "INTERACTION Dataset: An International, Adversarial and Cooperative motion Dataset in Interactive Driving Scenarios with Semantic Maps," 2019.
- [18] V. Alexiadis, J. Colyar, J. Halkias, *et al.*, "The next generation simulation program," *ITE Journal (Institute of Transportation Engineers)*, vol. 74, no. 8, pp. 22–26, 2004.
- [19] R. Krajewski, J. Bock, L. Kloeker, *et al.*, "The highD Dataset: A Drone Dataset of Naturalistic Vehicle Trajectories on German Highways for Validation of Highly Automated Driving Systems," *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, vol. 2018-Novem, pp. 2118–2125, 2018.
- [20] W. Zhan, A. L. de Fortelle, Y.-T. Chen, *et al.*, "Probabilistic Prediction from Planning Perspective: Problem Formulation, Representation Simplification and Evaluation Metric," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, IEEE, Jun. 2018, pp. 1150–1156.
- [21] K. Esterle, V. Aravatinos, and A. Knoll, "From Specifications to Behavior: Maneuver Verification in a Semantic State Space," in *IEEE Intelligent Vehicles Symposium*, May 2019, pp. 2140–2147.
- [22] T. Kessler and A. Knoll, "Cooperative Multi-Vehicle Behavior Coordination for Autonomous Driving," in *IEEE Intelligent Vehicles Symposium*, 2019, pp. 1696–1703.
- [23] P. Hart, L. Rychly, and A. Knoll, "Lane-Merging Using Policy-based Reinforcement Learning and Post-Optimization," *2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019*, pp. 3176–3181, 2019.
- [24] M. Bansal, A. Krizhevsky, and A. Ogale, "ChauffeurNet: Learning to Drive by Imitating the Best and Synthesizing the Worst," pp. 1–20, 2019.
- [25] F. Behbahani, K. Shiarlis, X. Chen, *et al.*, "Learning from demonstration in the wild," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2019-May, pp. 775–781, 2019.
- [26] J. Bernhard and A. Knoll, "Addressing inherent uncertainty: Risk-Sensitive behavior generation using Distributional Reinforcement Learning," in *IEEE Intelligent Vehicles Symposium*, 2019.
- [27] D. Ridel, E. Rehder, M. Lauer, *et al.*, "A Literature Review on the Prediction of Pedestrian Behavior in Urban Scenarios," in *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, 2018, pp. 3105–3112.
- [28] W. Damm, S. Kemper, E. Möhlmann, *et al.*, "Using Traffic Sequence Charts at the Development of HAVs," in *9th European Congress on Embedded Real Time Software and Systems (ERTS 2018)*, 2018.