

Safety Reinforced Model Predictive Control (SRMPC): Improving MPC with Reinforcement Learning for Motion Planning in Autonomous Driving

Johannes Fischer^{*1}, Marlon Steiner^{*1}, Ömer Şahin Taş² and Christoph Stiller¹

Abstract—Model predictive control (MPC) is widely used for motion planning, particularly in autonomous driving. Real-time capability of the planner requires utilizing convex approximation of optimal control problems (OCs) for the planner. However, such approximations confine the solution to a subspace, which might not contain the global optimum. To address this, we propose using safe reinforcement learning (SRL) to obtain a new and safe reference trajectory within MPC. By employing a learning-based approach, the MPC can explore solutions beyond the close neighborhood of the previous one, potentially finding global optima. We incorporate constrained reinforcement learning (CRL) to ensure safety in automated driving, using a handcrafted energy function-based safety index as the constraint objective to model safe and unsafe regions. Our approach utilizes a state-dependent Lagrangian multiplier, learned concurrently with the safe policy, to solve the CRL problem. Through experimentation in a highway scenario, we demonstrate the superiority of our approach over both MPC and SRL in terms of safety and performance measures.

I. INTRODUCTION

Autonomous vehicles have the potential to fundamentally change traffic by reducing risky driving behavior, congestion, carbon dioxide emissions, and transportation costs, as well as improving road safety [1]. An important area in the field of autonomous driving is motion planning, which aims to generate optimal trajectories in a given state that at the same time do not pose a threat to road traffic. Motion planning is often performed by model predictive control (MPC) [2]–[6]. For real-time applications, linear time-varying MPC (LTV-MPC) is often preferred because of the ability to faster solution calculation [7], [8]. Though, LTV-MPC only performs a local optimization which leads to the fact that only a locally optimal solution can be found.

Recent studies also focus on motion planning and control with reinforcement learning (RL). In RL applications it is a huge challenge to address constraint satisfaction, especially in

safety critical applications like autonomous driving. There is a lot of research trying to achieve safety in RL, which is called safe RL (SRL) [9]–[16]. Most of the latest approaches are not tested on driving tasks yet, and they often require further research. However, full autonomous driving is still confronted with the task to make feasible and especially safe decisions and planning in complex and uncertain dynamic traffic scenarios.

In order to address the existing issues, we present a novel approach which combines both methods, MPC and RL, to improve the motion planning in terms of safety and performance. This leads to a novel method that combines the mathematical optimization approach with a learning-based approach. Our main contributions are:

- The application of a SRL approach with state dependent Lagrange multiplier and safety index for autonomous driving tasks.
- Utilization of SRL for switching between local optima and ideally finding the global optimum, which serves as the reference for the locally optimizing MPC.
- Experimental results on a highway scenario, demonstrating that the proposed method achieves better results in terms of safety and performance than SRL and LTV-MPC.

II. RELATED WORK

To the best of our knowledge there are no other approaches, which use RL to improve the motion planning of an MPC controller by generating a reference trajectory. There are several approaches of so called *learning-based MPC*, which use machine learning to learn dynamic models or parameters of the MPC controller [17]–[21]. RL in combination with MPC is often used to imitate the control strategy of the MPC by *end-to-end learning* [22]–[24]. However, approaches where the MPC benefits from the learned value function of the RL algorithm can rarely be found in the literature. The authors in [25] integrate the RL value function into the objective function of the MPC. But they solved the optimal control problem (OCP) in a combinatorial way and not analytically. This makes the proposed approach only applicable for problems with a small and discrete action space. Another approach which is

^{*}Authors have equal contribution.

¹ Johannes Fischer, Marlon Steiner, Christoph Stiller are with the Institute of Measurement and Control Systems, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany. {johannes.fischer, marlon.steiner, stiller}@kit.edu

² Ömer Şahin Taş is with FZI Research Center for Information Technology, Karlsruhe, Germany. tas@fzi.de

closely related to ours is presented in [26]. In this paper the authors use RL to define a terminal area for the MPC, so that the controlled vehicle can be guided through an unknown environment. In contrast to [26], we do not define a terminal area but use RL to generate a whole trajectory as a reference respectively initial solution for the OCP. Furthermore, the learned trajectory in our work is not forced to be tracked by the MPC but can shift the optimization problem from one locally optimal solution to another.

III. FUNDAMENTALS

In this section, we formalize the definition of the MPC and of the RL problem as a constrained Markov decision process (CMDP) and present the algorithms that are used to solve it.

A. MPC and Linear Time-Varying MPC

MPC is a control strategy that predicts future system behavior and determines optimal control actions \mathbf{u} based on the modeled system dynamics. It computes control actions for a finite time horizon by solving an OCP, but only applying the first action to the system. The process is repeated at every time step, taking into account new measurements and updated predictions for the state \mathbf{x} . Due to often nonlinear constraints in the OCP within an autonomous driving task, a nonlinear MPC (NMPC) for the control of such a system is required. The NMPC solves an optimization problem

$$\min_{\mathbf{u}_{1:N}, \mathbf{x}_{1:N+1}} J(\mathbf{u}_{1:N}, \mathbf{x}_{1:N+1}) \quad (1a)$$

$$\text{s.t. } \mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k), \quad k = 1, \dots, N, \quad (1b)$$

$$\mathbf{x}_1 = \hat{\mathbf{x}}, \quad (1c)$$

$$\underline{\mathbf{g}}_{\text{col}} \leq \mathbf{g}_{\text{col}}(\mathbf{x}_k, \mathbf{u}_k) \leq \overline{\mathbf{g}}_{\text{col}}, \quad k = 2, \dots, N+1, \quad (1d)$$

$$\underline{\mathbf{u}} \leq \mathbf{u}_k \leq \overline{\mathbf{u}}, \quad k = 1, \dots, N, \quad (1e)$$

$$\underline{\mathbf{x}} \leq \mathbf{x}_k \leq \overline{\mathbf{x}}, \quad k = 1, \dots, N+1, \quad (1f)$$

in every time step specified by the objective function (1a), the system dynamics (1b), the initial state condition (1c), collision avoidance constraints (1d) denoted by general collision functions \mathbf{g}_{col} and box constraints for the control and state variables (1e)–(1f). In the OCP (1) the variable N represents the prediction horizon length.

Applying NMPC to an on-board system is challenging due to real-time constraints and the computational complexity of the optimization problem. LTV-MPC addresses this challenge by incorporating time-varying dynamics through the linearization of the optimization problem constraints at every time step. Together with a quadratic objective function to optimize, the problem becomes convex and a quadratic programming (QP) solver can be applied to solve it. However, this leads to the fact that only a locally optimal solution is calculated which depends on the trajectory used for linearization.

B. Constrained Markov Decision Process

Analogous to a Markov decision process (MDP) [27], a CMDP [28] is a model that describes a problem of sequential

decision-making between an agent and its environment. The difference is that the CMDP allows a clear separation of reward and safety signals. A CMDP can be defined as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}_0, r, c, d, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{P}_0: \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition probability kernel indicating the probability of state s' after taking action a in state s , $r: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, $c: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the cost function, d is the safety threshold and γ is the discount factor.

The goal in a CMDP is to find a policy π that maximizes the accumulated discounted reward

$$\max_{\pi} J_r(\pi) = \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right] \quad (2)$$

while keeping the accumulated discounted cost

$$J_c(\pi) = \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t c(s_t, a_t) \right] \leq d \quad (3)$$

bounded to d [29]. Here τ denotes a trajectory $\tau = (s_0, a_0, s_1, \dots)$ distributed according to the policy π . Equation (3) represents the constraint of the optimization problem. It is possible for the cost function to represent either a physically based function or an indicator function with discrete values for safety. In this context costs occur when the constraint is violated.

The set of feasible stationary policies that satisfy the constraint of the defined CMDP is then given by

$$\Pi_c \doteq \{\pi \in \Pi: J_c(\pi) \leq d\}. \quad (4)$$

C. Constrained Reinforcement Learning

Constrained RL (CRL) is a variant of classic RL which considers constraints in the MDP and can be classified as an approach for safe RL (SRL). The most commonly used CRL algorithm is the Lagrangian approach where the optimization criterion (2) is extended by the cost constraint (3) [30]. Examples for baseline CRL Lagrangian algorithms are PPO-Lagrangian and TRPO-Lagrangian [10]. For detailed information about the classification of SRL approaches we refer the reader to [9].

The goal of CRL is to learn a feasible optimal policy π^* regarding (4) that maximizes the objective J_r :

$$\pi^* = \arg \max_{\pi \in \Pi_c} J_r(\pi). \quad (5)$$

Within the Lagrangian approach for solving CMDPs, the objective function becomes a min-max optimization problem

$$\min_{\pi} \max_{\lambda \geq 0} \mathcal{L}(\pi, \lambda) = -J_r(\pi) + \lambda(J_c(\pi) - d) \quad (6)$$

to balance between reward and cost. The optimal solution to (6) is given by its saddle point.

D. Energy Function based Safety

In control theory a control strategy is considered as safe when unsafe regions of the state space are avoided and at the same time the safe region is forward invariant. The safety of a control can be evaluated by a safety certificate in form of a scalar energy function $\phi: \mathbb{R}^n \rightarrow \mathbb{R}$. A function like this is also called safety index (SI). In general those functions take on positive values to indicate unsafe states ($\phi > 0$) and take on negative values to indicate safe states ($\phi \leq 0$) [31].

To achieve safety, energy dissipation is required so that $\dot{\phi} < 0$ holds if the state is unsafe. According to this assumption a safe control follows the condition

$$\phi(s') < \max\{\phi(s) - \eta, 0\} \quad (7)$$

on the energy of the SI. In this context η is a small slack variable to force the energy to decrease if the current state s is unsafe. As soon as the state s gets safe the next state s' has to be safe, too. Safe states can be represented as a subset of all system states \mathcal{S} by a closed, connected safe set \mathcal{S}_s . The SI must be designed such that the set \mathcal{S}_s is a zero-sublevel set of $\phi: \mathcal{S} \rightarrow \mathbb{R}: \mathcal{S}_s = \{s \mid \phi(s) \leq 0\}$ [13].

Further, the following restrictions for the design of a feasible SI must hold:

- 1) The time derivative of the SI can be influenced by the control variables \mathbf{u} of the system: $\frac{\partial \phi}{\partial \mathbf{u}} \neq 0$
- 2) and the reachable set under ϕ is a subset of the safe set: $X(\phi) \subseteq \mathcal{X}_S$ [32].

IV. METHOD

In this section, we formulate the used MPC and RL methods independently to combine them afterwards into the proposed SRMPC approach.

A. MPC for Automated Driving

Subsequently, we define the relevant terms of the OCP (1). The used MPC scheme is state of the art, hence we only give a brief overview.

1) *Objective Function*: We use quadratic forms to define the objective function (1a) to be optimized. In this way, we obtain a convex function, which is a central property regarding the choice of the solution method. The objective function is defined such that a reference trajectory \mathbf{x}^{ref} is tracked and the control variables and their derivations are minimized:

$$J = \sum_{k=1}^{N+1} \|\mathbf{x}_k - \mathbf{x}_k^{\text{ref}}\|_{\mathbf{Q}}^2 + \sum_{k=1}^N \|\mathbf{u}_k\|_{\mathbf{R}}^2 + \|\mathbf{u}_1 - \hat{\mathbf{u}}\|_{\mathbf{S}}^2 + \sum_{k=2}^N \|\mathbf{u}_k - \mathbf{u}_{k-1}\|_{\mathbf{S}}^2 + J_{\text{slack}}, \quad (8)$$

where $\hat{\mathbf{u}}$ denotes the previously applied control. The term J_{slack} in equation (8) penalizes the solver when using slack variables. Slack variables are applied to the collision constraints (1d) to guarantee feasibility of the problem by relaxing the whole constraint. Additional slack variables are used to define a safety region around the vehicles which is only entered by the ego-vehicle if necessary [8].

2) *Vehicle Model*: For the system dynamics (1b) the kinematic bicycle model [33] is used, see Fig. 1. The purely kinematic

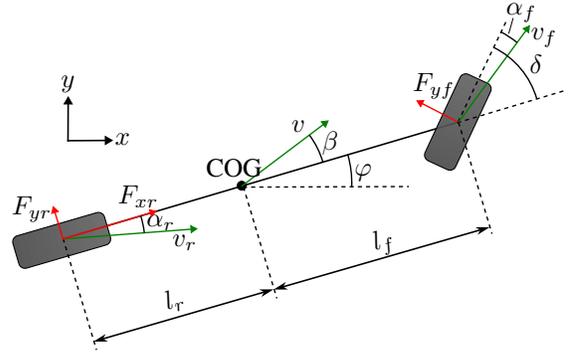


Figure 1: Bicycle model: The forces (red arrows) are not considered in the kinematic bicycle model and the sideslip angles are assumed to be $\alpha_{f,r} = 0$.

matic consideration of the vehicle dynamics can be expressed by equation (9):

$$\begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\varphi} \\ \dot{v} \end{pmatrix} = \begin{pmatrix} v \cos(\varphi + \beta) \\ v \sin(\varphi + \beta) \\ \frac{v}{l} \tan(\delta) \cos(\beta) \\ a \end{pmatrix} \quad (9)$$

with the slip angle β

$$\beta = \arctan\left(\frac{l_r}{l} \tan(\delta)\right). \quad (10)$$

Thereupon, the state variables \mathbf{x} are the global x - and y -coordinates, the heading angle φ and the velocity v . Concurrently the steering angle δ and the acceleration a function are the control variables \mathbf{u} .

3) *Collision constraints*: The collision constraints (1d) are the crucial part in the OCP concerning the safety. It turned out that over-approximating the ego-vehicle by a set of circles and other traffic participants by superellipses, leads to an easy to handle mathematical form as well as good results. The resulting collision constraint between the ego-vehicle and the j -th vehicle can be expressed by equation (11).

$$\begin{bmatrix} \Delta x_j \\ \Delta y_j \end{bmatrix}^T \mathbf{R}(\varphi^j)^T \begin{bmatrix} \frac{1}{a} & 0 \\ 0 & \frac{1}{b} \end{bmatrix}^2 \mathbf{R}(\varphi^j) \begin{bmatrix} \Delta x_j \\ \Delta y_j \end{bmatrix} > 1 \quad (11)$$

where \mathbf{R} is the rotation matrix in \mathbb{R}^2 , a and b are the sum of the semi-major respectively semi-minor axis and the radius of the circle [34].

4) *Approximate the nonlinear OCP*: We approximate the original OCP by a QP in every time step to make it applicable within the LTV-MPC method. Therefore, we use a second order Taylor series expansion to reformulate the objective function (1a) and a first order Taylor series expansion to linearize the constraints (1b)–(1f). The resulting QP has the form

$$\begin{aligned} \min_{\mathbf{z}} \quad & \frac{1}{2} \mathbf{z}^T \mathbf{H} \mathbf{z} + \mathbf{g}^T \mathbf{z} \\ \text{s.t.} \quad & \mathbf{l} \leq \mathbf{A} \mathbf{z} \leq \mathbf{u} \end{aligned} \quad (12)$$

with the new optimization variable \mathbf{z} . Assuming the weight matrices \mathbf{Q} , \mathbf{R} and \mathbf{S} are positive semidefinite and diagonal, the QP (12) is convex. The Taylor expansion requires a reference trajectory $\mathbf{z}^r = [\mathbf{x}_{1:N+1}^r; \mathbf{u}_{1:N}^r]$. To obtain such a reference, trajectory shifting [35] is applied to the previous solution. To clarify, this reference trajectory differs from \mathbf{x}^{ref} in equation (8). The QP is solved using the OSQP-Solver [36].

B. Defining the CMDP

In this section we define the elements of the constraint optimization task, the CMDP.

1) *Observation*: The observation approximates the state, and functions as the input of the policy and value functions. Fig. 2 shows an exemplary scene with multiple vehicles to explain which ones are considered in the observation.

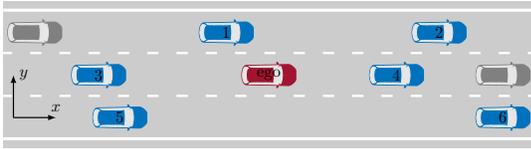


Figure 2: CMDP observation of the ego-vehicle (red): The ego-vehicle observes the vehicle in front of and the one behind itself on each lane (represented in blue). The vehicles shown in gray are not considered.

The matrices in (13) present all considered quantities within the observation.

$$\begin{bmatrix} \Delta x^1 & \Delta y^1 & \Delta v_x^1 & \Delta v_y^1 \\ \vdots & \vdots & \vdots & \vdots \\ \Delta x^6 & \Delta y^6 & \Delta v_x^6 & \Delta v_y^6 \end{bmatrix} \quad (13a)$$

$$\begin{bmatrix} v_x^{\text{ego}} & v_y^{\text{ego}} & \varphi^{\text{ego}} & \Delta y^{\text{rs,r}} & \Delta y^{\text{rs,l}} & \Delta y^{\text{target}} \end{bmatrix} \quad (13b)$$

Matrix (13a) poses the quantities for the six surrounded vehicles (see Fig. 2). The Δ expresses that the specific quantity is measured relative to the ego-vehicle. Information about the ego-vehicle state and relative coordinates to the left and right roadsides (rs) and the target lane can be found in Matrix (13b).

2) *Reward*: In our case the reward function r is based on the deviation from a longitudinal reference velocity v_x^{ref} and a reference lane y^{ref} . We define an auxiliary reward

$$r'(s, a) = -g(v_x^{\text{ref}} - v_x^{\text{ego}})^2 - (1 - g)(y^{\text{ref}} - y^{\text{ego}})^2 \quad (14)$$

In this equation g is a hyperparameter that balances the relation between speed reward and reward for driving on a specific lane. To achieve a better and more stable performance in training, the reward r is obtained by exponentiating the function r' and normalizing it to the interval $[0, 1]$.

3) *Actions*: The same control variables as in the MPC are used as actions to control the agent (see section IV-A2). But in contrast to the MPC, the RL algorithm outputs continuous actions due to the utilization of a stochastic policy.

C. Safety Index for Highway Driving

In this section we develop the cost function defining the CMDP. Therefore, we define SI based on the formulation from [11], [13], [31], [37]:

$$\phi(s) = (\sigma + d_{\text{min}})^n - d^n - k\dot{d} \quad (15)$$

with the tunable parameters σ , n and k , the minimum distance to the other vehicle d_{min} , the actual distance d and the velocity towards the other vehicle \dot{d} . We defined \dot{d} dependent to the lateral distance from the agent to another vehicle. When both are not on the same lane, the longitudinal velocity part of \dot{d} is set to zero to make the SI (15) applicable to the highway driving.

The time derivative of the chosen SI contains both action dimensions δ and a , so that the relative degree of the system is $r = 1$. Under the assumption that the parameters of the SI are chosen according to the safe set \mathcal{S}_s , the SI fulfills both requirements for a feasible SI (see section III-D). In this paper we do not formally proof the latter condition.

The cost function results from the condition for safe control (7):

$$c(s, a) = \max\left\{\phi(s') - \max\{\phi(s) - \eta, 0\}, -0.1\right\} \quad (16)$$

where cost occurs if the condition is not fulfilled. The outer max-function serves the purpose to prevent the algorithm from learning to become highly conservative.

D. State dependent Lagrangian Multiplier

As proposed in [38], we use a state dependent Lagrangian multiplier $\lambda(s)$ to handle the state dependent constraint arising from the SI. The value of the Lagrangian multiplier indicates the safety of the agent to be in a specific state. It gets approximated by a state value function which is learned through SI transitions by considering the *complementary slackness condition*. The condition comes from the Karush-Kuhn-Tucker (KKT) necessary conditions of optimality. Table I shows what

Table I: Relation between optimal multipliers λ^* and safety

$\lambda^*(s)$	Safety
Zero	Safe (inactive constraint)
Finite	On boundary (active constraint)
Infinite	Unsafe (infeasible constraint)

the value of the optimal multiplier implies for the safety. Numerically, the multiplier is considered infinite when exceeding a threshold we defined.

E. SRL Approach to Automated Driving

We use a state value function $v_\pi(s)$ to approximate the objective $J_r(\pi)$ (2) and another state value function $v_\pi^c(s)$ (referred to as cost value function) to approximate the associated constraint $J_c(\pi) - d$ (3). Together with the state dependent Lagrangian multiplier $\lambda(s)$ the loss function to be optimized can be rewritten as

$$\mathcal{L}(\pi, \lambda) = \mathbb{E}_s [-v_\pi(s) + \lambda(s)v_\pi^c(s)]. \quad (17)$$

To learn a policy that solves the optimization problem, the on-policy *proximal policy optimization* (PPO) [39] baseline algorithm is adapted by the new objective (17) and extended by the cost value function $v_{\pi}^c(s)$ and the state value function $\lambda(s)$ for the Lagrangian multiplier. We call the safe version of the PPO algorithm *PPO-Lagrangian-Safety Index* (PPO-L-SI).

F. Combining LTV-MPC and SRL

The limitation of LTV-MPC can be attributed to the approximating aspect of the solution computation [40]. Due to local optimization, LTV-MPC only identifies the optimal solution within the linearization range, thereby overlooking other locally optimal solutions. The linearization in LTV-MPC is based on the shifted solution from the previous time step, resulting in the dependency of subsequent time step solutions on the previous one. Although this approach is computationally efficient, it neglects other local optima and possibly the global optimum of the original optimization problem.

To address this limitation, an RL policy can be employed in LTV-MPC to generate a trajectory to be used for linearization. For the optimal use of the learned trajectories, it is crucial that the RL algorithm receives the same or a very similar optimization problem as the MPC for policy learning. Furthermore, both control algorithms should have access to the same control variables.

The trajectory generated by the policy for the planning horizon can be used to approximate the NLP as a QP. This provides an alternative reference trajectory compared to the shifted solution trajectory from the previous MPC step, allowing the LTV-MPC to discover a new local and potentially the global optimum. This approach also bypasses the need for a handcrafted initial trajectory for the first time step, $t = 0$, of the MPC. For example, a trajectory following the lane with constant velocity or a braking trajectory can be used as the initial trajectory. The proposed SRMPC approach is summarized in algorithm 1.

Fig. 3 shows an example of an RL trajectory based solution calculated by the LTV-MPC. The figure illustrates the state trajectories as well as the control trajectories.

V. RESULTS AND EVALUATION

A. Experimental Set-Up

For our evaluations we use the *highway-env* [41] framework, which provides environments for tactical decision-making in different automated driving tasks. Within this framework, the agent controls the ego vehicle, while the other vehicles follow the Intelligent Driver Model (IDM) and only react to the ego vehicle once it enters their lane. We focus on a three lane highway driving scenario as illustrated in Fig. 2. At the start of each episode, the agent and the other traffic participants are generated randomly on one of the lanes with a random velocity within a defined range. The lane changes of the IDM vehicles are disabled. Furthermore, we initialize the MPC controller with a trajectory with constant velocity.

Algorithm 1: Pseudocode for SRMPC

```

Input: policy  $\pi$ , prediction horizon length  $N$ 
1  $k \leftarrow 0$ ;
2 while goal is not reached do
3    $\mathbf{x}_k \leftarrow$  measure current state;
4   try calculate trajectory with  $\pi$  for linearization:
5      $\mathbf{x}_{k:k+N+1}^r, \mathbf{u}_{k:k+N}^r \leftarrow$  use  $\pi$  to simulate
       environment state forward and save state and
       control values;
6   catch error occurred:
7      $\mathbf{x}_{k:k+N+1}^r, \mathbf{u}_{k:k+N}^r \leftarrow$  use shifted trajectory
       from previous solution
8   end
9   NLP define and discretize the time-dependent
       OCP (1);
10   $\mathbf{H}, \mathbf{g}, \mathbf{A}, \mathbf{l}, \mathbf{u}$  calculate QP matrices and vectors
       (see equation (12)) from NLP with  $\mathbf{x}_k$ ,
        $\mathbf{x}_{k:k+N+1}^r, \mathbf{u}_{k:k+N}^r$  and  $\mathbf{x}_{k:k+N+1}^{\text{ref}}$ ;
11   $\mathbf{u}_{k+1}$  solve QP to get the optimal control
       value;
       Output:  $\mathbf{u}_{k+1}$ 
12   $k \leftarrow k + 1$ ;
13 end

```

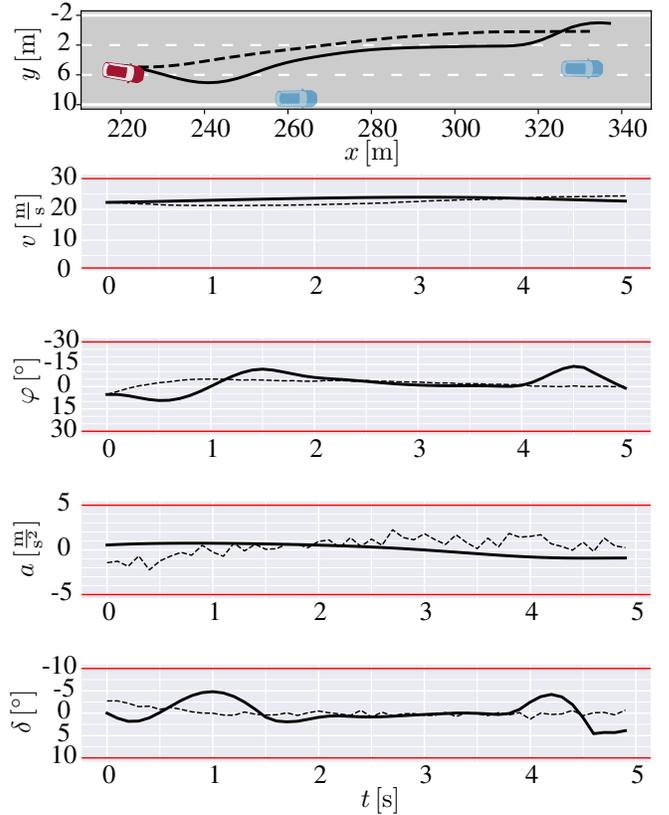


Figure 3: SRL trajectory (dashed lines) used for linearization within the LTV-MPC and the locally optimal solution (continuous lines) of the MPC based on the SRL trajectory.

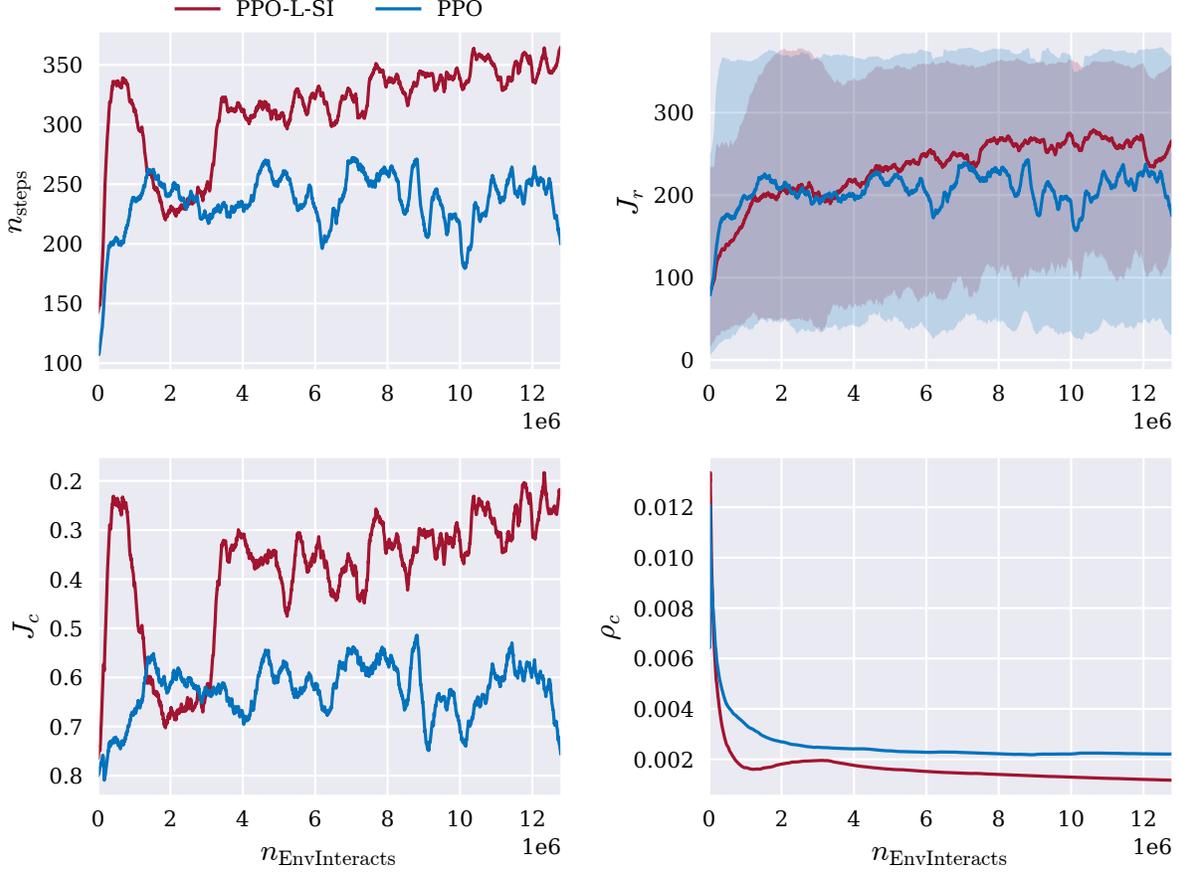


Figure 4: Training results for the baseline RL algorithm PPO and the SRL algorithm PPO-L-SI.

In both, MPC and RL, the middle of the rightmost lane is defined as optimal by means of the objective function and the reward function, respectively ($y^{\text{ref}} = \text{middle of right lane}$). Therefore, the agent will tend to drive on the right lane. Furthermore, the longitudinal velocity goal is defined as $v_x^{\text{ref}} = 25 \frac{\text{m}}{\text{s}}$.

B. Metrics

To evaluate and compare the described control methods with the proposed SRMPC, this section outlines the metrics. Regarding safety and performance, the following metrics are considered:

- Episodic return J_r : This metric uses the unprocessed reward function r' (14) to describe the performance of the agents.
- Episodic cost return J_c : For better comparability this metric considers only actual collisions and road exits as a cost ($J_c = 1$). This metric can only assume the values $J_c \in \{0, 1\}$, since the episode terminates when costs occur.
- Cost rate ρ_c : This metric describes the relationship between the total accumulated costs and the accumulated number of environment steps over all simulations.

- Episodic steps n_{steps} : The metric describes the number of steps taken by the agent without the occurrence of costs. n_{steps} can take on values in the interval $n_{\text{steps}} \in [0, 400]$ due to the maximum episode length.
- Average longitudinal velocity \bar{v}_x : A performance measure for comparison with the longitudinal reference velocity v_x^{ref} .

The episodic return J_r and average velocity \bar{v}_x are referred to as performance metrics and the episodic cost return J_c , the cost rate ρ_c and the episodic steps n_{steps} are referred to as safety metrics.

C. Results

First we evaluate the PPO-L-SI against the baseline PPO algorithm in training to demonstrate the effectiveness regarding safety. The training is shown in Fig. 4. Here, the values of the cost rate J_c are modified by a running mean filter to see a clearer trend of collisions respectively road exits. The filtered metric can be interpreted as the failure rate of the agent. It can be seen that the safety metrics of the PPO-L-SI take better values than those of the PPO algorithm almost over the entire training. Especially at the end of training the cost return J_c and cost rate ρ_c from PPO-L-SI are more than twice

as good as those from PPO. Also, the performance metric episodic return J_r of the PPO-L-SI lies above that of PPO almost during the entire training. Accordingly, the PPO-L-SI achieves much better results as the baseline PPO algorithm. In particular when looking at the safety metrics.

In the next step we evaluate the learning based PPO-L-SI and the optimization based LTV-MPC against the combined approach SRMPC. Therefore, we simulate all algorithms over 1000 episodes and track all described metrics from section V-B. In contrast to the RL training evaluation (see Fig. 4) we now use mean values for the comparison. The cost rate ρ_c already is a metric tracking a running mean. The metrics for each control method are reported from table II. All three methods are compared in a scenario with light and dense traffic volume. In light traffic the vehicles are randomly generated in one of the three lanes at a distance of around 25m from each other whereas in dense traffic the distance reduces to roughly 15m.

Table II: Evaluation results: Comparison between the SRL algorithm PPO-L-SI, LTV-MPC and SRMPC based on a highway driving scenario with light and dense traffic.

	Algorithm	\bar{J}_r (per step)	\bar{J}_c	ρ_c	\bar{n}_{steps}	\bar{v}_x
light	PPO-L-SI	-4669 (-13.4)	0.308	0.0009	348	23.11
	LTV-MPC	-1241 (-5.1)	0.625	0.0026	244	24.38
	SRMPC	-1108 (-3.1)	0.215	0.0006	357	24.63
dense	PPO-L-SI	-4499 (-16.4)	0.540	0.002	274	22.47
	LTV-MPC	-334 (-2.4)	0.788	0.0057	139	24.79
	SRMPC	-1211 (-4.0)	0.353	0.0012	305	24.47

When evaluating the light traffic case, the SRMPC outperforms both methods, the PPO-L-SI and the LTV-MPC in terms of safety and performance. The proposed SRMPC reaches a higher return and a higher average velocity compared to PPO-L-SI and LTV-MPC. Regarding the cost return and cost rate, the SRMPC achieves more than 30% better results compared to PPO-L-SI and more than 65% better results than LTV-MPC.

The evaluation results at dense traffic differ a little from those at light traffic. Nevertheless, the SRMPC outperforms the two methods on which it is based in terms of safety. But the LTV-MPC achieves slightly better results in terms of performance. However, the better results are not decisive so that the performance with respect to the return per step and average velocity is comparable to that of the SRMPC.

D. Computation Time

The computation time of the SRMPC increases compared to the LTV-MPC due to the calculation of the reference trajectory using the SRL algorithm. This step requires the evaluation of the policy and the forward simulation of the environment by the chosen horizon length. Since the used simulation environment was not optimized for simulation performance, the computation time of the forward simulation is relatively high. Hence, we did not run quantitative experiments to compare the computation time of the SRMPC with the LTV-MPC. To

achieve real-time capability, we propose using a small policy network architecture and an efficient prediction framework for the forward simulation.

VI. CONCLUSIONS AND FUTURE WORK

We have demonstrated the effective use of RL to enhance the local optimization of LTV-MPC. The resulting control method outperforms the underlying RL and LTV-MPC methods in terms of safety and performance for the highway driving task in light traffic. In dense traffic, the developed method achieves superior results with regard to safety and shows roughly similar results in the performance compared to the LTV-MPC.

The utilization of an energy-based function for expressing and learning safety has shown promising results, suggesting that further research should be done on the CRL approach with SI and a state dependent Lagrangian multiplier. However, a possible reformulation of the SI is worth considering, aligning it with the mathematical formulation of e.g. Responsibility-Sensitive Safety (RSS) [42] for autonomous vehicles. Furthermore, it is conceivable to extend the proposed method to NMPC by utilizing an RL trajectory as the initial solution of an NLP solver to shorten the computation time.

REFERENCES

- [1] Y. Li and H. Shi, Eds., *Advanced Driver Assistance Systems and Autonomous Vehicles : From Fundamentals to Applications*, 1st ed. 2022., ser. Springer eBook Collection. Singapore: Springer Nature Singapore, 2022, ISBN: 9789811950537.
- [2] M. A. Mousavi, Z. Heshmati, and B. Moshiri, "LTV-MPC based path planning of an autonomous vehicle via convex optimization," in *2013 21st Iranian Conference on Electrical Engineering (ICEE)*, 2013, pp. 1–7.
- [3] G. Inghilterra, S. Arrigoni, F. Braghin, and F. Cheli, "Firefly algorithm-based nonlinear mpc trajectory planner for autonomous driving," in *2018 International Conference of Electrical and Electronic Technologies for Automotive*, 2018, pp. 1–6.
- [4] X. Yu, W. Zhu, and L. Xu, "Real-time motion planning and trajectory tracking in complex environments based on bézier curves and nonlinear mpc controller," in *2020 Chinese Control And Decision Conference (CCDC)*, 2020, pp. 1540–1546.
- [5] L. Wang, G. Yin, L. Xu, F. Dong, M. Xu, and Z. a. He, "Motion planning framework for autonomous vehicle with surrounding vehicle motion prediction in highway scenarios," in *2022 6th CAA International Conference on Vehicular Control and Intelligence (CVCI)*, 2022, pp. 1–6.
- [6] Ö. Ş. Taş, "Motion planning for autonomous vehicles in partially observable environments," PhD thesis, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany, 2022.
- [7] P. F. Lima, J. Mårtensson, and B. Wahlberg, "Stability conditions for linear time-varying model predictive control in autonomous driving," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017, pp. 2775–2782.
- [8] V. Jain, U. Kolbe, G. Breuel, and C. Stiller, "Reacting to multi-obstacle emergency scenarios using linear time varying model predictive control," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1822–1829.
- [9] J. García and F. Fernández, "A comprehensive survey on safe reinforcement learning," *Journal of Machine Learning Research*, vol. 16, no. 42, pp. 1437–1480, 2015.

- [10] A. Ray, J. Achiam, and D. Amodei, *Benchmarking safe exploration in deep reinforcement learning*, 2019.
- [11] H. Ma, C. Liu, S. E. Li, S. Zheng, W. Sun, and J. Chen, *Learn zero-constraint-violation policy in model-free constrained reinforcement learning*, 2021.
- [12] Z. Qin, K. Zhang, Y. Chen, J. Chen, and C. Fan, *Learning safe multi-agent control with decentralized neural barrier certificates*, 2021.
- [13] H. Ma, C. Liu, S. E. Li, S. Zheng, and J. Chen, “Joint synthesis of safety certificate and safe control policy using constrained reinforcement learning,” in *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, et al., Eds., ser. Proceedings of Machine Learning Research, vol. 168, PMLR, Jun. 2022, pp. 97–109.
- [14] J. Fischer, C. Eyberg, M. Werling, and M. Lauer, “Sampling-based Inverse Reinforcement Learning Algorithms with Safety Constraints,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Prague, Czech Republic, Sep. 2021, pp. 791–798.
- [15] D. Kamran, T. Engelgeh, M. Busch, J. Fischer, and C. Stiller, “Minimizing Safety Interference for Safe and Comfortable Automated Driving with Distributional Reinforcement Learning,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Prague, Czech Republic, Sep. 2021, pp. 1236–1243.
- [16] D. Kamran, T. D. Simão, Q. Yang, C. T. Ponnambalam, J. Fischer, M. T. Spaan, et al., “A Modern Perspective on Safe Automated Driving for Different Traffic Dynamics Using Constrained Reinforcement Learning,” in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, Macau, China, Oct. 2022, pp. 4017–4023.
- [17] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, et al., “Information theoretic MPC for model-based reinforcement learning,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 1714–1721.
- [18] T. X. Nghiem, “Linearized gaussian processes for fast data-driven model predictive control,” in *2019 American Control Conference (ACC)*, 2019, pp. 1629–1634.
- [19] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, “Provably safe and robust learning-based model predictive control,” *Automatica*, vol. 49, no. 5, pp. 1216–1226, 2013, ISSN: 0005-1098.
- [20] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, “Robust constrained learning-based NMPC enabling reliable mobile robot path tracking,” *The International Journal of Robotics Research*, vol. 35, no. 13, pp. 1547–1563, 2016.
- [21] T. Koller, F. Berkenkamp, M. Turchetta, J. Boedecker, and A. Krause, *Learning-based model predictive control for safe exploration and reinforcement learning*, 2019.
- [22] B. Amos, I. D. J. Rodriguez, J. Sacks, B. Boots, and J. Z. Kolter, *Differentiable MPC for end-to-end planning and control*, 2018.
- [23] B. Chen, Z. Cai, and M. Bergés, “Gnu-RL: A practical and scalable reinforcement learning solution for building hvac control using a differentiable MPC policy,” *Frontiers in Built Environment*, vol. 6, 2020, ISSN: 2297-3362.
- [24] J. Drgoňa, A. Tuor, E. Skomski, S. Vasisht, and D. Vrabie, “Deep learning explicit differentiable predictive control laws for buildings,” *IFAC-PapersOnLine*, vol. 54, no. 6, pp. 14–19, 2021, 7th IFAC Conference on Nonlinear Model Predictive Control NMPC 2021, ISSN: 2405-8963.
- [25] J. Arroyo, C. Manna, F. Spiessens, and L. Helsen, “Reinforced model predictive control (RL-MPC) for building energy management,” *Applied Energy*, vol. 309, p. 118346, 2022, ISSN: 0306-2619.
- [26] C. Greatwood and A. G. Richards, “Reinforcement learning and model predictive control for robust embedded quadrotor guidance and control,” *Autonomous Robots*, vol. 43, no. 7, pp. 1681–1693, 2019.
- [27] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, Second, ser. Adaptive computation and machine learning. Cambridge, Massachusetts: Westchester Publishing Services and The MIT Press, 2018, ISBN: 9780262039246.
- [28] E. Altman, *Constrained Markov Decision Processes*. CRC Press, 1999, vol. 7.
- [29] J. Achiam, D. Held, A. Tamar, and P. Abbeel, *Constrained policy optimization*, 2017.
- [30] Z. Liu, Z. Cen, V. Isenbaev, W. Liu, Z. S. Wu, B. Li, et al., *Constrained variational policy optimization for safe reinforcement learning*, 2022.
- [31] T. Wei and C. Liu, *Safe control algorithms using energy functions: A unified framework, benchmark, and new directions*, 2019.
- [32] C. Noren, W. Zhao, and C. Liu, “Safe adaptation with multiplicative uncertainties using robust safe set algorithm,” *IFAC-PapersOnLine*, vol. 54, no. 20, pp. 360–365, 2021, Modeling, Estimation and Control Conference MECC 2021, ISSN: 2405-8963.
- [33] P. Polack, F. Alché, B. d’Andréa-Novel, and A. de La Fortelle, “The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles?” In *2017 IEEE Intelligent Vehicles Symposium (IV)*, 2017, pp. 812–818.
- [34] B. Brito, B. Floor, L. Ferranti, and J. Alonso-Mora, *Model predictive contouring control for collision avoidance in unstructured dynamic environments*, 2020.
- [35] M. Diehl, H. J. Ferreau, and N. Haverbeke, “Efficient numerical methods for nonlinear MPC and moving horizon estimation,” in *Nonlinear Model Predictive Control: Towards New Challenging Applications*, L. Magni, D. M. Raimondo, and F. Allgöwer, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 391–417, ISBN: 978-3-642-01094-1.
- [36] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, “OSQP: An operator splitting solver for quadratic programs,” *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020.
- [37] C. Liu and M. Tomizuka, “Control in a safe set: Addressing safety in human-robot interactions,” in *Proceedings of the ASME 2014 Dynamic Systems and Control Conference*, V003T42A003, vol. 3, Oct. 2014.
- [38] H. Ma, Y. Guan, S. E. Li, X. Zhang, S. Zheng, and J. Chen, *Feasible actor-critic: Constrained reinforcement learning for ensuring statewise safety*, 2021.
- [39] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, *Proximal policy optimization algorithms*, 2017.
- [40] P. Bender, Ö. Ş. Taş, J. Ziegler, and C. Stiller, “The combinatorial aspect of motion planning: Maneuver variants in structured environments,” in *2015 IEEE Intelligent Vehicles Symposium (IV)*, 2015, pp. 1386–1392.
- [41] E. Leurent, *An environment for autonomous driving decision-making*, <https://github.com/eleurent/highway-env>, 2018.
- [42] S. Shalev-Shwartz, S. Shammah, and A. Shashua, “On a formal model of safe and scalable self-driving cars,” *CoRR*, vol. abs/1708.06374, 2017. arXiv: 1708.06374.