Proceedings of Machine Learning Research vol 242:1058–1070, 2024 6th Annual Conference on Learning for Dynamics and Control

Convergence guarantees for adaptive model predictive control with kinky inference

Riccardo Zuliani RZULIANI@ETHZ.CH

Raffaele Soloperto SOLOPERR@ETHZ.CH

John Lygeros JLYGEROS@ETHZ.CH

All the authors are with the Automatic Control Laboratory, ETH Zürich

Editors: A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

Abstract

We analyze the convergence properties of a robust adaptive model predictive control algorithm used to control an unknown nonlinear system. We show that by employing a standard quadratic stabilizing cost function, and by recursively updating the nominal model through *kinky inference*, the resulting controller ensures convergence of the true system to the origin, despite the presence of model uncertainty. We illustrate our theoretical findings through a numerical simulation.

Keywords: Adaptive model predictive control, adaptive control, kinky inference.

1. Introduction

Motivation During the last few decades, model predictive control (MPC) has attracted large attention because of its efficiency in handling nonlinear systems subject to hard state and input constraints, while minimizing a user-defined cost function, Rawlings et al. (2017). An MPC scheme employs a nominal model of the system dynamics to predict future trajectories over a given prediction horizon. However, in several applications, obtaining an accurate model can be expensive in terms of money and resources, Darby and Nikolaou (2012). Model inaccuracies, combined with the presence of disturbances affecting the system, might lead to a deterioration of performance, constraint violation, or even instability, Forbes et al. (2015).

To address this issue, research has focused on adaptive MPC schemes where the identification of the model is performed online together with the computation of the input. Since excitation is not explicitly enforced, such approaches are often referred to as passive-learning controllers, Mesbah (2018). One of the main drawbacks of passive-learning approaches is that newly generated data might not be informative (for example, when the system reaches a steady state), and therefore an improvement in the model estimate is not guaranteed. This issue is addressed in active-learning approaches where a form of excitation is explicitly induced, often by incorporating a learning cost into the MPC cost function, Tanaskovic et al. (2019), Soloperto et al. (2019a). Even though excitation is beneficial for model adaption, it can be counterproductive in terms of stability. Motivated by this, we analyze the stability properties of an adaptive MPC scheme.

Related works A Gaussian process (GP) is a collection of random variables, any finite subset of which follows a multivariate Gaussian distribution, Rasmussen (2003). In the context of adaptive MPC, GPs are employed to construct a nominal model and sets that bound, within a certain probability, the evolution of the true system and the learned one, Hewing et al. (2019), Bradford et al. (2020). However, in the case where robust constraint satisfaction is needed, e.g., in safety-critical systems, GP-based adaptive MPC schemes cannot be employed as they fail to ensure hard constraint satisfaction.

Conversely, robust adaptive MPC approaches use a set-membership method to construct robust sets where the uncertainty lies, and then incorporate this knowledge into a robust tube-based MPC scheme, see, e.g., Lorenzen et al. (2019), Lu and Cannon (2019), Tanaskovic et al. (2019), Köhler et al. (2021), Soloperto et al. (2019a) for the case of systems subject to parametric uncertainty. If the system is subject to non-parametric uncertainty, then *kinky inference* is a learning method that is able to learn an unknown nonlinear system, while providing uncertainty sets that robustly bound the mismatch between the obtained estimate and the actual system, Calliess (2014). In Calliess et al. (2020), the authors studied the theoretical properties of kinky inference in the context of regression for simple adaptive control problems. Kinky inference has been successfully combined with model predictive control; for example, in Limon et al. (2017), a learning-based MPC scheme uses a kinky inference to model an unknown system while providing safety guarantees, and input-to-state stability. In Manzano et al. (2019), the authors propose a smoothed version of the method that is more suitable for real-time control thanks to its lower computation effort. In this case, only soft constraints on the outputs are imposed. The approach is subsequently improved in Manzano et al. (2022), where a kinky inference-based predictor that is guaranteed to be Lipschitz continuous is considered. Even though the authors show that the origin is input-to-state stable, convergence is not shown.

Contribution In this paper, we perform a stability analysis of a robust adaptive MPC scheme where the model is learned through a kinky inference. In particular, in addition to input-to-state stability, we show that standard robust adaptive MPC schemes can successfully ensure convergence of the closed-loop to the origin, despite the presence of model uncertainty. This theoretical result is obtained by only employing a tracking cost function, without the need to enforce any kind of excitation in the system, i.e., in a passive-learning fashion.

Outline Section 2 introduces the unknown system under consideration, shows how kinky inference can be used for modeling, and discusses the MPC scheme used for control. Section 3 demonstrates that the closed-loop dynamics are input-to-state stable, and that, in addition, the closed-loop system asymptotically converges to the origin. Simulations are presented in Section 4, while Section 5 concludes the paper.

Notation We use $\mathbb{R}_{\geq 0}$ and $\mathbb{R}_{>0}$ to represent the sets of non-negative and positive real numbers, respectively. A function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a class $\mathcal{K}_\infty$ function, i.e., $\alpha \in \mathcal{K}_\infty$ if α is strictly increasing, $\alpha(0) = 0$, and $\lim_{t\to\infty} \alpha(t) = \infty$. $\|\cdot\|$ denotes the Euclidean norm. $\mathbf{1}$ is a vector of appropriate dimension where each entry is equal to 1.

2. Problem Setup

System description: Consider the nonlinear, time-invariant system described, for each time $t \in \mathbb{N}$, as follows

$$x_{t+1} = f(x_t, u_t), \tag{1}$$

where $x_t \in \mathbb{R}^n$ and $u_t \in \mathbb{R}^m$ denote the state and the input of the system at time t, respectively. The state x_t is assumed to be fully available for measurement at each $t \in \mathbb{N}$. The system is subject to the following state and input constraints

$$(x_t, u_t) \in \mathcal{X} \times \mathcal{U} =: \mathcal{Z}, \tag{2}$$

for each $t \in \mathbb{N}$, where $\mathcal{X}$ and $\mathcal{U}$ are known convex and compact sets that contain the origin in their interior. To simplify the notation, we set $z := (x, u)$ when referring to a generic state-input pair, and $z_t := (x_t, u_t)$ for the state-input pair at time t. The function $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is unknown, but satisfies the following Assumption.

Assumption 1 *The function $f : \mathcal{Z} \to \mathbb{R}^n$ is Hölder continuous in all of its arguments, i.e., there exist known constants $q \in \mathbb{R}_{>0}$ and $\lambda \in \mathbb{R}$, with $0 < \lambda \leq 1$, such that*

$$\|f(z_1) - f(z_2)\| \leq q\|z_1 - z_2\|^\lambda,$$

holds for all $z_1, z_2 \in \mathcal{Z}$. Moreover, the origin is an equilibrium point for f, i.e., $f(0, 0) = 0$.

Note that the case where $\lambda > 1$ implies that the underlying function is constant on $\mathcal{Z}$ (see Proposition 1.1.16, Fiorenza (2017)). Estimating the Hölder constants q and λ from data can be done e.g., from first principles following the techniques outlined in Section 2.5 of Calliess (2014) or in Huang et al. (2023) for the Lipschitz case.

Kinky inference: *Kinky inference* is a learning technique that can be applied to Hölder continuous functions, Calliess (2014). We use kinky inference to construct a nominal model $f_t : \mathcal{Z} \to \mathbb{R}^n$ and an uncertainty function $\mathbb{W}_t : \mathcal{Z} \to 2^{\mathbb{R}^n}$. For each $t \in \mathbb{N}_{>0}$, we define the *data set* $\mathcal{D}_t \subset \mathbb{R}^{n+m}$ iteratively by $\mathcal{D}_t := \mathcal{D}_{t-1} \cup \{z_{t-1}\}$, with $\mathcal{D}_0 := \{0\}$. To construct the uncertainty functions $\mathbb{W}_t$, we define the confidence bounds $w_t^{\max}, w_t^{\min} : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$:

$$w_t^{\max}(z) := \min_{y \in \mathcal{D}_t} f(y) + \mathbf{1}q\|z - y\|^\lambda, \qquad w_t^{\min}(z) := \max_{y \in \mathcal{D}_t} f(y) - \mathbf{1}q\|z - y\|^\lambda. \tag{3}$$

Since the system f satisfies Assumption 1, it is possible to show that

$$w_t^{\min}(z) \leq f(z) \leq w_t^{\max}(z), \quad \forall z \in \mathcal{Z}, \tag{4}$$

where the inequalities are meant element-wise.

Assumption 2 *At each time $t \in \mathbb{N}$, the nominal model f_t is chosen as a Hölder continuous function satisfying the following conditions*

$$w_t^{\min}(z) \leq f_t(z) \leq w_t^{\max}(z), \quad \forall z \in \mathcal{Z}, \tag{5a}$$

$$\|f_t(z) - f_t(y)\| \leq \alpha_1(\|z - y\|), \quad \forall z, y \in \mathcal{Z}, \quad \forall t \in \mathbb{N}, \tag{5b}$$

$$\|f_t(z) - f_{t+1}(z)\| \leq \alpha_2(\|w_t^{\max}(z_t) - w_t^{\min}(z_t)\|), \quad \forall z \in \mathcal{Z}, \ \forall t \in \mathbb{N}. \tag{5c}$$

with $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$.

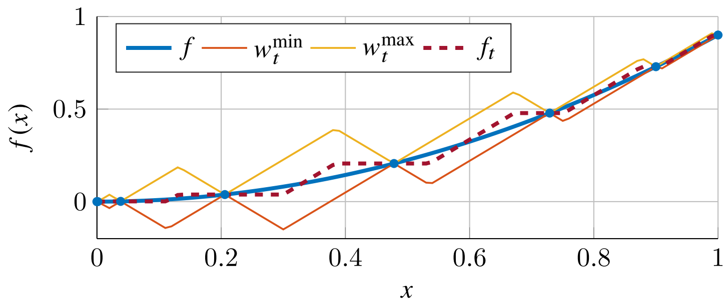

Figure 1: Kinky inference bounds for $\lambda = 1$ (thin continuous lines) and mean function (dashed line, in red) compared against the true function (thick and continuous line, in blue) for a one-dimensional system.

Assumption 2 is satisfied if the nominal model f_t is chosen as the average between the two confidence bounds, as depicted in Figure 1 for the case of a one-dimensional system. Note that condition (5c) implies that for any $z \in \mathcal{Z}$, the model update is bounded above by a function α_2 that solely depends on the most recently visited state-input pair z_t, and not the considered state z.

Lemma 1 *Condition* (5c) *can be equivalently re-stated as follows*

$$\|f_t(z) - f_{t+1}(z)\| \leq \alpha_2(\max_{y \in \mathcal{Z}} \|w_t^{\max}(y) - w_t^{\min}(y) - (w_{t+1}^{\max}(y) - w_{t+1}^{\min}(y))\|), \quad \forall z \in \mathcal{Z}, \ \forall t \in \mathbb{N}.$$

Proof In the following, we focus on the more elaborated case where $w_t^{\max}(z) \neq w_{t+1}^{\max}(z)$ and $w_t^{\min}(z) \neq w_{t+1}^{\min}(z)$. The alternative case can be trivially shown by following similar steps.

$$\begin{aligned}
&\|w_t^{\max}(z) - w_t^{\min}(z) - (w_{t+1}^{\max}(z) - w_{t+1}^{\min}(z))\| \\
\overset{(3)}{=} &\| \min_{y \in \mathcal{D}_t} [f(y) + \mathbf{1}q\|z - y\|^\lambda] - \max_{y \in \mathcal{D}_t} [f(y) - \mathbf{1}q\|z - y\|^\lambda] \\
&- \min\{w_t^{\max}(z), f(z_t) + \mathbf{1}q\|z - z_t\|^\lambda\} + \max\{w_t^{\min}(z), f(z_t) - \mathbf{1}q\|z - z_t\|^\lambda\}\| \\
\overset{(*)}{\leq} &\| \min_{y \in \mathcal{D}_t} [f(y) + \mathbf{1}q\|z_t - y\|^\lambda + \mathbf{1}q\|z - z_t\|^\lambda] - \max_{y \in \mathcal{D}_t} [f(y) - \mathbf{1}q\|z_t - y\|^\lambda - \mathbf{1}q\|z - z_t\|^\lambda] \\
&- f(z_t) - \mathbf{1}q\|z - z_t\|^\lambda + f(z_t) - \mathbf{1}q\|z - z_t\|^\lambda\| \\
= &\| \min_{y \in \mathcal{D}_t} [f(y) + \mathbf{1}q\|z_t - y\|^\lambda] - \max_{y \in \mathcal{D}_t} [f(y) - \mathbf{1}q\|z_t - y\|^\lambda] + 2 \cdot \mathbf{1}q\|z - z_t\|^\lambda - 2 \cdot \mathbf{1}q\|z - z_t\|^\lambda\| \\
\overset{(3)}{=} &\|w_t^{\max}(z_t) - w_t^{\min}(z_t)\|.
\end{aligned}$$

where in $(*)$ we use the subadditivity of the function $\|\cdot\|^\lambda$ and that $w_t^{\max}(z) \neq w_{t+1}^{\max}(z)$ and $w_t^{\min}(z) \neq w_{t+1}^{\min}(z)$. Since the bound above holds for any $z \in \mathcal{Z}$, then we have that

$$\max_{y \in \mathcal{Z}} \|w_t^{\max}(y) - w_t^{\min}(y) - (w_{t+1}^{\max}(y) - w_{t+1}^{\min}(y))\|) = \|w_t^{\max}(z_t) - w_t^{\min}(z_t)\|,$$

which concludes the proof. ■

Based on the confidence bounds, the uncertainty function $\mathbb{W}_t$ is chosen as follows

$$\mathbb{W}_t(z) := \left\{ w \in \mathbb{R}^n : w_t^{\min}(z) \leq w + f_t(z) \leq w_t^{\max}(z) \right\}. \tag{6}$$

In Calliess (2014), it is shown that

$$f_{t+1}(x,u) \oplus \mathbb{W}_{t+1}(x,u) \subseteq f_t(x,u) \oplus \mathbb{W}_t(x,u), \ \forall (x,u) \in \mathcal{Z}, \tag{7a}$$

$$\mathbb{W}_t(z) = \{0\}\,,\ \forall z \in \mathcal{D}_t,\ t \in \mathbb{N}, \tag{7b}$$

$$f(x,u) - f_t(x,u) \in \mathbb{W}_t(x,u),\ \forall (x,u) \in \mathcal{Z}, t \in \mathbb{N}. \tag{7c}$$

Note in particular that (7b) implies that the uncertainty is zero at all data-points that have been observed in the past.

Model Predictive Control scheme: Our goal is to steer system (1) from some initial condition $x_0 \in \mathcal{X}$ to the origin by choosing an appropriate input sequence u_t, for each $t \in \mathbb{N}$. To achieve this, we formulate a model predictive control scheme that employs the nominal dynamics f_t to predict the future evolution of the system, and leverages the uncertainty bounds $\mathbb{W}_t$ to guarantee robust constraint satisfaction at each time-step.

Consider the following finite horizon *cost function* $V : \mathbb{R}^{n\times N} \times \mathbb{R}^{m\times(N-1)} \to \mathbb{R}_{\geq 0}$ defined as

$$V(x_{\cdot|t}, u_{\cdot|t}) = V_f(x_{N|t}) + \sum_{k=0}^{N-1} \ell(x_{k|t}, u_{k|t}), \tag{8}$$

where $\ell : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}_{\geq 0}$ is the *stage cost*, while $V_f : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is the *terminal cost*, and $N \in \mathbb{N}_{>0}$ is the prediction horizon. The stage cost ℓ is chosen to satisfy the following.

Assumption 3 *The stage cost ℓ is continuous, and satisfies*

$$\alpha_3(\|x\|) \leq \inf_{u\in\mathcal{U}} \ell(x,u) \leq \alpha_4(\|x\|),$$

for some $\alpha_3, \alpha_4 \in \mathcal{K}_\infty$ and for all $x \in \mathbb{R}^n$.

Assumption 3 is satisfied if $\ell(x,u) := x^\top Q x + u^\top R u$, where $Q, R \succ 0$ are matrices of appropriate dimension, Soloperto et al. (2022a).

Assumption 4 *There exist a terminal cost $V_f : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$, terminal controller $\kappa_f : \mathbb{R}^n \to \mathbb{R}^m$, terminal region $\mathcal{X}_f \subseteq \mathcal{X}$, and functions $\alpha_5, \alpha_6 \in \mathcal{K}_\infty$, such that, for all $x \in \mathcal{X}_f$ and for all $w \in \mathbb{W}_0(x, \kappa_f(x))$, it holds that*

$$f_0(x, \kappa_f(x)) + w \subseteq \mathcal{X}_f,\ \ (x, \kappa_f(x)) \in \mathcal{Z}, \tag{9a}$$

$$\alpha_5(\|x\|) \leq V_f(x) \leq \alpha_6(\|x\|), \tag{9b}$$

$$V_f(f_0(x, \kappa_f(x)) + w) - V_f(x) \leq -\ell(x, \kappa_f(x)), \tag{9c}$$

Assumption 4 is satisfied if the origin is exponentially stable with a common Lyapunov function, Chen and Allgöwer (1998).

Assumption 5 *Given a nominal model f_t, an uncertainty function $\mathbb{W}_t$, and a state and input pair $(x,u) \in \mathcal{Z}$, it is possible to construct tubes $\mathcal{X}(x,u,t) \subseteq \mathcal{X}$ such that*

$$f_t(x,u) \oplus \mathbb{W}_t(x,u) \in \mathcal{X}(x,u,t), \quad \forall t \in \mathbb{N},$$

$$f_{t+1}(x,u) \oplus \mathbb{W}_{t+1} \subseteq f_t(x,u) \oplus \mathbb{W}_t \Rightarrow \mathcal{X}(x,u,t+1) \subseteq \mathcal{X}(x,u,t).$$

Assumption 5 is common in robust MPC approaches that only consider the initial knowledge of the system, Köhler et al. (2020). In our case, satisfying it becomes "easier" over time, due to the monotonicity in uncertainty implied by (7a).

Given a nominal model f_t, a description of the uncertainty $\mathbb{W}_t$, a cost function V, and an appropriately designed terminal set $\mathcal{X}_f$, we consider the following finite-horizon optimal control problem

$$\begin{aligned}
\underset{\mathcal{X}_{\cdot|t}, x_{\cdot|t}, u_{\cdot|t}}{\text{minimize}} \quad & V(x_{\cdot|t}, u_{\cdot|t}) \\
\text{subject to} \quad & \left.\begin{aligned} & x_{k+1|t} = f_t(x_{k|t}, u_{k|t}), \\ & x_{0|t} = x_t, \end{aligned}\right\} && \text{(10a)} \\
& \left.\begin{aligned} & f_t(\bar{x}_{k|t}, u_{k|t}) + w_{k|t} \in \mathcal{X}_{k+1|t}, \\ & (\bar{x}_{k|t}, u_{k|t}) \in \mathcal{Z}, \\ & \forall w_{k|t} \in \mathbb{W}_t(\bar{x}_{k|t}, u_{k|t}), \\ & \forall \bar{x}_{k|t} \in \mathcal{X}_{k|t}, \\ & k = 0, \ldots, N-1, \end{aligned}\right\} && \text{(10b)} \\
& x_t \in \mathcal{X}_{0|t}, \quad \mathcal{X}_{N|t} \subseteq \mathcal{X}_f, && \text{(10c)}
\end{aligned}$$

where $\mathcal{X}_{\cdot|t}$ are appropriately constructed tubes that satisfy Assumption 5. The optimization problem (10) is a general description of a robust MPC scheme, and is not implementable without a proper construction of the sets $\mathcal{X}_{\cdot|t}$. In Soloperto et al. (2019b), it is shown how (10) can describe several approaches, including MPC schemes for linear systems subject to bounded additive disturbance, Chisci et al. (2001) and Mayne et al. (2005), and nonlinear system subject to parametric uncertainties, Köhler et al. (2021).

We denote the optimizers of the Problem (10) by $(\mathcal{X}^*_{\cdot|t}, x^*_{\cdot|t}, u^*_{\cdot|t}) = (\mathcal{X}^*_{\cdot|t}, z^*_{\cdot|t})$, and the optimal value by

$$V_t^*(x_t) := V_f(x^*_{N|t}) + \sum_{k=0}^{N-1} \ell(x^*_{k|t}, u^*_{k|t}).$$

The equality constraints (10a) produce a nominal trajectory $x_{\cdot|t}$, obtained by propagating the initial condition x_t through the nominal model f_t with the input $u_{\cdot|t}$. The constraints in (10b) are designed to ensure that the closed-loop trajectory of the system satisfies the constraints $(x_t, u_t) \in \mathcal{Z}$. This is achieved thanks to the introduction of the tubes $\mathcal{X}_{\cdot|t} \subset \mathbb{R}^n$, which are guaranteed to contain the true state of the system at any time step if the inputs $u_{\cdot|t}$ were to be applied.

The closed-loop system can be obtained by combining the optimal control input $u^*_{0|t}$ with the dynamics (1) using the receding-horizon paradigm:

$$x_{t+1} = f(x_t, u_t), \;\; u_t = u^*_{0|t}. \tag{11}$$

3. Theoretical analysis

In this section, we show that the closed-loop system converges to the origin despite the presence of model uncertainty. Let

$$h_t(z) := \|w_t^{\max}(z) - w_t^{\min}(z)\|. \tag{12}$$

The function h_t can be interpreted as the level of uncertainty about the system $f(z)$ at a given time t, where a small value of $h_t(z)$ implies accurate knowledge of $f(z)$. Based on (3) and (7), we have that $h_t(z) = 0$ for all the previously visited points $z \in \mathcal{D}_t$.

We define the *uncertainty size* $C_t \in \mathbb{R}_{\geq 0}$ across the entire space $\mathcal{Z}$ as

$$C_t := \int_{z\in\mathcal{Z}} h_t(z)dz. \tag{13}$$

According to (7), it is easy to verify that $0 \leq C_{t+1} \leq C_t < \infty$ for all $t \in \mathbb{N}$. The upper-bound is ensured since $\mathcal{Z}$ is compact and $h_0(z) < \infty$ for all $z \in \mathcal{Z}$.

Lemma 2 *Under Assumption 1, there exists a function $\alpha_7 \in \mathcal{K}_\infty$ such that $C_{t+1} - C_t \leq -\alpha_7(h_t(x_t, u_t))$ for all $t \in \mathbb{N}$.*

Proof We begin by applying Lemma 3 of Soloperto et al. (2022b) to the function $h_t - h_{t+1}$. To this end, consider that $h_t - h_{t+1}$ is continuous in z (since both h_t and h_{t+1} are) and, in particular, it is uniformly continuous in z when restricted to $\mathcal{Z}$. As a result, there exists some $\delta_1 \in \mathcal{K}_\infty$ such that

$$|h_t(z) - h_{t+1}(z) - [h_t(y) - h_{t+1}(y)]| \leq \delta_1(\|z - y\|), \quad \forall z, y \in \mathcal{Z},$$

and we can therefore apply Lemma 3 of Soloperto et al. (2022b), which yields

$$h_t(\bar{z}) - h_{t+1}(\bar{z}) \leq \delta_2\left(\int_{z\in\mathcal{Z}} [h_t(z) - h_{t+1}(z)]dz\right),$$

for some $\delta_2 \in \mathcal{K}_\infty$ and for all $\bar{z} \in \mathcal{Z}$. The inequality holds if we apply the function $\delta_2^{-1} \in \mathcal{K}_\infty$ on both sides of the inequality (recall that the inverse of a $\mathcal{K}_\infty$ function exists and is itself $\mathcal{K}_\infty$, as stated in page 4, Kellett (2014)), obtaining for all $\bar{z} \in \mathcal{Z}$

$$\delta_2^{-1}(h_t(\bar{z}) - h_{t+1}(\bar{z})) \leq \int_{z\in\mathcal{Z}} [h_t(z) - h_{t+1}(z)]dz.$$

Choosing $\bar{z} = z_t$, we have $h_{t+1}(z_t) = 0$ from (7b), and therefore

$$C_{t+1} - C_t = \int_{z\in\mathcal{Z}} h_{t+1}(z)dz - \int_{z\in\mathcal{Z}} h_t(z)dz \leq -\delta_2^{-1}(h_t(z_t)) =: -\alpha_7(h_t(z_t)).$$

■

Theorem 3 *Let Assumptions 1, 3, 4, and 5 hold, and assume that the MPC problem* (10) *is feasible at time $t = 0$. Then, the MPC scheme* (10) *is feasible for all time $t \in \mathbb{N}_+$, and the closed-loop system* (11) *asymptotically converges to the origin while satisfying the state and input constraints* (2).

Proof The proof of Theorem 3 is divided into three parts: in part a) we start by showing recursive feasibility of (10), in part b) we then show that the origin is input-to-state stable, i.e., it holds that

$$V_{t+1}^*(x_{t+1}) - V_t^*(x_t) \leq -\ell(x_t, u_t) + \alpha_8(h_t(z_t)), \quad \forall t \in \mathbb{N}_{>0}, \tag{14}$$

for some $\alpha_8 \in \mathcal{K}_\infty$, and finally in c) we show that $h_t(z_t) \to 0$ for $t \to \infty$. The combination of the last two points implies that the system asymptotically converges to the origin.

a) Recursive feasibility: We introduce two different state and input trajectories. One, denoted by $(\hat{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1})$, considers the case where the nominal model is not updated, i.e., $f_{t+1} = f_t$, while the other, denoted by $(\bar{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1})$, the case where the nominal model is updated. Note that the latter is the only trajectory that is guaranteed to be feasible in Problem (10) at time $t+1$ and that both trajectories share the same input $\hat{u}_{\cdot|t+1}$.

Let define $(\hat{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1})$ as

$$\begin{gathered} \hat{u}_{k|t+1} = u^*_{k+1|t}, \quad k = 0, \dots, N-2, \quad \hat{u}_{N-1|t+1} = \kappa_f(x^*_{N|t}), \\ \hat{x}_{0|t+1} = f_t(x_t, u_t), \quad \hat{x}_{k|t+1} = f_t(\hat{x}_{k-1|t+1}, \hat{u}_{k-1|t+1}), \quad k = 1, \dots, N-1. \end{gathered} \tag{15}$$

The trajectory $(\hat{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1})$ satisfies the state and input constraints in (10) at time $t+1$, provided that the initial condition is $x_{t+1} = f_t(x_t, u_t)$. Likewise, we define $(\bar{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1})$ as

$$\bar{x}_{k+1|t+1} = f_{t+1}(\bar{x}_{k|t+1}, \hat{u}_{k|t+1}), \quad k = 0, \dots, N-1, \quad \bar{x}_{0|t+1} = f_{t+1}(x_t, u_t) = x_{t+1}. \tag{16}$$

Note that $(\bar{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1})$ is obtained by propagating the actual initial condition $x_{t+1} = f(x_t, u_t)$ through the updated dynamics f_{t+1} under the input $\hat{u}_{\cdot|t+1}$, defined in (15).

Thanks to (7) and (9), we have that the trajectory $(\bar{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1})$ is feasible at time $t+1$ for Problem (10). The existence of candidate tubes is ensured based on Assumption 5.

b) Input-to-state stability: Starting from $k = 0$, we have

$$\|\bar{x}_{0|t+1} - \hat{x}_{0|t+1}\| \overset{(15),(16)}{=} \|f_{t+1}(x_t, u_t) - f_t(x_t, u_t)\| \overset{(5c)}{\leq} \gamma_0(h_t(z_t)), \tag{17}$$

where $\gamma_0 := \alpha_2$. Then, for $k = 1, \dots, N$, we can use induction to show that

$$\begin{aligned} &\|\bar{x}_{k|t+1} - \hat{x}_{k|t+1}\| \overset{(15),(16)}{=} \|f_{t+1}(\bar{x}_{k-1|t+1}, \hat{u}_{k-1|t+1}) - f_t(\hat{x}_{k-1|t+1}, \hat{u}_{k-1|t+1})\| \\ \leq\ &\|f_{t+1}(\bar{x}_{k-1|t+1}, \hat{u}_{k-1|t+1}) - f_{t+1}(\hat{x}_{k-1|t+1}, \hat{u}_{k-1|t+1})\| \\ &+ \|f_{t+1}(\hat{x}_{k-1|t+1}, \hat{u}_{k-1|t+1}) - f_t(\hat{x}_{k-1|t+1}, \hat{u}_{k-1|t+1})\| \\ \overset{(5c),(5b)}{\leq}\ &\alpha_1(\|\bar{x}_{k-1|t+1} - \hat{x}_{k-1|t+1}\|) + \alpha_2(h_t(z_t)) \overset{(17)}{=} (\alpha_1 \circ \gamma_k + \alpha_2)(h_t(z_t)) =: \gamma_{k+1}(h_t(z_t)), \end{aligned} \tag{18}$$

where for $k \geq 0$ we recursively define $\gamma_{k+1} := \alpha_1 \circ \gamma_k + \alpha_2$. Since class $\mathcal{K}_\infty$ functions are closed under composition and summation, Kellett (2014), we have $\gamma_{k+1} \in \mathcal{K}_\infty$. Therefore, from (18), we have that there always exists some $\gamma_k \in \mathcal{K}_\infty$ such that for $k = 0, 1, \dots, N$ and for all $t \in \mathbb{N}$

$$\|\bar{x}_{k|t+1} - \hat{x}_{t|k+1}\| \leq \gamma_k(h_t(z_t)). \tag{19}$$

Next, since the stage cost ℓ is continuous and $\mathcal{X}$ is compact, there exists some $\alpha_9 \in \mathcal{K}_\infty$ such that

$$\|\ell(x,u) - \ell(y,u)\| \leq \alpha_9(\|x - y\|), \tag{20}$$

for all $x, y \in \mathcal{X}$, Limon et al. (2009), Lemma 1. We now upper-bound the following difference

$$\sum_{k=0}^{N-2} \ell(\bar{x}_{k|t+1}, \hat{u}_{k|t+1}) - \ell(\hat{x}_{k|t+1}, \hat{u}_{k|t+1}) \overset{(19),(20)}{\leq} \sum_{k=0}^{N-2} (\alpha_9 \circ \gamma_k)(h_t(z_t)) =: \alpha_{10}(h_t(z_t)), \tag{21}$$

where $\alpha_{10} := \sum_{k=0}^{N-2} \alpha_9 \circ \gamma_k \in \mathcal{K}_\infty$. Since by optimality we have that $V^*_{t+1}(x_{t+1}) \leq V(\bar{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1})$, we can make use of (21) to obtain

$$\begin{aligned} &V^*_{t+1}(x_{t+1}) - V^*_t(x_t) \leq V(\bar{x}_{\cdot|t+1}, \hat{u}_{\cdot|t+1}) - V^*_t(x_t) \\ &\overset{(21)}{\leq} -\ell(x_t, u_t) + \alpha_{10}(h_t(z_t)) + \ell(\bar{x}_{N-1|t+1}, \hat{u}_{N-1|t+1}) - V_f(x^*_{N|t}) + V_f(\bar{x}_{N|t+1}) \end{aligned} \tag{22}$$

Since V_f is continuous and $\mathcal{X}$ is compact, there exists a function $\alpha_{11} \in \mathcal{K}_\infty$ such that

$$\|V_f(x) - V_f(y)\| \leq \alpha_{11}(\|x - y\|), \tag{23}$$

for all $x, y \in \mathcal{X}$. In the following, we upper-bound the sum of the last three terms in (22):

$$\begin{aligned} &\ell(\bar{x}_{N-1|t+1}, \hat{u}_{N-1|t+1}) - V_f(x^*_{N|t}) + V_f(\bar{x}_{N|t+1}) \\ &= \ell(\bar{x}_{N-1|t+1}, \hat{u}_{N-1|t+1}) - \ell(\hat{x}_{N-1|t+1}, \hat{u}_{N-1|t+1}) + \ell(\hat{x}_{N-1|t+1}, \hat{u}_{N-1|t+1}) - V_f(x^*_{N|t}) \\ &\quad + V_f(\hat{x}_{N|t+1}) + V_f(\bar{x}_{N|t+1}) - V_f(\hat{x}_{N|t+1}) \\ &\overset{(20),(23)}{\leq} \alpha_9(\|\hat{x}_{N-1|t+1} - \bar{x}_{N-1|t+1}\|) + \alpha_{11}(\|\bar{x}_{N|t+1} - \hat{x}_{N|t+1}\|) + \ell(x^*_{N|t}, \kappa_f(x^*_{N|t})) \\ &\quad - V_f(x^*_{N|t}) + V_f(f_t(x^*_{N|t}, \kappa_f(x^*_{N|t}))) \\ &\overset{(9c),(19)}{\leq} (\alpha_9 \circ \gamma_{N-1} + \alpha_{11} \circ \gamma_N)(h_t(z_t)) =: \alpha_{12}(h_t(z_t)), \end{aligned} \tag{24}$$

where we defined $\alpha_{12} := \alpha_9 \circ \gamma_{N-1} + \alpha_{11} \circ \gamma_N \in \mathcal{K}_\infty$. Combining (22) and (24), and choosing $\alpha_8 := \alpha_{10} + \alpha_{12} \in \mathcal{K}_\infty$, we obtain (14). Since V^*_t is an ISS Lyapunov function, the closed-loop system is input-to-state stable (compare Definition 1 and Theorem 1 in Li et al. (2018)).

c) Convergence: Using Lemma 2, it holds that

$$\sum_{k=0}^{t} C_{k+1} - C_k \leq -\sum_{k=0}^{t} \alpha_7(h_k(z_k)) \Rightarrow \sum_{k=0}^{t} \alpha_7(h_k(z_k)) \leq C_0 - C_{t+1} \leq C_0.$$

Since the inequality above holds for all $t \in \mathbb{N}$, by taking the limit we have that

$$\lim_{t\to\infty} \sum_{k=0}^{t} \alpha_7(h_k(z_k)) \leq C_0 \Rightarrow \lim_{t\to\infty} \alpha_7(h_t(z_t)) = 0 \iff \lim_{t\to\infty} h_t(z_t) = 0.$$

Using the fact that the closed-loop system is ISS, and since the term h_t converges to zero for $t \to \infty$, we can leverage the converging-input converging-state property of ISS systems, as stated in Page 3, Jiang and Wang (2001), to conclude that $z_t \to 0$ as $t \to \infty$. ■

4. Numerical example

We consider a two-dimensional system described by the following dynamics

$$x_{t+1} = \underbrace{\begin{bmatrix} 1 & 0.4 \\ 0 & 0.56 + 0.1x_t^1 \end{bmatrix} x_t + \begin{bmatrix} 0 \\ 0.4 \end{bmatrix} u_t}_{:=\hat{f}(x_t, u_t)} + \begin{bmatrix} 0 \\ 0.9x_t^1 \exp(-x_t^1) \end{bmatrix},$$

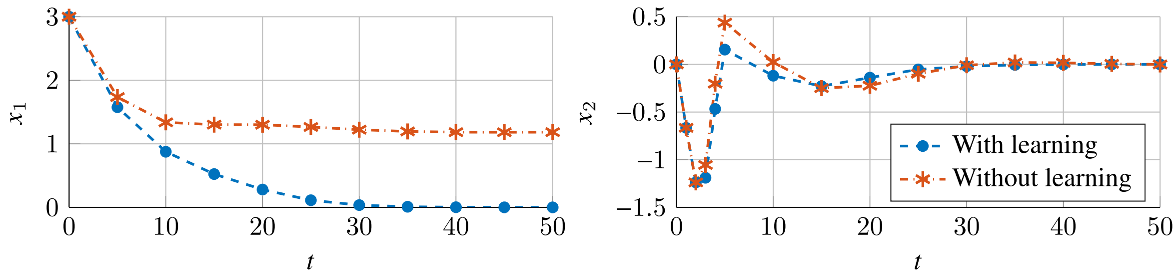

Figure 2: State evolution.

where $x_t = (x_t^1, x_t^2) \in \mathbb{R}^2$ and $u_t \in \mathbb{R}$. We initialize the system with $x_0 = (3, 0)$. The input constraints are defined as $\mathcal{U} = \{u \in \mathbb{R} : |u| \leq 2\}$. The goal is to reach the origin while minimizing the stage cost $\ell(x, u) = x^T Q x + u^T R u$, where $Q = \mathrm{diag}(1, 1)$, and $R = 1$.

For simplicity, we do not consider state constraints or a terminal region. It is possible to prove, using Lemma 2, Proposition 1 in Limon et al. (2009), and Theorem 4 in Boccia et al. (2014), that the system is ISS for a sufficiently large control horizon.

We considered the case where the function $\hat{f}$ is known, while we use kinky inference to learn the unknown function $x^1 \mapsto g(x^1) := 0.9x^1 \exp(-x^1)$. Our nominal model f_t is therefore defined as $f_t(x, u) = \hat{f}(x_t, u_t) + g_t(x^1)$, where $g_0(x^1) = 0$, for all $x^1 \in \mathbb{R}$ and g_t is given by the mean of the upper and lower bounds of the Kinky inference. We over-approximate the Lipschitz constant of g in the operating range of states by choosing $q = 1.5$, $\lambda = 1$.

The MPC problem described here requires solving a nonlinear optimization problem including nonsmooth equality constraints. This is generally challenging. Future work should focus on adapting the theoretical analysis to the smoothed Kinky inference introduced in Manzano et al. (2019), where the nonsmoothness is not present.

We compare the performance of our scheme against an MPC scheme where $f_t(x, u)$ is chosen as $f_t(x, u) = \hat{f}(x, u)$. In Figure 2, it is possible to see that the lack of learning prevents the second scheme from converging to the origin. On the contrary, our adaptive MPC scheme takes advantage of new data to successfully regulate the system to the origin.

5. Conclusion

We considered the problem of regulating an unknown nonlinear system using an adaptive model predictive control scheme. Specifically, we considered systems that are continuous and deterministic, while the online-updated nominal models are learned through a kinky inference. We showed that, for the considered class of systems, a standard adaptive MPC scheme, i.e., only with a standard tracking cost function, the system converges, in closed-loop, to the origin. We illustrated our findings in a simulation example.

Acknowledgments

The work of Riccardo Zuliani is supported by the Swiss National Science Foundation under NCCR Automation (grant agreement 51NF40_180545). The work of Raffaele Soloperto is supported by the European Research Council under the H2020 Advanced Grant no. 787845 (OCAL).

References

Andrea Boccia, Lars Grüne, and Karl Worthmann. Stability and feasibility of state constrained mpc without stabilizing terminal constraints. *Systems & control letters*, 72:14–21, 2014.

Eric Bradford, Lars Imsland, Dongda Zhang, and Ehecatl Antonio del Rio Chanona. Stochastic data-driven model predictive control using gaussian processes. *Computers & Chemical Engineering*, 139:106844, 2020.

Jan-Peter Calliess. *Conservative decision-making and inference in uncertain dynamical systems*. PhD thesis, University of Oxford Oxford, 2014.

Jan-Peter Calliess, Stephen J Roberts, Carl Edward Rasmussen, and Jan Maciejowski. Lazily adapted constant kinky inference for nonparametric regression and model-reference adaptive control. *Automatica*, 122:109216, 2020.

Hong Chen and Frank Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica*, 34(10):1205–1217, 1998.

Luigi Chisci, J Anthony Rossiter, and Giovanni Zappa. Systems with persistent disturbances: predictive control with restricted constraints. *Automatica*, 37(7):1019–1028, 2001.

Mark L Darby and Michael Nikolaou. Mpc: Current practice and challenges. *Control Engineering Practice*, 20(4):328–342, 2012.

Renato Fiorenza. *Hölder and locally Hölder Continuous Functions, and Open Sets of Class C^ k, C^{k, lambda}*. Birkhäuser, 2017.

Michael G Forbes, Rohit S Patwardhan, Hamza Hamadah, and R Bhushan Gopaluni. Model predictive control in industry: Challenges and opportunities. *IFAC-PapersOnLine*, 48(8):531–538, 2015.

Lukas Hewing, Juraj Kabzan, and Melanie N Zeilinger. Cautious model predictive control using gaussian process regression. *IEEE Transactions on Control Systems Technology*, 28(6):2736–2743, 2019.

Julien Walden Huang, Stephen J Roberts, and Jan-Peter Calliess. On the sample complexity of lipschitz constant estimation. *Transactions on Machine Learning Research*, 2023.

Zhong-Ping Jiang and Yuan Wang. Input-to-state stability for discrete-time nonlinear systems. *Automatica*, 37(6):857–869, 2001.

Christopher M Kellett. A compendium of comparison function results. *Mathematics of Control, Signals, and Systems*, 26:339–374, 2014.

Johannes Köhler, Raffaele Soloperto, Matthias A Müller, and Frank Allgöwer. A computationally efficient robust model predictive control framework for uncertain nonlinear systems. *IEEE Transactions on Automatic Control*, 66(2):794–801, 2020.

Johannes Köhler, Peter Kötting, Raffaele Soloperto, Frank Allgöwer, and Matthias A Müller. A robust adaptive model predictive control framework for nonlinear uncertain systems. *International Journal of Robust and Nonlinear Control*, 31(18):8725–8749, 2021.

Huijuan Li, Anping Liu, and Linli Zhang. Input-to-state stability of time-varying nonlinear discrete-time systems via indefinite difference lyapunov functions. *ISA transactions*, 77:71–76, 2018.

Daniel Limon, Teodoro Alamo, Davide M Raimondo, D Muñoz De La Peña, José Manuel Bravo, Antonio Ferramosca, and Eduardo F Camacho. Input-to-state stability: a unifying framework for robust model predictive control. *Nonlinear Model Predictive Control: Towards New Challenging Applications*, pages 1–26, 2009.

Daniel Limon, J Calliess, and Jan Marian Maciejowski. Learning-based nonlinear model predictive control. *IFAC-PapersOnLine*, 50(1):7769–7776, 2017.

Matthias Lorenzen, Mark Cannon, and Frank Allgöwer. Robust mpc with recursive model update. *Automatica*, 103:461–471, 2019.

Xiaonan Lu and Mark Cannon. Robust adaptive tube model predictive control. In *2019 American Control Conference (ACC)*, pages 3695–3701. IEEE, 2019.

Jose Maria Manzano, Daniel Limon, David Munoz de la Pena, and Jan Peter Calliess. Output feedback mpc based on smoothed projected kinky inference. *IET Control Theory & Applications*, 13(6):795–805, 2019.

Jose Maria Manzano, David Muñoz de la Peña, and Daniel Limon. Input-to-state stable predictive control based on continuous projected kinky inference. *International Journal of Robust and Nonlinear Control*, 2022.

David Q Mayne, María M Seron, and SV Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005.

Ali Mesbah. Stochastic model predictive control with active uncertainty learning: A survey on dual control. *Annual Reviews in Control*, 45:107–117, 2018.

Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.

James Blake Rawlings, David Q Mayne, and Moritz Diehl. *Model predictive control: theory, computation, and design*, volume 2. Nob Hill Publishing Madison, WI, 2017.

Raffaele Soloperto, Johannes Köhler, Matthias A Müller, and Frank Allgöwer. Dual adaptive mpc for output tracking of linear systems. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 1377–1382. IEEE, 2019a.

Raffaele Soloperto, Johannes Köhler, Frank Allgöwer, and Matthias A. Müller. Collision avoidance for uncertain nonlinear systems with moving obstacles using robust model predictive control. In *2019 18th European Control Conference (ECC)*, pages 811–817, 2019b. doi: 10.23919/ECC.2019.8796049.

Raffaele Soloperto, Johannes Köhler, and Frank Allgöwer. A nonlinear mpc scheme for output tracking without terminal ingredients. *IEEE Transactions on Automatic Control*, 68(4):2368–2375, 2022a.

Raffaele Soloperto, Matthias A Müller, and Frank Allgöwer. Guaranteed closed-loop learning in model predictive control. *IEEE Transactions on Automatic Control*, 68(2):991–1006, 2022b.

Marko Tanaskovic, Lorenzo Fagiano, and Vojislav Gligorovski. Adaptive model predictive control for linear time varying mimo systems. *Automatica*, 105:237–245, 2019.