

Learning-enhanced Nonlinear Model Predictive Control using Knowledge-based Neural Ordinary Differential Equations and Deep Ensembles

Kong Yao Chee

M. Ani Hsieh

Nikolai Matni

University of Pennsylvania, Philadelphia, PA 19104

CKONGYAO@SEAS.UPENN.EDU

M.HSIEH@SEAS.UPENN.EDU

NMATNI@SEAS.UPENN.EDU

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

Nonlinear model predictive control (MPC) is a flexible and increasingly popular framework used to synthesize feedback control strategies that can satisfy both state and control input constraints. In this framework, an optimization problem, subjected to a set of dynamics constraints characterized by a nonlinear dynamics model, is solved at each time step. Despite its versatility, the performance of nonlinear MPC often depends on the accuracy of the dynamics model. In this work, we leverage deep learning tools, namely knowledge-based neural ordinary differential equations (KNODE) and deep ensembles, to improve the prediction accuracy of this model. In particular, we learn an ensemble of KNODE models, which we refer to as the KNODE ensemble, to obtain an accurate prediction of the true system dynamics. This learned model is then integrated into a novel learning-enhanced nonlinear MPC framework. We provide sufficient conditions that guarantees asymptotic stability of the closed-loop system and show that these conditions can be implemented in practice. We show that the KNODE ensemble provides more accurate predictions and illustrate the efficacy and closed-loop performance of the proposed nonlinear MPC framework using two case studies.

Keywords: Nonlinear model predictive control, deep learning, neural ordinary differential equations, deep ensembles

1. Introduction

Advances in nonlinear optimization algorithms and improvements in hardware computational power, nonlinear model predictive control (MPC) is proliferating across a range of autonomous systems applications, *e.g.*, [Chowdhri et al. \(2021\)](#); [Yao and Shekhar \(2021\)](#). For similar reasons, deep learning tools are also becoming prevalent, particularly in the context of learning dynamical system representations. These tools can be used to extract meaningful information of the system from data, which is then incorporated into the dynamics models for MPC. Within the MPC framework, the models are deployed to predict the system dynamics, which consequently enables accurate closed-loop performance. Despite a surge in recent literature, the choice of deep learning tools and how they can be seamlessly integrated to learn dynamics models within the MPC framework remains a relatively open problem ([Mesbah et al. \(2022\)](#)). In this work, we present a novel learning-enhanced nonlinear MPC framework that uses a suite of deep learning techniques to enhance the accuracy of the dynamics model, which improves the performance of the closed-loop system.

Related Works. Learning-based model predictive control (LMPC) is a fast-growing research area and a variety of topics have been studied in recent literature. Here, we highlight recent developments, particularly in the context of learning nonlinear dynamics models, and refer interested readers to recent reviews by [Hewing et al. \(2020\)](#) and [Mesbah et al. \(2022\)](#). There are a number of

works, *e.g.*, [Kabzan et al. \(2019\)](#), [Maddalena and Jones \(2020\)](#), [Torrente et al. \(2021\)](#), that adopt a non-parametric approach to learning dynamics models, before incorporating them into an MPC framework. However, it is well known that non-parametric models such as Gaussian processes may not be amenable to large amounts of data. Hence, a data selection strategy may be required to distill the relevant data for training and deployment. On the other hand, parametric models such as feedforward neural networks (NNs) and its variants, have also been explored within the LMPC context. [Chen et al. \(2020\)](#) use NNs to model heating, ventilation and air-conditioning systems (HVAC) within smart buildings and integrate it into an MPC scheme. [Son et al. \(2022\)](#) utilized NNs to learn the model-plant mismatch and applied it into an offset-free MPC scheme. Architectural variants such as recurrent neural networks and long short-term memory networks have also been considered in LMPC, *e.g.*, [Terzi et al. \(2021\)](#); [Wu et al. \(2021\)](#); [Saviolo et al. \(2022\)](#). These networks are structurally more complex and incur significant computational costs when integrated with a model predictive controller. In general, literature that apply ensemble methods in the context of LMPC seem to be sparser. [Kusiak et al. \(2010\)](#) use NN ensembles to model HVAC systems and use it in a particle swarm optimization problem. However, these processes operate at slower time scales and may not be suitable for fast-paced applications. A related concept is scenario-based MPC, *e.g.*, [Bernardini and Bemporad \(2009\)](#); [Schildbach et al. \(2014\)](#), where scenarios are sampled from a given distribution and used in linear models within an MPC framework. One difference in this approach is that these scenarios collected from the system are in terms of the uncertainty and not of the states and control inputs. A second difference is in the formulation of the optimization problem, where the scenarios are considered simultaneously, instead of fusing them using a selected strategy, which is more common in ensemble methods. For a recent review on ensemble methods in deep learning, we refer the readers to [Ganaie et al. \(2021\)](#).

Contributions. Knowledge-based neural ordinary differential equations (KNODE), which comprises of a nominal first-principles model and a neural ordinary differential equation (NODE) component, are used to extract dynamics models which are then utilized within a MPC framework ([Chee et al. \(2022\)](#)). In this work, we leverage the fact that KNODE models are relatively lightweight and introduce deep ensembling to enhance the learned dynamics model. We show that this provides more accurate state predictions and improves closed-loop performance. Our contributions in this work are three-fold. First, we present a novel learning-enhanced nonlinear MPC framework, leveraging concepts in deep learning, namely KNODEs and deep ensembles. These tools allow us to assimilate data into the learned model in a systematic and amenable manner, which in turn provides an accurate representation of the system dynamics and improves closed-loop performance. Second, for our proposed framework, we provide sufficient conditions that render the closed-loop system under the proposed MPC framework asymptotically stable. We show that these conditions, when applied with our proposed framework, are practical and can be implemented readily. Third, using two case studies, we demonstrate the efficacy and versatility of our proposed framework.

2. Problem Framework and Preliminaries

Consider a discrete-time, nonlinear system with the following dynamics,

$$x^+ = f(x, u), \quad (1)$$

where $x, x^+ \in \mathbb{R}^n$ are the current and successor states, $u \in \mathbb{R}^m$ is the control input to the system and it is assumed that f is twice continuously differentiable. To control the system, we consider a nonlinear MPC framework that generates a sequence of optimal control inputs, which operates in a receding horizon manner. First, a dynamics model is constructed using first principles or prior

knowledge of the system. Using this model, together with information on the state and control input constraints, we formulate the following finite horizon constrained optimal control problem,

$$\underset{U}{\text{minimize}} \quad \sum_{i=0}^{N-1} q(x_i, u_i) + p(x_N) \quad (2a)$$

$$\text{subject to} \quad x_{i+1} = \hat{f}(x_i, u_i), \quad \forall i = 0, \dots, N-1, \quad (2b)$$

$$x_i \in \mathcal{X}, \quad u_i \in \mathcal{U}, \quad \forall i = 0, \dots, N-1, \quad (2c)$$

$$x_N \in \mathcal{X}_f, \quad x_0 = x(k), \quad (2d)$$

where N is the horizon, $q(x_i, u_i)$ and $p(x_N)$ are the state and terminal state costs, and the sets \mathcal{X} , \mathcal{U} and \mathcal{X}_f denote the constraint sets for the states, control inputs and the terminal state respectively. At each time step k , using state measurements $x(k)$, we solve (2) and obtain a sequence of optimal control inputs, $U^* := \{u_0^*, \dots, u_{N-1}^*\}$. We apply the first vector in this sequence $u_0^*(x(k))$ to the system as the control action. Hence, the control law is expressed as $u(k) = u_0^*(x(k))$ and the closed-loop system is given as $x^+ = f(x, u_0^*(x)) := f_{cl}(x)$.

If the dynamics model in (2b) represents the system (1) with sufficient accuracy, we can expect the MPC framework to perform well. However, like many models, it is unlikely for the dynamics model to capture the true dynamics perfectly. To improve upon this dynamics model, we utilize deep learning tools to enhance its accuracy within the MPC framework. We first adopt a NODE formulation to represent any unknown, poorly modeled, and/or residual dynamics. During training, we combine this NODE model with a nominal model derived from prior knowledge to learn a hybrid KNODE model. This allows us to get a more accurate representation of the true dynamics in (1). Next, instead of learning a single KNODE model, we learn an ensemble of KNODE models, which we denote as a KNODE ensemble. Through experiments, we show that the KNODE ensemble provides better accuracy and generalization in terms of state predictions and better closed-loop performance in terms of trajectory tracking.

Notation. \mathbb{R} and \mathbb{R}^n denote the set of real numbers and the n -dimensional vector space of real numbers. We use $\|\cdot\|$ to denote the Euclidean and spectral norm for vectors and matrices respectively. B_r denotes a closed ball in \mathbb{R}^n centered at the origin, *i.e.*, $B_r := \{x \in \mathbb{R}^n \mid \|x\| \leq r\}$. For a matrix $A \in \mathbb{R}^{n \times n}$, $A \succeq 0$ and $A \succ 0$ imply that A is positive semi-definite and positive definite respectively.

3. KNODE Ensembles

3.1. Knowledge-based Neural Ordinary Differential Equations

For many electro-mechanical and robotic systems, it is relatively straightforward to obtain a model of the system, using first principles and prior knowledge. However, in the presence of uncertainty and modeling errors, it is challenging to assess if this derived model is sufficiently accurate, especially if it is applied within a model-based control framework such as nonlinear MPC. On the other hand, when the true system is operating in its specified environment, we can collect useful data that inform us about its underlying dynamics. The KNODE framework (Jiahao et al. (2021)) provides a systematic approach to fuse the model obtained from first principles with the collected data. In this framework, we consider the following continuous-time representation of the true system,

$$\dot{x} = f_c(x, u) = \hat{f}_c(x, u) + d_c(x, u), \quad (3)$$

where the subscript c denotes the continuous-time nature and is used to distinguish these models from their discretized counterparts. The function $\hat{f}_c(x, u)$ is a nominal model derived from prior

knowledge. The dynamics of the true system $f_c(x, u)$ differs from that of the nominal model and this difference is represented by the unknown, residual dynamics, $d_c(x, u)$. We parameterize these unknown dynamics using a NODE with parameters θ , denoted by $d_{\theta, c}(x, u)$. The model $\hat{f}_c(x, u) + d_{\theta, c}(x, u)$ is jointly known as the KNODE model. To train this model, the first step is to collect a dataset $\mathcal{D}_{train} := \{(x_i, u_i)\}_{i=1}^M$ that consists of M data samples. Each of these samples contains the state measurement x_i and the control input u_i , sampled at times $\{t_i\}_{i=1}^M$. For each data sample in \mathcal{D}_{train} and with the nominal model \hat{f}_c , we generate one-step state predictions,

$$\hat{x}_{i+1} = x_i + \int_{t_i}^{t_{i+1}} \hat{f}_c(x_i, u_i) + d_{\theta, c}(x_i, u_i) dt, \quad (4)$$

where the integral is computed with a numerical method, *e.g.*, the explicit 4th order Runge-Kutta (RK4) method. Next, we define a loss function that characterizes the mean-squared errors across the predicted and measured states,

$$\mathcal{L}(\theta) := \frac{1}{M-1} \sum_{i=2}^M \|\hat{x}_i - x_i\|^2. \quad (5)$$

To obtain a set of optimal parameters θ^* , the model is put through a backpropagation procedure, where the gradients of the loss function with respect to the parameters are computed with automatic differentiation and applied in an optimization routine, Adam (Kingma and Ba (2015)). After training, we obtain a KNODE model, $\hat{f}_c(x, u) + d_{\theta^*, c}(x, u)$. It has been shown that a single KNODE model provides predictions of higher accuracy as compared to the nominal model, as described in Chee et al. (2022) and Jiahao et al. (2022). In this work, we propose an enhancement by learning an ensemble of KNODE models, instead of a single model.

3.2. Deep Ensembles

To improve upon the prediction accuracy and to reduce the variance of the predictions, we introduce a further enhancement to the KNODE model. Inspired by deep ensembling methods such as those described by Lakshminarayanan et al. (2017) and Wilson and Izmailov (2020), we construct a deep ensemble of KNODE models. It has been shown that ensembles provide improvements across a variety of classification and regression tasks in a number of application areas, *e.g.*, see reviews in Dietterich (2000), Ganaie et al. (2021). In our proposed framework, we adopt a randomized approach where each model in the KNODE ensemble is trained with the same training dataset, with the parameters in each model initialized at random. To promote diversity in the ensemble, we allow the models to have different network architectures. The ensemble residual dynamics is obtained by computing a weighted average across the individual models, *i.e.*, $d_{ens, c} := \sum_{j=1}^L \alpha_j d_{\theta_j^*, c}$, where L is the number of models in the ensemble and $\{\alpha_i\}_{j=1}^L$ are weights for the models in the ensemble, with $\sum_j \alpha_j = 1$. This formulation can also be interpreted as a form of stacked regression (Ganaie et al. (2021)). Two variants of the KNODE ensemble are proposed. In the first variant, equal weights are assigned to each model in the ensemble. For the second variant, we compute a set of weights for the models through an optimization problem. Specifically, we formulate and solve the following problem across a hold-out dataset, $\mathcal{D}_v := \{(x_i, u_i)\}_{i=1}^{M_v}$,

$$\underset{\alpha_1, \dots, \alpha_L}{\text{minimize}} \quad \frac{1}{M_v - 1} \sum_{i=2}^{M_v} \|x_i - \sum_{j=1}^L \alpha_j \hat{x}_{i,j}\|^2 \quad (6a)$$

$$\text{subject to} \quad \hat{x}_{i+1,j} = x_i + \int_{t_i}^{t_{i+1}} \hat{f}_c(x_i, u_i) + d_{\theta_j^*, c}(x_i, u_i) dt, \quad \forall i = 1, \dots, M_v - 1, \quad (6b)$$

$$\sum_{j=1}^L \alpha_j = 1. \quad (6c)$$

where $\hat{x}_{i,j}$ denotes the state predicted by the j^{th} model in the ensemble for the i^{th} data sample in \mathcal{D}_v . Similar to (4), the integral in (6b) is computed using a numerical solver. The hold-out set \mathcal{D}_v is obtained by splitting the collected dataset into $\mathcal{D}_{\text{train}}$ and \mathcal{D}_v .

4. Learning-based Nonlinear MPC with KNODE Ensembles

We integrate the trained KNODE ensemble into a learning-enhanced nonlinear MPC framework. In particular, we replace the dynamics model in (2b) with a discretized version of the KNODE ensemble. The resulting finite horizon optimal control problem is written as

$$J^*(x(k)) = \underset{\mathcal{U}}{\text{minimize}} \quad \sum_{i=0}^{N-1} q(x_i, u_i) + p(x_N) \quad (7a)$$

$$\text{subject to} \quad x_{i+1} = \hat{f}(x_i, u_i) + d_{\text{ens}}^*(x_i, u_i), \quad \forall i = 0, \dots, N-1, \quad (7b)$$

$$x_i \in \mathcal{X}, \quad u_i \in \mathcal{U}, \quad \forall i = 0, \dots, N-1, \quad (7c)$$

$$x_N \in \mathcal{X}_f, \quad x_0 = x(k), \quad (7d)$$

where $\hat{f}(x, u) + d_{\text{ens}}^*(x, u) := \hat{f}(x, u) + \sum_j \alpha_j^* d_{\theta_j^*}(x, u)$ denotes the discretized version of the trained KNODE ensemble with optimal parameters θ^* and α_j^* , $j \in \{1, \dots, L\}$. The control law is synthesized using a similar approach as described in Section 2. Upon solving (7), we obtain a sequence of optimal control inputs, $U^* := \{u_0^*, \dots, u_{N-1}^*\}$. The first element in this optimal control sequence $u_0^*(x(k))$ is applied to the system as the control action and this proceeds in a receding horizon manner. Next, we examine the stability properties of the learning-based MPC framework by showing that there exists a practical choice of $q(x, u)$, $p(x)$ and \mathcal{X}_f that guarantees asymptotic stability for the closed-loop system.

4.1. Conditions for Asymptotic Stability

For a compact set $\mathcal{X} \subset \mathbb{R}^n$ containing the origin, a function $V : \mathbb{R}^n \rightarrow \mathbb{R}_+$ is a Lyapunov function in $\mathcal{X} \subseteq \mathbb{R}^n$ for the system $f_{\text{cl}}(x)$ if

$$\begin{aligned} V(0) &= 0, \quad V(x) > 0, \quad \forall x \in \mathcal{X} \setminus \{0\}, \\ V(f_{\text{cl}}(x)) - V(x) &\leq -\alpha(\|x\|), \quad \forall x \in \mathcal{X} \setminus \{0\}, \end{aligned} \quad (8)$$

where $\alpha : \mathbb{R}^n \rightarrow \mathbb{R}$ is a continuous, positive definite function. It is well known that the existence of a Lyapunov function implies that the origin is asymptotically stable for the system in \mathcal{X} (Borrelli et al. (2017); Chellaboina and Haddad (2008)). In this work, we show that there exists a practical choice of the stage and terminal costs $q(x, u)$, $p(x)$, and terminal state constraint set \mathcal{X}_f such that the closed-loop system, under the proposed learning-based MPC scheme, is asymptotically stable. Before describing this design choice, we first state a set of conditions, conditions (a) through (e) in the following theorem, such that if they are fulfilled, then the value function in (7), $J^*(x(k))$, is a valid Lyapunov function and consequently, asymptotic stability is achieved. After which, we relate this design choice with these conditions. The following is a standard result that guarantees asymptotic stability for the closed-loop system.

Theorem 1 (Borrelli et al. (2017); Grüne and Pannek (2017)) *Suppose that*

- a) $q(x, u)$ and $p(x)$ are continuous and positive definite functions,
- b) The set $\mathcal{X}_f \subseteq \mathcal{X}$ is closed and contains the origin in its interior,
- c) $v(x) \in \mathcal{U}$, $\forall x \in \mathcal{X}_f$,
- d) $f(x, v(x)) \in \mathcal{X}_f$, $\forall x \in \mathcal{X}_f$,
- e) $p(f(x, v(x))) - p(x) + q(x, v(x)) \leq 0$, $\forall x \in \mathcal{X}_f$,

where the functions $q(x, u)$, $p(x)$, $f(x, u)$, sets \mathcal{X} , \mathcal{U} , \mathcal{X}_f are as described in (7) and $v(x)$ is a chosen local control law. Then, the value function $J^*(x(k))$ satisfies the conditions in (8) and the origin of the closed-loop system $f_{cl}(x)$ is asymptotically stable in \mathcal{X}_0 , where \mathcal{X}_0 is the set of initial states $x(k)$ for which the optimization problem (7) is feasible.

The proof of Theorem 1 is listed in the Appendix for completeness. Quadratic cost functions, i.e., $q(x, u) := x^\top Qx + u^\top Ru$, and $p(x) := x^\top Px$, are commonly used in the formulation of (7). If Q , R and $P \succ 0$, condition (a) is satisfied. While there is some flexibility in the choice of Q and R , we show that a particular choice of terminal cost matrix P , terminal state constraint set \mathcal{X}_f , and local control law $v(x)$ ensure that conditions (b) through (e) are fulfilled. The underlying idea is to utilize the fact that the dynamics of a nonlinear system and its linearized dynamics are not too different in some neighborhood of the origin. This concept is not new and has been studied in the form of receding horizon control, synonymous to MPC, for nonlinear constrained systems (see [Chen and Allgöwer \(1998\)](#); [Magni et al. \(2001\)](#); [Kwon and Han \(2005\)](#); [Rawlings et al. \(2017\)](#)). However, there are nuances in some of these approaches, which may make them prohibitive for implementation. For instance, in [Magni et al. \(2001\)](#), the terminal cost is a function of the parameters characterizing the exponential stability of the linearized system. In [Chen and Allgöwer \(1998\)](#), continuous-time systems are considered. In this work, we consider discrete-time nonlinear systems and provide an alternative perspective of the conditions in the Lyapunov equation, differing from those in [Kwon and Han \(2005\)](#) (cf. (7.23) and (7.28) therein). In Section 5, we further show that this choice of P and \mathcal{X}_f , together with our proposed framework, can be implemented and performs well in practice.

We consider a linearization of the system in (1) about the origin. Assuming that $f(0, 0) = 0$, the linear dynamics can be written as

$$x^+ = \nabla f_x(x, u)|_{(x, u)=(0, 0)} + \nabla f_u(x, u)|_{(x, u)=(0, 0)} := Ax + Bu, \quad (9)$$

where the pair (A, B) is assumed to be stabilizable. Next, we choose a linear feedback controller, $u_l(x) := Kx$, such that the linear closed-loop system, $x^+ = Ax + Bu_l(x) = (A + BK)x := A_{cl}x$, is exponentially stable. It is important to note that while this linear gain K is chosen and required for the computation of the terminal cost matrix P , it is not required for the control law in MPC. In particular, by considering the local control law $v(x)$ in Theorem 1 to be $u_l(x)$, and together with $q(x, u)$, $p(x)$ and \mathcal{X}_f as described below, it can be shown that condition (e) in Theorem 1 is fulfilled. This is given by the following proposition.

Proposition 2 Consider the stage and terminal costs to be $q(x, u) = x^\top Qx + u^\top Ru$, $p(x) := x^\top Px$, where Q , $R \succ 0$, and P is the solution to the following Lyapunov equation,

$$A_{cl}^\top PA_{cl} + \rho(Q + K^\top RK) - P = 0, \quad (10)$$

where $\rho \in (1, \infty)$. Let the terminal constraint set \mathcal{X}_f be a γ -sublevel set of $p(x)$, i.e., $\mathcal{X}_f := \{x \in \mathbb{R}^n \mid p(x) \leq \gamma\}$, with $\gamma > 0$. Then, it holds that

$$p(f(x, Kx)) - p(x) \leq -x^\top (Q + K^\top RK)x, \quad \forall x \in \mathcal{X}_f \quad (11)$$

To prove Proposition 2, we make use of a lemma that provides an upper bound on the ratio of the norm of a twice continuously differentiable function and the norm of its argument. This lemma with its proof, together with the proof of Proposition 2, are given in the Appendix.

Corollary 3 *Given $p(x)$ and \mathcal{X}_f defined in Proposition 2, together with the descent condition in (11), conditions (b) and (d) in Theorem 1 are satisfied.*

The proof for Corollary 3 is given in the Appendix. For condition (c) in Theorem 1, it is shown in Magni et al. (2001) that there exists a $\delta_1 \in (0, \infty)$ such that for any $x \in \mathcal{X}$, $Kx \in \mathcal{U}$, for all $x \in B_{\delta_1}$. Hence, to satisfy all the conditions in Theorem 1, we consider an additional criterion to the selection of ϵ in the proof of Proposition 2 such that $\epsilon \in (0, \min\{\delta, \delta_1\}]$ and choose the γ -sublevel set accordingly. Through the case studies in Section 5, we show that this choice of \mathcal{P} and \mathcal{X}_f given in Proposition 2 can be readily implemented, which guarantees asymptotic stability of the closed loop system under the proposed learning-enhanced MPC framework. We note that the conditions given above are not only valid for the proposed learning-based MPC scheme, but they can also be applied to other suitably formulated nonlinear MPC schemes. This approach for the selection of the terminal cost and constraint set may be more conservative than some recent results, e.g., Köhler et al. (2019); Eyüboğlu and Lazar (2022). Examining and comparing with these results is a direction for future work.

4.2. Enabling Efficient Implementation

Due to the nonlinear nominal dynamics and KNODE ensemble, (7) is a nonlinear constrained optimization problem. To ensure that it is solved in an efficient manner, we incorporate a number of measures in our implementation. We formulate (7) in CasADi (Andersson et al. (2019)) as a parametric optimization problem. This allows the solver to retain a fixed problem structure during deployment. A solver based on the interior-point method, IPOPT (Wächter and Biegler (2006)), is used to solve the optimization problem. It is warm-started at each time step by providing it with an initial guess, based on the solution from the previous time step. Since NODEs are used to learn only the residual and not the full dynamics, the trained neural networks are relatively lightweight in both width and depth. This further reduces the computational efforts required to solve (7). While the evaluation of the models in the KNODE ensemble are computationally cheaper as compared to training them, there is a trade-off between the number of models in the ensemble and the required computational efforts. In Section 5, we show that a reasonably small number of models improves the prediction accuracy of the dynamics model.

5. Case Studies

We consider two case studies to illustrate the flexibility and efficacy of the proposed learning-enhanced nonlinear MPC framework.

5.1. Inverted Pendulum

We first consider an inverted pendulum system with dynamics given by (Brockman et al. (2016)),

$$\ddot{\theta} = 3g \sin(\theta)/(2l) + 3\tau/(ml^2), \quad (12)$$

where θ is the angle made by the pendulum and the vertical, m and l are the mass and length of the pendulum, g is the gravitational force, and τ is the external torque acting on the pendulum. We define the state and control input to be $x := [\theta \ \dot{\theta}]^\top$ and $u := \tau$. To ascertain the learning ability of the KNODE ensemble, residual dynamics are introduced in the form of a mass difference. In particular, we consider the mass in the nominal model to be 1kg, while the mass of the true pendulum is 0.55kg. The inverted pendulum is required to track a time-varying reference trajectory defined by a sequence of step angle commands. The dynamics are simulated using an explicit RK45 method with a sampling time of 0.01s. To get the terminal cost matrix P , we linearize and discretize (12), and compute the matrix K that optimizes the discrete-time linear-quadratic regulator for the

chosen values of Q and R . We then solve the discrete-time Lyapunov equation as described in Proposition 2 to get P and subsequently, \mathcal{X}_f . For this system, we set the parameter γ to 0.01. More details on training, the ensemble architecture and controller parameters are provided in the Appendix.

5.2. Quadrotor System

Next, we consider a high-dimensional quadrotor system. The dynamics of the system are given as ([Mellinger and Kumar \(2011\)](#))

$$m\ddot{r} = mg + R\eta, \quad I\dot{\omega} = \tau - \omega \times I\omega, \quad (13)$$

where $r, \omega \in \mathbb{R}^3$ are the position and angular rate vectors of the quadrotor, $\eta \in \mathbb{R}$, $\tau \in \mathbb{R}^3$ are the thrust and moments generated by the motors of the quadrotor. The gravity vector is denoted by g and R is the transformation matrix that maps η to the acceleration of the quadrotor. m and $I \in \mathbb{R}^{3 \times 3}$ are the mass and inertia matrix of the quadrotor. The state and control inputs are defined as $x := [r^\top \dot{r}^\top q^\top \omega^\top]^\top \in \mathbb{R}^{13}$ and $u := [\eta \tau^\top]^\top \in \mathbb{R}^4$, where $q \in \mathbb{R}^4$ denotes the quaternions that represents the orientation of the quadrotor. The quadrotor is tasked to track circular trajectories of varying radii and speeds, which implies tracking time-varying position and velocity commands. It is assumed that there is an unknown disturbance force that is a function of the velocities acting on the quadrotor. The cost matrix P and terminal constraint set \mathcal{X}_f are obtained using a similar approach as the first case study. Details on training, the ensemble architecture and controller parameters are listed in the Appendix.

6. Results and Discussion

6.1. Predictions by the KNODE Ensemble

For the inverted pendulum, we evaluate the prediction accuracy of the KNODE ensemble by computing the statistics of the mean-squared errors (MSE) across 30 test trajectories. For each of these trajectories, both the initial states and the magnitude of the step commands are selected at random, under an uniform distribution. Statistics of the prediction accuracy are shown in Fig. 1. We compare the MSEs across three cases; (i) individual KNODE models, (ii) the first variant, *i.e.*, a KNODE ensemble of equal weights, and (iii) the second variant, *i.e.*, a KNODE ensemble with different weights. Observed from the left panel of Fig. 1, the median errors for the two variants are 18.1% and 56.2% lower than those given by the predictions across the individual models.

Next, we analyze the results by comparing the MSEs for each of the test runs, as depicted in the right panel of Fig. 1. The MSEs for the average model are computed by taking the average across the models in the ensemble. As observed, the MSEs for the second variant are lower than those for the first variant and the average model. These improvements indicate that the KNODE ensemble provides more accurate predictions and generalizes better to unseen trajectories, since these test runs are different from those in the training data. It is observed from the right panel of Fig. 1 that the first variant provides lesser MSE reduction, as compared to the second variant. These results illustrate the effectiveness of the weight optimization procedure described in Section 3.2.

For the quadrotor system, we compute MSEs of the states across 30 trajectories, where the commanded speed and radii, and initial states are selected at random. Statistics of the MSEs are given in Fig. 2. An improvement is observed for the KNODE ensemble with different weights, as compared to the other two schemes, in terms of the median errors. We further examine the MSEs for each of the 30 runs, as shown in the right panel of Fig. 2. While it is observed that both variants of the KNODE ensemble provide improved accuracy and lower MSEs for all runs over the

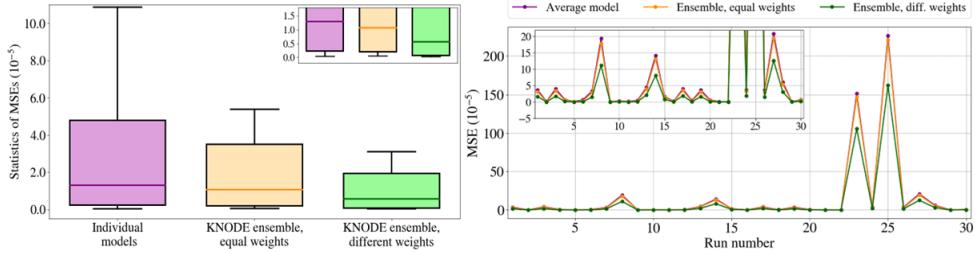


Figure 1: MSEs for individual KNODE models, KNODE ensembles with equal and different weights across 30 test trajectories, for the inverted pendulum. The left panel depicts the error statistics and the right panel shows the MSEs for each test trajectory.

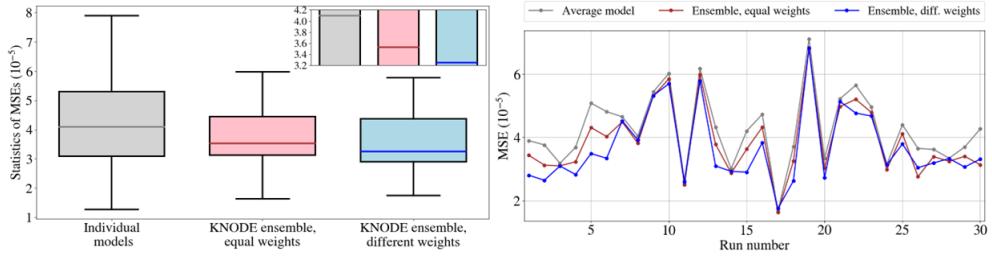


Figure 2: MSEs for individual KNODE models, KNODE ensembles with equal and different weights across 30 test trajectories, for the quadrotor. The left panel depicts the error statistics and the right panel shows the MSEs for each test trajectory.

average model, the application of different weights provides MSE reduction in 60% of the runs, when compared against using equal weights in the ensemble. This can be attributed to the challenge of finding suitable weights for the ensemble as the required complexity of the models increases, particularly for the quadrotor.

6.2. Closed-loop Tracking Performance

To evaluate the performance of the inverted pendulum, we consider the time histories of the tracking errors between the reference and closed-loop angle trajectories. We compare the errors under three schemes; learning-enhanced MPC with (i) individual KNODE models, (ii) a KNODE ensemble of equal weights and (iii) a KNODE ensemble with different weights. These schemes are denoted as *KNODE-MPC*, *EnKNODE-MPC*, *Equal* and *Different*. Fig. 3 depicts the time histories of the angle tracking errors across test trajectories. It is observed that using *KNODE-MPC* results in higher variance in the steady-state errors, as compared to the two ensemble approaches. Moreover, the median steady-state errors are smaller in magnitude for the MPC schemes that use a KNODE ensemble. This implies that using a KNODE ensemble leads to more consistent closed-loop performance. On the other hand, the number of models required for the ensemble would depend on the dynamics, model complexity and the amount of computational power available. While significant prediction improvements are obtained with the KNODE ensemble with different weights, as shown in Fig. 1, we only observe a slight improvement in the error bands for *EnKNODE-MPC, Opt*, as compared to *EnKNODE-MPC, Equal*. This implies that for this inverted pendulum, improvements in model accuracy may not necessarily lead to significant improvements in the closed-loop performance.

For the quadrotor, the task involves following a circular trajectory, which translates to the tracking of time-varying position and velocity commands. The time histories of the position errors for one test case are shown in the left panel of Fig. 4. The errors are computed by taking the difference

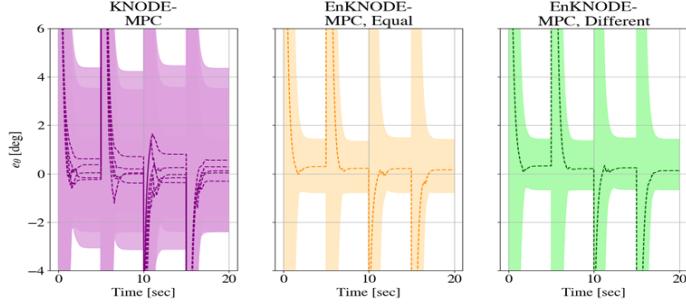


Figure 3: Time histories of the errors between the reference and angle trajectories for the pendulum, under the three learning-based MPC schemes. The dashed lines represent the time histories of the median across trajectories, while the colored bands indicate the maximum and minimum tracking errors across test runs.

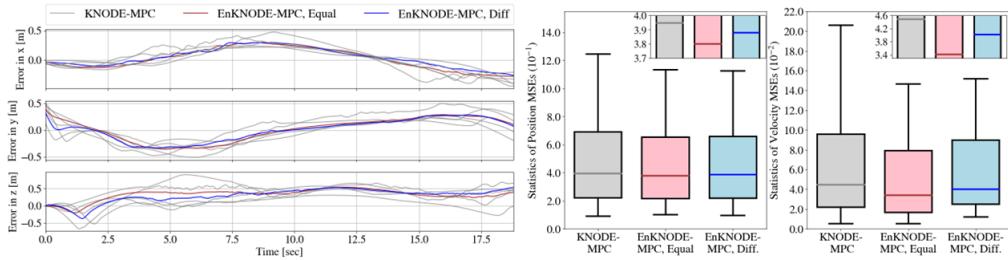


Figure 4: **Left:** Time histories of the position errors for the quadrotor for a single test case. For this test case, the MSEs for *EnKNODE-MPC, Equal* and *Diff.* are 1.125, while the MSEs for *KNODE-MPC* with the single models range from 1.128 to 1.260. **Right:** Statistics of the quadrotor position and velocity tracking MSEs under the three MPC schemes. The insets are zoomed-in plots of the median errors.

between the quadrotor and reference trajectories. The reference trajectories in the x and y axes are time-shifted for ease of comparison, as there is a time difference between the reference and quadrotor trajectories, due to the time-varying nature of the reference trajectories. As observed, the errors are slightly smaller with the KNODE ensemble in the MPC scheme, as compared to a single model. Next, we evaluate the overall closed-loop performance through the MSE statistics for the positions and velocities across 30 runs and they are shown in the right panel of Fig. 4. The position MSEs under *EnKNODE-MPC, Equal* and *EnKNODE-MPC, Different* are lower than *KNODE-MPC* by 3.8% and 1.7% and the velocity MSEs are lower by 24.1% and 10.4% respectively, illustrating the efficacy of the framework.

7. Conclusion

We present a learning-enhanced nonlinear MPC framework that incorporates concepts in deep learning, namely KNODEs and deep ensembles. For the proposed framework, sufficient conditions that render the closed-loop system asymptotically stable are provided. Through two case studies, we demonstrate that the proposed scheme not only provides more accurate predictions, but also improves the closed-loop performance. For future work, we aim to further reduce the computational efforts required, and improve performance for the framework, by combining ideas from multi-stage nonlinear MPC.

Appendix

An extended version of this paper can be found at: <https://arxiv.org/abs/2211.13829>.

Acknowledgments

This work is supported by DSO National Laboratories, 12 Science Park Drive, Singapore 118225, NSF IUCRC 1939132, and the Office of Naval Research (ONR) Award No. N00014-22-1-2157.

References

Joel A E Andersson, Joris Gillis, Greg Horn, James B Rawlings, and Moritz Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Math. Program. Computation*, 11(1):1–36, 2019.

Daniele Bernardini and Alberto Bemporad. Scenario-based model predictive control of stochastic constrained linear systems. In *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 6333–6338. IEEE, 2009.

Francesco Borrelli, Alberto Bemporad, and Manfred Morari. *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

Kong Yao Chee, Tom Z Jiahao, and M Ani Hsieh. KNODE-MPC: A Knowledge-Based Data-Driven Predictive Control Framework for Aerial Robots. *IEEE Robotics and Automation Letters*, 7(2):2819–2826, 2022.

Vijay Sekhar Chellaboina and Wassim M Haddad. *Nonlinear dynamical systems and control: A Lyapunov-based approach*. Princeton University Press, 2008.

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

Yujiao Chen, Zheming Tong, Yang Zheng, Holly Samuelson, and Leslie Norford. Transfer learning with deep neural networks for model predictive control of hvac and natural ventilation in smart buildings. *Journal of Cleaner Production*, 254:119866, 2020.

Nishant Chowdhri, Laura Ferranti, Felipe Santafé Iribarren, and Barys Shyrokau. Integrated nonlinear model predictive control for automated driving. *Control Engineering Practice*, 106:104654, 2021.

Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.

Mert Eyüboğlu and Mircea Lazar. snmpc: A matlab toolbox for computing stabilizing terminal costs and sets. *IFAC-PapersOnLine*, 55(30):19–24, 2022.

Mudasir A Ganaie, Minghui Hu, et al. Ensemble deep learning: A review. *arXiv preprint arXiv:2104.02395*, 2021.

Lars Grüne and Jürgen Pannek. *Nonlinear model predictive control*. Springer, 2017.

Lukas Hewing, Kim P Wabersich, Marcel Menner, and Melanie N Zeilinger. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:269–296, 2020.

Tom Z Jiahao, M Ani Hsieh, and Eric Forgoston. Knowledge-based learning of nonlinear dynamics and chaos. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(11):111101, 2021.

Tom Z Jiahao, Kong Yao Chee, and M Ani Hsieh. Online dynamics learning for predictive control with an application to aerial robots. *arXiv preprint arXiv:2207.09344*, 2022.

Juraj Kabzan, Lukas Hewing, Alexander Liniger, and Melanie N Zeilinger. Learning-based model predictive control for autonomous racing. *IEEE Robotics and Automation Letters*, 4(4):3363–3370, 2019.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.

Johannes Köhler, Matthias A Müller, and Frank Allgöwer. A nonlinear model predictive control framework using reference generic terminal ingredients. *IEEE Transactions on Automatic Control*, 65(8):3576–3583, 2019.

Andrew Kusiak, Mingyang Li, and Fan Tang. Modeling and optimization of hvac energy consumption. *Applied Energy*, 87(10):3092–3102, 2010.

Wook Hyun Kwon and Soo Hee Han. *Receding horizon control: model predictive control for state models*. Springer Science & Business Media, 2005.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

Emilio T Maddalena and Colin N Jones. Learning non-parametric models with guarantees: A smooth lipschitz regression approach. *IFAC-PapersOnLine*, 53(2):965–970, 2020.

Lalo Magni, Giuseppe De Nicolao, Lorenza Magnani, and Riccardo Scattolini. A stabilizing model-based predictive control algorithm for nonlinear systems. *Automatica*, 37(9):1351–1362, 2001.

Daniel Mellinger and Vijay Kumar. Minimum snap trajectory generation and control for quadrotors. In *2011 IEEE international conference on robotics and automation*, pages 2520–2525. IEEE, 2011.

Ali Mesbah, Kim P Wabersich, Angela P Schoellig, Melanie N Zeilinger, Sergio Lucia, Thomas A Badgwell, and Joel A Paulson. Fusion of machine learning and mpc under uncertainty: What advances are on the horizon? In *2022 American Control Conference (ACC)*, pages 342–357. IEEE, 2022.

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

Alessandro Saviolo, Guanrui Li, and Giuseppe Loianno. Physics-inspired temporal learning of quadrotor dynamics for accurate model predictive trajectory tracking. *arXiv preprint arXiv:2206.03305*, 2022.

Georg Schildbach, Lorenzo Fagiano, Christoph Frei, and Manfred Morari. The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, 50(12):3009–3018, 2014.

Sang Hwan Son, Jong Woo Kim, Tae Hoon Oh, Dong Hwi Jeong, and Jong Min Lee. Learning of model-plant mismatch map via neural network modeling and its application to offset-free model predictive control. *Journal of Process Control*, 115:112–122, 2022.

Enrico Terzi, Fabio Bonassi, Marcello Farina, and Riccardo Scattolini. Learning model predictive control with long short-term memory networks. *International Journal of Robust and Nonlinear Control*, 31(18):8877–8896, 2021.

Guillem Torrente, Elia Kaufmann, Philipp Föhn, and Davide Scaramuzza. Data-driven mpc for quadrotors. *IEEE Robotics and Automation Letters*, 6(2):3769–3776, 2021.

Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006.

Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708, 2020.

Zhe Wu, David Rincon, Quanquan Gu, and Panagiotis D Christofides. Statistical machine learning in model predictive control of nonlinear processes. *Mathematics*, 9(16):1912, 2021.

Ye Yao and Divyanshu Kumar Shekhar. State of the art review on model predictive control (mpc) in heating ventilation and air-conditioning (hvac) field. *Building and Environment*, 200:107952, 2021.