Delft University of Technology

Delft Center for Systems and Control

Technical report 23-031

Learning safety in model-based reinforcement learning using MPC and Gaussian processes*

F. Airaldi, B. De Schutter, and A. Dabiri

If you want to cite this report, please use the following reference instead: F. Airaldi, B. De Schutter, and A. Dabiri, "Learning safety in model-based reinforcement learning using MPC and Gaussian processes," *Proceedings of the 22nd IFAC World Congress*, Yokohama, Japan, pp. 5759–5764, July 2023. doi:10.1016/j.ifacol. 2023.10.563

Delft Center for Systems and Control Delft University of Technology Mekelweg 2, 2628 CD Delft The Netherlands phone: +31-15-278.24.73 (secretary) URL: https://www.dcsc.tudelft.nl

* This report can also be downloaded via https://pub.bartdeschutter.org/abs/23_031.html

Learning safety in model-based Reinforcement Learning using MPC and Gaussian Processes

Filippo Airaldi Bart De Schutter Azita Dabiri

Delft Center for Systems and Control, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands Emails: {f.airaldi,b.deschutter,a.dabiri}@tudelft.nl

Abstract: This paper proposes a method to encourage safety in Model Predictive Control (MPC)-based Reinforcement Learning (RL) via Gaussian Process (GP) regression. The framework consists of 1) a parametric MPC scheme that is employed as model-based controller with approximate knowledge on the real system's dynamics, 2) an episodic RL algorithm tasked with adjusting the MPC parametrization in order to increase its performance, and 3) GP regressors used to estimate, directly from data, constraints on the MPC parameters capable of predicting, up to some probability, whether the parametrization is likely to yield a safe or unsafe policy. These constraints are then enforced onto the RL updates in an effort to enhance the learning method with a probabilistic safety mechanism. Compared to other recent publications combining safe RL with MPC, our method does not require further assumptions on, e.g., the prediction model in order to retain computational tractability. We illustrate the results of our method in a numerical example on the control of a quadrotor drone in a safety-critical environment.

Keywords: Learning-based Model Predictive Control, Safe Reinforcement Learning, Gaussian Processes

1. INTRODUCTION

The modern proliferation of advances in Machine Learning has led to a growing interest in learning-based control methodologies that strive to learn different elements of the controller scheme from data in order to, e.g., reduce conservativeness and improve closed-loop performance, or encourage safety and robustness (Brunke et al., 2022). In particular, Model Predictive Control (MPC) is the subject of a large amount of the current research in this field thanks to its wide variety of applications and high versatility, especially in regard to multivariate and constrained systems (Hewing et al., 2020b). While robust and stochastic MPC schemes allow for a systematic handling of different sources of uncertainties affecting the MPC performance, they often follow a strict separation between the offline design phase and the closed-loop application, where in most cases the controller remains unresponsive to changes in, e.g., the system or task. Adaptive and learningbased methodologies aim at overcoming this paradigm by automating the controller design and adaptation based on data collected during operation (Mesbah et al., 2022).

To this end, several tools from ML have been successfully coupled with MPC in order to enhance its performance or augment its capabilities. Among all, we can refer to Gaussian Process (GP) regression, which has been widely employed to, e.g., automatically learn the unmodeled (often nonlinear) dynamics in the prediction model (Hewing et al., 2020a) and to construct provably accurate confidence intervals on predicted trajectories (Koller et al., 2018); and black-box Bayesian Optimization, which is a widespread tool for automatic tuning and optimization of the MPC controller hyperparameters (Piga et al., 2019), thus relieving some burden on the design phase. Equally interesting is the combination of MPC with Reinforcement Learning (RL), another paradigm of ML gaining more and more attention thanks to its recent developments, in which an agent learn how to interact with the environment and take actions so as to maximize/minimize a reward/cost signal (Sutton and Barto, 2018). The majority of RL algorithms that are nowadays popular in the ML community are model-free and rely purely on observed state transitions and realizations of stage cost to increase the performance of the control policy. Deep Neural Networks (DNN) are predominantly used as function approximators and, while they have proven effective, they lack interpretability, and their behavior is difficult to analyze, certify, and trust in regard to stability and constraint satisfaction.

Recently, Gros and Zanon (2020) proposed and justified the use of MPC as the function approximation of the optimal policy in model-based RL. In such a scheme, the MPC controller yields the policy and value functions via its optimization problem, while, at the same time, the learning algorithm is tasked with adjusting the parametrization of the controller in an effort to implicitly or explicitly discover the optimal policy, thus improving closed-loop performance in a data-driven fashion. Via sensitivity analysis (Büskens and Maurer, 2001), it is possible to apply RL algorithms such as Q-learning and stochastic/deterministic policy gradient by differentiating the MPC optimization problem with respect to its parameters and updating these accordingly. In contrast to DNN-based RL, MPC-based RL agents can ideally exploit the underlying optimization problem to provide stability requirements by construction and guarantee safety, which defines the ability of a control policy to yield state-action trajectories that do not violate system constraints. However, even in this novel paradigm, certificates on safety and stability remain an arduous challenge. If proper care is not taken, the RL updates can still potentially jeopardize the MPC scheme by updating its parameters to an undesirable configuration, resulting in a controller with unsatisfactory behavior, e.g., that does violate constraints and hence is unsafe.

In this paper, we propose and explore a learning method in the context of safe MPC-based RL where the estimation of the safety set (i.e., the set of MPC parameters that yields a controller able to attain constraint satisfaction) is carried out in a data-driven fashion. In particular, we borrow a technique for estimating unknown constraints from the fields of constrained Bayesian Optimization (Sorourifar et al., 2021) and safe exploration (Schreiter et al., 2015), which offer tools to tackle black-box optimization problems that are constrained by some unknown functions in the decision variables. As a consequence, we are able to learn in a probabilistic way which configurations of the MPC controller are safe or not directly from data.

The rest of this article is organized as follows. Background and motivations both for safe RL in combination with MPC and for unknown constraint modeling via GP are given in Section 2. The proposed algorithm and details of its implementation are discussed in Section 3. A numerical experiment showing the effectiveness of the proposed method in a safe-critical quadrotor control application is reported and analyzed in Section 4. Lastly, conclusions and directions for future research are presented in Section 5. More details on are available in our extended version (Airaldi et al., 2022).

2. BACKGROUND

2.1 MPC as Function Approximation in Safe RL

As in Sutton and Barto (2018), we describe discretetime system dynamics as a Markov Decision Process with continuous state $s \in \mathbb{R}^{n_s}$ and continuous action $a \in \mathbb{R}^{n_a}$, and state transitions $s, a \to s_+$ with the underlying conditional probability density

$$\mathbb{P}\left[s_{+}|s,a\right]:\mathbb{R}^{n_{s}}\times\mathbb{R}^{n_{s}}\times\mathbb{R}^{n_{a}}\to\left[0,1\right].$$
(1)

Considering a deterministic policy $\pi_{\theta}(s) : \mathbb{R}^{n_s} \to \mathbb{R}^{n_a}$ parametrized in $\theta \in \mathbb{R}^{n_{\theta}}$ and resulting in the state distribution $\tau_{\pi_{\theta}}$. The performance of such policy is defined as

$$J(\pi_{\theta}) \coloneqq \mathbb{E}_{\tau_{\pi_{\theta}}} \left[\sum_{k=0}^{\infty} \gamma^{k} L(s_{k}, \pi_{\theta}(s_{k})) \right], \qquad (2)$$

where s_k is the state at time step k, $L(s, a) : \mathbb{R}^{n_s} \times \mathbb{R}^{n_a} \to \mathbb{R}$ is the stage cost and $\gamma \in (0, 1]$ is the discount factor. The RL problem is then to find the optimal policy π_{θ}^{\star} as

$$\pi_{\theta}^{\star} = \underset{\theta}{\arg\min} \quad J(\pi_{\theta}).$$
(3)

While DNNs are possibly the most common choice for representing such policy, an MPC scheme can be exploited as function approximation as proposed in Gros and Zanon (2020). Consider the following MPC problem approximating the value function $V_{\theta} : \mathbb{R}^{n_s} \to \mathbb{R}$ as

$$V_{\theta}(s) = \min_{u, x, \sigma} \quad \lambda_{\theta}(x_0) + \sum_{k=0}^{N-1} \gamma^k \left(L_{\theta}(x_k, u_k) + w^{\top} \sigma_k \right) \\ + \gamma^N \left(V_{\theta}^{\mathsf{f}}(x_N) + w_{\mathsf{f}}^{\top} \sigma_N \right)$$
(4a)

s.t.

x

$$x_{k+1} = f_{\theta}(x_k, u_k), \tag{4b}$$

$$h_{\theta}(x_k, u_k) \le \sigma_k, \tag{4c}$$

$$h_{\theta}^{\mathrm{f}}(x_N) \le \sigma_N,\tag{4d}$$

$$_0 = s$$
 (4e)

where vectors x, u and σ respectively collect states, actions and slack variables over the horizon N. In (4a), $\lambda_{\theta}(x_0)$ is an initial cost term, $L_{\theta}(x, u)$ is the stage cost, and $V_{\theta}^{\rm f}(x)$ is a terminal cost approximation. $f_{\theta}(x, u)$ is the model approximation, and $h_{\theta}(x, u), h_{\theta}^{\rm f}(x)$ are inequality constraints. Lastly, w and $w_{\rm f}$ are the weights of the slack variable in the objective. The value function (4) satisfies the fundamental equalities of the Bellman equations, so $Q_{\theta}(s, a)$ and $\pi_{\theta}(s)$ can be defined similarly.

Therefore, in RL terms, the MPC parametric scheme acts as policy provider for the learning agent, whose goal is to modify the parameters θ of the controller in an effort to minimize (3). Various forms of RL exist that solve this problem directly or indirectly via iterative updates

$$\theta \leftarrow \theta - \alpha \nabla_{\theta} \sum_{k=0}^{m} \psi(s_k, a_k, s_{k+1}, \theta), \tag{5}$$

where $\alpha \in \mathbb{R}_+$ is the learning rate, *m* denotes the number of observations used in the update (i.e., a batch of observations), and ψ captures the controller's performance and varies with the specific RL algorithm.

While in most learning algorithms a safe policy is usually implicitly achieved at convergence assuming violations are appropriately penalized, naively performing updates (5) does not guarantee such property during the learning process itself. Thus, we would like to restrict the search of θ to a set S that ensures the safety of the corresponding parametric policy. More specifically, consider a generic episodic task of duration T with n_c safety-critical constraints, gathered as the function $h : \mathbb{R}^{n_s} \times \mathbb{R}^{n_a} \to \mathbb{R}^{n_c}$ (e.g., bounds on the temperature of a chemical reactor, or obstacles in autonomous navigation). By deploying the MPC policy π_{θ} to accomplish the task, safety amounts to satisfying

$$h(x_t, \pi_\theta(x_t)) \le 0, \ t = 1, \dots, T.$$
(6)

Thus, we can define the safe set \mathcal{S} as

$$S \coloneqq \{\theta \mid (6) \text{ holds}\},\tag{7}$$

leading to the following reformulation of (5) as a constrained update

$$\min_{\theta_{+}} \quad \frac{1}{2} \|\theta_{+} - \theta\|_{2}^{2} + \alpha \nabla_{\theta} \sum_{k=0}^{m} \psi(s_{k}, a_{k}, s_{k+1}, \theta)^{\top} (\theta_{+} - \theta)$$
(8a)

s.t.
$$\theta_+ \in \mathcal{S},$$
 (8b)

by then setting $\theta \leftarrow \theta_+^*$, where θ_+^* is the optimal point of (8). Nonetheless, the characterization of S is in general difficult, even more so in the case of function approximators with vast parametrizations. Fortunately, an MPC approximator can be leveraged here: as shown in Zanon and Gros (2021) and further analyzed in terms of feasibility and stability in Gros and Zanon (2022), a robust MPC scheme coupled with an online set-membership system identification approach allows to guarantee that the MPC-RL algorithm is robustly safe with respect to the estimated disturbance set throughout the whole training, despite the fact that the dynamics of the real system are unknown to the learning agent. However, even though an adaptive robust scheme is deployed, safety is only guaranteed up to a probability smaller than one, due to the lack of perfect knowledge of the underlying system (1). In this context, safety can be viewed in the light of Bayesian inference, where it is probabilistically conditioned on our limited knowledge of the system, i.e., prior information and actual data, labeled as \mathcal{D} , with probability β . Henceforth, we will refer to safety in probabilistic terms and define $\mathcal{S}_{\mathcal{D}}$ as

$$\mathcal{S}_{\mathcal{D}} \coloneqq \{\theta \mid \mathbb{P}\left[(6) | \mathcal{D}\right] \ge \beta\}.$$
(9)

A limitation in the approach proposed by Zanon and Gros (2021) is that, in order to retain computational tractability in the algorithm, the MPC prediction model f_{θ} has to be linear and kept fixed during learning. As acknowledged in the paper (Zanon and Gros, 2021), while linear MPC applied to nonlinear systems can still be successful, strong nonlinearities are likely to deteriorate its performance, compromise safety, and prevent the agent from tuning the prediction model, which might severely hinder learning and lead to convergence to suboptimal policies.

In this paper, we propose a method that estimates the safe set $S_{\mathcal{D}}$ directly from episodic data in a way that bypasses the aforementioned limitation and that conforms with the probabilistic nature of the issue of safety itself. As detailed next, we propose to leverage GP regression tools that deliver surrogate models that are able to provide uncertainty estimates whether a specific θ belongs to $S_{\mathcal{D}}$ or not. This will then allow us to formulate an RL algorithm that can learn while being safe, up to some probability.

2.2 Estimation of Unknown Constraints

In essence, the matter of the estimation of set $S_{\mathcal{D}}$ in order to impose safety coincides with estimating whether a given state-action trajectory, produced by an MPC controller parametrized in θ , will satisfy constraints $h(x_t, u_t) \leq$ $0, t = 1, \ldots, T$, with $u_t = \pi_{\theta}(x_t)$. However, while on the one hand the relationship h(x, u) is given by the task, on the other hand we cannot predict beforehand whether $\pi_{\theta}(x)$ will yield safe control actions and hence safe trajectories. Therefore, we introduce the following safety constraint function $z : \mathbb{R}^{n_{\theta}} \to \mathbb{R}^{n_c}$

$$z(\theta) = \max_{t=1,\dots,T} h\big(x_t, \pi_\theta(x_t)\big) \tag{10}$$

that directly maps parametrizations to positive values, in case θ resulted in trajectories with violations, or to negative values, in case of safe trajectories. Analytically, this function is unknown to us, but each time a stateaction trajectory for the given task is computed, it can be evaluated for the current θ . Therefore, by leveraging past observed data, an approximation of (10) can be learned and used to predict safety at any previously unobserved query point θ^* . Thus, we could use this surrogate model of $z(\theta)$ to constrain learning only to those parameters that are estimated to be safe. Assumption 1. We assume the trajectory to be dependent only on θ , and, consequently, that $z(\theta)$ wholly captures safety (i.e., no other factor has an impact).

Remark 1. In general, this is not the case as also the initial and final conditions of the task, as well as the initialization of the optimization solver in case of nonlinear MPC, will affect constraint satisfaction. To mitigate these dependencies, we assume the episodic task to be repetitive, e.g., to have the same initial and final conditions, and the MPC problem to be convex. Analogously, we could also assume these factors to be readily included into the MPC parametrization θ so that, in principle, all parameters affecting safety may be included in $z(\theta)$ and made learnable.

Remark 2. In (10), the maximum over the time horizon T is taken in order to reduce violations along a trajectory to a point-wise maximum violation, but this does not invalidate (9) since $z(\theta) \leq 0 \Leftrightarrow h(x_t, \pi_\theta(x_t)) \leq 0, t = 1, \dots, T$. Likewise, functions other than max could be crafted to quantify safety of a trajectory while still encoding (9).

Unknown constraint estimation has been already tackled in literature. In constrained Bayesian Optimization, Sorourifar et al. (2021) propose a method for black-box optimization with unknown constraints that can only be evaluated at specific query points, where GP regression is exploited to achieve surrogate models of both the objective and the constraint functions, the latter being then used to weigh the exploration strategy by the probability of constraint satisfaction; conversely, Krishnamoorthy and Doyle (2022) include the GP models as barrier penalty terms. Safe exploration algorithms also make extensive use of GPs, either for classification (Schreiter et al., 2015) or in combination with regularity tools (Turchetta et al., 2019).

In this paper, GPs are used to model the safety-critical constraints (10) since these processes allow to inherently address the probabilistic nature of (9). Training data for the GPs are collected each time the given task is solved, yielding the tuple $\langle \theta, \tilde{z} \rangle$, where θ is the current MPC parametrization and \tilde{z} is a noisy observation of (10). Observations are generally regarded as noisy due to, e.g., measurement noise or disturbances in the dynamics. Therefore, the basic idea is that, given the dataset $\mathcal{D}_n = \{(\theta^i, \tilde{z}^i)\}_{i=1}^n$, one independent GP regressor per constraint is used to model the corresponding safety constraint function z_j , j = $1, \ldots, n_c$. In particular, we assume the function to have a GP prior with mean μ_0 : $\mathbb{R}^{n_{\theta}} \to \mathbb{R}$ and covariance kernel ν : $\mathbb{R}^{n_{\theta}} \times \mathbb{R}^{n_{\theta}} \to \mathbb{R}$. Under the GP prior and assumed i.i.d. Gaussian noise, observations \tilde{z}_i^i are jointly Gaussian distributed, so that, at any query point θ^* , the corresponding constraint violation $z_j(\theta^*)$ must be jointly Gaussian too, i.e., $z_j(\theta^*) \sim \mathcal{N}(\mu_j^n(\theta^*), \sigma_j^{n2}(\theta^*))$, where $\mu_j^n(\theta^*)$ is the estimate mean and $\sigma_j^n(\theta^*)$ its uncertainty (Rasmussen and Williams, 2005, Algorithm 2.1). In this context, (9) can be approximated from data as

$$\mathcal{S}_{\mathcal{D}_n} = \{ \theta \mid \mathbb{P} \left[z_j(\theta) \le 0 | \mathcal{D}_n \right] \ge \beta, \ j = 1, \dots, n_c \}, \quad (11)$$

where z_j is represented by a GP trained on data in \mathcal{D}_n .

In the next section, we illustrate how to combine MPCbased RL with GP constraint modeling into a safety-aware algorithm.

3. DATA-DRIVEN SAFE MPC-BASED RL

Algorithm 1 illustrates the procedure to promote safety, up to some probability β , upon an MPC-based RL agent. Starting from an initial parametrization of the MPC controller and a (possibly empty) dataset of initial observations, experiment *i* is carried out by completing the given task with an MPC scheme parametrized in θ^i . At the end of such experiment, the full state-action trajectory can be collected and the maximum constraint violation \tilde{z}^i computed. Thanks to the collected data, constraints $z_j(\theta), j = 1, \ldots, n_c$ are then approximated via GP regressors. The RL update is then enhanced with the trained GP-based constraints, which help the algorithm in selecting the next parameters θ_+ that satisfy these constraints and thus are likely to yield a safe policy.

| Algorithm 1: Data-driven Safe MPC-based RL | |
|--|---|
| | Input: Initial MPC parameters θ_0 ; initial |
| | observations \mathcal{D}_0 ; $\varrho \in (0, 1)$ |
| 1 | for $i = 1, \ldots, n_{\max} \operatorname{do}$ |
| 2 | Perform MPC closed-loop task with θ^i by |
| | solving (4) at each time step $t = 1, \ldots, T$ |
| 3 | Observe \tilde{z}_j^i , $j = 1, \ldots, n_c$ as per (10) |
| 4 | Augment dataset $\mathcal{D}_i \leftarrow \mathcal{D}_{i-1} \cup \{(\theta^i, \tilde{z}^i)\}$ |
| 5 | Train GPs for $z_j(\theta), \ j = 1, \ldots, n_c$ on \mathcal{D}_i |
| 6 | do |
| 7 | Try performing safe RL update |
| | $\theta_{\perp}^{\star} \leftarrow \arg \min (8a) $ (12a) |
| | θ_+ |
| | s.t. $\theta_+ \in \mathcal{S}_{\mathcal{D}_i}$ (12b) |
| 8 | Reduce safety probability $\beta \leftarrow \rho\beta$ |
| 9 | while (12) is not feasible |
| 10 | Update parametrization $\theta_{i+1} \leftarrow \theta_+^{\star}$ |
| 11 end | |

Remark 3. Gaussianity of the GPs can be conveniently leveraged to express constraint (12b) deterministically as

$$\mu_j^i(\theta_+) + \Phi^{-1}(\beta)\sigma_j^i(\theta_+) \le 0, \ j = 1, \dots, n_{\rm c}, \tag{13}$$

where Φ^{-1} is the normal distribution quantile function. However, since the GP terms $\mu_i^i(\theta_+)$ and $\sigma_i^i(\theta_+)$ are in practice always nonlinear, the overall RL update optimization problem (12) becomes nonlinear too. Moreover, there exist no guarantees on the feasibility of (12), i.e., the update might get stuck as it is unable to find a safe θ^{\star}_{\perp} . This is especially relevant at the beginning of the simulation in case \mathcal{D}_0 is empty, i.e., data is scarce and the GP approximation is poor. Currently, one can mitigate this issue by backtracking (i.e., iteratively reducing) the requested safety probability β to smaller and smaller values till feasibility is recovered, as shown in lines 6-9. Of course, this trades off safety with feasibility, which might be undesirable in cases where the former must be guaranteed with high enough certainty. However, one could easily envision methods for alleviating this issue by, e.g., injecting prior knowledge via the GP prior mean or the initial dataset \mathcal{D}_0 , or, alternatively, assuming the initial condition θ_0 to be safe with unitary probability (thus acting as a sort of backup configuration).

4. NUMERICAL EXPERIMENT

The proposed algorithm was implemented and simulated in a numerical experiment of a quadrotor drone application adapted from Wabersich and Zeilinger (2022). In the following sections, we detail the system, the controller and the safe learning algorithm, and present the final results.

4.1 System Description

The aim of the task is for the drone to reach a final destination in the fastest and most accurate way, while avoiding violations of safe-critical boxing constraints on the states and control actions, with wind disturbances at different altitudes influencing the drone's dynamics. The goal of the RL agent is, on one side, to tune the parameters of the MPC controller to achieve stable flight and better performance in terms of the objective function and, on the other side, to learn to avoid constraint violations as much as possible. The discrete-time dynamics of the drone are

$$x_{t+1} = Ax_t + Bu_t + C\Psi(x_t)w_t + G,$$
 (14)

where the state $x \in \mathbb{R}^{10}$ contains the pose (position, velocity, attitude, and its rate), the control action $u \in$ \mathbb{R}^3 dictates the desired pitch/roll attitude and vertical acceleration, and system matrices A and B describe the dynamics of the quadrotor around the hovering state, and vector G represents the constant gravity pull. The term $C\Psi(x_t)$ models nonlinear wind disturbances at time step $t: \Psi(x) \in \mathbb{R}^3$ is a state-dependent vector simulating three different wind currents at different altitudes, each implemented as a squared exponential basis function with unknown hyperparameters, whereas matrix C modulates the influence of each gust on each state, and $w_t \sim \mathcal{U}(0,1)$ induces randomness. Box constraints h are imposed on the state and control action. The RL stage cost takes into account, at each step, the error from the target position $s_{\rm f}$, the control action usage, as well as constraint violations

 $L(s,a) = \|s - s_{\rm f}\|_2^2 + c_1 \|a\|_2^2 + c_2^{\top} \max\left(0, h\left(s,a\right)\right), \quad (15)$ where scalar c_1 and vector c_2 are weights.

4.2 MPC Function Approximation

The controller chosen for this task consists of an MPC scheme based on (4): the stage/final costs are tracking costs (in both states and actions) but their cost matrices are fixed/non-learnable, and no initial cost is used, i.e., $\lambda_{\theta}(x_0) = 0$; the parametric prediction model is $f_{\theta}(x, u) =$ $A_{\theta}x + B_{\theta}u + G_{\theta}$; the constraints are parametrized by means of a backoff parameter $h_{\theta}(x, u) = (1 + \delta)h(x, u)$ and are made soft in the MPC to avoid infeasibilities. Eight different system parameters in A_{θ} , B_{θ} and G_{θ} affect the dynamics of the quadrotor drone but, out of these, the learning is restricted exclusively to the gravitational pull constant q and the vertical thruster coefficient K_z . Additionally, the constraint backoff δ is made learnable, i.e., $\theta = [g K_z \delta]^{\top}$, while all the other parameters in the drone model and optimization are kept fixed during learning. To better test the viability of the proposed approach, all eight dynamics parameters in the MPC scheme are initialized wrongly (with an error $\approx \pm 20\%$), so as to lead to a controller with initial poor performance and constraint violations. Therefore, the agent has to achieve

better performance and safety by only tuning θ , a small subset of all the dynamics parameters.

4.3 Safe RL Algorithm

For learning safety constraints (10), one independent GP is employed to model each constraint. A zero prior mean function is assumed, whereas the kernel function is selected as the sum of the squared exponential kernel and a white noise kernel. As per Algorithm 1, the GP regressors are trained on past observed data before each update of the MPC parameters. For the RL updates, a secondorder LSTD Q-learning algorithm (Lagoudakis et al., 2002) is employed to find the parameter vector θ that minimizes closed-loop performance under the policy $\pi_{\theta}(s)$. Q-learning is a classical RL algorithm that tries to find the parametrization which best fits the action-value function to the observed data, thus indirectly finding the optimal policy, and has shown promising results also with MPC (Esfahani et al., 2021). Additionally, the second-order Newton's method ensures faster convergence by exploiting an approximation of the Hessian. The gradient $\nabla_{\theta} Q_{\theta}$ is computed via nonlinear programming sensitivity analysis (Büskens and Maurer, 2001). A more in-depth explanation of the RL algorithm can be found in Airaldi et al. (2022). Lastly, in the constrained update (12), the target safety level β is set to 0.9, but backtracking on this level is enabled in case of infeasibility.

4.4 Results

The experiment was implemented in Python leveraging the symbolic framework CasADi (Andersson et al., 2019) and its interface to the IPOPT solver (Wächter and Biegler, 2006). Source code and simulation results can be found in the following repository: https://github.com/ FilippoAiraldi/learning-safety-in-mpc-based-rl.

To investigate the effectiveness of the proposed algorithm, simulations were carried out for the LSTD Q-learning algorithm both with and without our GP-based safety enhancement. The safe variant of the algorithm was also tested with zero initial knowledge (i.e., \mathcal{D}_0 empty), and with limited prior knowledge in the form of 5 datapoints picked randomly from previous simulations. Where relevant, we also compare the algorithm with a baseline, in which the MPC controller is initialized with perfect knowledge on the system model. Each algorithm is averaged over 100 simulations for 50 learning episodes. In the figures, average results are plotted with one standard deviation.

As shown in Figure 1, despite starting from an unfavorable initial configuration θ that yields large constraint violations, both safe and unsafe algorithms are able to recover and converge to a performance that is just $\approx 20\%$ worse than the baseline. Note that they achieve so by adjusting only a limited subset of the system dynamics parameters. However, the convergence speed of the two substantially differs, the safe variants being much faster. This is explained by the fact that already in the first few iterations the safety mechanism is able to approximate the safe set $S_{\mathcal{D}}$, even if only locally, and to steer the learning towards parameters θ that are likely to yield safer policies. However, we note that at the very beginning of the



Fig. 1. Comparison in performance between the original LSTD Q-learning algorithm and its GP-based safe variant (with and without prior knowledge). In dashed black, the baseline cumulative cost when exact knowledge of the system is given to the MPC controller.



Fig. 2. Comparison in (top) violations of the altitude constraint, where positive values imply violation, and (bottom) the cumulative number of unsafe episodes.

learning process the zero-knowledge safe algorithm suffers from much higher variance in the results compared to its counterpart initialized with some prior information. This confirms that at the beginning the ability to predict safety is extremely poor and can even damage performance.

Figure 2 shows the effectiveness of our method in encouraging safety. At the end of the learning process, we witnessed a 57% and 64% reduction in the total number of unsafe episodes (episodes in which constraint violation occurs) for the two safe algorithms compared to the unsafe LSTD Q-learning variant. These results are achieved in a purely data-driven fashion and with no or little previous information on safety. Furthermore, constraint violations are recovered much faster, especially in the quadrotor altitude, which is the state variable most prone to violations. This explains why the performance of the safe algorithms shows faster convergence to better suboptimal policies.

Lastly, Figure 3 shows the impact of the backtracking on β , as explained in Remark 3. In particular, it shows on average during learning at which probability Algorithm 1 had to decrease β in order to obtain the feasibility of the constrained update (12). In the case of zero initial information, at the beginning β has to be decreased to achieve feasible updates, while at convergence, when enough data for the GPs has been collected, the required



Fig. 3. Backtracked safety probability β during learning.

probability of 0.9 is achieved. Conversely, the simulations with initial knowledge seems to confirm that providing a small \mathcal{D}_0 to pre-train the GP-based safety constraints (12b) is an effective way to avoid backtracking on β .

5. CONCLUSIONS

In this paper, we propose a method that enables safety in the context of MPC-based RL via GP regression. It learns and encourages safety exclusively from observed past trajectories, without the need of, e.g., a priori information on the system or estimation of disturbances affecting its dynamics. The algorithm is straightforward and can be implemented without extensive customization on top of several RL agents, such as Q-learning and policy gradient methods. The numerical experiment for the quadrotor control application demonstrated that the safety mechanism successfully reduces the number of unsafe episodes during learning and also promotes faster convergence by avoiding unnecessary constraint violations.

ACKNOWLEDGEMENTS

This research is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101018826 - CLariNet).

REFERENCES

- Airaldi, F., De Schutter, B., and Dabiri, A. (2022). Learning safety in model-based reinforcement learning using mpc and gaussian processes. arXiv preprint arXiv:2211.01860.
- Andersson, J.A.E., Gillis, J., Horn, G., Rawlings, J.B., and Diehl, M. (2019). CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1), 1–36.
- Brunke, L., Greeff, M., Hall, A.W., Yuan, Z., Zhou, S., Panerati, J., and Schoellig, A.P. (2022). Safe learning in robotics: from learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5(1), 411–444.
- Büskens, C. and Maurer, H. (2001). Sensitivity analysis and real-time optimization of parametric nonlinear programming problems. In M. Grötschel, S.O. Krumke, and J. Rambau (eds.), Online Optimization of Large Scale Systems, 3–16. Springer, Berlin, Heidelberg.
- Esfahani, H.N., Kordabad, A.B., and Gros, S. (2021). Approximate robust NMPC using reinforcement learning. In 2021 European Control Conference (ECC), 132–137.
- Gros, S. and Zanon, M. (2022). Learning for MPC with stability & safety guarantees. *Automatica*, 146, 110598.

- Gros, S. and Zanon, M. (2020). Data-driven economic NMPC using reinforcement learning. *IEEE Transac*tions on Automatic Control, 65(2), 636–648.
- Hewing, L., Kabzan, J., and Zeilinger, M.N. (2020a). Cautious model predictive control using gaussian process regression. *IEEE Transactions on Control Systems Technology*, 28(6), 2736–2743.
- Hewing, L., Wabersich, K.P., Menner, M., and Zeilinger, M.N. (2020b). Learning-based model predictive control: toward safe learning in control. Annual Review of Control, Robotics, and Autonomous Systems, 3(1).
- Koller, T., Berkenkamp, F., Turchetta, M., and Krause, A. (2018). Learning-based model predictive control for safe exploration. In 2018 IEEE Conference on Decision and Control (CDC), 6059–6066.
- Krishnamoorthy, D. and Doyle, F.J. (2022). Safe bayesian optimization using interior-point methods—applied to personalized insulin dose guidance. *IEEE Control Sys*tems Letters, 6, 2834–2839.
- Lagoudakis, M.G., Parr, R., and Littman, M.L. (2002). Least-squares methods in reinforcement learning for control. In *Methods and Applications of Artificial Intelligence*, 249–260. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Mesbah, A., Wabersich, K.P., Schoellig, A.P., Zeilinger, M.N., Lucia, S., Badgwell, T.A., and Paulson, J.A. (2022). Fusion of machine learning and MPC under uncertainty: What advances are on the horizon? In 2022 American Control Conference (ACC), 342–357.
- Piga, D., Forgione, M., Formentin, S., and Bemporad, A. (2019). Performance-oriented model learning for datadriven MPC design. *IEEE Control Systems Letters*, 3(3), 577–582.
- Rasmussen, C.E. and Williams, C.K.I. (2005). *Gaussian Processes for Machine Learning*. MIT Press.
- Schreiter, J., Nguyen-Tuong, D., Eberts, M., Bischoff, B., Markert, H., and Toussaint, M. (2015). Safe exploration for active learning with gaussian processes. In *Machine Learning and Knowledge Discovery in Databases*, 133– 149. Springer International Publishing.
- Sorourifar, F., Makrygirgos, G., Mesbah, A., and Paulson, J.A. (2021). A data-driven automatic tuning method for MPC under uncertainty using constrained bayesian optimization. *IFAC-PapersOnLine*, 54(3), 243–250. 16th IFAC Symposium on Advanced Control of Chemical Processes ADCHEM 2021.
- Sutton, R.S. and Barto, A.G. (2018). Reinforcement Learning: An Introduction. MIT Press.
- Turchetta, M., Berkenkamp, F., and Krause, A. (2019). Safe exploration for interactive machine learning. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc.
- Wabersich, K.P. and Zeilinger, M.N. (2022). Cautious bayesian MPC: regret analysis and bounds on the number of unsafe learning episodes. *IEEE Transactions on Automatic Control*, 1–8.
- Wächter, A. and Biegler, L.T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1), 25–57.
- Zanon, M. and Gros, S. (2021). Safe reinforcement learning using robust MPC. *IEEE Transactions on Automatic Control*, 66(8), 3638–3652.