
UMBRELLA: Uncertainty-Aware Model-Based Offline Reinforcement Learning Leveraging Planning

Christopher Diehl, Timo Sievernich, Martin Krüger, Frank Hoffmann, Torsten Bertram*
TU Dortmund University
christopher.diehl@tu-dortmund.de

Abstract

Offline reinforcement learning (RL) provides a framework for learning decision-making from offline data and therefore constitutes a promising approach for real-world applications as automated driving. Self-driving vehicles (SDV) learn a policy, which potentially even outperforms the behavior in the sub-optimal data set. Especially in safety-critical applications as automated driving, explainability and transferability are key to success. This motivates the use of model-based offline RL approaches, which leverage planning. However, current state-of-the-art methods often neglect the influence of aleatoric uncertainty arising from the stochastic behavior of multi-agent systems. This work proposes a novel approach for Uncertainty-aware Model-Based Offline REinforcement Learning Leveraging plAnning (UMBRELLA), which solves the prediction, planning, and control problem of the SDV jointly in an interpretable learning-based fashion. A trained action-conditioned stochastic dynamics model captures distinctively different future evolutions of the traffic scene. The analysis provides empirical evidence for the effectiveness of our approach in challenging automated driving simulations and based on a real-world public dataset.

1 Introduction

Safe and interpretable motion prediction, planning, and control are crucial components for automated driving. Here, current decision-making systems are challenged by modeling multi-agent interactions. Often engineers manually design driving policies for specific scenarios, which does not scale well for more complex task. In contrast, RL [Sutton and Barto, 2018] learns a policy with little design and engineering effort but comes with the downside of trial-and-error learning. Hence, offline RL [Levine et al., 2020] is a promising research direction as it takes advantage of previously collected datasets and is applicable in safety-critical systems as robotics or automated driving. Moreover, in contrast to imitation learning approaches [Ross et al., 2011; Codevilla et al., 2018, 2019], which can only recover the (sub-optimal) behavior exhibited by the dataset, offline RL methods can leverage and improve on different types of data quality.

Applying offline RL approaches in real-world domains faces some practical challenges. Model-free Kumar et al. [2020]; Yin et al. [2021] and model-based policy learning methods [Kidambi et al., 2020; Yu et al., 2020] have low interpretability, which is important for designing safety layers [Pek et al., 2020] in automated driving. Moreover, they lack control flexibility. For example, adding new constraints often requires an expensive retraining of these model.

In contrast, the model-based offline planning (MBOP) [Argenson and Dulac-Arnold, 2021] framework addresses the above challenges in an comprehensive manner. The algorithm utilizes different offline

*The authors are with the Institute of Control Theory and Systems Engineering, TU Dortmund University, D-44227, Germany.

learned models to plan an optimal action using model-predictive control (MPC). It allows for a simple extension of the reward function and the incorporation of state constraints. Moreover, planning with the learned dynamics model enhances interpretability, which is crucial for the testing and deployment of SDV’s. While showing promising results on single-agent control tasks, MBOP only uses a simple deterministic dynamics model. However, this neglects the stochasticity of the underlying process, also known as *aleatoric* uncertainty [Kiureghian and Ditlevsen, 2009] arising from uncertainty about the behavior of other traffic agents in automated driving. In addition, MBOP operates in a fully observable setting. However, this assumption does not hold in the real-world driving setting. For example, the intent of human drivers is not directly observable and can only be estimated indirectly from observations.

This work tackles the previously described problems and contributes in the following way: First, it proposes UMBRELLA, a model-based offline planning approach, which learns from offline data and plans considering both *epistemic* and *aleatoric* uncertainty while operating in the partially observable setting. The proposed method solves the prediction, planning, and control problem of the SDV using interpretable representations. Second, we introduce an ablation of UMBRELLA, which optimizes for the worst-case model. Third, we demonstrate that our approach consistently outperforms behavior cloning (BC) and the state-of-the-art model-based offline planning algorithm MBOP in urban and highway automated driving scenarios with dense traffic.

2 Related Work

The proposed approach is situated within the broader literature on model-based offline RL and interaction-aware motion prediction and planning. This section provides a brief description of the underlying concepts and how they relate to the work presented here.

Model-based Offline Reinforcement Learning. As in the offline RL setting, the agents do not interact with the environment during learning, one major challenge arises: the *distributional shift* Levine et al. [2020]. Model-based offline RL methods like MOREL and MOPO [Kidambi et al., 2020; Yu et al., 2020] address this issue by incorporating an epistemic uncertainty estimate of the dynamics model into the reward function to penalize states which are not covered by the behavior distribution. [Yu et al., 2021] penalizes the value function on out-of-distribution states, making the challenging uncertainty estimation of MOREL and MOPO obsolete. MBOP uses planning based on bootstrap ensembles [Lakshminarayanan et al., 2017] of offline learned models [Argenson and Dulac-Arnold, 2021]. MOPP improves upon MBOP by pruning trajectories to avoid potential out-of-distribution samples [Zhan et al., 2021]. Kahn et al. [2021] introduce BADGR, a system for self-supervised robot navigation learned from off-policy data. All these model-based approaches are evaluated in single-agent environments. Hence, they lack the modeling of aleatoric uncertainty, which is caused by uncertain human behavior in the interactive automated driving setting. [Henaff et al., 2019] address this issue by a stochastic dynamics model represented by a conditional variational autoencoder (CVAE) [Kingma and Welling, 2014]. However, their approach relies on policy learning, which in contrast to model-based offline planning comes with a reduced interpretability and control flexibility.

Interaction-aware Motion Prediction and Planning. Traditional automated driving pipelines follow the concept of planning a safe motion based on the prediction of all other agents, which neglects the interaction between both planning and prediction. Game-theoretic approaches [Cleac’h et al., 2020, 2021] model the multi-agent dynamics as games but come with the challenge of computing (Nash) equilibria, and their computation times scale poorly with the number of agents. In contrast, learning-based approaches as [Liu et al., 2021; Rhinehart et al., 2021] have the potential to generalize to a higher number of scenarios and learn from high-dimensional sensor data. However, in these works, epistemic uncertainty is not explicitly considered, which could lead to failure in out-of-distribution states.

3 The UMBRELLA Framework

The proposed algorithm is a model-based offline RL algorithm which learns from previously recorded datasets. UMBRELLA learns a stochastic dynamics model, a BC policy, and a truncated value function as shown in Figure 1 a). UMBRELLA is an extension of the MBOP [Argenson and Dulac-Arnold, 2021] method and plans for different future evolutions. Each model is a bootstrap ensemble

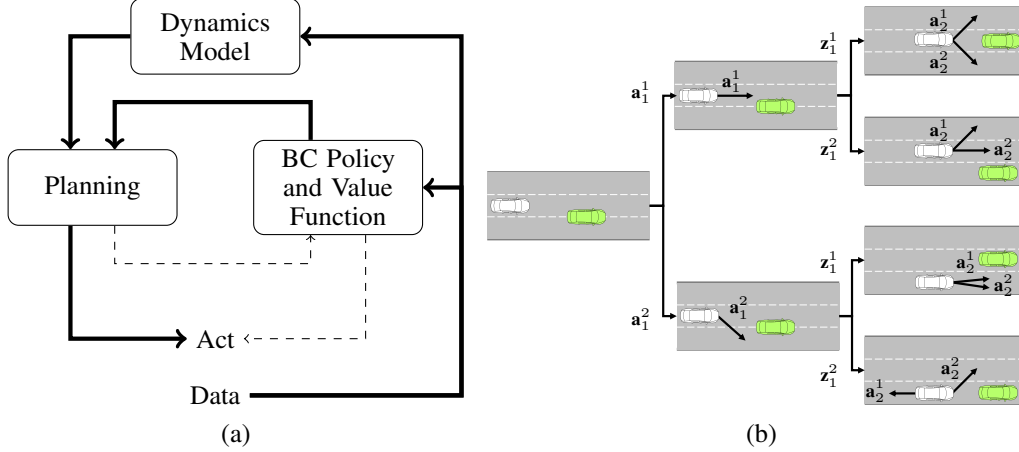


Figure 1: (a) UMBRELLA’s integration of learning and planning following the general visualization of Moerland et al. [2021]: Model learning includes the dynamics, BC policy, and value function. The thick arrows indicate which parts of the planning/learning loop are used here. There is no arrow from *Act* to *Data* due to offline learning. b) One step of the planning procedure. For the sake of clarity, the number of samples equals $N = 2$. Actions and latent variables are indexed by time step t and sample n , resulting in \mathbf{a}_t^n and \mathbf{z}_t^n , respectively.

[Lakshminarayanan et al., 2017] of K neural networks. The weights of the each ensemble head f^i with $i \in [1, \dots, K]$ are initialized differently but get trained on the same dataset D .

3.1 Problem Formulation

The RL task is to control a dynamical control, described by a *Markov Decision Process* (MDP). The MDP is defined by the tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p, r, \gamma)$, where \mathcal{S} is a set of states $\mathbf{s} \in \mathcal{S}$ and \mathcal{A} is a set of actions $\mathbf{a} \in \mathcal{A}$. Let $r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$ denote the *reward*, which the agent receives when it takes action \mathbf{a}_t in state \mathbf{s}_t at timestep t and arrives in state \mathbf{s}_{t+1} with a probability described by the *transition dynamics* $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$. The *policy* $\pi : \mathcal{S} \rightarrow \mathcal{A}$ maps from states to a probability distribution over actions. The objective is to estimate an *optimal policy* function $\pi^* = \arg \max_{\pi} \sum_{t=1}^H \gamma^t r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$ that maximizes the finite-horizon cumulative reward over the horizon H . $\gamma = 1$ is the scalar time-wise discount factor. In the offline setting, the agent does not interact with the environment but rather learns from a static dataset \mathcal{D} generated by a behavior policy π_d . When the sampled observations $\mathbf{o}_t \in \mathcal{O}$ do not fully capture the ground truth state, the MDP becomes a partially observable MDP (POMDP). It is described by the tuple $\mathcal{M}_{PO} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, p, r, \gamma)$. A common method to provide a solution in the partially observable setting is the *n*-th-order history [Sutton and Barto, 2018] approach. Here the state estimation is approximated by a sequence of the $n_c \in \mathbb{N}_0^+$ last observations and actions. The result is a large but finite Markov decision process (MDP) and standard RL approaches can then be applied.

3.2 General Approach

Forecasting other agent responses to the actions of the SDV vehicle is critical for automated driving, and faces the challenge of the inherent uncertainty of human behavior. As an example, let us consider the situation illustrated in Figure 1 b). Here another agent (green) starts to cut in onto the lane of the SDV (white). The SDV has to estimate the probability of each future outcome (i.e. the other agent aborts or continues the maneuver) and should plan its action \mathbf{a}_t accordingly. UMBRELLA models the stochasticity of other agents’ behavior at time t with a continuous latent variable $\mathbf{z}_t \in \mathcal{Z}$. As enumerating all possible actions of the SDV during planning is intractable, we sample actions based on a learned BC policy (Section 3.4). That leads to efficient, expert-like rollouts of $N \in \mathbb{N}_0^+$ potential state-trajectories over a planning horizon H . After trajectory sampling, the approach applies a return-weighted trajectory optimizer. This work follows the *n*-th-order history approach to account for states, which are not fully observable (e.g. human drivers intent) and are merely estimated from the last observations $\mathbf{o}_{t-n_c:t}$ up to timestep t .

3.3 Stochastic Dynamics Model

To model different futures, we learn stochastic forward dynamics models $f_{m,\theta} : \mathcal{S} \times \mathcal{A} \times \mathcal{Z} \rightarrow \mathcal{S} \times \mathbb{R}$ parameterized by θ based on the work of [Henaff et al., 2019]. As the evolution of the traffic scene also depends on the action \mathbf{a}_t of the SDV, a dynamics model is action-conditioned and answers the question: *If the SDV takes this action, how will the other agents in the scene respond?* UMBRELLA’s augmented models captures the interaction between prediction and planning in the context of automated driving. The model, which is a CVAE [Kingma and Welling, 2014], outputs the prediction for the next state $\hat{\mathbf{s}}_{t+1} = f_{m,\theta}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{z}_t)_s$ and the corresponding reward prediction $\hat{r}_t = f_{m,\theta}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{z}_t)_r$. As the model outputs two predictions, we minimize a multi-task mean squared error loss during training.

The latent variable \mathbf{z}_t models different future predictions and makes sure the output is non-deterministic to the input. During training, the latent variable is sampled from the posterior distribution $q_\phi(\mathbf{z} | \mathbf{s}_t, \mathbf{s}_{t+1})$ parameterized by ϕ . As we can only sample from the prior distribution during inference, the Kullback-Leibler (KL) divergence between the posterior and the prior distribution $p(\mathbf{z})$ is also minimized, following the definition of the Evidence Lower Bound (ELBO) objective used when training VAEs. Using a weighting factor ζ the per-sample loss is then given by

$$\mathcal{L}(\theta, \phi; \mathbf{s}_t, \mathbf{s}_{t+1}, \mathbf{a}_t, r_t) = \|\mathbf{s}_{t+1} - f_{m,\theta}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{z}_t)_s\|_2^2 + \|r_t - f_{m,\theta}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{z}_t)_r\|_2^2 + \zeta D_{\text{KL}}(q_\phi(\mathbf{z}_t | \mathbf{s}_t, \mathbf{s}_{t+1}) || p(\mathbf{z}_t)) \quad (1)$$

3.4 Behavior Cloned Policy and Truncated Value Function

During inference, the algorithm rolls out potential state trajectories according to the stochastic forward dynamics model. For this, an action-sampling procedure guided by a BC policy is applied. UMBRELLA learns a bootstrap ensemble of BC policies $f_{b,\psi} : \mathcal{S} \times \mathcal{A}^{n_c} \rightarrow \mathcal{A}$ parameterized by $\psi \in \mathbb{R}$. The model $f_b(\mathbf{s}_t, \mathbf{a}_{(t-n_c):(t-1)})$ takes the current state and the n_c previous actions as input and outputs the action \mathbf{a}_t . By concatenating the previous actions the learned action \mathbf{a}_t is supposed to be more smoothly.

Computational resources limit the planning horizon length. Therefore prior works [Lowrey et al., 2019; Argenson and Dulac-Arnold, 2021; Zhan et al., 2021] successfully propose a value function to extend the planning horizon. UMBRELLA also learn a truncated value function $f_{R,\xi} : \mathcal{S} \times \mathcal{A}^{n_c} \rightarrow \mathbb{R}$ parameterized by ξ to estimate the expected return of the next H episodes \hat{R}_H when being in state \mathbf{s}_t and executed $\mathbf{a}_{(t-n_c):(t-1)}$ as n_c previously taken actions. That effectively extends the planning horizon without increasing the number of rollouts of the dynamics model. Appendix A and B report additional details of the models’ architecture and training procedure.

3.5 UMBRELLA-Planning

Model Predictive Control Formulation. UMBRELLA uses model-predictive control (MPC) [Richalet et al., 1978], which has a long history in control and automation, to plan it’s actions. Hence in every planning step, the algorithm solves a finite-horizon optimal control problem resulting in an optimal trajectory T of length H . Then the first action of the optimal control sequence is executed. The repetitive solution of an optimal control problem reduces the influence of modeling errors.

UMBRELLA Trajectory Optimizer. Next, we will describe the UMBRELLA planning algorithm 1. It is used in every MPC planning cycle to receive an optimal action trajectory. Note that for the sake of clarity, we omit indexing of the models by the parameters ψ, ϕ, ξ . The algorithm generates N trajectories, whereas $M \in \mathbb{N}_0^+$ trajectories are planned per ensemble head using the sampled latent variables \mathbf{z}_m from the prior distribution with $m \in [1, \dots, M]$. As human drivers are assumed to not switch their driving style erratically, the latent variable is fixed across the entire trajectory indexed by $n \in [1, \dots, N]$ (Line 6), which results in a consistent prediction. Moreover, as [Nagabandi et al., 2020] and [Argenson and Dulac-Arnold, 2021], the same l^{th} ensemble head of the BC policy f_b^l and the dynamics model f_m^l are employed consistently across the entire trajectory.

The BC policy guides the expansion of the trajectory by sampling an action \mathbf{a}_t from the BC policy with added Gaussian noise (Line 14). Afterward, the action is averaged together with the trajectory of the previous time-step (Line 15) using the mixture coefficient β [Argenson and Dulac-Arnold, 2021].

UMBRELLA then rolls out state-trajectories using the dynamics model f_m^l (line 16) and calculates the average reward over all ensemble members (line 17). At the end of the trajectory, it calculates an average over all ensemble members of the truncated value function. Averaging over all models of the same ensemble family is inspired by prior work [Nagabandi et al., 2020; Argenson and Dulac-Arnold, 2021].

Last, an optimal action-trajectory is calculated using the model predictive path integral (MPPI) framework [Williams et al., 2017]. Prior works also employ MPPI in the online and offline planning setting [Nagabandi et al., 2020; Argenson and Dulac-Arnold, 2021; Kahn et al., 2021]. Let $\mathbf{A}_{N,H}$ denote set of N generated action trajectories and \mathbf{R}_N the corresponding rewards. Then the optimal trajectory is obtained by re-weighting each action based on the associated reward:

$$\mathbf{T}_t^* = \frac{\sum_{n=1}^N e^{\kappa \mathbf{R}_n} \mathbf{A}_{n,t+1}}{\sum_{n=1}^N e^{\kappa \mathbf{R}_n}}, \forall t \in [0, \dots, H-1] \quad (2)$$

Algorithm 1 UMBRELLA Planning

```

1: procedure UMBRELLA-PLANNING( $\mathbf{s}, \mathbf{T}, f_m, f_b, f_R, H, N, \sigma^2, \beta, \kappa$ )
2:   Set  $\mathbf{R}_N = \vec{0}_N, \mathbf{A}_{N,H} = \vec{0}_{N,H}$  ▷ Returns and actions of  $N$  sampled trajectories.
3:    $M = N/K$  ▷ Plan  $M$  different trajectories per head.
4:   for  $m = 1, \dots, M$  do
5:      $\mathbf{z}_m \sim p(\mathbf{z})$  ▷ Sample latent variable  $\mathbf{z}_m$  from prior  $p(\mathbf{z})$ .
6:   end for
7:   for  $n = 1, \dots, N$  do ▷ Sample  $N$  trajectories over horizon  $H$ 
8:      $l = n \bmod K$  ▷ Use  $K$  different ensemble heads.
9:      $m = n \bmod M$ 
10:     $\mathbf{s}_1 = \mathbf{s}, \mathbf{a}_0 = \mathbf{T}_0, R = 0$ 
11:    for  $t = 1, \dots, H$  do
12:       $\epsilon \sim \mathcal{N}(0, \sigma^2)$ 
13:       $\mathbf{a}_t = f_b^l(\mathbf{s}_t, \mathbf{a}_{(t-n_c):(t-1)}) + \epsilon$  ▷ BC policy prior to guide the action sampling.
14:       $\mathbf{A}_{n,t} = (1 - \beta)\mathbf{a}_t + \beta\mathbf{T}_{i=\min(t,H-1)}$  ▷ Beta-mixture with prev. trajectory  $\mathbf{T}$ .
15:       $\mathbf{s}_{t+1} = f_m^l(\mathbf{s}_t, \mathbf{A}_{n,t}, \mathbf{z}_m)_s$  ▷ Sample next state from stochastic dyn. model.
16:       $R = R + \frac{1}{K} \sum_{i=1}^K f_m^i(\mathbf{s}_t, \mathbf{A}_{n,t}, \mathbf{z}_m)_r$  ▷ Take avg. reward over heads.
17:    end for
18:     $\mathbf{R}_n = R + \frac{1}{K} \sum_{i=1}^K f_R^i(\mathbf{s}_{H+1}, \mathbf{A}_{n,(H-n_c+1):(H)})$  ▷ Append pred. truncated value.
19:  end for
20:  Compute  $\mathbf{T}_t^*$  according to Eq. (2) ▷ Compute return-weighted trajectory
21:  return  $\mathbf{T}_t^*$ 
22: end procedure

```

Pessimistic Trajectory Optimizer. This work also proposes UMBRELLA-P, a pessimistic trajectory optimizer. UMBRELLA and MBOP [Argenson and Dulac-Arnold, 2021] use all sampled trajectories for the computation (equation 2) of the weighted trajectory. UMBRELLA-P only aggregates those trajectories of the ensemble head with the lowest sum over returns. Therefore, the algorithm first calculates $\mathbf{R}_{k,\text{sum}} \forall k \in [1, \dots, K]$, which is the sum over all rewards belonging to the ensemble indexed by k . Then, it picks the index of the ensemble with the lowest sum over returns by $k^* = \arg \min_{k \in [1, \dots, K]} \mathbf{R}_{k,\text{sum}}$. Lastly, equation 2 is only applied to the trajectories belonging to k^* . Therefore, UMBRELLA-P optimizes in the face of epistemic uncertainty for the worst-case outcome and acts pessimistically.

4 Experimental Evaluation

This section evaluates the proposed method in two interactive automated driving environments. The goal of this section is to answer the following research questions: *Q1*: Does the modeling of aleatoric uncertainty enhance the performance of model-based offline planning methods in interactive automated driving scenarios? *Q2*: Does the algorithm improve upon simple behavior cloned policies? *Q3*: Does the pessimistic variant UMBRELLA-P improve the planning performance? *Q4*: What are the limitations of model-based offline planning methods in the context of automated driving?

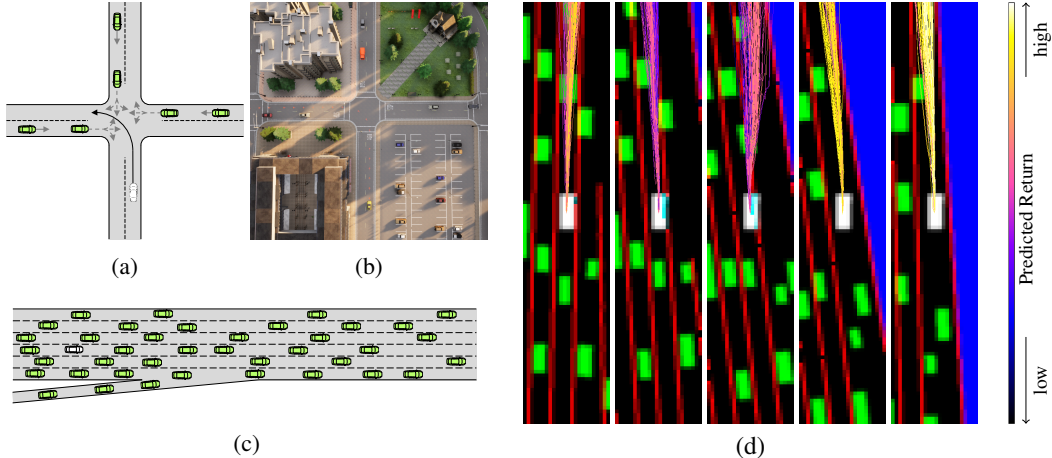


Figure 2: (a) Schematic representation of the left turn task. (b) Intersection in CARLA. (c) NGSIM environment. (d) Context images overlaid with the planned UMBRELLA trajectories in the time interval $t \in [0s, 8s]$ in steps of 2s during a lane change maneuver. The trajectories are colored concerning their predicted return colors. The background describes the context image: The SDV is denoted in white, other agents are visualized in green. Black denotes the road and red lane markings.

Environments. *NGSIM*: The first environment is a challenging multi-agent automated driving environment from the work of [Henaff et al., 2019] based on the Next Generation Simulation program’s Interstate 80 (NGSIM I-80) dataset [Halkias and Colyar, 2006]. The goal of the SDV is to stay in the middle of a lane while avoiding collisions. *CARLA*: The second environment includes a urban multi-agent scenario implemented in the CARLA simulator [Dosovitskiy et al., 2017]. The SDV is supposed to make progress along the route and execute an unprotected left turn at a busy intersection while avoiding collisions. Figure 2 (a)-(c) illustrates both experiments.

Baselines. This work benchmarks against the following methods: (i) *1-step IL*: A learned policy imitating the expert driver using the BC policy of Section 3.4 (ii) *MBOP* [Argenson and Dulac-Arnold, 2021]: A current state-of-the-art model-based offline RL method that uses the deterministic dynamics model of Henaff et al. [2019]. Note, that for a fair comparison all other components of MBOP are identical to the UMBRELLA approach. In the NGSIM environment we further compare against (iii) *MPUR*: [Henaff et al., 2019] A state-of-the-art model-based policy learning method, which accounts for epistemic and aleatoric uncertainty (iv) *Human*: The ground truth actions of the human. (v) *No action*: A policy that always applies actions of zeros.

Metrics. The comparative analysis uses the following metrics. *Success rate* (SR): The rate of collision-free episodes. In the CARLA experiment, also an episode is considered as failure when the vehicle does not reach the target location within the allocated time. *Mean distance* (MD): The distance traveled longitudinal direction along the NGSIM-highway averaged over all episodes. The NGSIM analysis further evaluated the individual parts of the reward function (i.e. *mean proximity reward* \bar{r}_{prox} and *mean lane reward* \bar{r}_{lane} and the mean final reward \bar{r}). The CARLA experiment determines the *mean successful time* (MST): The average time required for successfully navigating through the intersection averaged over all successful episodes. Appendix C provides further information about the state, action, and reward representation.

4.1 NGSIM Experiments

The NGSIM I-80 dataset consists of 45 minutes of real-world driving data on a highway with dense traffic recorded with cameras. The dataset has high environment (or aleatoric) uncertainty as all agents exhibit different driving maneuvers as lane changes, merges, braking, and acceleration maneuvers. We follow the preprocessing and experimental setup of Henaff et al. [2019] utilizing the OpenAI Gym [Brockman et al., 2016]. In each evaluation episode, the SDV is initialized at the beginning of a random trajectory of the test set. Other agents drive along their fixed trajectories from the dataset. Fixed trajectories avoid the need of hand-designing (unrealistic) driving policies of neighboring cars. However, this log-replay comes with the downside that the other agents do not react to the SDV, which enhances the difficulty of the task as multiple agents can trap the SDV.

Table 1: Performance of the different evaluated methods on the highway (NGSIM) and urban (CARLA) control task by diverse metrics. CARLA (ID) denotes evaluation using simulations with parameters from the same distribution as used during training. CARLA (OOD) defines experiments, where a distribution shift is simulated. The unit of the metric MD is meter. The values of the MST metric are denoted in seconds. Bold numbers mark the best result and an underlined number the second best.

Method	NGSIM					CARLA ID		CARLA OOD	
	SR	MD	\bar{r}_{prox}	\bar{r}_{lane}	\bar{r}	SR	MST	SR	MST
UMBRELLA	<u>0.60</u>	<u>320.06</u>	0.817	<u>0.902</u>	1.718	0.63	11.62	<u>0.43</u>	13.10
UMBRELLA-P	0.53	300.43	<u>0.820</u>	0.904	<u>1.724</u>	<u>0.50</u>	11.97	0.53	13.55
MBOP	0.50	290.62	0.775	0.893	1.669	0.27	<u>11.78</u>	0.20	<u>13.37</u>
1-step IL	0.00	67.65	0.580	0.605	1.185	0.00	-	0.00	-
MPUR	0.70	351.38	0.886	0.896	1.781	-	-	-	-
No action	0.27	225.98	0.665	0.668	1.332	-	-	-	-
Human	1.00	378.04	0.821	0.642	1.463	-	-	-	-

UMBRELLA is compared against the previously described baselines. Table 1 illustrates the results. In terms of success rate and mean traveled distance, MPUR performs best followed by UMBRELLA. However, MPUR does not exhibit the control flexibility as UMBRELLA and requires retraining whenever new parts are added to the reward function. In contrast, model-based offline planning methods like MBOP and UMBRELLA can easily be modified, which was shown in prior work [Argenson and Dulac-Arnold, 2021]. Moreover, in contrast to MPUR, these methods do not require a differentiable reward function.

Regarding questions $Q1$ and $Q2$, UMBRELLA outperforms both MBOP and the 1-step IL method. We can conclude that modeling aleatoric uncertainty is critical when applying model-based offline RL approaches in automated driving scenarios. Prediction with a deterministic dynamics model leads to an averaging over multiple futures. Hence, UMBRELLA outperforms MBOP w.r.t the considered metrics. Interestingly, within the IL approach the SDV fails on every episode. However, UMBRELLA improves drastically upon this policy. Considering question $Q3$, UMBRELLA-P is inferior compared to UMBRELLA on the metrics SR and MD, while achieving slightly higher mean rewards. Our experiments revealed, that UMBRELLA-P performs fewer lane changes and it gets more frequently trapped by the non-reactive agents. We hypothesize that this occurs due to the more conservative planning. As the pessimistic variant plans for the worst-case, it stays on the same lane most of the time to avoid potential risk of lane changes. As a result, the follower vehicle causes a rear collision, as its original trajectory is ignorant of the phantom SDV.

Figure 2 (d) provides an exemplary visualization of the planning result in a challenging lane change scenario with dense traffic. The SDV predicts a safety-critical future evolution of the scenario (visualized with low return trajectories) as the follower vehicle starts to accelerate and potentially traps the SDV. However, the SDV reacts accordingly and successfully changes its lane in dense traffic.

4.2 CARLA Experiments

The second experiment describes an unprotected left-turning task at a busy intersection in Town04 of the CARLA simulator (Version 0.9.10) [Dosovitskiy et al., 2017], as visualized in Figures 2 (a) and (b). Six other agents, parameterized with different random driving style, are controlled by CARLA’s traffic manager and approach the intersection from three directions. The other agents perform random maneuvers at the intersection. That makes the scenario challenging, as the prediction of uncertain multi-agent interaction is required to successfully navigate to the goal location. For training, we generate a dataset using different sub-optimal behavior policies. We construct two different test settings. In the first one, CARLA in-distribution (ID), the SDV is spawned at the same locations as during training. In the second setting, CARLA out-of-distribution (OOD), initial configurations are changed to generate novel unseen situations for decision-making. We evaluate on 30 randomly generated scenarios per test setting, and the results are visualized on the right side of Table 1. Appendix C provides additional details about the experimental setup and the data generation.

Regarding question $Q1$, again UMBRELLA improves the success rate of MBOP by a factor of 2.33 (CARLA ID) and 2.15 (CARLA OOD) while also navigating the intersection faster. *Our insight is that the superior performance can be attributed to the use of a stochastic dynamics model.* That underlines the importance of modeling different future outcomes in the context of model-based offline RL for automated driving. Appendix D provides further visualizations of the model’s prediction performance. 1-step IL consistently fails and does not re-accelerate once it stops at the intersection. This behavior is attributed to the *inertia problem* [Codevilla, Eder, Lopez, and Gaidon, 2019], a common problem in behavior cloning due to causal confusion [de Haan et al., 2019]. Interestingly, while using the 1-step IL as a prior for rolling out state-trajectories, UMBRELLA and MBOP can overcome this issue, which answers question $Q2$.

Considering question $Q3$, UMBRELLA achieves a higher success rate and lower MST than its ablation in the CARLA ID experiment. However, in the CARLA OOD experiment, the performance of UMBRELLA drops, whereas the pessimistic variant performs nearly identical. Using UMBRELLA-P, the SDV plans for the worst-case model, i.e. the model, which predicts that all other agents decline the right of way to the SDV. The result is a more conservative and reactive behavior in the intersection causing fewer or no collisions. That is also reflected by the higher MST value.

4.3 Limitations

This section approaches question $Q4$ and outlines the limitation of current model-based offline planning methods in the context of automated driving.

Reward Function Mismatch. While UMBRELLA, UMBRELLA-P, MBOP, and MPUR learn to perform different challenging maneuvers in dense highway traffic, none of these methods reaches human level performance. To further investigate the cause, let us consider the mean reward \bar{r} in Table 1. Notice, that all approaches achieve a higher reward than the human baseline. In the context of offline RL, the task of the agent is to outperform the dataset policy in terms of the reward. We conclude that optimizing the reward function proposed by Henaff et al. [2019] does not necessarily lead to optimization in terms of driving performance. That is caused by the fact that the reward function does not exactly represent the human driving style. As other agents do not react to the SDV, the most likely way to get a success rate of 1.00 is to drive exactly like the human in the dataset. One remedy, is to learn the reward function with inverse RL (IRL) methods [Osa et al., 2018]. However, this raises the almost philosophical question whether a human-like reward function is the ultimate goal for automated driving, as humans sometimes also tend to drive in a safety-critical way and violate traffic rules. In general, reward (mis)design is a common problem in automated driving contributions [Knox et al., 2021] and needs to be addressed in future work. We hypothesize that this effect also occurs in our CARLA experiments.

Dependency on BC Policy. While UMBRELLA and MBOP can improve upon a simple BC policy, we observed that their performance is limited by the unimodal BC policy. That often leads to sampled trajectories, which correspond to one possible maneuver (see Figure 2 (d)), whereas in reality a human driver plans, for multiple possible maneuvers. Therefore, future work should investigate the integration of multi-modal BC policies. In addition, the CARLA dataset contains suboptimal expert policies, which drive through the intersection aggressively and are mostly non-reactive. We observed that this driving style is partially also present in MBOP’s and UMBRELLA’s policies, which shows a downside of both approaches when learning from non-expert-like datasets.

5 Conclusion

This work proposed an approach for model-based offline reinforcement learning, which considers aleatoric and epistemic uncertainty. Experiments in multiple challenging automated driving scenarios with dense traffic showed that incorporating aleatoric uncertainty improves the performance of planning-based approaches. Moreover, the proposed method improves upon the poor performance of simple BC policies. A pessimistic ablation is inferior in scenarios of the training distribution but seems to offer advantages in out-of-distribution scenarios. We also showed the limitations of the proposed approach and current state-of-the-art planning-based offline RL in the context of automated driving. Future work should focus on improving the BC policy priors, use graph-based representations instead of images and investigate the use of other optimization techniques.

Acknowledgments and Disclosure of Funding

This research was funded by the Federal Ministry for Economic Affairs and Energy on the basis of a decision by the German Bundestag in the project “KISSaF – AI-based Situation Interpretation for Automated Driving”.

References

- Arthur Argenson and Gabriel Dulac-Arnold. Model-based offline planning. In *International Conference on Learning Representations*, 2021.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016. arxiv:1606.01540.
- Simon Le Cleac’h, Mac Schwager, and Zachary Manchester. Algames: A fast solver for constrained dynamic games. In *Robotics: Science and Systems*, 2020.
- Simon Le Cleac’h, Mac Schwager, and Zachary Manchester. Lucidgames: Online unscented inverse dynamic games for adaptive trajectory prediction and planning. *IEEE Robotics and Automation Letters*, 6(3):5485–5492, 2021.
- Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4693–4700, 2018.
- Felipe Codevilla, Santana Eder, Antonio M. Lopez, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *IEEE/CVF International Conference on Computer Vision*, pages 9328–9337, 2019.
- Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proceedings of the Conference on Robot Learning*, pages 1–16, 2017.
- Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectornet: Encoding hd maps and agent dynamics from vectorized representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- John Halkias and James Colyar. Ngsim interstate 80 freeway dataset. Technical report, US Federal Highway Administration, FHWA-HRT-06-137, 2006.
- Mikael Henaff, Alfredo Canziani, and Yann LeCun. Model-predictive policy learning with uncertainty regularization for driving in dense traffic. In *International Conference on Learning Representations*, 2019.
- Gregory Kahn, Pieter Abbeel, and Sergey Levine. Badgr: An autonomous self-supervised learning-based navigation system. *IEEE Robotics and Automation Letters*, 6(2):1312–1319, 2021.
- Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. Morel: Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21810–21823, 2020.
- Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 2014.
- Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? does it matter? *Structural Safety*, 31(2):105–112, 2009.
- W. Bradley Knox, Alessandro Allievi, Holger Banzhaf, Felix Schmitt, and Peter Stone. Reward (mis)design for autonomous driving. 2021. arxiv:2104.13906.

- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *CoRR*, 2020. arXiv:2005.01643.
- Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *Proceedings of the European Conference on Computer Vision*, 2020.
- Jerry Liu, Wenyuan Zeng, Raquel Urtasun, and Ersin Yumer. Deep structured reactive planning. In *IEEE International Conference on Robotics and Automation*, 2021.
- Kendall Lowrey, Aravind Rajeswaran, Sham Kakade, Emanuel Todorov, and Igor Mordatch. Plan online, learn offline: Efficient learning and exploration via model-based control. In *International Conference on Learning Representations*, 2019.
- Thomas M. Moerland, Joost Broekens, and Catholijn M. Jonker. Model-based reinforcement learning: A survey, 2021. arXiv:2006.16712.
- Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Anusha Nagabandi, Kurt Konolige, Sergey Levine, and Vikash Kumar. Deep dynamics models for learning dexterous manipulation. In *Proceedings of the Conference on Robot Learning*, pages 1101–1112, 2020.
- Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J. Bagnell, Pieter Abbeel, and Jan Peters. An algorithmic perspective on imitation learning. *Foundations and Trends in Robotics*, 7:1–179, 11 2018.
- Tom Le Paine, Cosmin Paduraru, Andrea Michi, Çağlar Gülçehre, Konrad Zolna, Alexander Novikov, Ziyu Wang, and Nando de Freitas. Hyperparameter selection for offline reinforcement learning. 2020. arxiv:2007.09055.
- Christian Pék, Stefanie Manzingger, Markus Koschi, and Matthias Althoff. Using online verification to prevent autonomous vehicles from causing accidents. *Nature Machine Intelligence*, 2:518–528, 2020.
- Nicholas Rhinehart, Jeff He, Charles Packer, Matthew A. Wright, Rowan McAllister, Joseph E. Gonzalez, and Sergey Levine. Contingencies from observations: Tractable contingency planning with learned behavior models. In *IEEE International Conference on Robotics and Automation*, 2021.
- J. Richalet, A. Rault, J.L. Testud, and J. Papon. Model predictive heuristic control: Applications to industrial processes. *Automatica*, 14(5):413–428, 1978.
- Stephane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 627–635, 2011.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.
- Grady Williams, Andrew Aldrich, and Evangelos A. Theodorou. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 40(2): 344–357, 2017.
- Ming Yin, Yu Bai, and Yu-Xiang Wang. Near-optimal offline reinforcement learning via double variance reduction. 2021. arxiv:2102.01748.

- Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. Mopo: Model-based offline policy optimization. In *Advances in Neural Information Processing Systems*, volume 33, pages 14129–14142, 2020.
- Tianhe Yu, Aviral Kumar, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, and Chelsea Finn. Combo: Conservative offline model-based policy optimization, 2021. arXiv:2102.08363.
- Xianyuan Zhan, Xiangyu Zhu, and Haoran Xu. Model-based offline planning with trajectory pruning, 2021. arXiv:2105.07351.
- Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.

A Model Architecture

A.1 Dynamics Model

The dynamics model is inspired by the work of Henaff et al. [2019] and its architecture is visualized in Figure 3. The network consists of the three main components: the encoder f_{enc} , the posterior network q_ϕ , and the decoder f_{dec} . For reasons of clarity, we do not include latent dropout (see Appendix B.1) in the figure. The illustration is further simplified as in reality, the posterior network outputs the parameters of a Gaussian distribution, which are then used to estimate \mathbf{z}_t by

$$(\boldsymbol{\mu}_\phi, \boldsymbol{\sigma}_\phi) = q_\phi(\mathbf{s}_t, \mathbf{s}_{t+1}) \quad (3)$$

$$\boldsymbol{\epsilon} \sim \mathcal{N}(0, I) \quad (4)$$

$$\mathbf{z}_t = \boldsymbol{\mu}_\phi + \boldsymbol{\sigma}_\phi * \boldsymbol{\epsilon}. \quad (5)$$

The work of Henaff et al. [2019] gives further information about the layer sizes and different parts of the model.

A.2 BC Policy and Truncated Value Function

The network architecture of the BC policy and the Truncated Value Function are visualized in Figure 4 and 5. They encode the sequence of images with a 3-layer convolutional neural network with (64 -128- 256) feature maps. The sequences of measurement vectors \mathbf{u}_t and actions \mathbf{a}_t are encoded using a 2-layer fully connected network with 256 units. The addition of all encoding is then further processed by two 2-layer MLPs. The result is the predicted action $\hat{\mathbf{a}}_t$ or return $\hat{\mathbf{R}}_t$, respectively.

B Training Details

B.1 Dynamics Model

The signal flow during training and testing is visualized in Figure 6. The network f_{enc} encodes the current state \mathbf{s}_t and action \mathbf{a}_t . Moreover, a latent variable \mathbf{z}_t is added. f_{dec} decodes the result leading to the prediction of the successor state $\hat{\mathbf{s}}_{t+1}$ and reward \hat{r}_{t+1} . During training, the mean squared error (MSE) between the predicted successor state $\hat{\mathbf{s}}_{t+1}$ and the true successor state \mathbf{s}_{t+1} is minimized. The same holds for the reward. The network further encodes the current state \mathbf{s}_t and the true successor state \mathbf{s}_{t+1} with f_{enc} to the parameterized posterior distribution $q_\phi(\mathbf{z} | \mathbf{s}_t, \mathbf{s}_{t+1})$. Besides the MSE losses, the Kullback-Leibler divergence (D_{KL}) between the posterior and the prior distribution $p(\mathbf{z})$ is minimized during training. To make the network sensitive to the action input, we apply *latent dropout* [Henaff et al., 2019] during training. First, a random variable x is sampled from a Bernoulli distribution $\mathcal{B}(p_x)$ with probability p_x . Hence we sample the latent variable during training according to

$$\mathbf{z}_t \sim (1 - x) * q_\phi(\mathbf{z} | \mathbf{s}_t, \mathbf{s}_{t+1}) + x * p(\mathbf{z}) \quad (6)$$

That forces the model to receive as much information as possible from the current state and action to predict the successor state. To mitigate aggregated multi-step prediction errors [Moerland et al., 2021] at test time, we unroll the model for five timesteps and train using the corresponding loss. The dynamics model is first trained in deterministic mode for 50 000 steps. Deterministic mode denotes the training configuration, where no latent variable is added. Afterward, the dynamics model is again trained for 50 000 steps with the added latent variable. We denote this model as the stochastic dynamics model. The *deterministic dynamics model* used for MBOP is trained without the latent variable addition.

B.2 BC Policy and Truncated Value Function

We train the BC policy network by minimizing the loss

$$\mathcal{L}(\boldsymbol{\psi}; \mathbf{s}_t, \mathbf{a}_{(t-n_c):(t-1)}, \mathbf{a}_t) = \|\mathbf{a}_t - f_{\text{b},\boldsymbol{\psi}}(\mathbf{s}_t, \mathbf{a}_{(t-n_c):(t-1)})\|_2^2, \quad (7)$$




Figure 3: Simplified network architecture of the stochastic forward dynamics model during training




Figure 4: Behavior-clone policy network architecture.




Figure 5: Truncated value function network architecture.

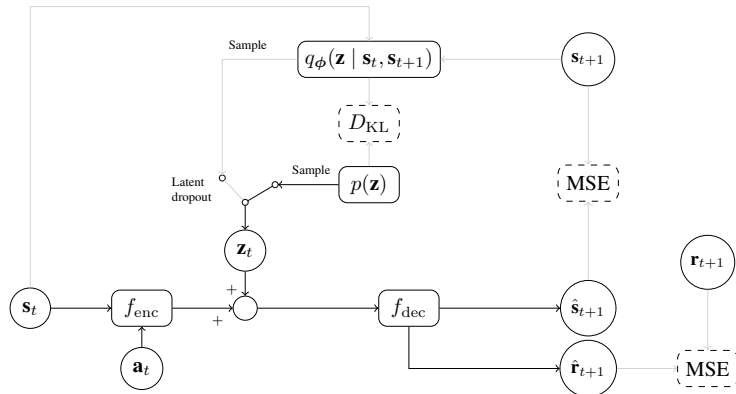


Figure 6: Signal flow of the stochastic forward dynamics model (inspired by Henaff et al. [2019]). Black color denotes the signal flow during training and inference. Grey denotes the signal flow, which is only used during training.

whereas parameter vector of one behavior-cloned policy model is denoted by ψ . The parameterized truncated value function model $f_{R,\xi}(\mathbf{s}_t, \mathbf{a}_{(t-n_c):(t-1)})$ is learned by minimizing the following per-sample loss:

$$\mathcal{L}(\xi; R_H, \mathbf{s}_t, \mathbf{a}_{(t-n_c):(t-1)}) = \|R_H - f_{R,\xi}(\mathbf{s}_t, \mathbf{a}_{(t-n_c):(t-1)})\|_2^2 \quad (8)$$

ξ is the parameter vector. The value function target $R_H \in \mathbb{R}$ is defined as the sum over the next H rewards:

$$R_H = \sum_{t=\tau}^{t+H-1} r_t \quad (9)$$

Table 2 lists the hyperparameters for training the stochastic dynamics model, the BC policy model, and the value function model.

Table 2: Training hyperparameters used for the stochastic forward dynamics model, the BC policy model, and the value function model. n_p denotes the number of time steps the model is unrolled during training. η is the initial learning rate. p_{dropout} describes the dropout probability. p_x denotes the latent dropout probability

Model	Type	n_c	n_p	Opt.	Steps	Batchsize	η	p_{dropout}	p_x
Dynamics	stochastic	20	5	Adam	2 * 50000	32	0.0001	0.1	0.5
BC Policy	deterministic	20	1	Adam	50000	32	0.0001	0.0	-
Value Func.	deterministic	20	30	Adam	50000	32	0.0001	0.1	-

C Experimental Setup

C.1 NGSIM

State Space. The state \mathbf{s}_t at time t is approximated by a set of n_c observations \mathbf{o}_t . One observation consists of a context image in bird’s-eye-view $\mathbf{i}_t \in \mathbb{R}^{4 \times H \times W}$ and a measurement vector $\mathbf{u}_t \in \mathbb{R}^4$. The context image has sizes $H = 117$ pixel and $W = 24$ pixel and is centered around the SDV. It contains four channels with information about the SDV’s size, the drivable area, lane markings, and the surrounding vehicles. The resolution is approximately 0.5 meter per pixel. The measurement vector contains information about the 2D position $\mathbf{p}_t \in \mathbb{R}^2$ and velocity $\mathbf{v}_t \in \mathbb{R}^2$ of the SDV.

Action Space. The action $\mathbf{a}_t = (\Delta v_t, \Delta \delta_t)$ describes a change of velocity $\Delta v_t \in \mathbb{R}$ and a change of steering angle $\Delta \delta_t \in \mathbb{R}$ by

$$\Delta v_t = \|\mathbf{v}_{t+1}\|_2 - \|\mathbf{v}_t\|_2 \quad (10)$$

$$\Delta \delta_t = (\mathbf{v}_{t+1} - \mathbf{v}_t)^T \frac{(\mathbf{v}_t)_\perp}{\|\mathbf{v}_t\|_2} \quad (11)$$

Reward. The SDV optimizes the following reward function, which was proposed by Henaff et al. [2019]:

$$r_t = r_{\text{prox},t} + r_{\text{lane},t} \quad (12)$$

It penalizes small distances to other objects scaled by the SDV’s velocity (proximity reward $r_{\text{prox},t}$) and small distances to the lane markings (lane reward $r_{\text{lane},t}$).

Dataset. A total of 5596 vehicle trajectories sampled at 10Hz are included. Note that one trajectory corresponds to one episode. We split the trajectories into training (80%), validation (10%) and testing sets (10%).

C.2 CARLA

State Space. The state space is the same as in the NGSIM experiment besides that $H = 126$, $W = 74$ and $\mathbf{i}_t \in \mathbb{R}^{3 \times H \times W}$. In this experiment, we excluded the ego-layer as the SDV always has the same size during training. Note that this is not the case in the NGSIM experiment due to different vehicle types and sizes.

Action Space. The action $\mathbf{a}_t = (d_t, \delta_t)$ describes a normalized lateral steering wheel command $\delta_t \in [-1, 1]$ and longitudinal driving command $d_t \in [-1, 1]$ (acceleration or braking), which is converted to a throttle and brake command using a low-level PID controller.

Reward. The SDV optimizes the following reward function:

$$r_t = w_1 * r_{\text{prog},t} + w_2 * r_{\text{lane},t} + w_3 * r_{\text{coll},t}. \quad (13)$$

$r_{\text{prog},t}$ rewards the SDV’s progress towards its goal while obeying the speed limit. The lane reward $r_{\text{lane},t}$ penalizes deviations from the reference lane center. $r_{\text{coll},t}$ is a collision reward encouraging the agent to avoid collisions. Note that $w_1, w_2, w_3 \in \mathbb{R}$ are weighting factors. Further, the three reward terms are defined as:

$$r_{\text{prog},t} = \begin{cases} \gamma^t * (e_t - e_{t-1}), & \text{if } \|\mathbf{v}_t\|_2^2 \leq v_{\text{limit}} \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

$$r_{\text{lane},t} = -\|\mathbf{p}_t - \mathbf{p}_{\text{center}}\|_2^2 \quad (15)$$

$$r_{\text{coll},t} = \begin{cases} -2 + 0.2(t_{\text{coll}} - t), & \text{if } t \in [t_{\text{coll}} - 9, t_{\text{coll}}] \\ 0, & \text{otherwise} \end{cases} \quad (16)$$

$e_t \in \mathbb{R}$ is the distance to the target location and $\gamma = 0.99$ a discount factor. $v_{\text{limit}} \in \mathbb{R}$ describes the speed limit. The lane reward is defined as the negative Euclidean distance to the lane center defines the lane reward. The global position of the lane center closest to the position of the SDV is given by $\mathbf{p}_{\text{center}}$. If a collision occurs in time step t the collision reward also penalizes the prior 9 time steps according to a linear function. $t_{\text{coll}} \in \mathbb{R}$ denotes the collision timestamp. That should make the agent learn dangerous pre-collision states. This work chooses $w_1 = 1, w_2 = 0.1, w_3 = 1$.

Dataset. The models are trained based on trajectories corresponding to 1200 episodes, which are recorded by driving the intersection with different policies. We intentionally make the policies sub-optimal to see how the different algorithms can handle non-expert-like data. This work uses a total of three different behavior policies because drivers in real-world datasets also follow different driving styles. The first policy controls the SDV using PID-controllers, which track the path given by a route planning module. All agents are spawned according to the parameters in Table 3. To generate a diverse set of scenarios the parameters are randomly drawn. Other agents are controlled using the autopilot of the CARLA traffic manager. At the intersection they, follow a First-In-First-Out (FIFO) order. The SDV’s policy can be randomized by three different driving styles: cautions, normal and aggressive. This behavior affects the maximum speed and maintained safety distance to other agents on its planned path. The other agent’s behavior is further randomized by changing their traffic sign, other vehicles, and speed limit ignorance rate. When the SDV reaches the intersection, it waits for a random amount of time $t_{\text{wait}} \in [0\text{s}, 15\text{s}]$ before starting the turn. Note, that this often causes collisions, as the SDV does not react to vehicles that approach from the side. Hence, we filter out collision-prone episodes during data generation. However, the recorded driving behavior is still sub-optimal as the SDV assumes it always has the right of way, which results in a safety-critical behavior. The second policy drives like the first one, but random noise is applied to the controls while maneuvering the intersection. For this purpose, a random offset is injected for 5 time steps. This work uses both policies to generate 400 collision-free episodes each. Lastly, another 400 collision-free episodes are constructed by a human driver, who controls the car with the keyboard. Due to difficult control the policy is also sub-optimal.

C.3 Planning Hyperparameters

We use the hyperparameter of Table 4 during our experiments, which were found using a grid search. Note that, Off Policy Evaluation [Nachum et al., 2019] and Offline Hyperparameter Selection [Paine et al., 2020] are outside the scope of this work. For a fair comparison we performed a search for all used models during evaluation.

D Additional Experimental Results

D.1 Video Prediction Results

As the work of Henaff et al. [2019] already showed prediction results of a similar dynamics model in the NGSIM environment, this section only shows video predictions results from the CARLA

Table 3: Randomized parameters for data generation.

	SDV	Other agents
Spawn position [m]	[25, 45]	first vehicle of each direction: [25, 35], second vehicle of each direction: [45, 50]
Agent behavior	[cautious, normal, aggressive]	safety distance [m]: [0.1, 10] sign ignorance: [0, 0.2], vehicle ignorance: [0, 0.2] speed difference: [-0.3, 0.3]
Waiting time at intersection [s]	[0, 15]	FIFO
Noise injection offset	throttle: [-0.5, 0.5] brake: [-0.5, 0.5] steering: [-0.05, 0.05]	-

Table 4: Used hyperparameters of UMBRELLA for the NGSIM and CARLA experiments.

Environment	K	H	N	σ^2	β	κ
NGSIM	2	30	300	1.2	0.6	0.5
CARLA	2	30	100	1.5	0.9	0.5

experiments in Figure 7. The prediction horizon is 20s. We observe that the stochastic dynamics model can model different future outcomes over a long time horizon. The predictions of the deterministic model are more blurred (last column) as it is averaging over multiple futures. Both models have learned what the road geometry looks like after the intersection. Note that, the prediction results on NGSIM Henaff et al. [2019] are more blurred as different locations, vehicle types and more maneuvers occur during training.

D.2 Runtime

This work conducts runtime experiments using an Nvidia Titan X and an Intel i5-3550@3.3 GHz. In its current version, the implementations of MBOP and UMBRELLA are not real-time capable using the models described in Appendix A, as shown in Figure 8. However, the runtime could be improved as the current implementation is not optimized for performance. Moreover, both methods encode and decode high-dimensional images using CNN’s leading to a high computation cost. Note, that high computation times also occur when using MBOP with the proposed context image-based network. Hence, future work should focus on using more efficient networks. For example, the observation space could be represented by a graph [Gao et al., 2020; Liang et al., 2020] instead of an image. Then using graph neural networks [Zhou et al., 2020] would lead to a more efficient implementation regarding computation time and memory consumption.

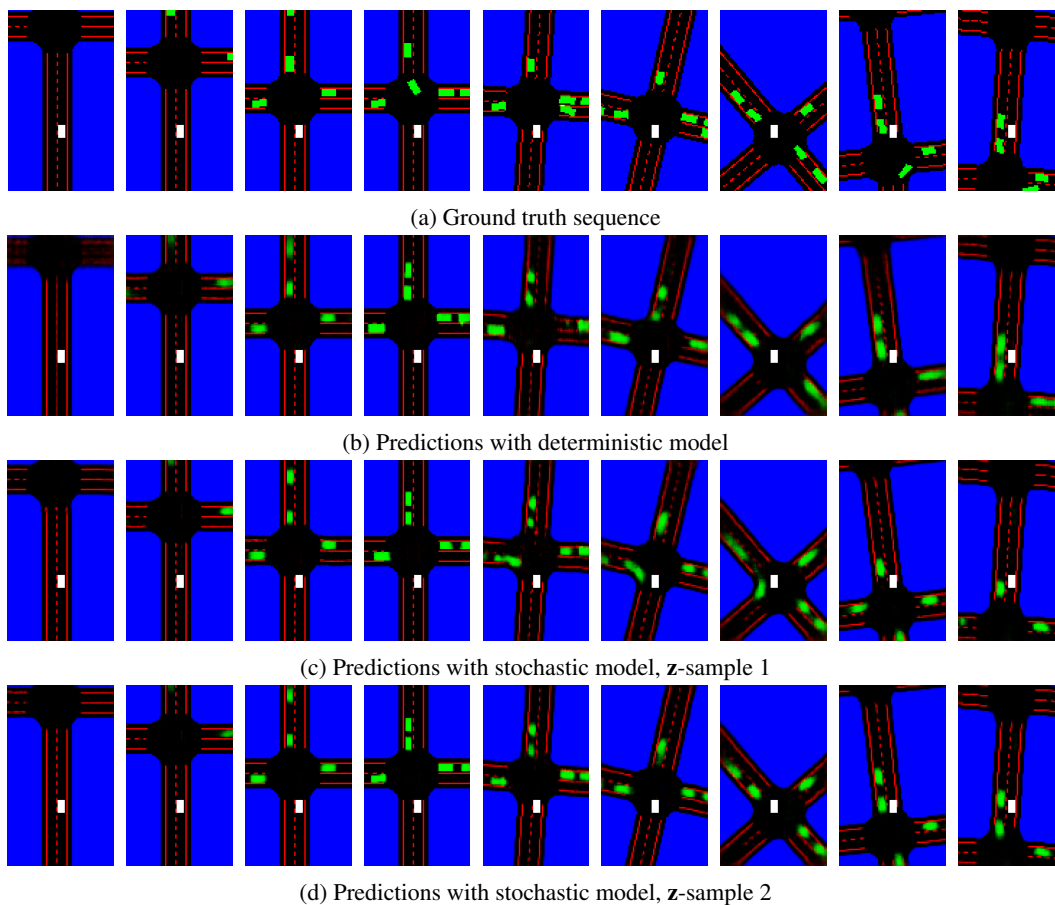


Figure 7: Video prediction results using the different models. (a) Ground truth sequence of the test dataset. (b) Predicted sequence made using the deterministic model (c),(d) Predicted sequence using the stochastic dynamics model. Sampling two different latent variables \mathbf{z} from the prior distribution results in two different predicted sequences.




Figure 8: Average computation time of UMBRELLA as function of N and K with a horizon of $H = 30$.

Table 5: Performance of UMBRELLA in the NGSIM environment for varying parameter combinations. Only one hyperparameter is changed at a time and the rest of the configuration from Table 4

Parameter	SR	MD
$K = 1$	0.43	266.71
$K = 4$	0.43	284.46
$H = 25$	0.57	240.57
$H = 35$	0.60	293.37
$H = 40$	0.50	280.30
$N = 10$	0.20	221.72
$N = 50$	0.47	290.65
$N = 100$	0.50	281.62
$N = 200$	0.53	297.44
$N = 500$	0.50	299.71
$\sigma^2 = 0.3$	0.43	273.21
$\sigma^2 = 0.5$	0.50	282.07
$\sigma^2 = 0.7$	0.53	296.21
$\sigma^2 = 1.0$	0.53	283.80
$\sigma^2 = 1.5$	0.50	303.52
$\sigma^2 = 3.0$	0.27	202.44
$\beta = 0.0$	0.03	116.06
$\beta = 0.1$	0.07	128.12
$\beta = 0.3$	0.13	223.59
$\beta = 0.5$	0.53	307.63
$\beta = 0.7$	0.40	284.99
$\beta = 0.9$	0.20	202.57
$\kappa = 0.01$	0.10	182.77
$\kappa = 0.1$	0.40	272.43
$\kappa = 0.3$	0.53	339.97
$\kappa = 0.5$	0.60	320.06
$\kappa = 1.0$	0.47	292.17
$\kappa = 1.5$	0.57	287.37
$\kappa = 3.0$	0.20	166.66

D.3 Hyperparameter Sensitivity

Additional experiments to research the sensitivity to hyperparameters are conducted. Table 5 shows the performance of UMBRELLA in the NGSIM environment for varying individual hyperparameters based on the configuration from Table 4.

While the most sensitive hyperparameter is β , the algorithm remains stable for different planning horizons H . β can be interpreted as a way to introduce momentum to the algorithm. A too low value makes the actions independent of previous actions, leading to behaviors with high jerk. However, a too high value makes the SDV less reactive. Future, work could also investigate the effect of a context dependent β value. As the work of Argenson and Dulac-Arnold [2021], UMBRELLA performance mostly degrades near extreme values. We note that the prior work evaluates the sensitivity based on rewards. However as discussed in Section 4.3 and the work Knox et al. [2021], this often does not necessarily correlate with a good driving performance.