Received: Added at production | Revised: Added at production | Accepted: Added at production
DOI: xxx/xxxx

ARTICLE TYPE

Nonlinear learning-based model predictive control supporting state and input dependent model uncertainty estimates

Kim P. Wabersich* | Melanie N. Zeilinger

Institute for Dynamic Systems and Control, ETH Zurich, Switzerland

Correspondence
*Kim P. Wabersich, Institute for Dynamic Systems and Control, ETH Zurich, Switzerland.
Email: wkim@ethz.ch

Funding Information
This research was supported by the Swiss National Science Foundation under grant no. PP00P2 157601/1.

Summary

While model predictive control (MPC) methods have proven their efficacy when applied to systems with safety specifications and physical limitations, their performance heavily relies on an accurate prediction model. As a consequence, a significant effort in the design of MPC controllers is dedicated to the modeling part and often requires advanced physical expertise. In order to facilitate the controller design, we present an MPC scheme supporting nonlinear learning-based prediction models, i.e. data-driven models with probabilistic parameter uncertainties. A tube-based MPC formulation in combination with an additional implicit state and input constraint forces the closed-loop system to be operated in domains of sufficient model confidence, thereby ensuring asymptotic stability and constraint satisfaction at a pre-specified level of probability. Furthermore, by relying on tube-based MPC concepts, the proposed learning-based MPC formulation offers a general framework for addressing different problem classes, such as economic MPC, while providing a general interface to probabilistic prediction models based, e.g., on Bayesian regression or Gaussian processes. A design procedure is proposed for approximately linear systems and the efficiency of the method is illustrated using numerical examples.

KEYWORDS:
model predictive control, safe learning-based control, control of constrained systems, robust control of nonlinear systems

1 | INTRODUCTION

Over the past decades, model predictive control methods have been established as the standard for controlling safety-critical dynamical systems with limited input authority across various industries.[1,2] One of its main benefits is the possibility to ensure closed-loop constraint satisfaction in a principled way while approximately maximizing a given objective that specifies the control task at hand. This is achieved through solving a finite-time optimal control problem at each sampling time in a receding horizon fashion, using a model of the true system dynamics to simultaneously predict and optimize the performance of the system in real-time. As a consequence, the quality of model predictive controllers heavily relies on the accuracy of the prediction model. At the same time, the effort needed for manually obtaining a prediction model by first-principles and previous experiments increases drastically with the desired level of prediction accuracy. In industrial applications, this often results in high cost for the development and maintenance of model predictive controllers or in the use of suboptimal prediction models that are often subject to a significant amount of model uncertainty, which can cause performance degeneration and safety violations.

To explicitly cope with these uncertainties, model predictive control techniques were extended to support uncertain and stochastic prediction models, resulting in a robust or probabilistic closed-loop analysis. While these approaches can cope with model inaccuracies, they typically still require a control engineer to manually derive a sound uncertainty description of the prediction model and to select and execute a suitable model predictive control synthesis. As a consequence, the resulting closed-loop performance depends largely on the offline design phase, in which a-priori assumptions on the model uncertainty can lead to conservatism.

Increasing availability of system data, computation and sensing lead to a growing interest in machine learning techniques, which offer sophisticated tools for highly automated inference of prediction models from data, resulting in novel learning-based model predictive controllers.[3,4,5,6,7] On the one hand, even though some of these controllers have shown to perform well in practice,[8,5,9] they do not provide rigorous safety or closed-loop performance certificates. On the other hand, there are attempts to define novel prediction models[10,11] compatible with existing robust model predictive control methods that offer desired closed-loop properties, although their performance have not yet been demonstrated in practice or compared with established machine learning methods. The goal of this paper is to combine the advantages of both research directions by providing rigorous probabilistic safety guarantees using prediction models that are obtained through successful and highly automated machine learning tools.

1.1 | Contributions

We extend a previously proposed framework[12] for safe learning-based model predictive control in two ways to significantly reduce conservatism. Firstly, we avoid the robust treatment of all possible model uncertainties by using nonlinear probabilistic prediction models. Secondly, a large portion of conservatism originates from a global treatment of the worst-case model uncertainty across the entire state and input space, which we reduce by leveraging probabilistic state and input dependent uncertainty information. The basic mechanism is to extend well-known tube-based model predictive control concepts with an additional constraint that forces predictions to a subset of the admissible state and input space with high model confidence at a desired level of probability while maintaining computational efficiency of the resulting MPC problem. The additional constraint is defined through a so-called set-valued model confidence map, describing state and input dependent model uncertainties, such that the proposed model predictive controller can be interfaced with any probabilistic prediction model and provides closed-loop chance constraint satisfaction. At the same time, by using tube-based model predictive control concepts at the core of the method, the scheme allows for direct adaptations to different problem settings including standard set-point or advanced economic control as two important examples. While the prediction model can be computed based on available data using machine learning tools, we present an automated procedure to compute the required components of the proposed model predictive control method based on only a few tuning parameters for the relevant case of approximately linear systems. The method and its design procedure are illustrated using numerical simulations for an approximately linear 10-dimensional quadrotor system and a nonlinear economic control problem.

1.2 | Related work

The use of learning-based prediction models in model predictive control was first proposed by Aswani et al. (2013)[12] using a tube-based approach, which requires a global bound on the prediction error and provides provable guarantees in terms of constraint satisfaction. Practical demonstrations of the method were presented in a number of subsequent results.[13,14,15,16] In addition, learning-based prediction models were developed that provide the required global model error bounds, for example based on Lipschitz interpolation[17] and Kinky inference,[10,18] leading to variations of the original learning-based model predictive control scheme.[19] One significant limitation of these approaches originates from the underlying assumption that the prediction models cover all possible model errors with probability one. Furthermore, even if these models provide state and input dependent uncertainty estimates, most of the learning-based model predictive control frameworks presented so far did not take advantage of them and conservatively have considered a global worst-case model error over the whole state and input space. We overcome these limitations by supporting probabilistic model estimates, allowing to consider model errors at a desired probability level and by enabling the use of state and input dependent uncertainty information in the model predictive controller.

Another class of methods that originate from robust model predictive control ideas is based on parametric set-membership model estimation techniques.[20,21,22,23,24,11,25,26] The idea is to iteratively rule out impossible model parameters successively over time and to explicitly take care of the remaining uncertainty in the online model predictive control problem. While these methods explicitly take online updates of the model into account in terms of recursive feasibility, they perform a conservative robust treatment of uncertainties with probability one and do not allow for using established stochastic methods from machine learning.

Different from robust approaches, stochastic model predictive control techniques that use probabilistic uncertainty information to reduce conservatism were investigated,[27] which among others can support learning-based prediction models. However, their theoretical analysis tends to be rather challenging, even for simple linear stochastic prediction models, while they are at the same time often difficult to implement, resulting in approximate implementations[3] that work well in practice[8,4] but are lacking theoretical statements about closed-loop properties. A relatively simple alternative is based on Monte Carlo simulations or scenario-based optimization techniques[28,29], allowing to iteratively select parameters of the MPC problem to achieve a desired probability of chance-constraint satisfaction. These results are, however, limited to linear systems or fixed initial conditions.

The goal of this paper is therefore to combine the benefits of purely robust and fully stochastic techniques into a novel model predictive control scheme, that treats probabilistic model uncertainties *robust in probability by design*, thereby being less conservative than robust approaches, while providing rigorous theoretical guarantees. Different from previous work[7,6] using a similar concept, our approach is scalable to larger dimensional systems by avoiding the need for expensive explicit off-line computations and it can be easily adopted to different problem classes such as economic control by relying on a static tube-based prediction mechanism.

Conceptually, the proposed techniques are related to stochastic and learning-based predictive control techniques[30,31,32] that are used for enhancing reinforcement learning algorithms with safety certificates and which are also centred around the idea of treating model uncertainties robustly in probability to provide rigorous closed-loop guarantees in terms of constraint satisfaction. Different from these approaches, we rely on a standard tube-based model predictive control formulation, enabling the use of established model predictive control techniques to obtain closed-loop guarantees together with a principled design procedure. Finally, note that while stochastic safety verification techniques[33,34] can be used to certify, e.g., nominal or robust MPC methods in the presence of stochastic uncertainties, our goal is to provide a principled controller design that guarantees the desired level of safety by construction.

Notation and definitions

The distance between a vector $x \in \mathbb{R}^n$ and a set $\mathcal{A} \subseteq \mathbb{R}^n$ is defined as $\|x\|_{\mathcal{A}} := \inf_{a\in\mathcal{A}} \|x - a\|$. A ball with radius $\epsilon > 0$ is denoted as $\mathcal{B}(\epsilon) = \{x \in \mathbb{R} | \, \|x\| \leq \epsilon\}$ and the radius of an outer bounding ball of a bounded set $\mathcal{A} \subset \mathbb{R}^n$ as $\operatorname{diam}(\mathcal{A}) := \max_{x\in\mathcal{A}} \|x\|$. The Minkowski sum of two sets $\mathcal{A}_1, \mathcal{A}_2 \subset \mathbb{R}^n$ is denoted by $\mathcal{A}_1 \oplus \mathcal{A}_2 := \{a_1 + a_2 | a_1 \in \mathcal{A}_1, a_2 \in \mathcal{A}_2\}$ and the Pontryagin set difference by $\mathcal{A}_1 \ominus \mathcal{A}_2 := \{a_1 \in \mathbb{R}^n | a_1 + a_2 \in \mathcal{A}_1, \, \forall a_2 \in \mathcal{A}_2\}$. An affine image of a set $\mathcal{A}_1 \subseteq \mathbb{R}^n$ under $x \mapsto Kx$ is defined as $K\mathcal{A}_1 := \{Kx | x \in \mathcal{A}_1\}$. The convex hull of a set of vectors $\{x_i\}_{i=0}^N$ with $x_i \in \mathbb{R}^n$ is denoted by $\operatorname{conv}(\{x_i\}_{i=0}^N)$. We use $\mathbb{P}(E)$ for the probability of an event E and indicate with $x \sim \mathcal{Q}_x$ a random variable x of distribution $\mathcal{Q}_x$. For time-dependent quantities $x(k)$ with time $k = 0, 1, ..$, we use the notation $x^+ := x(k+1)$ and $x := x(k)$ when convenient. Optimal solutions to an optimization problem will be denoted by an asterisk, e.g., $x^* = \operatorname{argmin}_{x\in\mathbb{R}} f(x)$.

2 | PROBLEM STATEMENT

We consider the problem of controlling deterministic non-linear systems that can be modelled as

$$x(k+1) = f(x(k), u(k); \theta), \; k \in \mathbb{N}^+, \tag{1}$$

with $x(k) \in \mathbb{R}^n$, $u(k) \in \mathbb{R}^m$, initial condition $x(0) = x_0$, and true system parameters $\theta \in \mathbb{R}^{n_\theta}$ that are assumed to be unknown but constant over time. We assume a distributional belief over the true system parameters θ of the form

$$\theta \sim \mathcal{Q}_\theta, \text{ with mean } \bar{\theta} := \mathbb{E}[\theta], \tag{2}$$

inferred for example using machine learning techniques based on imperfect system data given as

$$\mathbb{D} := \{(x(i+1) + \epsilon(i), x(i), u(i))\}_{i=1}^{N_{\mathbb{D}}} \text{ with } \epsilon(i) \sim \mathcal{Q}_\epsilon \text{ zero mean i.i.d. noise.} \tag{3}$$

Remark 1. System models of the form (1) can be obtained using, e.g., Bayesian regression[35]. While we focus on parametric models for the sake of clarity, the presented methods can directly be applied together with non-parametric estimation techniques such as Gaussian Processes[36] if the true system model is sampled from a reproducing kernel hilbert space or the state and input domain is bounded,[37,38] see also Section 4.1 for further details.

The goal is to develop a learning-based controller that is based on a model of the form (1), (2) and that minimizes the expected sum over stage cost functions up to a possibly infinite horizon $\bar{N}$ given by

$$J = \mathbb{E}_\theta \left[\sum_{k=0}^{\bar{N}} \ell(x(k), u(k)) \right], \tag{4}$$

where the expected value is taken with respect to the current distributional belief over the true system parameter θ. While minimizing (4), the controller additionally needs to ensure that the closed-loop system satisfies chance constraints on the system states and inputs, i.e.

$$\mathbb{P}(\text{For all } k = 1, 2, .., \bar{N} : \ x(k) \in \mathbb{X}, \ u(k) \in \mathbb{U}) \geq p_S \tag{5}$$

where p_S denotes the desired probability of safety, i.e. constraint satisfaction, $\mathbb{X} \subseteq \mathbb{R}^n$ typically denotes safety-critical state constraints, and $\mathbb{U} \subseteq \mathbb{R}^m$ represents physical input limitations.

3 | NONLINEAR LEARNING-BASED MODEL PREDICTIVE CONTROL

In this section, we show how conventional tube-based MPC formulations that provide robustness against additive disturbances can be used to leverage learning-based models in a safe manner. This is achieved by exploiting probabilistic state and input dependent uncertainty information to reduce the overall conservatism of the controller compared with a classical formulation. To this end we briefly introduce the tube-based MPC formulation in Section 3.1 and highlight its main limitation in case of probabilistic parameter uncertainties according to (1) and (2) in Section 3.2. We then introduce an efficient state and input dependent uncertainty description supporting probabilistic parameter uncertainties in Section 3.3 together with an extension of standard tube-based MPC methods that allows to exploit nonlinear learned system models.

3.1 | Background: Tube-based model predictive control

The system model and the nominal prediction model typically considered in tube-based MPC are given by

$$x(k+1) = f(x(k), u(k)) + w(k) \quad \text{and} \quad z(k+1) = f(z(k), v(k)), \tag{6}$$

where the disturbance $w(k)$ is assumed to lie in a compact disturbance set $\mathcal{W}$, which is constant for all times $k \in \mathbb{N}$ and $z(k)$ and $v(k)$ are the nominal system states and inputs. A common task is to steer an initial state $x(0)$ to a neighborhood of a desired equilibrium state z_s of the nominal system. We consider the origin as a set point, i.e. $z_s = 0$ with corresponding zero nominal input $v_s = 0$, for notational simplicity, but the results directly extend to non-zero set points. The resulting trajectory is required to robustly satisfy state and input constraints of the form $x(k) \in \mathbb{X}$ and $u(k) \in \mathbb{U}$, i.e. for all admissible disturbance sequences $\{w(k)\}$. In tube-based MPC, the MPC problem is formulated using the nominal system dynamics with respect to tightened state and input constraints and by compensating resulting closed-loop errors defined as $e(k) := x(k) - z(k)$ by a trajectory tracking feedback controller of the form $\kappa_\Omega : \mathbb{U} \times \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{U}$ in addition to the nominal system inputs $v(k)$, i.e. $u(k) = \kappa_\Omega(v(k), x(k), z(k))$. An appropriate tightening of the state and input constraints for the nominal system is determined in a principled way, by defining a tube of the form $z(k) \oplus \Omega$ with $\Omega \subset \mathbb{R}^n$, which robustly contains the true system state $x(k)$ around $z(k)$ under application of κ_Ω as follows.

Definition 1 (Robust positive invariant error set[39]). A set Ω is robust positive invariant (RPI) for the error system $e(k+1) = x(k+1) - z(k+1)$ if for all $x(k), z(k) \in \mathbb{X}$ with $e(k) \in \Omega$, $v(k) \in \mathbb{U}$, and $w(k) \in \mathcal{W}$ it holds that $e(k+1) \in \Omega$ with $u(k) = \kappa_\Omega(v(k), x(k), z(k)) \in \mathbb{U}$.

Using the set Ω, we can explicitly specify the required state and input constraint tightening as

$$\bar{\mathbb{Z}} := \{z \in \mathbb{X}, v \in \mathbb{U} \mid (x, \kappa_\Omega(v, x, z)) \in \mathbb{X} \times \mathbb{U} \ \forall x \in z \oplus \Omega\}, \tag{7}$$

ensuring that for all $k \geq 0$ the condition $(z(k), v(k)) \in \bar{\mathbb{Z}}$ implies by construction that $x(k) \in \mathbb{X}$ and $u(k) \in \mathbb{U}$ holds under application of $u(k) = \kappa_\Omega(v(k), x(k), z(k))$.

Taking into account the tightened constraints, the resulting tube-based MPC problem takes the form

$$\min_{\{v_{i|k}\}} \ \ell_f(z_{N|k}) + \sum_{i=0}^{N-1} \ell(z_{i|k}, v_{i|k}) \tag{8a}$$

$$\text{s.t. for all } i = 1, 2, .., N-1 :$$

$$x(k) \in z_{0|k} \oplus \Omega \tag{8b}$$

$$z_{i+1|k} = f(z_{i|k}, v_{i|k}), \tag{8c}$$

$$(z_{i|k}, v_{i|k}) \in \bar{\mathbb{Z}}, \tag{8d}$$

$$z_{N|k} \in \mathcal{X}_f, \tag{8e}$$

where we denote nominal states and inputs predicted at time k for i time steps into the future as $z_{i|k}$ and $v_{i|k}$. The additional terminal cost ℓ_f on the last predicted state together with the imposed terminal constraint (8e) w.r.t. the target set $\mathcal{X}_f$ is selected to ensure convergence of the nominal system to the origin[40]. The initial condition (8b) enforces that the tube at the initial time step $z_{0|k} \oplus \Omega$ contains the system state $x(k)$, thereby introducing feedback w.r.t the real system state in closed-loop. The tube-based MPC control input is defined as

$$\kappa_{\text{MPC}}(x(k)) = \kappa_\Omega \left(v^*_{0|k}, x(k), z^*_{0|k} \right) \tag{9}$$

where $v^*_{i|k}$ and $z^*_{i|k}$ denote the optimal solution of (8) at time step k. In the following, we summarize common assumptions under which we obtain recursive feasibility of (8) together with robust closed-loop constraint satisfaction as well as convergence guarantees of the system state to a neighborbood around the origin.[41,42,39]

Assumption 1. A control law (9) with a corresponding RPI set $\Omega \subset \mathbb{X}$ according to Definition 1 is available.

Assumption 2. The stage cost function ℓ is positive definite w.r.t. the origin and continuous, $\mathbb{X}$ and $\mathbb{U}$ are compact sets, and there exists a set $\mathcal{X}_f \subseteq \mathbb{R}^n$, a feedback law $\kappa_f : \mathcal{X}_f \to \mathbb{R}^m$, and a positive definite terminal cost function $\ell_f : \mathcal{X}_f \to \mathbb{R}^+$, such that for all $z \in \mathcal{X}_f$ it holds i) $(f(z, \kappa_f(z)), \kappa_f(z)) \in \bar{\mathbb{Z}}$, ii) $f(z, \kappa_f(z)) \in \mathcal{X}_f$, and iii) $\ell_f(f(z, \kappa_f(z))) - \ell_f(z) \le -\ell(z, \kappa_f(z))$.

Proposition 1. Consider system (6) and let Assumption 1 and 2 hold. If (8) is feasible for $x(0)$, then under application of the tube-based model predictive control law (9) it follows for all future time steps $k \in \mathbb{N}$ that (8) is feasible, $x(k) \in \mathbb{X}$, $u(k) \in \mathbb{U}$, as well as $\lim_{k\to\infty} \|x(k)\|_\Omega = 0$.

3.2 | Conservatism of global uncertainty bounds in case of learning-based prediction models

While tube-based MPC approaches provide an efficient and simple approach to compensate model uncertainties, learning-based model inference that results in parametric uncertainties such as (1), (2) is often difficult to handle in this framework due to possibly large or unbounded state and input dependent uncertainties $w(z, u) = f(z, v; \theta) - f(z, v; \bar{\theta})$, see Figure 1. The resulting large uniform uncertainty bounds $\mathcal{W}$ over the state and input space typically lead to a larger diameter of the tube Ω and therefore often render tube-based MPC infeasible. This is of particular relevance if some regions of the state and input space are poorly covered by available data, as it is often the case for larger-scale systems. An approach to keep the additive disturbance reasonably small is to manually introduce auxiliary state constraints that keep the system state in a domain of low model uncertainty. While the computation of such auxiliary state constraint sets can be automated for small scale systems[7], drawbacks are the need for re-designing the auxiliary state constraint every time the model is updated with new data and their limited scalability, similarly as for explicit MPC techniques[43].

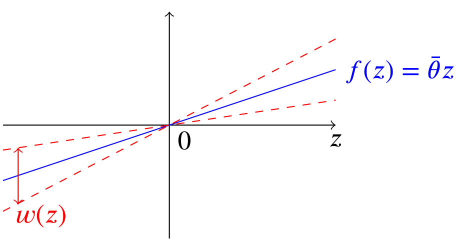

FIGURE 1 Parametric uncertainty according to (1) and (2) in case of a scalar linear system $f(x; \theta) = \theta x$ with expected linear model $f(z; \bar{\theta}) = \bar{\theta} z$ in blue together with state dependent uncertainties $w(z) = (\theta - \bar{\theta})z$ between the two dotted red lines.

This paper presents an alternative approach to limit planning to confident subsets of the state space, which mitigates these limitations.

3.3 | Learning-based MPC through efficient planning in confident subsets

In this section we present a computationally efficient approach to restrict predicted state and input sequences in (8) to subsets with high model confidence, i.e. regions for which the admissible model error $w(z_{i|k}, v_{i|k}) = f(z_{i|k}, v_{i|k}; \theta) - f(z_{i|k}, v_{i|k}; \bar{\theta})$ is contained in a pre-specified additive disturbance set $\mathcal{D} \subset \mathcal{W}$ at desired probability level p_S according to (5) without explicitly pre-computing the allowed domain. Since the size of the RPI set Ω is strongly related to the size of $\mathcal{D}$, we thus enable to trade off conservatism induced by the constraint tightening (7) via the choice of $\mathcal{D}$, against the locally admissible uncertainty magnitude and probability of safety. This trade-off can be efficiently incorporated into a design procedure as shown in Section 4.2.

Because $w(z_{i|k}, v_{i|k})$ is uncertain and we cannot simply impose $w(z_{i|k}, v_{i|k}) \in \mathcal{D}$ as an additional constraint to (8a), we utilize analytic state and input dependent model confidence estimates $\mathcal{D}_{p_S}(x, u)$ such that $w(x, u) \in \mathcal{D}_{p_S}(x, u)$ holds at probability level p_S. Such a model confidence estimate is available for relevant system descriptions of the form (1), (2) and can be seen as a possibly general interface between different machine learning techniques and tube-based MPC schemes by incorporating an additional constraint of the form '$\mathcal{D}_{p_S}(z_{i|k}, v_{i|k}) \subseteq \mathcal{D}$' into the tube-based MPC problem (8). More formally:

Definition 2 (Set-valued model confidence map)**.** A set-valued map $\mathcal{D}_{p_S}(.,.)$ mapping states and inputs from $\mathbb{X} \times \mathbb{U}$ to subsets of $\mathcal{D}$ with $\mathcal{D} \subset \mathbb{R}^n$ is said to be a *set-valued model confidence map* associated with (1), (2), for a given $\bar{\theta} \in \mathbb{R}^q$ at probability level $p_S > 0$ if, at probability greater or equal to p_S, it holds for all $k \in \mathbb{N}$, $x(k) \in \mathbb{X}$, and $u(k) \in \mathbb{U}$ that

$$x(k+1) - f(x(k), u(k), \bar{\theta}) \in \mathcal{D}_{p_S}(x(k), u(k)). \tag{10}$$

A discussion together with an example for obtaining such confidence maps $\mathcal{D}_p(x, u)$ is provided in Section 4.1. While Definition 2 principally allows to guide the closed-loop system within confident regions of the state and input space, it does not suffice to require '$\mathcal{D}_{p_S}(z_{i|k}, v_{i|k}) \subseteq \mathcal{D}$' as an additional constraint in the nominal MPC problem (8), since $\mathcal{D}_{p_S}(z^*_{1|k}, v^*_{1|k}) \subseteq \mathcal{D} \not\Rightarrow \mathcal{D}_{p_S}(x(k+1), u(k+1)) \subseteq \mathcal{D}$ under application of (9) because $(z^*_{0|k}, v^*_{0|k}) \neq (x(k), u(k))$ in general. Similarly as for the case of state and input constraint satisfaction, we therefore also tighten the set-valued confidence constraint as follows, which will play a key role for recursive feasibility.

Definition 3 (Tightened confidence set)**.** Consider system (1), (2) and a control law (9) with a corresponding RPI set $\Omega \subset \mathbb{X}$ according to Definition 1. A set $\bar{\mathcal{D}} \subseteq \mathcal{D}$ is a *tightened confidence set* associated with a set-valued model confidence map $\mathcal{D}_{p_S}(.,.)$ if for all $(z, v) \in \bar{\mathbb{Z}}$, $x \in z \oplus \Omega$ it holds

$$\mathcal{D}_{p_S}(z, v) \subseteq \bar{\mathcal{D}} \Rightarrow \mathcal{D}_{p_S}(x, \kappa_\Omega(v, x, z)) \subseteq \mathcal{D}. \tag{11}$$

In Section 4 we show that for a sufficiently small local model uncertainty together with stabilizability of the linearization of $x(k+1) = f(x(k), u(k); \bar{\theta})$ around $x = 0$ and $u = 0$ one can always find a model error set $\mathcal{D}$ such that a non-empty tightened model error set $\bar{\mathcal{D}}$ (11) exists locally. Using the notion of set-valued model confidence maps from Definition 2 together with a tightened confidence set according to Definition 3 we can state the resulting additional constraint to conventional tube-based MPC schemes, supporting probabilistic state and input dependent uncertainty estimates as

$$\min_{\{v_{i|k}\}} \ \ell_f(z_{N|k}) + \sum_{i=0}^{N-1} \ell(z_{i|k}, v_{i|k}) \tag{12a}$$

$$\text{s.t. for all } i \in \mathcal{I}_{[0,N-1]}:$$

$$(8b) - (8e) \tag{12b}$$

$$\mathcal{D}_{p_S}(z_{i|k}, v_{i|k}) \subseteq \bar{\mathcal{D}}, \tag{12c}$$

where (12c) can be efficiently implemented for several learning-based prediction models, see Section 4.1. Application of the resulting MPC control law (9) based on the modified problem (12) provides the following closed-loop system properties.

Theorem 1. Consider system (1), (2), together with a set-valued model confidence map according to Definition 2 and let Assumption 1 hold w.r.t. a pre-specified additive disturbance set $\mathcal{W} = \mathcal{D}$. If $\bar{\mathcal{D}}$ is a tightened confidence set according to Definition 3 and if Assumption 2 holds with the additional requirement that for all $z \in \mathbb{X}_f$ it follows that $\mathcal{D}_{p_S}(z, \kappa_f(z)) \subseteq \bar{\mathcal{D}}$, then under application of the tube-based model predictive control law (9), it holds jointly for all future time steps $k \in \mathbb{N}$ at probability level p_S that $x(k) \in \mathbb{X}$, $u(k) \in \mathbb{U}$ and $\lim_{k\to\infty} \|x(k)\|_\Omega = 0$.

Proof. To streamline notation, define $d(k) := f(x(k), u(k); \theta) - f(x(k), u(k); \bar{\theta})$. We show recursive feasibility based on the relation

$$\begin{aligned} &\mathbb{P}(\forall k \in \mathbb{N}^+ : \text{(12) is feasible}) && (13)\\ \geq &\mathbb{P}(\forall k \in \mathbb{N}^+ : \text{(12) is feasible}, d(k) \in \mathcal{D}_{p_S}(x(k), u(k))) \\ \geq &\mathbb{P}(\forall k \in \mathbb{N}^+ : \text{(12) is feasible} | d(k) \in \mathcal{D}_{p_S}(x(k), u(k)))\mathbb{P}(\forall k \in \mathbb{N}^+ : d(k) \in \mathcal{D}_{p_S}(x(k), u(k))). && (14) \end{aligned}$$

Since $\mathbb{P}(\forall k \in \mathbb{N}^+ : d(k) \in \mathcal{D}_{p_S}(x(k), u(k))) \geq p_S$ by assumption due to Definition 2, relation (14) enables us to prove (13) by establishing

$$\mathbb{P}(\forall k \in \mathbb{N}^+ : \text{(12) is feasible} | d(k) \in \mathcal{D}_{p_S}(x(k), u(k))) = 1. \tag{15}$$

Similar to the standard tube-based model predictive control proof, we show (15) recursively by induction, i.e. if (12) is feasible at time k, it will be feasible at time $k+1$. From feasibility at time k we have from (8b) that $x(k) \in z^*_{0|k} \oplus \Omega$ and together with (12c) it follows from Definition 3 that $\mathcal{D}_{p_S}(x, \kappa_\Omega(v^*_{0|k}, x(k), z^*_{0|k})) \subseteq \mathcal{D}$. Under the condition in (15) we therefore conclude that $d(k) \in \mathcal{D}$. Since Ω is by assumption RPI for the error dynamics $e(k+1) = f(x(k), \kappa_\Omega(v^*_{0|k}, x(k), z^*_{0|k})) - z^*_{1|k}$, we can further establish that $e(k+1) \in \Omega$ and $x(k+1) \in z^*_{1|k} \oplus \Omega$. From here it follows from Assumption 2 in combination with the additional assumption $\forall z \in \mathbb{X}_f \Rightarrow \mathcal{D}_{p_S}(z, \kappa_f(z)) \subseteq \bar{\mathcal{D}}$ that we can leverage the standard candidate input sequence given by $V(k+1) := (v^*_{1|k}, .., v^*_{N|k}, \kappa_f(z^*_{N|k}))$ to construct a feasible nominal candidate sequence for (12) at time $k+1$. Hence, we have shown (15), implying recursive feasibility at probability level p_S, which immediately implies satisfaction of chance constraints (5) and asymptotic convergence $\lim_{k\to\infty} \|x(k)\|_\Omega = 0$ as in standard tube-based MPC[41]. □

Model update without re-design of the MPC problem

One of the main advantages of learning-based control is the possibility to leverage available closed-loop system data in order to obtain a more accurate system model, leading to improved closed-loop system performance. In particular, it would be desirable to perform updates of the system model based on collected data during closed-loop system operation in order to subsequently reduce conservatism of the MPC controller. The proposed implicit formulation of the requirement to stay within confident subsets according to (12c) allows us to derive online verifiable sufficient conditions on an updated set-valued model confidence map according to Definition 2 that preserve the closed-loop properties of Theorem 1. Thereby the tube-MPC formulation does not have to be re-designed, if collected data leads to a more confident model estimate in the sense that for all $x \in \mathbb{X}$, $u \in \mathbb{U}$ we have that $\mathcal{D}^+_{p_S}(x, u) \subseteq \mathcal{D}_{p_S}(x, u)$ and that the tightened confidence set $\bar{\mathcal{D}}$ is still valid for an updated set-valued model confidence map $\mathcal{D}^+_{p_S}$.

Corollary 1. Let $\mathcal{D}^+_{p_S}$ be an updated set-valued model confidence map according to Definition 2 and let the assumptions in Theorem 1 be satisfied. If it holds for all $(z, v) \in \bar{\mathbb{Z}}$, $x \in z \oplus \Omega$ that i) $\mathcal{D}^+_{p_S}(z, v) \subseteq \mathcal{D}_{p_S}(z, v)$ and ii) $\mathcal{D}^+_{p_S}(z, v) \subseteq \bar{\mathcal{D}} \Rightarrow \mathcal{D}^+_{p_S}(x, \kappa_\Omega(x, z, v)) \subseteq \mathcal{D}$, then choosing $\mathcal{D}_{p_S} := \mathcal{D}^+_{p_S}$ satisfies the assumptions of Theorem 1.

Note that the conditions of assumptions of Corollary 1 can be verified by solving a nonlinear optimization problem, i.e. by searching for particular z, v, x that violate i) or ii), see Appendix B.

Remark 2. For substantial model updates that also require a different mean estimated parameter $\bar{\theta}$, one needs to recompute the RPI set, terminal set, and terminal cost computations. The updated model predictive control problem would then have to be solved in parallel until initial feasibility, allowing to switch to the updated controller.

A framework for uncertain parametric prediction models and diverse control tasks

By relying on a deterministic nominal prediction model (e.g. expected estimate or maximum-a-posteriori estimate $\bar{\theta}$ of θ) in combination with the set-valued model confidence map as uncertainty measure for planning in confident subsets of the state and input space, the proposed approach can directly make use of different classes of probabilistic prediction models as illustrated in Figure 2 and as explained for the special case of Bayesian regression in Section 4.1 in more detail. Furthermore, the close relation to standard tube-based MPC formulations offers easy adaptation of the concept to different problem settings such as trajectory tracking and economic control[44,39] via the choice of the stage cost, the terminal cost, and the terminal constraint. It is important to note that these compontents only require the nominal model and the set-valued model confidence set as illustrated in Figure 2 and in the numerical example in Section 5.2, where economic steady-state control is considered.

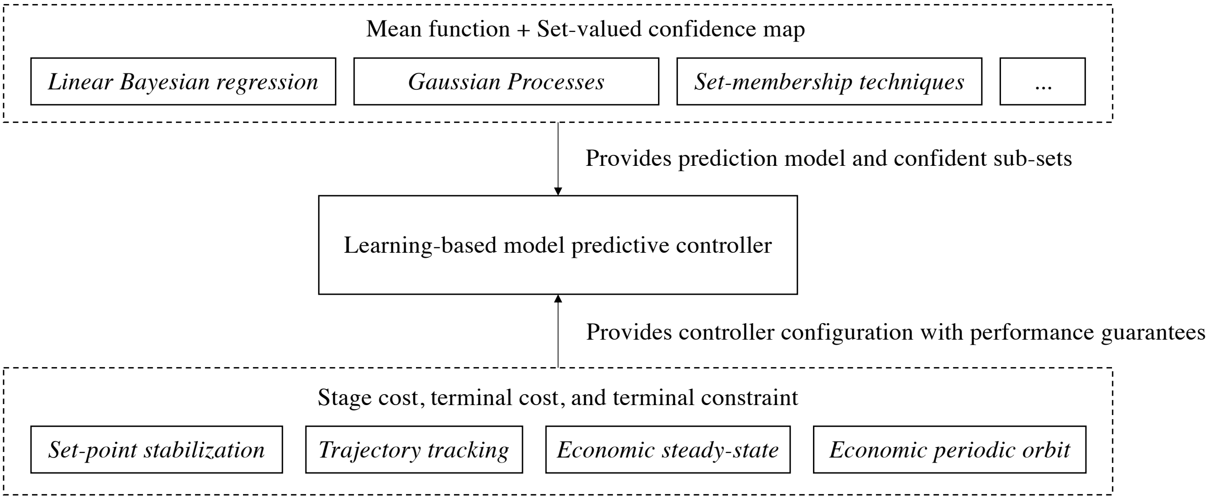

FIGURE 2 The proposed learning-based model predictive controller can be seen as a modular framework in terms of different prediction models (top) and available model predictive control configurations (bottom), each specialized to a different application.

Remark 3. There are a number of recently proposed robust, nonlinear model predictive control schemes[45,46,47] that could be adopted to support state and input dependent uncertainty estimates robustly in probability by using the set-valued confidence set according to Definition 2 as well. However, since these methods differ from the standard tube-based model predictive control framework, more effort would generally be needed to adapt these schemes beyond set-point stabilization towards, e.g., economic steady-state or periodic operation, which build on existing tube-based approaches[39,48].

4 | LEARNING-BASED MODEL PREDICTIVE CONTROLLER SYNTHESIS

In this section, we provide principled synthesis techniques for the proposed MPC scheme through a learning-based computation of the prediction model (1), (2) together with a set-valued model confidence map according to Definition 2 for the important case of Bayesian regression in Section 4.1. Furthermore, we derive a corresponding tube and tube controller computation as required by Assumption 1 with a confidence set tightening in compliance with Definition 3 in Section 4.2, and outline terminal cost and terminal set computation methods in Section 4.3.

4.1 | Learning-based predictions and set-valued model confidence maps using Bayesian inference

A simple, yet powerful Bayesian regression approach is based on parametrized models of the form

$$f(x,u;\theta) = \theta^\top \phi(x,u), \quad \theta \in \mathbb{R}^{n_\theta \times n} \tag{16}$$

consisting of a nonlinear transformation into feature space via $\phi : \mathbb{R}^{n+m} \to \mathbb{R}^{n_\theta}$, which are linearly combined via the parameters θ. Bayesian regression offers great model flexibility by choosing ϕ for example[49,50] to be a collection of higher-order polynomials, spline functions, radial basis functions centered around data points, or even feed-forward neural nets.

Due to being linear in the unknown parameters, one can efficiently obtain a parameter estimate of the form (2) together with a corresponding set-valued model confidence map according to Definition 2 based on a Gaussian parameter prior distribution $\text{col}_i(\theta) \sim \mathbb{N}(\mu_i^\theta, \Sigma_i^\theta)$ together with data $\mathbb{D}$ (3) affected by Gaussian i.i.d. noise $\mathcal{Q}_\epsilon = \mathbb{N}(0, I_n\sigma_s^2)$. Specifically, the posterior distribution $\mathcal{Q}_{\theta|\mathbb{D}}$ over parameters θ conditioned on data $\mathbb{D}$ can be obtained as[36,49]

$$\mathbb{P}(\text{col}_i(\theta)|\mathbb{D}) = \mathbb{N}\left(\mu_i^{\theta|\mathbb{D}},\ \Sigma_i^{\theta|\mathbb{D}}\right), \quad \mu_i^{\theta|\mathbb{D}} = \Sigma_i^{\theta|\mathbb{D}}\left(\sigma_s^{-2}X^\top \text{col}_i(Y) + (\Sigma_i^\theta)^{-1}\mu_i^\theta\right), \quad \Sigma_i^{\theta|\mathbb{D}} = \left(\sigma_s^{-2}X^\top X + (\Sigma_i^\theta)^{-1}\right)^{-1} \tag{17}$$

with data matrices X and Y defined as $\text{row}_j(X) = \phi(x_j, u_j)^\top$ and $\text{row}_j(Y) = x_{j+1}^\top$. Propagation of the posterior distribution $\mathcal{Q}_{\theta|\mathbb{D}}$ through the model (16) then yields the posterior predictive distribution of the model f at locations (x, u) given by[36,49]

$$f_i(x, u; \theta) \sim \mathbb{N}(\phi(x, u)^\top \mu_i^{\theta|\mathbb{D}},\ \phi(x, u)^\top \Sigma_i^{\theta|\mathbb{D}} \phi(x, u)). \tag{18}$$

The corresponding nominal mean prediction model can then, e.g., be selected as $f_i(x, u; \bar{\theta}) = \phi(x, u)^\top \mu_i^{\theta|\mathbb{D}}$ and the set-valued model confidence map according to Definition 2, describing deviations from $f(x, u; \bar{\theta})$ at probability level p_S, can be defined using the chi-squared distribution[51] $\chi_n^2(p_S)$ as

$$\mathcal{D}_{p_S}(x, u) = \{d \in \mathbb{R}^n | d^\top \Sigma^{-1}(x, u) d \leq \chi_n^2(p_S)\} \text{ with } \Sigma(x, u) = \text{diag}([\phi(x, u)^\top \Sigma_i^{\theta|\mathbb{D}} \phi(x, u)]_{i=1,..,n}). \tag{19}$$

It is important to note that in case of a normally distributed prior distribution as considered here, selecting $p_S = 1$ yields an unbounded set-valued model confidence map, meaning that a robust treatment is not possible for unbounded priors on θ.

Given the closed-form expressions for the mean function and set-valued model confidence map (19), we turn our attention to efficient ways of implementing the tightened set-valued model confidence map constraint (12c). We exemplify the procedure by selecting a hyper rectangular admissible disturbance set $\mathcal{D} := \text{conv}(\{d_j e_j\}_{j=1,..,n} \cup \{-d_j e_j\}_{j=1,..,n})$ where e_j are the basis vectors, i.e. $[e_j]_j = 1$, $[e_j]_{i,i\neq j} = 0$, and d_j are the corresponding scalings. We define a corresponding tightened version as $\bar{\mathcal{D}} = (1 - \gamma_{\mathcal{D}})\mathcal{D}$ where $\gamma_{\mathcal{D}} > 0$ is selected sufficiently small according to Definition 3 as described in Section 4.2. In order to implement $\mathcal{D}_{p_S}(z, v) \subseteq \bar{\mathcal{D}}$ in the MPC problem (12), we enforce the semi-axis of $\mathcal{D}_{p_S}(z, v)$ to be smaller or equal to the respective unit vector length $(1 - \gamma_{\mathcal{D}})d_i$ as illustrated in Figure 3, resulting in

$$\phi(z_{i|k}, v_{i|k})^\top \Sigma_i^{\theta|\mathbb{D}} \phi(z_{i|k}, v_{i|k}) \chi_n^2(p_S) \leq ((1 - \gamma_{\mathcal{D}})d_j)^2, \quad j = 1, .., n, \tag{20}$$

representing the resulting set-valued model confidence map constraint (12c). Note that (20) is convex, if $[\phi(.,.)]_i$ are convex, e.g. $\phi(x, u) = [x^\top, u^\top]^\top$.

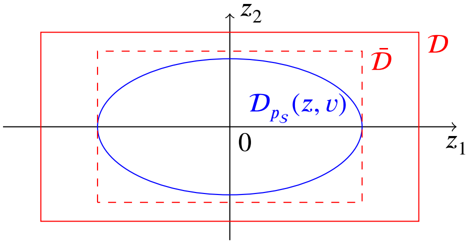

FIGURE 3 Illustration of the ellipsoidal set-valued model confidence set $\mathcal{D}_{p_S}(z, v)$ in blue for a system with two dimensional state space, which is constrained inside the tightened version $\bar{\mathcal{D}}$ (dashed red) of the admissible model error $\mathcal{D}$ (red), i.e. $\mathcal{D}_{p_S}(z, v) \subseteq \bar{\mathcal{D}}$.

Remark 4. While we focus on parametric models of the form (1), it is similarly possible to derive a set-valued model confidence map for non-parametric learning-based predictions. In particular, in case of Gaussian process regression[36], the corresponding set-valued confidence map takes the same form as (19), where $\chi^2(p_S)$ needs to be scaled by an additional factor as described by Srinivas et al. (2012)[37], which was as used before in the learning-based model predictive control context.[6]

Remark 5. While we focus on being robust against parametric uncertainties in this paper, bounded additive disturbances of the form $x(k+1) = f(x(k), u(k); \theta) + w(k)$ with $w(k) \in \mathcal{W}$ can additionally be considered in the set-valued model confidence map by re-defining $\mathcal{D}_{p_S}^{\mathcal{W}}(x, u) := \mathcal{D}_{p_S}(x, u) \oplus \mathcal{W}$

4.2 | Tube synthesis and tightened confidence sets for approximately linear systems

We exemplify a tube synthesis procedure by considering systems of the form (1), which can be approximately described by a linear system model

$$x(k + 1) = Ax(k) + Bu(k) + d(k) \tag{21}$$

with $A := \frac{\partial f}{\partial x}\Big|_0$, $B := \frac{\partial f}{\partial u}\Big|_0$, and $d(k) \in \mathcal{D}_{p_s}(x(k), u(k))$ where $\mathcal{D}_{p_s}$ is the corresponding set-valued model confidence map at probability level p_s as specified in Definition 2. Note that models of the form (16) can easily be turned into (21) by defining $\mathcal{D}_{p_S}^{A,B}(z, v) := (f(z, u; \bar{\theta}) - Az - Bv) \oplus \mathcal{D}_{p_S}(z, v)$ as corresponding set-valued model confidence map.

To enable an efficient synthesis for the auxiliary controller κ_Ω, the RPI set Ω, and the locally admissible model uncertainty set $\mathcal{D}$, as required for the model predictive controller (12), we restrict ourselves to a linear auxiliary tracking controller that results in the well-known standard tube-based model predictive control input[40] $u(k) = \kappa_\Omega(v_{0|k}^*, x(k), z_{0|k}^*) = v_{0|k}^* + K(x(k) - z_{0|k}^*)$ with $K \in \mathbb{R}^{m\times n}$ such that all eigenvalues of $A + BK$ are strictly inside the unit circle and with $v_{0|k}^*$, $z_{0|k}^*$ being the optimal solution to the learning-based model predictive control problem (12) at time k. The resulting error dynamics in case of a linear nominal model (21) reads[41]

$$e(k + 1) = (A + BK)e(k) + d(k). \tag{22}$$

Next, we exploit that constraint (12c) enables us to adjust the considered model uncertainty by planning in confident subsets for controller synthesis. In particular, by extending the design procedure for robust tube MPC by Limon et al. (2008),[52] we present *one* optimization problem that allows to simultaneously synthesize the RPI set Ω, auxiliary control law K, and locally admissible model error set $\mathcal{D}$ for (22) using convex optimization. Specifically, for polytopic state and input constraints $\mathbb{X} := \{x \in \mathbb{R}^n | A_x x \leq b_x\}$ and $\mathbb{U} := \{u \in \mathbb{R}^m | A_u u \leq b_u\}$ in (5), the goal is to impose a maximum amount of constraint tightening of the form $\bar{\mathbb{Z}} = \{\mathbb{X} \ominus \Omega\} \times \{\mathbb{U} \ominus K\Omega\}$. By parametrizing the resulting tightened state and input constraints as

$$\mathbb{X} \ominus \Omega \supseteq \{x \in \mathbb{R}^n | A_x x \leq (1 - \bar{\gamma}_x) b_x\} \text{ and } \mathbb{U} \ominus K\Omega \supseteq \{u \in \mathbb{R}^m | A_u u \leq (1 - \bar{\gamma}_u) b_u\}, \tag{23}$$

we can impose a minimum size of the tightened constraints through bounds on γ_x and γ_u, which directly affects the size of the feasible set of the model predictive control problem (12) and admissible disturbance set $\mathcal{D}$. More precisely, given the minimum size of the tightened constraint sets (23), we compute Ω and K to maximize the locally admissible model uncertainty set $\mathcal{D}$ and thereby make constraint (12c) as least restrictive as possible.

In order to enable scalablity to larger scale systems, we restrict Ω to an ellipsoidal set of the form $\Omega := \{e \in \mathbb{R}^n | e^\top P e \leq 1\}$ with optimization variable $P \in \mathbb{R}^{n \times n}$ that must be positive definite. The form of the admissible uncertainty set $\mathcal{D}$ is chosen to be a scaled polytope $\mathcal{D} = \alpha\tilde{\mathcal{D}}$ with pre-defined shape given by $\tilde{\mathcal{D}} := \operatorname{conv}(\{d_j\}_{i=1}^{n_{\tilde{\mathcal{D}}}})$ and variable scaling $\alpha \geq 0$. A natural choice for stabilization tasks is to select $\tilde{\mathcal{D}}$ such that constraint (12c) is at least approximately fulfilled at the origin, i.e.

$$\tilde{\mathcal{D}} \supseteq \mathcal{D}_{p_s}(0, 0), \quad \alpha \geq 1. \tag{24}$$

The resulting synthesis problem is given by the following optimization problem with $E = P^{-1} \in \mathbb{R}^n$ and $Y = KE \in \mathbb{R}^{m \times n}$

$$\max_{E, Y, \tau, \alpha} \quad \alpha \tag{25a}$$

$$\text{s.t.} \quad \alpha \geq 1, \ \tau \geq 0, E \succeq 0 \tag{25b}$$

$$\begin{pmatrix} \tau E & 0 & EA^\top + Y^\top B \\ 0 & 1 - \tau & \alpha d_j^\top \\ AE + BY & \alpha d_j & E \end{pmatrix} \succeq 0, \quad \forall j = 1, .., n_{\tilde{\mathcal{D}}} \tag{25c}$$

$$\begin{pmatrix} \bar{\gamma}_x^2 b_{x,j}^2 & A_{x,j} E \\ E A_{x,j}^\top & E \end{pmatrix} \succeq 0, \quad \forall j = 1, .., n_x \tag{25d}$$

$$\begin{pmatrix} \bar{\gamma}_u^2 b_{u,j}^2 & A_{u,j} Y \\ Y^\top A_{u,j}^\top & E \end{pmatrix} \succeq 0, \quad \forall j = 1, .., n_u, \tag{25e}$$

the solution of which provides an RPI set for the error system (22) given by $\Omega^* := \{e | e^\top P^* e \leq 1\}$, $P^* = E^{*-1}$, a state-feedback gain $K^* = Y^* E^{*-1}$ with corresponding constraint tightening, which fulfills (23), as well as a possibly large admissible disturbance set $\mathcal{D}^* = \alpha^* \tilde{\mathcal{D}}$. In (25), the previously introduced representations of Ω and $\mathcal{D}$ are used to reformulate the robust invariance condition $\forall e(k) \in \Omega, d(k) \in \mathcal{D} \Rightarrow e(k+1) \in \Omega$ via constraint (25c), obtained by repeated application of the Schur complement[53] and the S-Lemma[54]. Note that constraint (25c) is bilinear only in the scalar parameter $\tau \geq 0$ introduced by the S-Lemma, which needs to be contained in the interval $[0, 1]$. This allows for efficient gridding of τ, resulting in a convex linear matrix inequality for each grid point. The minimum size requirement on the tightened constraints (23) is represented by the convex linear matrix inequalities (25d) and (25e), see e.g. Boyd (1994)[55], Section 5.2.2. While we omit a detailed proof, similar derivations can also be found in Limon et. al. (2008).[52]

Finally, it is important to note that even if (24) holds, $\mathcal{D} = \alpha^* \tilde{\mathcal{D}}$ with $\alpha^* \geq 1$ does not yet guarantee a valid tightening according to Definition 3. This is due to the fact that even for $z_{0|k}^* = 0$ and $v_{0|k}^* = 0$ we have $\mathcal{D}_{p_s}(e, Ke) \neq \mathcal{D}_{p_s}(0, 0)$ for $e \in \Omega$. However, if $\mathcal{D}_{p_s}$ is Lipschitz continuous under the Hausdorff metric (Appendix A, Definition 4) with Lipschitz constant $L_{\mathcal{D}_{p_s}}$, a sufficient condition for a feasible tightened set-valued model confidence map constraint (12c) for $z_{0|k}^* = 0$ and $v_{0|k}^* = 0$ is given by

$$\bar{\mathcal{D}} \oplus \mathcal{B}\left(L_{\mathcal{D}_{p_s}} \operatorname{diam}(\Omega)\right) \subseteq \alpha^* \tilde{\mathcal{D}}. \tag{26}$$

Alternatively, one can verify the requirements in Definition 3 by solving the nonlinear program associated with Corollary 1 condition ii), see Appendix B, (B2). To recap, by focussing on approximately linear system model representations of the form (21), we were able to set up the synthesis problem (25), in which the only required design parameters are given by $\tilde{\mathcal{D}}$ and the maximum constraint tightening factors γ_x, γ_u.

Sufficient conditions to ensure local feasibility of the model predictive control problem

Different from the common case in robust control, where the disturbance set $\mathcal{W}$ is given by the problem formulation, we were able to scale $\mathcal{D}$ as large as possible through the design problem (25) to increase the chance of fulfilling (26) and to obtain a non-empty tightened set-valued model confidence map constraint according to Definition 3. In the following, we show that this additional degree of freedom allows us to pose sufficient conditions in terms of stabilizability and model accuracy to ensure existence of a solution to the design problem (25) in turn implying a non-empty feasible set of the model predictive control problem (12) around the set-point. To this end, the following intermediate step establishes a linear relation between the maximum diameter of the RPI set Ω resulting from (25) and the maximum diameter of the disturbance set $\mathcal{D}^*$.

Lemma 1. Consider a model uncertainty set $\tilde{\mathcal{D}}$ around the origin subject to (24) and let $(A, B) := \left(\frac{\partial f}{\partial x}\Big|_{0,0}, \frac{\partial f}{\partial u}\Big|_{0,0}\right)$ be stabilizable. For a given non-negative maximum RPI set diameter $\bar{e} > 0$ such that $\mathrm{diam}(\Omega) \leq \bar{e}$ and a state-feedback matrix K such that all eigenvalues of $A + BK$ are strictly inside the unit circle and that satisfies

$$\mathcal{B}(\bar{e}) \subseteq \{e \in \mathbb{R}^n | A_x e \leq \bar{\gamma}_x\} \text{ and } K\mathcal{B}(\bar{e}) \subseteq \{u \in \mathbb{R}^m | A_u u \leq \bar{\gamma}_u\}, \tag{27}$$

there exists a corresponding linear bound

$$\mathrm{diam}\left(\tilde{\mathcal{D}}\right) \leq c_\Omega \bar{e}, \tag{28}$$

on the model uncertainty set diameter with $c_\Omega \geq 0$ that ensures feasibility of (25).

Proof. The proof can be found in Appendix C. □

Lemma 1 enables us to state sufficient conditions in terms of the model accuracy via the set-valued model confidence map such that the design procedure (25) leads to a learning-based MPC problem (12) with non-empty region of attraction (=feasible set).

Proposition 2. Consider a system model of the form (21), let $(A, B) := \left(\frac{\partial f}{\partial x}\Big|_{0,0}, \frac{\partial f}{\partial u}\Big|_{0,0}\right)$ be stabilizable, and assume that $\mathbb{X}$ and $\mathbb{U}$ contain 0 in their interior. If $\mathcal{D}_{p_s}(0,0)$ is Lipschitz continuous under the Hausdorff metric with Lipschitz constant $L_{\mathcal{D}_{p_s}} > 0$ and $\mathrm{diam}\left(\tilde{\mathcal{D}}\right) \geq \mathrm{diam}\left(\mathcal{D}_{p_s}(0,0)\right) > 0$ sufficiently small, then the MPC Problem (12) with $\ell_f = 0$, $\mathcal{X}_f = \{0\}$ has a non-empty feasible set.

Proof. Since $\mathbb{X}$ and $\mathbb{U}$ contain 0 in their interior, we can conclude from Lemma 1 that there exists a maximum diameter $\mathrm{diam}(\Omega) > 0$ and a constant $c_\Omega > 0$ such that $\alpha^* \mathrm{diam}\left(\tilde{\mathcal{D}}\right) \leq c_\Omega \mathrm{diam}(\Omega)$, for an $\alpha^* \geq 1$, which implies feasibility of the synthesis problem (25). Let $\mathrm{diam}\left(\tilde{\mathcal{D}}\right) \leq \alpha^{*-1} c_\Omega \mathrm{diam}(\Omega) > 0$ and note that

$$L_{\mathcal{D}_{p_s}} \mathrm{diam}(\Omega) \leq (\alpha^* - 1)\mathrm{diam}\left(\tilde{\mathcal{D}}\right) \tag{29}$$

is a sufficient condition for (26). Feasibility of the tightened set-valued model confidence map constraint (12c) for $\{z_i\}_{i=0}^N = 0$ and $\{v_i\}_{i=0}^{N-1} = 0$ can therefore be obtained by inserting the linear bound $\alpha^* \mathrm{diam}\left(\tilde{\mathcal{D}}\right) \leq c_\Omega \mathrm{diam}(\Omega)$ in (29) yielding $L_{\mathcal{D}_{p_s}} \leq (1 - \alpha^{*-1}) c_\Omega > 0$. Lastly, we have by assumption that $(0, 0)$ is a steady state for system (21) and therefore $\{z_i\}_{i=0}^N = 0$ and $\{v_i\}_{i=0}^{N-1} = 0$ is a feasible solution to (12) with the given terminal configuration, completing the proof. □

From Proposition 2 we can conclude that we can always find a solution to (25) by selecting $\mathcal{D}$, i.e. d_i according to Section 4.1, sufficiently small. Furthermore, it implies that a corresponding tightened confidence map (Definition 3) exists such that feasibility of the resulting MPC problem can be guaranteed around the origin, if the Lipschitz constant associated with the set-valued model confidence map and the predicted model error at the origin are sufficiently small. For practical applications, we can therefore conclude that infeasibility of the MPC problem due to the set-valued model confidence map constraint (12c) can be fixed by either collecting more data around the set point or by lowering the chance constraint satisfaction probability p_S.

Remark 6. While we have shown that for sufficiently accurate model estimates the tube design based on (25) leads to a well-defined MPC problem, the RPI set Ω might be overly conservative, since it is restricted to ellipsoidal shapes. The set can be improved by approximating the true minimal polytopic robust invariant set, e.g., using the method described by Rakovic et. al. (2005),[56] which can also be applied to higher dimensional systems, see Darup and Dieter (2019)[57].

4.3 | Terminal ingredients for learning-based prediction models

As presented for example by Rawlings and Mayne (2008)[40] in Section 2.5 for setpoint stabilization control tasks and in Amrit et al. (2011)[58] for economic control tasks, there exist principled procedures to compute a corresponding terminal cost ℓ_f together with a terminal set $\mathcal{X}_f$, satisfying Assumption 2. To this end, any nominal model of the form $z(k+1) = f(z(k), v(k); \bar{\theta})$ can be used long as the linearization (A, B) around a target steady-state z_s, v_s given by $A := \frac{\partial f}{\partial x}(z_s, v_s; \bar{\theta})$ and $B := \frac{\partial f}{\partial u}(z_s, v_s; \bar{\theta})$ is stabilizable and the model is sufficiently confident meaning that $z \in \mathcal{X}_f \Rightarrow \mathcal{D}_{p_S}(z, \kappa_f(z)) \subseteq \bar{\mathcal{D}}$ according to Theorem 1. The latter requirement can be verified through the condition

$$
\begin{aligned}
0 = \max_{z,v,d} \quad & \|d\|_{\bar{\mathcal{D}}} && \text{(30a)} \\
\text{s.t.} \quad & z \in \mathcal{X}_f && \text{(30b)} \\
& v = \kappa_f(z) && \text{(30c)} \\
& d \in \mathcal{D}_{p_S}(z, v). && \text{(30d)}
\end{aligned}
$$

In case (30) does not hold, one can often either shrink the terminal set iteratively until (30) is fulfilled, with $\mathcal{X}_f = 0$ reducing to the case discussed in Proposition 2, or gather more data to improve the model confidence. If more data needs to be collected, note that (30) provides the location (z^*, v^*) in state and input space where the magnitude of model uncertainty causes problems and where data should be obtained first. Safe collection of additional data around the desired steady state can, e.g., be achieved using a linear stabilizing control law together with probabilistic safety verification techniques[33,34,59]. Existing literature therefore provides principled ways for designing terminal components ℓ_f and $\mathcal{X}_f$ providing recursive feasibility and performance guarantees based on the nominal system dynamics.

5 | NUMERICAL EXAMPLES OF PROVABLY SAFE LEARNING-BASED MPC

5.1 | Approximately linear 10-dimensional quadrotor system

Similar to recent work on robust model predictive techniques,[60,47] we consider the problem of controlling a simplified quadrotor system, which can be described by

$$
\dot{x}(t) = \begin{pmatrix} x_4(t) \\ x_5(t) \\ x_6(t) \\ g\tan(x_7(t)) \\ g\tan(x_8(t)) \\ -g \\ -d_1 x_7(t) - x_9(t) \\ -d_1 x_8(t) + x_{10}(t) \\ -d_0 x_3(t) \\ -d_0 x_7(t) \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ k_T u_3(t) \\ 0 \\ 0 \\ n_0 u_1(t) \\ n_0 u_2(t) \end{pmatrix}, \quad \text{with states} \quad \begin{matrix} x_1: & \text{x rel. position} \\ x_2: & \text{y rel. position} \\ x_3: & \text{z rel. position} \\ x_4: & \text{x velocity} \\ x_5: & \text{y velocity} \\ x_6: & \text{z velocity} \\ x_7: & \text{pitch angle} \\ x_8: & \text{roll angle} \\ x_9: & \text{pitch rate} \\ x_{10}: & \text{roll rate} \end{matrix}, \quad \text{inputs} \quad \begin{matrix} u_1: & \text{pitch rate change} \\ u_2: & \text{roll rate change} \\ u_3: & \text{vertical thrust} \end{matrix},
$$

with relative position coordinates around any target position and parameters $d_0 = 10$, $d_1 = 8$, $k_T = 0.91$, $g = 9.81$, and $n_0 = 10$. The input authority is limited by $|u_1| \leq \pi/2$, $|u_2| \leq \pi/2$ and $-g/k_T \leq u_3 \leq 2g$. The only constraint on states is given by $x_1 \leq 1$, representing a wall and therefore a safety-critical specification to keep the quadrotor from crashing. The control objective is to reach and control the steady-state $x_r = 0_n$ without crashing into the wall. The simulation is implemented using a Euler-forward discretization scheme with sampling time $h = 0.3$. We generate two data sets $\mathbb{D}_{40}$ and $\mathbb{D}_{200}$ by sampling uniformly random points (x_i, u_i) around the set-point, tainted with i.i.d. normally distributed noise with zero mean and standard deviation given by $\sigma_s = 0.001$.

We leverage an approximately linear system representation for small angles and select $|x_7| \leq 0.75$ and $|x_8| \leq 0.75$ as additional state constraints together with a linear model of the form $f(x, u) = \theta^\top \phi(x, u)$ with $\phi(x, u) = [x^\top, u^\top]^\top$ and θ representing the linear system dynamics with unknown parameters d_0, d_1, g, n_0 ,k_T. To infer a distribution over the unknown parameters θ from available data using Bayesian regression as described in Section 4.1 we pick a normal prior distribution over the parameters θ with $\mu_i^\theta = 0$ and $\Sigma_i^\theta = I_n$.

For the computation of the tube Ω, tube auxiliary controller κ_Ω, and locally admissible model uncertainty set $\mathcal{D}$, we apply the synthesis procedure for approximately linear systems as presented in Section 4. We select a maximum tightening factor $\gamma_x = \gamma_u = 0.1$ together with an initial shape $\tilde{\mathcal{D}} := \text{conv}(\{e_j, -e_j\}_{j=1,..,10})$ with e_j unit vectors in $j-th$ direction, leading to a feasible solution with maximum admissible disturbance set $\mathcal{D} = 0.01365\tilde{\mathcal{D}}$. While we directly use the tightened state and input constraint according to (23), we derive the tightened confidence set $\bar{\mathcal{D}}$ using the Lipschitz constant of the right hand side of (20) for each state dimension. The terminal set is selected equal to the target steady-state $x_s = 0_n$, $u_s = [0, 0, g/k_T]^\top$, fulfilling Assumption 2. To demonstrate the effect of the set-valued model confidence map constraint presented in Figure 4, we compare the approximate feasible sets of the MPC problems (12) resulting from the data set $\mathbb{D}_{40}$ with the one based on $\mathbb{D}_{200}$ using a prediction horizon of $N = 13$ time steps.

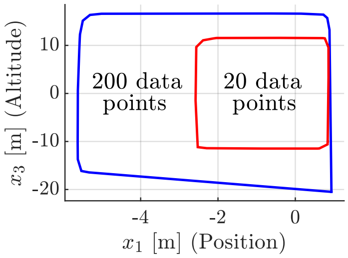

FIGURE 4 Approximate feasible sets of the model predictive control problems projected on the $x-z$ plane based on 20 (red) and 200 (blue) data points.

As expected, an increased number of data points leads to an increased model confidence that, in turn, render constraint (22c) less conservative, yielding a larger feasible set. Note that there does not exist a uniform error bound, which prohibits application of standard tube-based MPC, see Section 3.1. This is due to the state and input dependent uncertainty estimate that grows unbounded with the unbounded system states as described in Section 3.2.

5.2 | Economic operation of a nonlinear system

To demonstrate applicability of the proposed method to a nonlinear system and the extension to economic control tasks, we consider the following numerical example[61,39] with bilinear system dynamics given as

$$x(k+1) = \begin{bmatrix} 0.55 & 0.12 \\ 0 & 0.67 \end{bmatrix} x(k) + \left(\begin{bmatrix} -0.6 & 1 \\ 1 & -0.8 \end{bmatrix} x(k) + \begin{bmatrix} 0.01 \\ 0.15 \end{bmatrix} \right) u(k). \tag{31}$$

The system is subject to state and input constraints $\mathbb{X} = \{x \in \mathbb{R}^2 : \|x\|_\infty \leq 1\}$ and $\mathbb{U} = \{u \in \mathbb{R} : |u| \leq 0.15\}$. The objective is to minimize the value of state x_2 while maintaining a positive value of x_1 which is given by

$$\ell(x) = x_2 + \begin{cases} -100x_1^3, & x_1 < 0 \\ 0, & \text{else.} \end{cases} \tag{32}$$

The corresponding model with nonlinear features according to Section 4.1 is selected as $f(x,u) = \theta^\top \phi(x,u)$ with $\phi(x,u) = [x^\top, u^\top, x_1 u, x_2 u]^\top$, $\theta \in \mathbb{R}^{2\times 5}$ with 10 unknown system parameters. For illustration, we infer these parameters using two different data sets $\mathbb{D}_{30}$ and $\mathbb{D}_{60}$, each tainted with i.i.d. normally distributed noise with zero mean and standard deviation $\sigma_s = 0.1$ and using a normal prior distribution over the parameters θ with $\mu_i^\theta = 0$ and $\Sigma_i^\theta = 10I_5$. For controller design we begin by defining the admissible disturbance set as $\mathcal{D} := \{w \in \mathbb{R}^2 : \|w\|_\infty \leq 0.05\}$ and follow the corresponding tube derivation as described in Bayer et. al. (2013)[61] that is based on incremental input-to-state stability of (31). we can then select $\kappa_\Omega(v, x, z) = v$ and verify that $\Omega = \{e \in \mathbb{R}^2 : \|e\|_\infty \leq 1/3\}$ is an RPI set according to Definition 1 for both nominal models, i.e. for expected parameters $\mathbb{E}[\theta]$ infered from both data sets $\mathbb{D}_{30}$ and $\mathbb{D}_{60}$. In Figure 5 we illustrate a simplification of the resulting set-valued model confidence map constraint (12c) for $p_S = 0.9$ by displaying the maximum model uncertainty w.r.t. the infinity norm for different states and inputs. We deduce $\bar{\mathcal{D}}_{60} = \mathcal{D}$ and $\bar{\mathcal{D}}_{30} = 0.01\mathcal{D}$ as tightened confidence sets according to Definition 3.

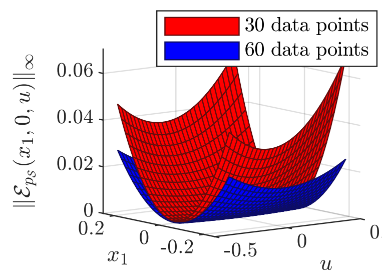

FIGURE 5 Illustration of the set-valued model confidence map resulting from $\mathbb{D}_{30}$ and $\mathbb{D}_{60}$. On the z-axis we display with $\|\mathcal{D}_{p_S}(x_1, 0, u)\|_\infty$ the maximum model uncertainty w.r.t. the infinity norm at x_1 and u for $x_2 = 0$.

From Figure 5 we can also conclude that the specified disturbance set $\mathcal{D}$ would not suffices to uniformly describe all possible state and input dependent uncertainties. Following a standard tube-based MPC approach as described in Section 3.1 would therefore require to increase the additive disturbance set $\mathcal{D}$, rendering the resulting controller overly conservative.

To demonstrate the effect of the set-valued model confidence map constraint on the economic performance we perform a stochastic simulation analysis, where we randomly generate 100 different data sets $\mathbb{D}_{30}$, $\mathbb{D}_{60}$ and apply the resulting learning based controller with prediction horizon $N = 70$ to system (31) for each data set over 100 time steps starting from the initial condition $x(0) = 0$. While all generated controllers remain recursively feasible during each experiment, the average cost in case of $\mathbb{D}_{30}$ is 0.045 ± 0.03, whereas we get -0.18 ± 0.11 for $\mathbb{D}_{60}$. This result illustrates that limited availability of data can lead to cautious behavior, and the performance can be improved through additional data.

6 | CONCLUSION

We presented a model predictive control approach that supports learning-based probabilistic prediction models and thereby enables efficient learning-based controller design. Using a tube-based model predictive control mechanism at the core of the method in combination with the idea of planning in confident subsets, we provide rigorous recursive feasibility and closed-loop constraint satisfaction and performance. The presented formulation provides a flexible interface to different classes of probabilistic prediction models and allows easy adaptation to different control problems such as trajectory tracking and economic model predictive control. To facilitate application of the controller, we proposed an efficient design procedure for obtaining the required tube and tube controller for linear prediction models. Using an approximately linear 10-dimensional quadrotor system to achieve set-point stabilization together with a nonlinear economic control task, we demonstrated the application of the design procedure and the behavior of the learning-based model predictive controller.

How to cite this article: K. P. Wabersich and M. N. Zeilinger (2019), Nonlinear learning-based model predictive control supporting stochastic state and input dependent model uncertainty estimates.

APPENDIX

A HAUSDORFF METRIC

Definition 4. The *Hausdorff metric* between two sets $\mathcal{A}$ and $\mathcal{B}$ in a metric space (M, d_M) is defined as

$$d_\mathrm{H}(\mathcal{A}, \mathcal{B}) := \max\left\{ \sup_{a\in\mathcal{A}} \inf_{b\in\mathcal{B}} d_\mathrm{M}(a,b), \inf_{a\in\mathcal{A}} \sup_{b\in\mathcal{B}} d_\mathrm{M}(a,b) \right\}.$$

B MODEL UPDATE VERIFICATION PROBLEMS ACCORDING TO COROLLARY 1

Conditions i) and ii) in Corollary 1 can be verified through the conditions

$$\begin{aligned} 0 = \max_{z,v} \; & \|d\|_{\mathcal{D}_{p_S}(z,v)} && \text{(B1a)}\\ \text{s.t. } & (z,v) \in \bar{\mathbb{Z}} && \text{(B1b)}\\ & d \in \mathcal{D}^+_{p_S}(z,v) && \text{(B1c)} \end{aligned}$$

$$\begin{aligned} 0 = \max_{z,v,x} \; & \|d\|_{\mathcal{D}} && \text{(B2a)}\\ \text{s.t. } & (z,v) \in \bar{\mathbb{Z}},\; x \in z \oplus \Omega && \text{(B2b)}\\ & \mathcal{D}^+_{p_S}(z,v) \subseteq \bar{\mathcal{D}} && \text{(B2c)}\\ & d \in \mathcal{D}^+_{p_S}(x, \kappa(v,x,z)) && \text{(B2d)} \end{aligned}$$

using nonlinear optimization methods for, e.g., the case of Bayesian regression as described in Section 4.1, rendering (B2c) into a set of nonlinear inequality constraints. If either (B1) or (B2) has an optimal value greater than zero, then i) or ii) in Corollary 1 does not hold.

C PROOF OF LEMMA 1

Proof. The proof of Lemma 1 is structured as follows. We start by solving the standard Lyapunov equation based on the stabilizing feedback gain K to obtain an ellipsiodal nominal invariant set Ω (RPI set for disturbances $\tilde{\mathcal{D}} = \{0\}$) and scale it such that $\mathrm{diam}\,(\Omega) \le \bar{e}$. Based on the set Ω we then argue that there exists the desired linear relation of the form (28) that ensures that Ω is RPI for $\tilde{\mathcal{D}} \neq \{0\}$ small enough.

Since all eigenvalues of $A + BK$ are strictly inside the unit circle, for every symmetric positive definite matrix Q there exists a symmetric positive definite matrix P such that the equation

$$e^\top (A_c^\top P A_c - P)e \le -e^\top Q e \tag{C3}$$

holds with $A_c = A + BK$. Using (C3) we define the ellipsoidal set of the form $\Omega := \{e | e^\top P e \le c_P\}$ and select $c_P := \lambda^-(P)\bar{e}^2$ with $\lambda^-(P)$ the smallest eigenvalue of P. The specific choice of Ω implies for $e \in \Omega$ that $e^\top P e \ge \lambda^-(P) e^\top e$ and $e^\top P e \le c_P = \lambda^-(P)\bar{e}^2$ yielding the relation $\max_{e\in\Omega} \|e\| \le \bar{e} \Rightarrow \mathrm{diam}(\Omega) \le \bar{e}$.

Next, we derive an auxiliary result, which we will be needed later to show robust invariance of Ω for small disturbances. Therefore we first need to find a possibly large 2-norm ball contained in Ω. Note that for all e such that $e^\top P e = c_P$ we have $\|e\|^2 \lambda^+(P) \ge c_P$ resulting in $\|e\|^2 \ge \frac{\lambda^-(P)}{\lambda^+(P)}\bar{e}^2$, with $\lambda^+(P)$ the largest eigenvalue of P, implying that the desired ball is given by $\mathcal{B}\left(\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}\right) \subseteq \Omega$. The auxiliary result we want to establish is a linear bound of the form (28) such that for all $e \in \mathcal{B}\left(\frac{0.5}{\max(\|A_c\|,1)}\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}\right)$ it follows $e^+ \in \mathcal{B}\left(\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}\right)$. To this end we select for all $e \in \mathcal{B}\left(\frac{0.5}{\max(\|A_c\|,1)}\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}\right)$ through the bound

$$e^+ \le \left\|A_c e + d\right\| \le 0.5\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e} + \bar{d}$$

the maximum disturbance diameter as $\bar{d} \le (1 - 0.5)\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}$, which provides us the desired relation of the form $\bar{d} \le c_1\bar{e}$ such that $\|e^+\| \le \sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}$ and therefore $e^+ \in \mathcal{B}\left(\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}\right) \subseteq \Omega$.

Using the previously derived auxiliary result, we will now establish the desired bound (28) such that Ω is RPI. To this end it remains to show that

$$\forall e \in \Omega \text{ such that } \|e\| \ge \frac{0.5}{\max(\|A_c\|,1)}\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e} \text{ it holds } e^+ \in \Omega, \tag{C4}$$

because for all $\|e\| \le \frac{0.5}{\max(\|A_c\|,1)}\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}$ we already know that $e^+ \in \Omega$ by the auxiliary result. From (C4) we continue by requiring for all $e \in \Omega$ with $\|e\| \ge \frac{0.5}{\max(\|A_c\|,1)}\sqrt{\frac{\lambda^-(P)}{\lambda^+(P)}}\bar{e}$ that $e^{+\top} P e^{+\top} \le e^\top P e$, yielding a sufficient condition for (C4) equal to

$$e^\top (A_c^\top P A_c - P)e + 2d^\top P A_c e + d^\top P d \le 0. \tag{C5}$$

The left hand side of (C5) can be bounded from above as

$$\begin{aligned}
&e^\top (A_c^\top P A_c - P)e + 2d^\top P A_c e + d^\top P d \\
\le& - e^\top Q e + 2d^\top P A_c e + d^\top P d \\
\le& - \lambda^-(Q)\frac{0.25}{\max(\|A_c\|, 1)^2}\frac{\lambda^-(P)}{\lambda^+(P)}\bar{e}^2 + 2\left\|PA_c\right\|\bar{e}\bar{d} + \lambda^+(P)\bar{d}^2,
\end{aligned}$$

providing us a sufficient choice for $\bar{d} > 0$ by solving

$$-\lambda^-(Q)\frac{0.25}{\max(\|A_c\|, 1)^2}\frac{\lambda^-(P)}{\lambda^+(P)}\bar{e}^2 + 2\left\|PA_c\right\|\bar{e}\bar{d} + \lambda^+(P)\bar{d}^2 = 0.$$

Utilizing the quadratic formula one obtains

$$\bar{d} = \frac{-2\left\|PA_c\right\|\bar{e} + (-)\sqrt{4\left\|PA_c\right\|^2\bar{e}^2 + \frac{\lambda^-(Q)}{\max(\|A_c\|,1)^2}\frac{\lambda^-(P)}{\lambda^+(P)}\lambda^+(P)\bar{e}^2}}{2\lambda^+(P)} = \underbrace{\frac{-2\left\|PA_c\right\| + \sqrt{4\left\|PA_c\right\|^2 + \frac{\lambda^-(Q)\lambda^-(P)}{\max(\|A_c\|,1)^2}}}{2\lambda^+(P)}}_{c_2>0}\bar{e}.$$

Finally, define the required $c_\Omega := \min\{c_1, c_2\}$ for which the previously constructed set Ω is RPI under condition (28). The resulting solution candidate $E^* := (Pc_P^{-1})^{-1}$ and $Y^* := KE^*$ satisfy by assumption (27) the constraints (25d) and (25e). Furthermore, because Ω is RPI, it follows due to necessity of the S-Lemma[54] that there exists a $\tau^* \ge 0$ such that (25c) holds, implying together with $\alpha^* = 1$ overall feasibility of (25), which is was we wanted to show. □

Data Availability Statement

Data sharing is not applicable to this article as no new real world data were created or analyzed in this article.

References

1. Morari M, Lee JH. Model predictive control: Past, present and future. *Comput. & Chem. Eng.* 1999; 23(4): 667–682.
2. Qin SJ, Badgwell TA. An Overview of Nonlinear Model Predictive Control Applications. In: Nonlinear Model Predictive Control. Birkhäuser Basel; 2000: 369–392.
3. Hewing L, Kabzan J, Zeilinger MN. Cautious Model Predictive Control Using Gaussian Process Regression. *IEEE Transactions on Control Systems Technology* 2019: 1–12.
4. Carron A, Arcari E, Wermelinger M, Hewing L, Hutter M, Zeilinger MN. Data-Driven Model Predictive Control for Trajectory Tracking With a Robotic Arm. *IEEE Robotics and Automation Letters* 2019; 4(4): 3758-3765.
5. Kamthe S, Deisenroth MP. Data-efficient reinforcement learning with probabilistic model predictive control. *International Conference on Artificial Intelligence and Statistics, AISTATS 2018* 2018; 84: 1701–1710.
6. Koller T, Berkenkamp F, Turchetta M, Krause A. Learning-Based Model Predictive Control for Safe Exploration. In: 2018 IEEE Conference on Decision and Control (CDC). ; 2018: 6059-6066.
7. Soloperto R, Müller MA, Trimpe S, Allgöwer F. Learning-Based Robust Model Predictive Control with State-Dependent Uncertainty. In: Proc. 6th IFAC Conf. Nonlinear Model Predictive Control. ; 2018: 442–447
8. Hewing L, Liniger A, Zeilinger MN. Cautious NMPC with Gaussian Process Dynamics for Autonomous Miniature Race Cars. In: 2018 European Control Conference (ECC). ; 2018: 1341-1348.
9. Ostafew CJ, Schoellig AP, Barfoot TD. Robust Constrained Learning-based NMPC enabling reliable mobile robot path tracking. *International Journal of Robotics Research* 2016; 35(13): 1547–1563.
10. Calliess JP. *Conservative decision-making and inference in uncertain dynamical systems*. PhD thesis. University of Oxford, Oxford, England; 2014.
11. Terzi E, Fagiano L, Farina M, Scattolini R. Learning multi-step prediction models for receding horizon control. In: 2018 European Control Conference, ECC 2018. ; 2018: 1335–1340
12. Aswani A, Gonzalez H, Sastry SS, Tomlin CJ. Provably safe and robust learning-based model predictive control. *Automatica* 2013; 49(5): 1216–1226.
13. Bouffard P, Aswani A, Tomlin CJ. Learning-based model predictive control on a quadrotor: Onboard implementation and experimental results. In: Proc. IEEE Int. Conf. Robot. and Automat. ; 2012: 279–284
14. Aswani A, Bouffard P, Tomlin C. Extensions of learning-based model predictive control for real-time application to a quadrotor helicopter. In: Proceedings of the American Control Conference. ; 2012: 4661–4666
15. Aswani A, Master N, Taneja J, Culler D, Tomlin C. Reducing Transient and Steady State Electricity Consumption in HVAC Using Learning-Based Model-Predictive Control. *Proceedings of the IEEE* 2012; 100(1): 240-253.
16. Aswani A, Master N, Taneja J, Krioukov A, Culler D, Tomlin CJ. Energy-Efficient Building HVAC Control Using Hybrid System LBMPC. In: Proc. 4th IFAC Conf. Nonlinear Model Predictive Control. ; 2012: 496–501
17. Beliakov G. Interpolation of Lipschitz functions. *Journal of Computational and Applied Mathematics* 2006; 196(1): 20–44.
18. Calliess JP, Roberts S, Rasmussen C, Maciejowski J. Nonlinear Set Membership Regression with Adaptive Hyper-Parameter Estimation for Online Learning and Control. In: 2018 European Control Conference, ECC 2018. ; 2018: 3167–3172.
19. Limon D, Calliess J, Maciejowski JM. Learning-based Nonlinear Model Predictive Control. In: Proc. 20th IFAC World Congress. ; 2017: 7769–7776.
20. Fukushima H, Kim TH, Sugie T. Adaptive model predictive control for a class of constrained linear systems based on the comparison model. *Automatica* 2007; 43(2): 301–308.

21. Tanaskovic M, Fagiano L, Smith R, Morari M. Adaptive receding horizon control for constrained MIMO systems. *Automatica* 2014; 50(12): 3019–3029.

22. Adetola V, DeHaan D, Guay M. Adaptive model predictive control for constrained nonlinear systems. *System & Control Letters* 2009; 58(5): 320–326.

23. Adetola V, Guay M. Robust adaptive MPC for constrained uncertain nonlinear systems. *International Journal Adaptive Control and Signal Processing* 2011; 25(2): 155–167.

24. Gonçalves GAA, Guay M. Robust discrete-time set-based adaptive predictive control for nonlinear systems. *J. Process Control* 2016; 39: 111–122.

25. Di Cairano S. Indirect adaptive model predictive control for linear systems with polytopic uncertainty. In: Proceedings of the American Control Conference. ; 2016: 3570–3575.

26. Lorenzen M, Cannon M, Allgöwer F. Robust MPC with recursive model update. *Automatica* 2019; 103: 461–471.

27. Mesbah A. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine* 2016; 36(6): 30–44.

28. Hewing L, Zeilinger MN. Scenario-based probabilistic reachable sets for recursively feasible stochastic model predictive control. *IEEE Control Systems Letters* 2019; 4(2): 450–455.

29. Mammarella M, Alamo T, Lucia S, Dabbene F. A probabilistic validation approach for penalty function design in Stochastic Model Predictive Control. *arXiv preprint arXiv:2003.07241* 2020.

30. Wabersich KP, Zeilinger MN. Linear model predictive safety certification for learning-based control. In: Proc. 57th IEEE Conf. Decision and Control. ; 2018: 7130–7135.

31. Wabersich KP, Hewing L, Carron A, Zeilinger MN. Probabilistic model predictive safety certification for learning-based control. *arXiv:1906.10417* 2019.

32. Wabersich KP, Zeilinger MN. Safe exploration of nonlinear dynamical systems: A predictive safety filter for reinforcement learning. *arXiv:1812.05506* 2018.

33. Prajna S, Jadbabaie A, Pappas GJ. Stochastic safety verification using barrier certificates. In: 43rd IEEE Conference on Decision and Control (CDC). ; 2004: 929–934.

34. Huang C, Chen X, Lin W, Yang Z, Li X. Probabilistic safety verification of stochastic hybrid systems using barrier certificates. *ACM Transactions on Embedded Computing Systems (TECS)* 2017; 16(5s): 1–19.

35. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB. *Bayesian data analysis*. Chapman and Hall/CRC . 2013.

36. Rasmussen CE, Williams CKI. *Gaussian processes for machine learning*. The MIT Press . 2006.

37. Srinivas N, Krause A, Kakade SM, Seeger MW. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory* 2012; 58(5): 3250–3265.

38. Chowdhury SR, Gopalan A. On kernelized multi-armed bandits. In: International Conference on Machine Learning (ICML). ; 2017: 844–853.

39. Bayer FA, Müller MA, Allgöwer F. Tube-based robust economic model predictive control. *Journal of Process Control* 2014; 24(8): 1237–1246.

40. Rawlings JB, Mayne DQ, Diehl MM. *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing. 2 ed. 2017.

41. Mayne DQ, Seron MM, Raković S. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica* 2005; 41(2): 219–224.

42. Yu S, Maier C, Chen H, Allgöwer F. Tube MPC scheme based on robust control invariant set with application to Lipschitz nonlinear systems. *Systems & Control Letters* 2013; 62(2): 194–200.

43. Borrelli F. *Constrained optimal control of linear and hybrid systems*. 290. Springer . 2003.

44. Limon D, Alvarado I, Alamo T, Camacho E. Robust tube-based MPC for tracking of constrained linear systems with additive disturbances. *Journal of Process Control* 2010; 20(3): 248–260.

45. Villanueva ME, Quirynen R, Diehl M, Chachuat B, Houska B. Robust MPC via min–max differential inequalities. *Automatica* 2017; 77: 311–321.

46. Singh S, Majumdar A, Slotine JJ, Pavone M. Robust online motion planning via contraction theory and convex optimization. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). ; 2017: 5883–5890.

47. Köhler J, Soloperto R, Müller MA, Allgöwer F. A computationally efficient robust model predictive control framework for uncertain nonlinear systems.

48. Wabersich KP, Bayer FA, Müller MA, Allgöwer F. Economic Model Predictive Control for Robust Periodic Operation with Guaranteed Closed-Loop Performance. In: ; 2018: 507-513

49. Bishop CM. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag . 2006.

50. Hastie T, Tibshirani R, Friedman J. *The Elements of Statistical Learning*. New York, NY, USA: Springer New York Inc. . 2001.

51. Slotani M. Tolerance regions for a multivariate normal population. *Annals of the Institute of Statistical Mathematics* 1964; 16(1): 135–153.

52. Limon D, Alvarado I, Alamo T, Camacho E. On the design of Robust tube-based MPC for tracking. *IFAC Proceedings Volumes* 2008; 41(2): 15333–15338.

53. Boyd S, Vandenberghe L. *Convex Optimization*. New York, NY, USA: Cambridge University Press . 2004.

54. Derinkuyu K, Pınar MÇ. On the S-procedure and some variants. *Mathematical Methods of Operations Research* 2006; 64(1): 55–77.

55. Boyd S, El Ghaoui L, Feron E, Balakrishnan V. *Linear Matrix Inequalities in System and Control Theory*. SIAM studies in applied mathematics: 15 . 1994.

56. Rakovic SV, Kerrigan EC, Kouramas KI, Mayne DQ. Invariant approximations of the minimal robust positively invariant set. In: . 50 of *IEEE Transactions on Automatic Control*. IEEE; 2005: 406–410

57. Darup MS, Teichrib D. Efficient computation of RPI sets for tube-based robust MPC. In: 18th European Control Conference (ECC). IEEE; 2019: 325–330.

58. Amrit R, Rawlings JB, Angeli D. Economic optimization using model predictive control with a terminal cost. *Annual Reviews in Control* 2011; 35(2): 178–186.

59. Wabersich KP, Zeilinger MN. Scalable synthesis of safety certificates from data with application to learning-based control. In: 2018 European Control Conference, ECC 2018. ; 2018: 1691–1697

60. Hu H, Feng X, Quirynen R, Villanueva ME, Houska B. Real-time tube MPC applied to a 10-state quadrotor model. In: 2018 Annual American Control Conference (ACC). ; 2018: 3135–3140.

61. Bayer F, Bürger M, Allgöwer F. Discrete-time incremental ISS: A framework for robust NMPC. In: 2013 European Control Conference (ECC). ; 2013: 2068–2073.