



Brief paper

A neural network-based approach to hybrid systems identification for control[☆]

Filippo Fabiani^{a,*}, Bartolomeo Stellato^b, Daniele Masti^c, Paul J. Goulart^d

^a IMT School for Advanced Studies Lucca, Piazza San Francesco 19, 55100, Lucca, Italy

^b Operations Research and Financial Engineering, Princeton University, Princeton, NJ, 08540, USA

^c Gran Sasso Science Institute, Viale F. Crispi 7, 67100, L'Aquila, Italy

^d Department of Engineering Science, University of Oxford, OX1 3PJ, Oxford, United Kingdom

ARTICLE INFO

Article history:

Received 2 April 2024

Received in revised form 9 October 2024

Accepted 22 November 2024

Available online xxxx

ABSTRACT

We consider the problem of designing a machine learning-based model of an unknown dynamical system from a finite number of (state-input)-successor state data points, such that the model obtained is also suitable for optimal control design. We adopt a neural network (NN) architecture that, once suitably trained, yields a hybrid system with continuous piecewise-affine (PWA) dynamics that is differentiable with respect to the network's parameters, thereby enabling the use of derivative-based training procedures. We show that a careful choice of our NN's weights produces a hybrid system model with structural properties that are highly favorable when used as part of a finite horizon optimal control problem (OCP). Specifically, we rely on available results to establish that optimal solutions with strong local optimality guarantees can be computed via nonlinear programming (NLP), in contrast to classical OCPs for general hybrid systems which typically require mixed-integer optimization. Besides being well-suited for optimal control design, numerical simulations illustrate that our NN-based technique enjoys very similar performance to state-of-the-art system identification methods for hybrid systems and it is competitive on nonlinear benchmarks.

© 2025 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Within the systems-and-control community, *system identification* methods have greatly benefited from powerful tools originating in the machine learning literature (Bemporad, 2023; Breschi, Piga and Bemporad, 2016; Pillonetto, Dinuzzo, Chen, De Nicola, & Ljung, 2014). In particular, NNs (Goodfellow, Bengio, & Courville, 2016) have been used since the 1980s (Werbos, 1989) to produce data-driven system model surrogates. Their flexibility – stemming from their universal approximation capability (Hornik, Stinchcombe, & White, 1989) and the wide range of available architectures – has motivated research into NN-based techniques for (non)linear system identification adopting various structures (Andersson, Ribeiro, Tiels, Wahlström, & Schön, 2019;

Forgione, Muni, Piga, & Gallieri, 2023; Forgione & Piga, 2021; Masti & Bemporad, 2021; Mavkov, Forgione, & Piga, 2020).

Despite a few exceptions (Hoekstra et al., 2023; Liu & Kadiramanathan, 1998), available learning-based system identification techniques do not consider whether the resulting models possess an internal structure favorable for use within an OCP. In fact, unless one assumes a certain simplified structure, deep NNs are generally hard to analyze due to their nonlinear and large-scale structure (Goodfellow et al., 2016). As an undesirable consequence, adopting the resulting system model surrogates into OCPs may make the latter extremely hard to solve efficiently (wherever a solution is guaranteed to exist). In this paper we propose a NN-based technique to identify a hybrid system representation, in the form of continuous PWA dynamics, with specific structure suitable for optimal control design.

1.1. Related work

Although hybrid models represent a broad class of systems (Garulli, Paoletti, & Vicino, 2012) whose identification is known to be NP-hard in general (Lauer, 2015; Roll, Bemporad, & Ljung, 2004), to the best of our knowledge little attention has focused on the derivation of such models through (deep) NN-based surrogates.

[☆] This work has been partially funded by the European Union - NextGenerationEU under the Italian Ministry of University and Research (MUR) National Innovation Ecosystem grant ECS00000041 - VITALITY – CUP: D13C21000430001. Daniele Masti is also part of INDAM/GNAMPA. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Dario Piga under the direction of Editor Alessandro Chiuso.

* Corresponding author.

E-mail addresses: filippo.fabiani@imtlucca.it (F. Fabiani), bstellato@princeton.edu (B. Stellato), daniele.masti@gssi.it (D. Masti), paul.goulart@eng.ox.ac.uk (P.J. Goulart).

Starting from the pioneering work in [Batruni \(1991\)](#), where a continuous PWA representation was obtained through a NN with a specific two-layer structure by relying on [Choi and Choi \(1994\)](#), [Chua and Deng \(1988\)](#) constructed a PWA local model by introducing an input space tessellation via hyperpolyhedral convex cells, associating to them NN granules with a local interpolation capability. In [Hush and Horne \(1998\)](#), instead, the authors introduced a constructive algorithm that builds a NN with PWA sigmoidal nodes and one hidden layer. In particular, one node at a time was added by fitting the residual, a task that was accomplished by searching for the best fit through a sequence of quadratic programs (QPs). A single hidden layer network was also proposed in [Gad, Atiya, Shaheen, and El-Dessouki \(2000\)](#), which constructed a continuous PWA error function and developed an efficient, two-step algorithm working in the weight space to minimize it. Recently, [Fält and Giselsson \(2019\)](#) proposed a series of experiments in which novel libraries are employed to identify dynamical models with NNs for complicated hybrid systems, while [de Avila Belbute-Peres, Smith, Allen, Tenenbaum, and Kolter \(2018\)](#) and [Li, Wu, Tedrake, Tenenbaum, and Torralba \(2018\)](#) have analyzed the effects of exploiting structured knowledge to NN surrogates in describing the system multi-modal behavior.

Similar to our approach, [Jin, Aydinoglu, Halm, and Posa \(2022\)](#) proposed to infer a control-oriented linear complementarity (LC) dynamics (equivalent to a PWA one [Heemels, De Schutter, & Bemporad, 2001](#)) directly from data. In particular, the approach in [Jin et al. \(2022\)](#) built upon the minimization of a tailored violation-based loss function, which allowed one to learn LC dynamics via standard gradient-based methods. In view of our results, we also mention the approach in [Yang, Balaprakash, and Leyffer \(2022\)](#). There, a stationarity condition was identified as necessary and sufficient to characterize local optima of a program with a rectified linear unit (ReLU) NN entering both in the cost and constraints, and the corresponding LC-based reformulation yielding an MPCC.

1.2. Summary of contribution

Following our recent contributions ([Fabiani & Goulart, 2023, 2024](#)), in this paper we take a NN-based approach to the problem of system identification for control. In particular, we wish to answer the following question: given N (state-input)-successor state samples $\{(x^{(i)}, u^{(i)}, x^{+, (i)})\}_{i=1}^N$, how can we obtain a descriptive system model that is also suitable for optimal control design? While the meaning of “descriptive” and “suitable” will be clarified later in Section 2, we summarize the contribution as follows:

- (1) We employ a NN-based method to obtain a hybrid system with continuous PWA dynamics from available data. The simple NN architecture we adopt combines an OptNet layer ([Amos, Jimenez, Sacks, Boots, & Kolter, 2018; Amos & Kolter, 2017](#)) and an affine one, and results in a differentiable output with respect to (w.r.t.) the NN’s parameters (Sections 3, 4);
- (2) Given its end-to-end trainability, we show that a careful choice of some weights of the NN allows us to produce a hybrid dynamics model with a specific structure. By relying on ([Hempel, Goulart, & Lygeros, 2017](#)), we then establish that the latter can be controlled (locally) optimally by solving the Karush–Kuhn–Tucker (KKT) conditions of the underlying OCP (Section 5);
- (3) Extensive numerical simulations show that, as a NN-based identification procedure, our technique has competitive performance compared to available hybrid system identification methods (Section 6).

Our method thus requires a simple identification step, represented by a careful training of a NN with specific structure via standard tools, which yields a hybrid system model that is well-suited to optimal control design. In particular, solving an OCP involving a hybrid system with PWA dynamics as an NLP has been proven to require shorter computation times and feature better scaling in the problem dimensions than standard approaches based on mixed-integer optimization ([Hall, Nurkanović, Messerer, & Diehl, 2021; Hempel et al., 2017](#)). In addition, we recall that PWA regression is NP-hard in general ([Lauer, 2015](#)), since it requires simultaneous classification of the N samples $\{(x^{(i)}, u^{(i)}, x^{+, (i)})\}_{i=1}^N$ into modes, thereby calling for a regression of a submodel for each mode. By taking a NN-based perspective, we circumvent such a potentially challenging classification issue, thereby reducing the identification step to the training of the adopted NN architecture, a task that can be accomplished through standard gradient-based methods in view of the network’s output differentiability.

Remarkably, we obtain an LC model suitable for control as a direct consequence of a careful choice of a NN with specific architecture, thereby circumventing the requirement of strict complementarity to recover differentiability w.r.t. the main parameters as in [de Avila Belbute-Peres et al. \(2018\)](#) and [Jin et al. \(2022\)](#). In contrast to [Yang et al. \(2022\)](#), we give an explicit structure of a NN for which local stationarity conditions, coinciding with the standard KKT system, are known to hold for the MPCC obtained by embedding the NN as a hybrid model in an OCP.

2. Problem formulation

We will assume that we have available a finite collection of (state-input)-successor state measured triplets, (x, u, x^+) , $x, x^+ \in \mathbb{R}^n$, $u \in \mathbb{R}^m$ for an unknown but deterministic dynamic system. Our aim is to produce a data-driven model of this unknown system without running further experiments. Such a model shall be “descriptive enough”, i.e., it shall belong to a model class capable of capturing a wide variety of system behaviors, while at the same time being suited for optimal control.

To this end, we will consider throughout the paper the following finite horizon optimal control problem (OCP):

$$\begin{aligned} \min_{\mathbf{u}, \mathbf{x}} \quad & \ell_T(x_T) + \sum_{t \in \mathcal{T}} \ell_t(x_t, u_t) \\ \text{s.t.} \quad & x_{t+1} = \mathcal{N}_\theta(x_t, u_t), \text{ for all } t \in \mathcal{T}, \\ & x_0 = x(0), \end{aligned} \quad (1)$$

where x_t , $t \in \mathcal{T} := \{0, \dots, T-1\}$, denotes the predicted system state t timesteps into the future through the (possibly nonlinear) dynamical model $\mathcal{N}_\theta : \mathbb{R}^p \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$, which is to be inferred from data and is characterized by parameters $\theta \in \mathbb{R}^p$. We denote by $\mathbf{u} := \text{col}((u_t)_{t \in \mathcal{T}}) \in \mathbb{R}^{mT}$ the collection of control inputs to be chosen and by $\mathbf{x} := \text{col}((x_t)_{t \in \mathcal{T} \cup \{T\}}) \in \mathbb{R}^{n(T+1)}$ the resulting state trajectory starting from the measured initial state $x_0 = x(0)$. In addition, let ℓ_t and ℓ_T denote the stage and terminal costs respectively, which are assumed to meet the following standard condition:

Standing Assumption 1. The stage cost $\ell_t : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ and the terminal cost $\ell_T : \mathbb{R}^n \rightarrow \mathbb{R}$ are convex functions.

Thus, given N samples $\{(x^{(i)}, u^{(i)}, x^{+, (i)})\}_{i=1}^N$, we aim at developing a system model $\mathcal{N}_\theta : \mathbb{R}^p \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that, once \mathcal{N}_θ is used within the OCP in (1), an optimal solution to the latter can be characterized via standard optimality conditions. To this end, we design \mathcal{N}_θ as a neural network (NN)-based representation of a hybrid system, ultimately leading to an OCP equivalent to (1) in the form of a mathematical program

with complementarity constraints (MPCC) (Luo, Pang, & Ralph, 1996), whose Karush–Kuhn–Tucker (KKT) conditions (Bertsekas, Nedic, & Ozdaglar, 2003, §5.1) will be necessary and sufficient to characterize an optimal solution.

For technical reasons, we will further assume that the state-input pair takes values in some bounded set:

Standing Assumption 2. $(x, u) \in \Omega$, where $\Omega \subset \mathbb{R}^n \times \mathbb{R}^m$ is a convex polytope.

3. A neural network-based representation of hybrid systems

To address the problem introduced in Section 2, we will design \mathcal{N}_θ as a two-layer NN with a simplified architecture, and then make use of available data $\{(x^{(i)}, u^{(i)}, x^{+, (i)})\}_{i=1}^N$ to train the associated parameters, generically described by $\theta \in \mathbb{R}^p$. In particular, the proposed NN consists of an OptNet layer (Amos et al., 2018; Amos & Kolter, 2017), which takes the input pair $\text{col}(x, u) =: y$ as a parameter to solve the following generic quadratic program (QP):

$$\begin{aligned} \min_z \quad & \frac{1}{2}z^\top Q(y)z + q(y)^\top z \\ \text{s.t.} \quad & A(y)z = b(y), \\ & G(y)z \leq h(y), \end{aligned} \quad (2)$$

thus returning an optimal solution $z^* : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^s$ as output. This output is then passed through a linear layer with tailored matrix $W \in \mathbb{R}^{n \times s}$ and affine term $c \in \mathbb{R}^n$:

$$x^+ = Wz^*(x, u) + c =: \mathcal{N}_\theta(x, u). \quad (3)$$

To ensure a unique optimizer for (2) so that there is no ambiguity in the definition of the dynamics in (3), during the training phase we will impose $Q \in \mathbb{S}_{>0}^s$. This will be made possible in view of the main technical features possessed by the proposed NN, established next:

Proposition 1. *The NN $\mathcal{N}_\theta : \mathbb{R}^p \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ in (3) enjoys the following properties:*

- (i) In (2), assume that $Q \in \mathbb{S}_{>0}^s$ and that A has full row rank. Then the output of \mathcal{N}_θ in (3) is differentiable with respect to (w.r.t.) the whole set of parameters θ ;
- (ii) \mathcal{N}_θ can represent any continuous piecewise-affine (PWA) mapping defined over Ω . Specifically, in case the mapping to be modeled is defined over a regular partition of Ω with r pieces, then the number of constraints that we require to reproduce it is no more than $2nr$, i.e., $l \leq 2nr$, and $s \leq 2n$.

Proof. (i) Since the composition of differentiable mappings remains differentiable, the proof follows by combining (Amos & Kolter, 2017, Th. 1), which proves differentiability of $z^*(x, u)$ w.r.t. the OptNet parameters under the postulated conditions on matrices Q and A , and the fact that any successor state x^+ is defined by an affine combination of $z^*(x, u)$ for any state-input pair (x, u) , which is hence differentiable w.r.t. both W and c . (ii) This part follows from Assumption 2 and Hempel, Goulart, and Lygeros (2014, Th. 2) directly. In fact, if one chooses $c = 0$, and parameters Q, q, A, b, G, h so that (2) becomes:

$$\begin{aligned} \min_z \quad & \frac{1}{2}z^\top Qz + (p + Ry)^\top z \\ \text{s.t.} \quad & Fz + Gy \leq h, \end{aligned} \quad (4)$$

with $Q \in \mathbb{S}_{>0}^s$, then the structure in Hempel et al. (2014, Eq. 4) is immediately recovered. Specifically, the inequality constraints in (4) are obtained from those in (2) by simply imposing $G(y) = F$ and $h(y) = h - Gy$, for appropriate matrices F, G and vector h

with row dimension $l \leq 2nr$ and $s \leq 2n$ columns (notice the slight abuse of notation). Remarkably, these choices are always possible since differentiability of the NN output w.r.t. θ , which follows from the first property shown in this statement, implies that \mathcal{N}_θ is an end-to-end trainable NN. \square

In the rest of the paper, when we will refer to the OptNet layer we will tacitly consider the structure in (4). The set of parameters characterizing \mathcal{N}_θ is then $\theta := \{(Q, R, p, F, G, h), W, c\} \in \mathbb{R}^p$, which shall be determined during an offline training procedure using the available dataset $\{(x^{(i)}, u^{(i)}, x^{+, (i)})\}_{i=1}^N$. With this regard, Proposition 1.(i) makes \mathcal{N}_θ in (3) an end-to-end trainable NN, meaning that one is actually able to set values for the parameters θ exploiting available samples, as commonly happens with any other NN. A differentiable output enables the use of a typical backpropagation strategy to find the gradient of the training loss function. Thus, one can train \mathcal{N}_θ with standard gradient descent methods for finding a minimum of that function.

As a main consequence of the second property in Proposition 1, instead, we note that (1) turns out to coincide with an OCP involving a continuous PWA dynamics. The latter is known to be equivalent to a number of hybrid system model classes (Heemels et al., 2001), such as linear complementarity (LC) (Jin et al., 2022) or mixed-logical dynamical (MLD) models (Bemporad & Morari, 1999). For this reason, the family of continuous PWA dynamics excels in capturing a wide variety of real-world system behaviors. In particular, since \mathcal{N}_θ will be trained on the basis of available data $\{(x^{(i)}, u^{(i)}, x^{+, (i)})\}_{i=1}^N$, (1) actually amounts to an OCP involving a data-based hybrid system representation.

Finally, note that the architecture expressed in (3)–(4) represents a totally valid NN in the machine learning realm, featuring attractive computational properties (Blondel et al., 2022). In particular, it incorporates an *implicit layer* (Amos & Kolter, 2017; Blondel et al., 2022), which requires an iterative process to compute its output, unlike traditional *explicit layers*. Since the proposed method involves not only solving the QP in (4) but also learning its parameters, referring to it as a NN is consistent with the terminology used in the literature.

4. The OCP in (1) is a data-based MPCC

We derive next an equivalent formulation for the OCP in (1). Successively, we will then present in Section 4.2 available results that are applicable to MPCCs models with the particular structure we obtain, before applying those results to our problem in Section 5.

4.1. Mathematical derivation

Suppose that we have trained \mathcal{N}_θ in (3) to obtain the structure of the OptNet layer as in (4) with $Q \in \mathbb{S}_{>0}^s$, and with $c = 0$. The following KKT system is then necessary and sufficient to characterize the optimal solution z^* :

$$\begin{cases} Qz^* + Ry + p + F^\top \lambda = 0, \\ 0 \leq (h - Fz^* - Gy) \perp \lambda \geq 0, \end{cases} \quad (5)$$

where $\lambda \in \mathbb{R}_{\geq 0}^l$ represents the vector of Lagrange multipliers associated to the linear constraints. By recalling that $y = \text{col}(x, u)$, from the KKT system above we obtain a so-called LC model of the system dynamics (Heemels & Brogliato, 2003) associated with the NN \mathcal{N}_θ , with architecture as in (3):

$$\begin{cases} x^+ = -WQ^{-1}R \begin{bmatrix} x \\ u \end{bmatrix} - WQ^{-1}F^\top \lambda - WQ^{-1}p, \\ 0 \leq (FQ^{-1}R - G) \begin{bmatrix} x \\ u \end{bmatrix} + FQ^{-1}F^\top \lambda \\ \quad \quad \quad + FQ^{-1}p + h \perp \lambda \geq 0. \end{cases} \quad (6)$$

Note that all of the terms in (6) except (x, u) and λ are determined through a training procedure for \mathcal{N}_θ in (3). Hence the terminology “data-based representation”.

Substituting our LC model from (6) into (1) yields:

$$\begin{aligned} \min_{\mathbf{u}, \mathbf{x}, \mathbf{w}} \quad & \ell_T(x_T) + \sum_{t \in \mathcal{T}} \ell_t(x_t, u_t) \\ \text{s.t.} \quad & x_{t+1} = Ax_t + B_u u_t + B_w w_t + d, \quad \forall t \in \mathcal{T}, \\ & 0 \leq E_w w_t + E_x x_t + E_u u_t + e \perp w_t \geq 0, \quad \forall t \in \mathcal{T}, \\ & x_0 = x(0), \end{aligned} \quad (7)$$

where the matrices $A, B_u, B_w, E_w, E_x, E_u$, and vectors d and e can be obtained by rearranging the terms and definition of the complementarity variable $w := \lambda$. Specifically, $B_w := -WQ^{-1}F^\top$, $d := -WQ^{-1}p$, $E_w := FQ^{-1}F^\top$, $e := FQ^{-1}p + h$. The matrices A and B_u (respectively, E_x and E_u) follow by partitioning $-WQ^{-1}R$ (resp., $FQ^{-1}R - G$) with $n + m$ columns into two matrices with n and m columns, respectively. The notation is then analogous to (1), and $\mathbf{w} := \text{col}((w_k)_{t \in \mathcal{T}})$ denotes the trajectory of complementarity variables corresponding to \mathbf{u} and \mathbf{x} . A distinct feature of the LC dynamics are the complementarity constraints inherited from (6), which hence turns the OCP (7), equivalent to (1) under our choice of \mathcal{N}_θ in (3), into a data-based MPCC.

Remark 1. Unlike Jin et al. (2022), where an LC model shall be inferred from data directly, we obtain it as a consequence of choosing to train a NN \mathcal{N}_θ with architecture as in (3). For this reason, we do not require strict complementarity of the resulting LC model to recover differentiability w.r.t. the main parameters, as postulated instead in de Avila Belbute-Peres et al. (2018) and Jin et al. (2022).

4.2. Prior results on a class of MPCCs

In general, MPCCs amount to nonlinear, nonconvex optimization problems that can be very challenging to solve (Ralph, 2008). Indeed, for such problems the standard constraint qualifications for nonlinear programming (NLP) (e.g., the classical Mangasarian-Fromovitz one (Bertsekas et al., 2003, §5.5.4)) typically fail to hold at any feasible point (Chen & Florian, 1995).

On the other hand, it has been recently proven in Hempel et al. (2017) that if the MPCC has a specific structure then one is able to establish strong stationarity conditions characterizing a (local) optimal solution. In particular, by referring to the data-based MPCC in (7) the following three requirements on the LC model are sufficient to recover the strong stationarity conditions derived in Hempel et al. (2017, Th. 1):

Condition 1. Given any state-input pair (x, u) , every complementarity variable w that solves $0 \leq E_w w + E_x x + E_u u + e \perp w \geq 0$ results in the same successor state x^+ .

Condition 2. The complementarity problem $0 \leq E_w w + E_x x + E_u u + e \perp w \geq 0$ can be decomposed elementwise w.r.t. the complementarity variable w , i.e.,

$$\forall i \in \{1, \dots, l\} : 0 \leq m_i w_i + D_i x + G_i u + e_i \perp w_i \geq 0, \quad (8)$$

for $m_i > 0$ so that $E_w = \text{diag}((m_i)_{i=1}^l)$, $E_x = \text{col}((D_i)_{i=1}^l)$, $E_u = \text{col}((G_i)_{i=1}^l)$ and $e = \text{col}((e_i)_{i=1}^l)$.

Condition 3. Given any state-input pair (x, u) , there exists some $i \in \{1, \dots, l\}$ such that $D_i x + G_i u + e_i < 0$.

While Condition 1 guarantees the well-posedness of the LC model we consider, entailing a deterministic behavior for the resulting dynamics, Conditions 2 and 3 are rather technical and partially limit the LC models we are allowed to consider. Specifically, Condition 2 means that the solution set of the complementarity problem $0 \leq E_w w + E_x x + E_u u + e \perp w \geq 0$ is given by the cartesian product of the solution sets of (8), while Condition 3 requires the existence of a solution $w_i \neq 0$ to (8), for any fixed pair (x, u) .

Remark 2. We consider an elementwise decomposition of the complementarity problem $0 \leq E_w w + E_x x + E_u u + e \perp w \geq 0$ for simplicity, although a generalization of Condition 2 allowing for a block-diagonal decomposition is also possible – see Hempel et al. (2017, Ass. 2).

Armed with these requirements, our next result provides necessary and sufficient conditions to characterize a local solution to the optimal control MPCC in (7):

Lemma 1 (Hempel et al., 2017, Th. 1). Let $(\mathbf{x}^*, \mathbf{u}^*, \mathbf{w}^*)$ be feasible for the MPCC in (7) with an LC model satisfying Conditions 1–3. Then, $(\mathbf{x}^*, \mathbf{u}^*, \mathbf{w}^*)$ is locally optimal if and only if the standard KKT conditions for (7) admit a primal–dual solution pair.

The statement in Lemma 1 enables one to seek a local solution to the optimal control MPCC (7) through the solution of a classical KKT system, i.e., as an NLP. This is typically computationally advantageous with significantly better scaling features compared to more traditional mixed-integer approaches to solving OCPs based on generic PWA or hybrid models¹ (Hall et al., 2021; Hempel et al., 2017).

Then, if we manage to train our NN \mathcal{N}_θ in (3) so that Conditions 1–3 are satisfied, \mathcal{N}_θ would also bring major computational advantages when used as part of an OCP.

5. Main results

We now show how the NN \mathcal{N}_θ in (3), where $z^*(x, u)$ minimizes (2) (or (4)) for a given pair $(x, u) \in \Omega$, can be trained so that Conditions 1–3 are met, thereby enabling us to solve the data-based OCP (7) as an NLP.

Specifically, in Hempel et al. (2017) it was shown that the following inverse optimization model of PWA dynamics leads to an LC model so that, once it is embedded into an OCP as in (1), we can satisfy all three conditions collectively:

$$x^+ = \alpha^*(x, u) - \beta^*(x, u) = \alpha^* - \beta^*, \quad \text{with} \quad (9a)$$

$$\alpha^* = \begin{cases} \text{argmin}_{\alpha \in \mathbb{R}^n} & \frac{1}{2} \|\alpha - A_\psi x - B_\psi u - c_\psi\|_{Q_\alpha}^2 \\ \text{s.t.} & \alpha \geq A_{\alpha,i} x + B_{\alpha,i} u + c_{\alpha,i}, \\ & \text{for all } i \in \{1, \dots, n_r^\alpha\} \end{cases} \quad (9b)$$

$$\beta^* = \begin{cases} \text{argmin}_{\beta \in \mathbb{R}^n} & \frac{1}{2} \|\beta - A_\phi x - B_\phi u - c_\phi\|_{Q_\beta}^2 \\ \text{s.t.} & \beta \geq A_{\beta,j} x + B_{\beta,j} u + c_{\beta,j}, \\ & \text{for all } j \in \{1, \dots, n_r^\beta\}, \end{cases} \quad (9c)$$

where $Q_\alpha, Q_\beta \in \mathbb{S}_{>0}^n$ are diagonal, but otherwise arbitrary, matrices. Moreover, the elements $\{(A_{\alpha,i}, B_{\alpha,i}, c_{\alpha,i})\}_{i=1}^{n_r^\alpha}$ and $\{(A_{\beta,j}, B_{\beta,j}, c_{\beta,j})\}_{j=1}^{n_r^\beta}$ were determined in Hempel et al. (2017) on the basis of a PWA partitions, computed through available methods as,

¹ For NP-hard problems as (7) one should not expect globally optimal solutions to be efficiently computable in general (Daafouz, Di Benedetto, Blondel, & Hetel, 2009), regardless of the method employed.

e.g., Delaunay triangulations and Voronoi diagrams, of the state-input space $\Omega \subset \mathbb{R}^n \times \mathbb{R}^m$ introduced in [Assumption 2](#) – see, for instance, the discussion in [Hempel et al. \(2014\)](#) and [Hempel, Goulart, and Lygeros \(2013\)](#). In addition, the affine functions in the costs originating from (A_ψ, B_ψ, c_ψ) (respectively, (A_ϕ, B_ϕ, c_ϕ)), instead, were required to lower bound $\{(A_{\alpha,i}, B_{\alpha,i}, c_{\alpha,i})\}_{i=1}^{n_r^\alpha}$ (resp., $\{(A_{\beta,j}, B_{\beta,j}, c_{\beta,j})\}_{j=1}^{n_r^\beta}$) in each region of the underlying partition. Specifically, one has to satisfy the following:

$$\begin{aligned} A_\psi x + B_\psi u + c_\psi &< \max_{i \in \{1, \dots, n_r^\alpha\}} \{A_{\alpha,i}x + B_{\alpha,i}u + c_{\alpha,i}\}, \quad \forall (x, u) \in \Omega, \\ A_\phi x + B_\phi u + c_\phi &< \max_{j \in \{1, \dots, n_r^\beta\}} \{A_{\beta,j}x + B_{\beta,j}u + c_{\beta,j}\}, \quad \forall (x, u) \in \Omega. \end{aligned} \quad (10)$$

As suggested in [Hempel et al. \(2017\)](#), note that the above can always be satisfied for a given PWA system, e.g., by choosing any index $i \in \{1, \dots, n_r^\alpha\}$ and $j \in \{1, \dots, n_r^\beta\}$ and by setting

$$\begin{aligned} A_\psi x + B_\psi u + c_\psi &= A_{\alpha,i}x + B_{\alpha,i}u + c_{\alpha,i} - \eta, \\ A_\phi x + B_\phi u + c_\phi &= A_{\beta,j}x + B_{\beta,j}u + c_{\beta,j} - \zeta, \end{aligned} \quad (11)$$

with arbitrary terms $\eta, \zeta > 0$.

By taking the NN-based perspective of in this work, however, one does not need to compute all of the aforementioned elements explicitly. As will be made clear in our main result, once the number of regions determining each partition is fixed, i.e., $n_r^\alpha, n_r^\beta \geq 1$ (which will hence coincide with the hyperparameters of our data-driven approach, along with Q_α, Q_β), we will obtain matrices and vectors $\{(A_{\alpha,i}, B_{\alpha,i}, c_{\alpha,i})\}_{i=1}^{n_r^\alpha}, \{(A_{\beta,j}, B_{\beta,j}, c_{\beta,j})\}_{j=1}^{n_r^\beta}, (A_\psi, B_\psi, c_\psi)$ and (A_ϕ, B_ϕ, c_ϕ) , as weights of \mathcal{N}_θ in (3). Specifically, we can claim the following:

Theorem 1. *Let $n_r^\alpha, n_r^\beta \geq 1$ be fixed. The NN $\mathcal{N}_\theta : \mathbb{R}^p \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ in (3) can be trained to produce a hybrid system model satisfying [Conditions 1–3](#).*

Proof. To start we note that, for given $n_r^\alpha, n_r^\beta \geq 1$ and any state-input pair $(x, u) \in \Omega$, (9b) and (9c) can be lumped together to obtain the following separable QP:

$$\begin{aligned} \min_{(\alpha, \beta) \in \mathbb{R}^{2n}} & \frac{1}{2} \left\| \begin{bmatrix} \alpha \\ \beta \end{bmatrix} - \begin{bmatrix} A_\psi \\ A_\phi \end{bmatrix} x - \begin{bmatrix} B_\psi \\ B_\phi \end{bmatrix} u - \begin{bmatrix} c_\psi \\ c_\phi \end{bmatrix} \right\|_H^2 \\ \text{s.t.} & \begin{bmatrix} I_n \otimes \mathbf{1}_{n_r^\alpha} & \mathbf{0}_{n n_r^\alpha \times n} \\ \mathbf{0}_{n n_r^\beta \times n} & I_n \otimes \mathbf{1}_{n_r^\beta} \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \\ & \geq \begin{bmatrix} A_{\alpha,1} \\ \vdots \\ A_{\alpha, n_r^\alpha} \\ A_{\beta,1} \\ \vdots \\ A_{\beta, n_r^\beta} \end{bmatrix} x + \begin{bmatrix} B_{\alpha,1} \\ \vdots \\ B_{\alpha, n_r^\alpha} \\ B_{\beta,1} \\ \vdots \\ B_{\beta, n_r^\beta} \end{bmatrix} u + \begin{bmatrix} c_{\alpha,1} \\ \vdots \\ c_{\alpha, n_r^\alpha} \\ c_{\beta,1} \\ \vdots \\ c_{\beta, n_r^\beta} \end{bmatrix}, \end{aligned}$$

where, for simplicity, we have imposed $\text{diag}(Q_\alpha, Q_\beta) = H \in \mathbb{S}_{>0}^{2n}$. Then, by introducing the decision variable $\xi := \text{col}(\alpha, \beta)$ and after suitably redefining all matrices/vectors, the QP above admits a compact form as:

$$\begin{aligned} \min_{\xi \in \mathbb{R}^{2n}} & \frac{1}{2} \|\xi - A_\gamma x - B_\gamma u - c_\gamma\|_H^2 \\ \text{s.t.} & S_\xi \xi \geq A_\xi x + B_\xi u + c_\xi. \\ \implies & \min_{\xi \in \mathbb{R}^{2n}} \frac{1}{2} \xi^\top H \xi - (A_\gamma x + B_\gamma u + c_\gamma)^\top H \xi \\ & \text{s.t.} \quad S_\xi \xi \geq A_\xi x + B_\xi u + c_\xi. \end{aligned} \quad (12)$$

In particular, $A_\gamma \in \mathbb{R}^{2n \times n}, B_\gamma \in \mathbb{R}^{2n \times m}, c_\gamma \in \mathbb{R}^{2n}$, while $A_\xi \in \mathbb{R}^{n(n_r^\alpha + n_r^\beta) \times n}, B_\xi \in \mathbb{R}^{n(n_r^\alpha + n_r^\beta) \times m}$ and $c_\xi \in \mathbb{R}^{n(n_r^\alpha + n_r^\beta)}$. Note that, in addition, the matrix $S_\xi \in \mathbb{R}^{n(n_r^\alpha + n_r^\beta) \times 2n}$ is a full column rank matrix for any $n_r^\alpha, n_r^\beta \geq 1$. The statement then follows by making a one-to-one correspondence between the model in (3), where $z^*(x, u)$ minimizes (4) for some pair $(x, u) \in \Omega$ with affine layer in cascade, and (12). In particular, for given $n_r^\alpha, n_r^\beta \geq 1$, and diagonal $H \in \mathbb{S}_{>0}^{2n}$, to fully recover (12) while training \mathcal{N}_θ one can treat z in (4) as a decision variable living in \mathbb{R}^{2n} , additionally imposing $Q = H \in \mathbb{S}_{>0}^{2n}, F = -S_\xi, W = [I_n \quad -I_n]$, and $c = 0$. Thus, once trained the NN \mathcal{N}_θ , we obtain equivalences: $-H^{-1}R = [A_\gamma \quad B_\gamma], -H^{-1}p = c_\gamma, G = [A_\xi \quad B_\xi]$, and $-h = c_\xi$. To conclude, note that the required conditions on parameters θ , specifically, on the fixed matrices/vector Q, F, W , and c can be easily imposed during the training as a direct consequence of [Proposition 1](#)(i). \square

Our NN-based hybrid system identification approach for control can therefore be summarized as follows (where the symbol “ \rightarrow ” denotes “obtain”):

- (1) Set $n_r^\alpha, n_r^\beta \geq 1$, diagonal $H \in \mathbb{S}_{>0}^{2n} \rightarrow \{Q, F, W, c\}$
- (2) Train \mathcal{N}_θ in (3) and (4) with $\{Q, F, W, c\}$ and available data $\{(x^{(i)}, u^{(i)}, x^{+, (i)})\}_{i=1}^N \rightarrow \{R, p, G, h\}$;
- (3) Plug the elements in (7), solve the KKT system associated $\rightarrow (x^*, u^*, w^*)$.

In item (2) above, one may also decide to train the structure in (9) directly, and then successively recover the weights in (3)–(4). Specifically, learning the model in (9) amounts to determine the elements $\{(A_{\alpha,i}, B_{\alpha,i}, c_{\alpha,i})\}_{i=1}^{n_r^\alpha}, \{(A_{\beta,j}, B_{\beta,j}, c_{\beta,j})\}_{j=1}^{n_r^\beta}, (A_\psi, B_\psi, c_\psi)$, and (A_ϕ, B_ϕ, c_ϕ) , with hyperparameters n_r^α and n_r^β . To this end, note that also the choice of the Hessian matrix in (9b) and (9c), i.e., Q_α and Q_β , is arbitrary.

To conclude, [Corollary 1](#) essentially particularizes the result in [Lemma 1](#) to the optimal control of the data-based hybrid model designed through the NN \mathcal{N}_θ in (3):

Corollary 1. *Let the NN $\mathcal{N}_\theta : \mathbb{R}^p \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ in (3) be trained to produce a hybrid system model satisfying [Conditions 1–3](#), along with the inequalities in (10). Let (x^*, u^*, w^*) be feasible for the resulting OCP in the form of (7). Then, (x^*, u^*, w^*) is locally optimal if and only if the standard KKT conditions for (7) admit a primal–dual solution pair.*

[Corollary 1](#) requires that if our NN model \mathcal{N}_θ is to be included in the OCP (7), its parameters must also satisfy the strict inequalities in (10) in order for the local optimality results of [Lemma 1](#) to apply. However, as common to any type of NN, these inequalities cannot be enforced uniformly over Ω during the training phase, i.e., strict inequalities are guaranteed to hold true for the considered (state-input)-successor state samples only, $\{(x^{(i)}, u^{(i)}, x^{+, (i)})\}_{i=1}^N$, and not for all $(x, u) \in \Omega$. Therefore, some a-posteriori verification procedure shall be applied to make sure they are satisfied for all $(x, u) \in \Omega$. Conversely, the relations in (11) provide us with a possible means of resolution by setting (A_ψ, B_ψ, c_ψ) and (A_ϕ, B_ϕ, c_ϕ) to meet (10). In particular, imposing equalities in (11) can be done either before or after the training process of \mathcal{N}_θ . Also in this case, however, applying some a-posteriori verification procedure is inevitable.

Remark 3. In contrast to [Yang et al. \(2022\)](#), we give an explicit expression of a NN \mathcal{N}_θ for which local stationarity conditions, coinciding with the standard KKT system, are known to hold true for the MPCC obtained once employed the NN \mathcal{N}_θ as a data-based hybrid system model in the OCP (1).

6. Numerical experiments

We now verify the performance of \mathcal{N}_θ in (3)–(4) to learning continuous PWA models from dataset produced by PWA and nonlinear dynamics.² All simulations are run on a laptop equipped with a AMD R7 8840U processor.

To this end, we focus on the structure reported in (9) directly, since [Theorem 1](#) establishes that the two representations are equivalent, and we set $Q_\alpha = Q_\beta = I$, and $n_r^\alpha = n_r^\beta = 7$ unless otherwise stated. We implement our method in Python, utilizing JAX ([Bradbury et al., 2018](#)), JaxOPT ([Blondel et al., 2022](#)) and Scipy ([Virtanen & Contributors, 2020](#)). The models have been trained to solve the following standard regularized training problem:

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^N \|x^{+, (i)} - \hat{x}^{+, (i)}\|_2^2 + \lambda \|\theta\|_2^2 \quad (13)$$

$$\text{s.t. } \hat{x}^{+, (i)} = \alpha_\theta(x^{(i)}, u^{(i)}) - \beta_\theta(x^{(i)}, u^{(i)}), \quad \forall i = 1, \dots, N,$$

where $\alpha_\theta(x, u)$ and $\beta_\theta(x, u)$ reminds the explicit dependence from $\theta = \{(A_{\alpha,i}, B_{\alpha,i}, c_{\alpha,i})\}_{i=1}^7 \cup \{(A_{\beta,j}, B_{\beta,j}, c_{\beta,j})\}_{j=1}^7 \cup \{(A_\psi, B_\psi, c_\psi), (A_\phi, B_\phi, c_\phi)\}$ of (9b) and (9c), respectively. The training problem was solved using the SLSQP algorithm ([Kraft, 1988](#)) with $\lambda = 0.01$, and BoxCDQP from JaxOPT as internal QP solver.

We then evaluate the approximation quality of the models obtained on the basis of open-loop predictions, produced by exciting the trained \mathcal{N}_θ with an unseen input sequence for 1000 time steps, and employing standard metrics ([Breschi, Bemporad and Piga, 2016](#); [Breschi, Piga et al., 2016](#); [Masti & Bemporad, 2021](#); [Mejari, 2018](#)) such as best fit ratio (BFR), defined as $\text{BFR} = \frac{\|\hat{X}^+ - X^+\|_F}{\|X^+ - \bar{X}^+\|_F}$, and the root mean square (RMS) error, $\text{RMS} = \sqrt{\frac{1}{N} \sum_{i=1}^N \|x^{+, (i)} - \hat{x}^{+, (i)}\|_2^2}$. Referring to BFR, all matrices have dimension $n \times N$, with X^+ stacking the true data, \bar{X}^+ is the componentwise average vector of x , while \hat{X}^+ is the predicted value of \hat{X}^+ obtained in open-loop prediction. For system identification purposes, a high (respectively, low) value for BFR (RMS) is desirable.

Unless otherwise stated, all tests have been conducted by using 5000 normalized samples, adding white noise $N(0, 0.01)$ to the sequence of state measurements.

6.1. Comparison with PARC

We first contrast the performance achieved by our method with that obtained by piecewise-affine regression and classification (PARC), a state-of-the-art PWA regression method that has been configured to perform up to 25 iterations and consider 20 clusters. To take into account the nonconvex nature of both learning problems, each test has been repeated for 10 different initial conditions of the optimizer.

6.1.1. Piecewise system benchmarks

To compare the two approaches on learning hybrid systems, we consider the following benchmarks:

- A system described by the following equation:

$$\Sigma_{\text{PWA}} : x^+ = Ax + Bu + W_A \text{clip}_{[0,2]}(W_B x),$$

where $A, W_A, W_B \in \mathbb{R}^{4 \times 4}$, $B \in \mathbb{R}^{4 \times 2}$, and $\text{clip}(\cdot)$ is a mapping that rounds the value of its argument in the range defined by the subscript. The entries of A, B have been drawn from uniform distributions $U(0, 1)$ and $U(0, 1/3)$, respectively, while the entries of W_A, W_B follow a normal distribution $N(0, 1/2)$;

- The problem adopted in [Breschi, Bemporad et al. \(2016, §IV.A\)](#), which we simulate from a random initial condition using an excitation signal with the features described in [Breschi, Bemporad et al. \(2016\)](#). We will refer to this system, which has also been used as benchmark in [Breschi, Piga et al. \(2016\)](#) and [Mejari \(2018\)](#), as $\Sigma_{\text{B-PWA}}$;
- The queuing network described in [Garbi, Incerto, and Tribastone \(2020, Ex. 5\)](#), in which we assume the second and third components of the vector μ are each restricted to the interval $[10, 100]$, while its first component remains set to 30 and the vector s set to $[1000, 11, 11]$. The system is excited using a random white sequence drawn from a uniform distribution defined on the range above. The resulting ODEs are integrated using the LSODA suite ([Hindmarsh, 1983](#)) via Scipy. The sampling time is set to 5 s. We will refer to this system as Σ_{QN} .

The numerical results for the examples above are reported in [Table 1](#), which clearly show how our methodology works well, achieving performance comparable with PARC and, in turn, with the current state-of-the-art. From our numerical experience we note that solving the training problem in (13) on, e.g., the Σ_2 benchmark, required around 15 [s], while PARC approximately 30 [s].

6.1.2. Linear parameter-varying and nonlinear benchmarks

We now compare the performance of our NN-based approach and PARC on the following nonlinear or parameter-varying system models:

- The two-tank system available with the MATLAB system identification toolbox ([MathWorks Inc, 0000](#)). As the dataset contains input–output data only, we define the state as $x_t = [y_{t-2} \ y_{t-1} \ y_t]^\top$. For this example, we use 2000 samples for training and 1000 for validation without adding any white noise to the measurements. We will refer to this system as Σ_{Tank} ;
- The linear parameter-varying dynamics in [Breschi, Bemporad et al. \(2016, §IV.B\)](#), excited according to the related discussion in that paper. We will refer to this system, which has also been used as benchmark in [Breschi, Piga et al. \(2016\)](#) and [Mejari \(2018\)](#), as $\Sigma_{\text{B-LPV}}$;
- The tank system “ Σ_2 ” described in [Masti and Bemporad \(2021\)](#), which exhibits a strong nonlinear input–output behavior. For this example, the noise treatment and system excitation have been performed as in [Masti and Bemporad \(2021\)](#).

[Table 2](#) reports the numerical values obtained for these examples, where we can see that, on the $\Sigma_{\text{B-LPV}}$ benchmark, our method greatly outperforms PARC. On the reference hardware, for Σ_{tank} PARC takes 25 [s], while our training procedure around 40 [s]. On this benchmark, while the identification performance of our approach is comparable to that of [Masti and Bemporad \(2021\)](#), our training time is significantly lower than that reported in [Masti and Bemporad \(2021\)](#) (around 20 minutes). This is mostly related to the fact that our method needs to learn a few hundreds of coefficients only. Finally, the results on Σ_2 are also very competitive w.r.t. those in [Masti and Bemporad \(2021\)](#), especially in view of the larger dataset employed in the latter paper (i.e., 20 000 against 5000 samples).

6.2. Sensitivity analysis – parameters n_r^α, n_r^β

We now investigate the sensitivity of our NN-based approach w.r.t. the main parameters characterizing our technique, i.e., n_r^α and n_r^β . Note that also Q_α, Q_β represent further possible hyperparameters to tune, however, we note that according to both the

² Python code is available upon request (email at daniele.masti@gssi.it).

Table 1
BFR and RMS obtained by our method and PARC (within parentheses) on PWA benchmarks.

	Median BFR	Best BFR	Variance BFR	Median RMS	Best RMS	Variance RMS
Σ_{PWA}	0.952 (0.937)	0.959 (0.944)	$2.88 \cdot 10^{-4}$ ($2.66 \cdot 10^{-5}$)	0.048 (0.064)	0.042 (0.057)	$2.95 \cdot 10^{-4}$ ($3.12 \cdot 10^{-5}$)
Σ_{B-PWA}	0.883 (0.855)	0.932 (0.891)	$7.28 \cdot 10^{-3}$ ($7.48 \cdot 10^{-4}$)	0.084 (0.109)	0.050 (0.082)	$4.20 \cdot 10^{-3}$ ($4.58 \cdot 10^{-4}$)
Σ_{QN}	0.838 (0.81)	0.851 (0.829)	$2.20 \cdot 10^{-4}$ ($4.41 \cdot 10^{-4}$)	0.193 (0.224)	0.175 (0.205)	$3.40 \cdot 10^{-4}$ ($5.90 \cdot 10^{-4}$)

Table 2
BFR and RMS obtained by our method and PARC (within parentheses) on the nonlinear and LPV benchmarks.

	Median BFR	Best BFR	Variance BFR	Median RMS	Best RMS	Variance RMS
Σ_{Tank}	0.901 (0.900)	0.932 (0.909)	$4 \cdot 10^{-4}$ ($1.64 \cdot 10^{-3}$)	0.108 (0.109)	0.075 (0.09)	$4.76 \cdot 10^{-4}$ ($1.95 \cdot 10^{-3}$)
Σ_{B-PLV}	0.696 (0.511)	0.731 (0.528)	$4.47 \cdot 10^{-4}$ ($1.62 \cdot 10^{-4}$)	0.287 (0.46)	0.255 (0.44)	$3.96 \cdot 10^{-4}$ ($1.30 \cdot 10^{-4}$)
Σ_2	0.924 (0.921)	0.947 (0.931)	$1.69 \cdot 10^{-4}$ ($3.47 \cdot 10^{-3}$)	0.061 (0.063)	0.042 (0.056)	$1.09 \cdot 10^{-4}$ ($1.96 \cdot 10^{-5}$)

Table 3
BFR obtained by our method in learning the behavior of Σ_2 and Σ_{B-PWA} for different values of n_r^α .

	Σ_{B-PWA}			Σ_2		
	Median	Best	Variance	Median	Best	Variance
$n_r^\alpha = 2$	0.915	0.943	$1.06 \cdot 10^{-3}$	0.863	0.863	$1.21 \cdot 10^{-3}$
$n_r^\alpha = 3$	0.897	0.925	$3.51 \cdot 10^{-4}$	0.889	0.913	$8.92 \cdot 10^{-5}$
$n_r^\alpha = 5$	0.904	0.932	$3.23 \cdot 10^{-3}$	0.915	0.924	$7.05 \cdot 10^{-5}$
$n_r^\alpha = 7$	0.883	0.932	$7.28 \cdot 10^{-3}$	0.924	0.947	$1.69 \cdot 10^{-4}$
$n_r^\alpha = 11$	0.845	0.936	$5.99 \cdot 10^{-3}$	0.939	0.956	$2.85 \cdot 10^{-4}$
$n_r^\alpha = 17$	0.875	0.910	$5.87 \cdot 10^{-3}$	0.946	0.952	$2.49 \cdot 10^{-4}$




Fig. 1. Comparison between an optimal controller based on SQP solver and one based on the DIRECT global optimization solver when tracking the reference sine-sweep r_k . The y -axis represents the (normalized) magnitude of the quantities involved, while the x -axis denotes the time steps. Both controllers have been parametrized to solve the same optimization problem based on a learnt model of Σ_2 .

discussion in Hempel et al. (2014, 2017) and our own observation, the performance of our method is not particularly affected by their values. We therefore omit the related sensitivity analysis for these terms.

We report in Table 3 the median, worst and best BFR achieved on the Σ_2 and Σ_{B-PWA} for different values of $n_r^\alpha = n_r^\beta$. In general we can observe that, while a higher n_r^α is beneficial to improve the best-case scenario performance, it also makes the training problem more difficult. This is clearly shown by the obtained median BFR, especially when dealing with Σ_2 . From our numerical experience, the difficulties of training models featuring higher values of n_r^α (and therefore the capability of representing richer continuous PWA mappings) can be often overcome by making use of larger dataset. These results also stress that, in case the data are fixed in advance, choosing appropriate values for n_r^α and n_r^β is key.

6.3. Performance for predictive control

The computational advantages of employing a model as in (9) for optimal control purposes have been already analyzed

in Hempel et al. (2017). Here, we have shown that the strong stationarity conditions offered by Lemma 1 can be solved through standard NLP solvers such as, e.g., IPOPT (Wächter & Biegler, 2006), outperforming classical mixed-integer programming approaches to hybrid system optimal control.

Nevertheless, one may wonder how the adopted NN model can be used within an OCP in conjunction with a sensitivity-based NLP solver. This is relevant for systems as those introduced in Section 6.1.2, which do not exhibit behaviors involving logic and dynamic at the same time. We then design a simple OCP for Σ_2 , mimicking the setup of Masti and Bemporad (2021), with prediction horizon equal $T = 7$. We solve, at each step, the following optimization problem:

$$\begin{aligned} \min_{u_1, \dots, u_7} \quad & \sum_{i=1}^7 (x_{i,2}^+ - r_{i+1})^2 + 0.01 u_i^2 + 0.001 \delta_i^2 \\ \text{s.t.} \quad & \hat{x}_i^+ = \alpha_\theta(x_i, u_i) - \beta_\theta(x_i, u_i), \quad i = 1, \dots, 7, \\ & \delta_i = u_i - u_{i-1}, \quad i = 1, \dots, 7, \\ & 0.95 \leq u_i \leq 1.2, \quad i = 1, \dots, 7, \end{aligned}$$

where with $x_{i,2}$ we mean the second component of x_i .

We then compare the behavior of a closed-loop trajectory achieved when using SLSQP solver with the one obtained with the global optimization solver DIRECT_L (Gablonsky & Kelley, 2001). The results are shown in Fig. 1 where, although the selected input sequence is slightly different, the closed-loop trajectories achieved by the two controllers are almost indistinguishable, thereby showing that derivative-based solver may exploit the adopted NN-based models.

We note that on our reference hardware, the whole closed-loop simulation with the SLSQP-based controller required approximately one millisecond per time-step. Overall, the SLSQP-based approach requires around 4 [s], in contrast to the 25 [s] needed by that based on DIRECT_L. While we opted for “off-the-shelf” solvers, adopting dedicated methodologies such as, e.g., Fletcher and Leyffer (2004) and Nurkanović, Pozharskiy, and Diehl (2024), would yield even better computational performance.

7. Conclusion and outlook

We have proposed a NN-based methodology to system identification for control that only requires the training of a NN via standard tools as simple identification step, and yields a hybrid system model suited for optimal control design. Specifically, we have employed a NN with specific, yet simple, architecture, which turns out to be end-to-end trainable and produces a hybrid system with PWA dynamics from available data. Extensive numerical simulations have illustrated that, as a NN-based identification procedure, our technique has very similar performance compared to the state-of-the-art of hybrid system identification methods.

By relying on available results (Hempel et al., 2017), we have also shown that, under a careful choice of some weights of the NN, the resulting hybrid dynamics can be controlled locally optimally by just solving the KKT system associated to the underlying finite horizon OCP. This is computationally advantageous compared to traditional approaches to optimal control of hybrid systems, usually requiring mixed-integer optimization.

Future work will concentrate on the integration of OptNet layers together with standard explicit layers characterized by different activation functions. In view of the requirement in Corollary 1, it is also key to investigate the practical impact the various strategies to impose the conditions in (10) have on the resulting NN model.

References

- Amos, B., Jimenez, I., Sacks, J., Boots, B., & Kolter, J. Z. (2018). Differentiable MPC for end-to-end planning and control. *Advances in Neural Information Processing Systems*, 31.
- Amos, B., & Kolter, J. Z. (2017). OptNet: Differentiable optimization as a layer in neural networks. In *International conference on machine learning* (pp. 136–145). PMLR.
- Andersson, C., Ribeiro, A. H., Tiels, K., Wahlström, N., & Schön, T. B. (2019). Deep convolutional networks in system identification. In *2019 IEEE 58th conference on decision and control CDC*, (pp. 3670–3676).
- Batruni, R. (1991). A multilayer neural network with piecewise-linear structure and back-propagation learning. *IEEE Transactions on Neural Networks*, 2(3), 395–403.
- Bemporad, A. (2023). A piecewise linear regression and classification algorithm with application to learning and model predictive control of hybrid systems. *IEEE Transactions on Automatic Control*, 68(6), 3194–3209.
- Bemporad, A., & Morari, M. (1999). Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3), 407–427.
- Bertsekas, D., Nedic, A., & Ozdaglar, A. (2003). *Convex analysis and optimization: vol. 1*, Athena Scientific.
- Blondel, M., Berthet, Q., Cuturi, M., Frostig, R., Hoyer, S., Llinares-López, F., et al. (2022). Efficient and modular implicit differentiation. *Advances in Neural Information Processing Systems*, 35, 5230–5242.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., et al. (2018). JAX: composable transformations of Python+NumPy programs.
- Breschi, V., Bemporad, A., & Piga, D. (2016). Identification of hybrid and linear parameter varying models via recursive piecewise affine regression and discrimination. In *2016 European control conference ECC*, (pp. 2632–2637). IEEE.
- Breschi, V., Piga, D., & Bemporad, A. (2016). Piecewise affine regression via recursive multiple least squares and multicategory discrimination. *Automatica*, 73, 155–162.
- Chen, Y., & Florian, M. (1995). The nonlinear bilevel programming problem: Formulations, regularity and optimality conditions. *Optimization*, 32(3), 193–209.
- Choi, C.-H., & Choi, J. Y. (1994). Constructive neural networks with piecewise interpolation capabilities for function approximations. *IEEE Transactions on Neural Networks*, 5(6), 936–944.
- Chua, L. O., & Deng, A.-C. (1988). Canonical piecewise-linear representation. *IEEE Transactions on Circuits and Systems*, 35(1), 101–111.
- Daafouz, J., Di Benedetto, M. D., Blondel, V. D., & Hetel, L. (2009). Switched and piecewise affine systems. In *Handbook of hybrid systems control: theory, tools, applications* (pp. 87–138). Cambridge University Press.
- de Avila Belbute-Peres, F., Smith, K., Allen, K., Tenenbaum, J., & Kolter, J. Z. (2018). End-to-end differentiable physics for learning and control. *Advances in Neural Information Processing Systems*, 31.
- Fabiani, F., & Goulart, P. J. (2023). Reliably-stabilizing piecewise-affine neural network controllers. *IEEE Transactions on Automatic Control*, 68(9), 5201–5215.
- Fabiani, F., & Goulart, P. J. (2024). Robust stabilization of polytopic systems via fast and reliable neural network-based approximations. *International Journal of Robust and Nonlinear Control*, 34(9), 6180–6201.
- Fält, M., & Giselsson, P. (2019). System identification for hybrid systems using neural networks. arXiv preprint arXiv:1911.12663.
- Fletcher, R., & Leyffer, S. (2004). Solving mathematical programs with complementarity constraints as nonlinear programs. *Optimization Methods & Software*, <http://dx.doi.org/10.1080/10556780410001654241>.
- Forgione, M., Muni, A., Piga, D., & Gallieri, M. (2023). On the adaptation of recurrent neural networks for system identification. *Automatica*, 155, Article 111092.
- Forgione, M., & Piga, D. (2021). Continuous-time system identification with neural networks: Model structures and fitting criteria. *European Journal of Control*, 59, 69–81.
- Gablonsky, J. M., & Kelley, C. T. (2001). A locally-biased form of the DIRECT algorithm. *Journal of Global Optimization*, 21, 27–37.
- Gad, E. F., Atiya, A. F., Shaheen, S., & El-Dessouki, A. (2000). A new algorithm for learning in piecewise-linear neural networks. *Neural Networks*, 13(4), 485–505.
- Garbi, G., Incerto, E., & Tribastone, M. (2020). Learning queuing networks by recurrent neural networks. In *Proceedings of the ACM/SPEC international conference on performance engineering* (pp. 56–66).
- Garulli, A., Paoletti, S., & Vicino, A. (2012). A survey on switched and piecewise affine system identification. *IFAC Proceedings Volumes*, 45(16), 344–355, 16th IFAC Symposium on System Identification.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Hall, J., Nurkanović, A., Messerer, F., & Diehl, M. (2021). A sequential convex programming approach to solving quadratic programs and optimal control problems with linear complementarity constraints. *IEEE Control Systems Letters*, 6, 536–541.
- Heemels, M., & Brogliato, B. (2003). The complementarity class of hybrid dynamical systems. *European Journal of Control*, 9(2–3), 322–360.
- Heemels, M., De Schutter, B., & Bemporad, A. (2001). Equivalence of hybrid dynamical models. *Automatica*, 37(7), 1085–1091.
- Hempel, A. B., Goulart, P. J., & Lygeros, J. (2013). Every continuous piecewise affine function can be obtained by solving a parametric linear program. In *2013 European control conference ECC*, (pp. 2657–2662). IEEE.
- Hempel, A. B., Goulart, P. J., & Lygeros, J. (2014). Inverse parametric optimization with an application to hybrid system control. *IEEE Transactions on Automatic Control*, 60(4), 1064–1069.
- Hempel, A. B., Goulart, P. J., & Lygeros, J. (2017). Strong stationarity conditions for optimal control of hybrid systems. *IEEE Transactions on Automatic Control*, 62(9), 4512–4526.
- Hindmarsh, A. C. (1983). ODEPACK, a systemized collection of ODE solvers. *Scientific Computing*.
- Hoekstra, J. H., Cseppento, B., Beintema, G. I., Schoukens, M., Kollár, Z., & Tóth, R. (2023). Computationally efficient predictive control based on ANN state-space models. In *2023 62nd IEEE conference on decision and control CDC*, (pp. 6336–6341). IEEE.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5), 359–366.
- Hush, D. R., & Horne, B. (1998). Efficient algorithms for function approximation with piecewise linear sigmoidal networks. *IEEE Transactions on Neural Networks*, 9(6), 1129–1141.
- Jin, W., Aydinoglu, A., Halm, M., & Posa, M. (2022). Learning linear complementarity systems. In *Learning for dynamics and control conference* (pp. 1137–1149). PMLR.
- Kraft, D. (1988). A software package for sequential quadratic programming. *Forschungsbericht-Deutsche Forschungs-und Versuchsanstalt fur Luft-und Raumfahrt*.
- Lauer, F. (2015). On the complexity of piecewise affine system identification. *Automatica*, 62, 148–153.
- Li, Y., Wu, J., Tedrake, R., Tenenbaum, J. B., & Torralba, A. (2018). Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In *International conference on learning representations*.
- Liu, G. P., & Kadiramanathan, V. (1998). Predictive control for non-linear systems using neural networks. *International Journal of Control*, 71(6), 1119–1132.
- Luo, Z.-Q., Pang, J.-S., & Ralph, D. (1996). *Mathematical programs with equilibrium constraints*. Cambridge University Press.
- Masti, D., & Bemporad, A. (2021). Learning nonlinear state-space models using autoencoders. *Automatica*, 129, Article 109666.
- MathWorks Inc Control of a Two-Tank System, Available: <https://it.mathworks.com/help/robust/ug/control-of-a-two-tank-system.html>.
- Mavkov, B., Forgione, M., & Piga, D. (2020). Integrated neural networks for nonlinear continuous-time system identification. *IEEE Control Systems Letters*, 4(4), 851–856.
- Mejari, M. D. (2018). *Towards automated data-driven modeling of linear parameter-varying systems* (Ph.D. thesis), IMT School for Advanced Studies Lucca.
- Nurkanović, A., Pozharskiy, A., & Diehl, M. (2024). Solving mathematical programs with complementarity constraints arising in nonsmooth optimal control. *Vietnam Journal of Mathematics*, <http://dx.doi.org/10.1007/s10013-024-00704-z>.
- Pillonetto, G., Dinuzzo, F., Chen, T., De Nicolao, G., & Ljung, L. (2014). Kernel methods in system identification, machine learning and function estimation: A survey. *Automatica*, 50(3), 657–682.
- Ralph, D. (2008). Mathematical programs with complementarity constraints in traffic and telecommunications networks. *Philosophical Transactions of the Royal Society of London A (Mathematical and Physical Sciences)*, 366(1872), 1973–1987.
- Roll, J., Bemporad, A., & Ljung, L. (2004). Identification of piecewise affine systems via mixed-integer programming. *Automatica*, 40(1), 37–50.

Virtanen, P., & Contributors, S. . (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17, 261–272. <http://dx.doi.org/10.1038/s41592-019-0686-2>.

Wächter, A., & Biegler, L. T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106, 25–57.

Werbos, P. J. (1989). Neural networks for control and system identification. In *Proceedings of the 28th IEEE conference on decision and control*, (pp. 260–265).

Yang, D., Balaprakash, P., & Leyffer, S. (2022). Modeling design and control problems involving neural network surrogates. *Computational Optimization and Applications*, 1–42.



Filippo Fabiani is an Assistant Professor at the IMT School for Advanced Studies Lucca, Italy. He received the B.Sc. degree in Bio-Engineering, the M.Sc. degree in Automatic Control Engineering, and the Ph.D. degree in Automatic Control, all from the University of Pisa, in 2012, 2015, and 2019 respectively.

In 2018–2019 he was post-doctoral Research Fellow in the Delft Center for Systems and Control at TU Delft, the Netherlands, while in 2019–2022 he was a post-doctoral Research Assistant in the Control Group at the Department of Engineering Science, University of Oxford, United Kingdom. His research interests include game theory, data-driven and robust control of uncertain systems, and machine learning, with applications in energy systems, smart grids and smart cities.



Bartolomeo Stellato is an Assistant Professor in the Department of Operations Research and Financial Engineering at Princeton University. Previously, he was a Postdoctoral Associate at the MIT Sloan School of Management and Operations Research Center. He holds a DPhil (Ph.D.) in Engineering Science from the University of Oxford, a M.Sc. in Robotics, Systems and Control from ETH Zürich, and a B.Sc. in Automation Engineering from Politecnico di Milano. His awards include the Beale – Orchard-Hays Prize, the ONR Young Investigator Award, the NSF CAREER Award, the Princeton SEAS

Howard B. Wentz Jr. Faculty Award, the Franco Strazzabosco Young Investigator Award from ISSNAF, the Princeton SEAS Innovation Award in Data Science, the Best Paper Award in Mathematical Programming Computation, and the First Place Prize Paper Award in IEEE Transactions on Power Electronics. His research focuses on data-driven computational tools for mathematical optimization, machine learning, and optimal control.



Daniele Masti received his Ph.D. in Systems Science in 2021 at IMT School for Advanced Studies Lucca, Italy. From 2022 to March 2024, he was a researcher in Cyber Security at IMT School for Advanced Studies in Lucca, Italy. He currently serves as PostDoc at Gran Sasso Science Institute, Italy. His main research interest lies at the intersection between control theory and machine learning, with the overall aim of bridging the gap between the two disciplines.



Paul J. Goulart received the B.Sc. and M.Sc. degrees in aeronautics and astronautics from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1998 and 2001, respectively, and the Ph.D. degree in control engineering in 2007 from the University of Cambridge, Cambridge, U.K., where he was selected as a Gates Scholar.

From 2007 to 2011, he was a Lecturer in control systems with the Department of Aeronautics, Imperial College London, and from 2011 to 2014, a Senior Researcher with the Automatic Control Laboratory, ETH Zürich. He joined the Department of Engineering Science, University of Oxford, Oxford, U.K., in 2014 as an Associate Professor, where he became Full Professor in 2024. He is also a Tutorial Fellow in Engineering at St. Edmund Hall. His research interests include high speed optimization and control for embedded systems, data-driven and robust optimization, and machine learning, with applications ranging from energy systems to fluid mechanics.