Proceedings of Machine Learning Research vol 168:1–14, 2022 4th Annual Conference on Learning for Dynamics and Control

A Piecewise Learning Framework for Control of Unknown Nonlinear Systems with Stability Guarantees

Milad Farsi MFARSI@UWATERLOO.CA
Department of Applied Mathematics, University of Waterloo, Canada

Yinan Li YINAN.LI@UWATERLOO.CA
Clearpath Robotics, Kitchener, Canada

Ye Yuan YYE@HUST.EDU.CN
School of Artificial Intelligence & Automation, Huazhong University of Science and Technology, China

Jun Liu J.LIU@UWATERLOO.CA
Department of Applied Mathematics, University of Waterloo, Canada

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

We propose a piecewise learning framework for controlling nonlinear systems with unknown dynamics. While model-based reinforcement learning techniques in terms of some basis functions are well known in the literature, when it comes to more complex dynamics, only a local approximation of the model can be obtained using a limited number of bases. The complexity of the identifier and the controller can be considerably high if obtaining an approximation over a larger domain is desired. To overcome this limitation, we propose a general piecewise nonlinear framework where each piece is responsible for locally learning and controlling over some region of the domain. We obtain rigorous uncertainty bounds for the learned piecewise models. The piecewise affine (PWA) model is then studied as a special case, for which we propose an optimization-based verification technique for stability analysis of the closed-loop system. Accordingly, given a time-discretization of the learned PWA system, we iteratively search for a common piecewise Lyapunov function in a set of positive definite functions, where a non-monotonic convergence is allowed. This Lyapunov candidate is verified on the uncertain system to either provide a certificate for stability or find a counter-example when it fails. This counter-example is added to a set of samples to facilitate the further learning of a Lyapunov function. We demonstrate the results on two examples and show that the proposed approach yields a less conservative region of attraction (ROA) compared with alternative state-of-the-art approaches. Moreover, we provide the runtime results to demonstrate potentials of the proposed framework in real-world implementations.

Keywords: Nonlinear systems, learning-based control, system identification, piecewise affine, Lyapunov analysis

1. Introduction

The flexibility of piecewise affine (PWA) systems makes them suitable for different approaches in control. Hence, the control problem of piecewise systems has been extensively studied in the literature (see, e.g., Marcucci and Tedrake (2019); Zou and Li (2007); Rodrigues and Boyd (2005); Baotic (2005); Strijbosch et al. (2020); Christophersen et al. (2005); Rodrigues and How (2003)). Moreover, various applications can be found for PWA systems, including robotics (Andrikopoulos et al., 2013; Marcucci et al., 2017), automotive control (Borrelli et al., 2006; Sun et al., 2019), and power electronics (Geyer et al., 2008; Vlad et al., 2012). Since PWA systems are highly adaptable to different problems, various techniques are presented to efficiently fit a piecewise model to data (see, e.g., Toriello and Vielma (2012); Breschi et al. (2016); Ferrari-Trecate et al. (2003); Amaldi et al. (2016); Rebennack and Krasko (2020); Du et al. (2021)). A review of some of the techniques can be found in Gambella et al. (2021); Garulli et al. (2012).

Solving the optimal control problem for PWA systems has also been the topic of various works. In Zou and Li (2007), the robust model predictive control (MPC) strategy is extended to PWA systems with polytopic uncertainty, where multiple PWA quadratic Lyapunov functions are employed for different vertices of the uncertainty polytope in different partitions. In another work by Marcucci and Tedrake (2019), hybrid MPC is formulated as a mixed-integer program to solve the optimal control problem for PWA systems. However, these techniques are only available in the open-loop form, which decreases their applicability for real-time control.

Deep neural networks (DNN) offer an efficient technique for control in the closed loop. However, one drawback of DNN-based control is the difficulty in stability analysis. This becomes even more challenging when PWA are considered. Chen et al. (2020) suggested a sample-efficient technique for synthesizing a Lyapunov function for the PWA system controlled through a DNN in the closed loop. In this approach, the analytic center cutting plane method (ACCPM) (Goffin and Vial, 1993; Nesterov, 1995; Boyd et al., 2004) is first used for searching for a Lyapunov function. Then, this Lyapunov candidate is verified on the closed-loop system using mixed-integer quadratic programming (MIQP). This approach relies on the exact model of the system and therefore cannot be directly implemented on an identified PWA system with uncertainty.

Although learning with guarantees has motivated plenty of recent research Chang et al. (2020); Chen et al. (2020); Dai et al. (2021), fewer works have explicitly considered the model uncertainty. In Berkenkamp et al. (2016), the authors consider an approach that learns the region of attraction (ROA) from experiments on a partially unknown system. Based on regularity assumptions on the model errors in terms of a Gaussian process (GP) prior, they employ an underlying Lyapunov function to determine an ROA from which the system is asymptotically stable with high probability. Even though a partially unknown model including uncertainty is considered in this approach, the approach may not computationally scale well with the number of samples considering the use of a GP learner.

Model-based learning approaches in the literature are mainly categorized under reinforcement learning (RL) in two groups: value function and policy search methods. Approximate/adaptive dynamic programming techniques (Wang et al., 2009; Lewis and Vrabie, 2009; Balakrishnan et al., 2008) as a well-known value-based approach can efficiently approximate the solution to the optimal control problem. Even though a set of polynomial bases, for instance, is known to be sufficient as a universal approximator over a compact domain (Kamalapurkar et al., 2018), the number of bases required for a tight approximation of the dynamics over a given domain may be exceedingly high. This highly impedes implementations, especially in an online learning and control setting.

On the other hand, employing a piecewise approach can improve the applicability of model-based learning greatly by keeping the online computations needed for updating the model and control in a tractable size, since at any instance only a particular mode of the system is involved. Hence, in this paper, we consider a partition of the domain consisting of different pieces, where for each piece, we run a local learner in terms of a limited number of bases using a structured online learning (SOL) approach (Farsi and Liu, 2020, 2021, 2022) to obtain a piecewise feedback control. Then, by considering the special case of PWA systems, an optimization-based technique is employed to verify a variant of the piecewise model to obtain rigorous stability guarantees. Two examples are shown to demonstrate the advantages of the proposed framework in terms of providing less conservative stability guarantees in terms of the size of the verified ROA, compared with Lyapunov functions learned using neural networks for systems with known dynamics (Chang et al., 2020).

Due to the page limit, some of the theoretical and simulation results are omitted in this paper. However, the extended version of this work can be accessed on arXiv (Farsi et al. (2022)).

2. Problem Formulation

Consider the nonlinear system in control-affine form

$$\dot{x} = F(x,u) = f(x) + g(x)u = f(x) + \sum_{j=1}^{m} g_j(x)u_j, \tag{1}$$

where $x \in D \subset \mathbb{R}^n$, $u \in \Omega \subset \mathbb{R}^m$, $f:\ D \to \mathbb{R}^n$, and $g:\ D \to \mathbb{R}^{n\times m}$.

The cost functional to be minimized along the trajectory, starting from the initial condition $x(0) = x_0$, is considered to be in the following linear quadratic form

$$J(x_0,u) = \lim_{T\to\infty} \int_0^T \mathrm{e}^{-\gamma t} \left(x^T Q x + u^T R u\right) \mathrm{d}t, \tag{2}$$

where $Q \in \mathbb{R}^{n\times n}$ is positive semi-definite, $\gamma \geq 0$ is the discount factor, and $R \in \mathbb{R}^{m\times m}$ is a diagonal matrix with only positive values, given by the design criteria.

3. The Piecewise Learning and Control Framework

We approximate the nonlinear system (1) by a piecewise model with a bounded uncertainty

$$\dot{x} = W_\sigma \Phi(x) + \sum_{j=1}^{m} W_{j\sigma}\Phi(x)u_j + d_\sigma, \tag{3}$$

where $d_\sigma \in \mathbb{R}^n$ is a time-varying uncertainty, W_σ and $W_{j\sigma} \in \mathbb{R}^{n\times p}$ are the matrices of the coefficients for $\sigma \in \{1,2,\ldots,n_\sigma\}$ and $j \in \{1,2,\ldots,m\}$, with a set of differentiable bases $\Phi(x) = [\phi_1(x) \quad \ldots \quad \phi_p(x)]^T$, and n_σ denoting the total number of pieces. Moreover, any piece of the system is defined over a convex set given by a set of linear inequalities as $\Upsilon_\sigma = \{x \in D | Z_\sigma x \leq z_\sigma\}$, where $\sigma \in \{1,\ldots,n_\sigma\}$ and Z_σ and z_σ are a matrix and a vector, respectively, of appropriate dimensions.

We assume that the set $\{\Upsilon_\sigma\}$ forms a partition of the domain and its elements do not share any interior points, i.e. $\bigcup_{\sigma=1}^{n_\sigma} \Upsilon_\sigma = D$ and $\mathrm{int}[\Upsilon_\sigma] \bigcap \mathrm{int}[\Upsilon_l] = \emptyset$ for $\sigma \neq l$ and $\sigma, l \in \{1,2,\ldots,n_\sigma\}$. Furthermore, the piecewise model is assumed to be continuous across the boundaries of $\{\Upsilon_\sigma\}$ (see Appendix A.1 of Farsi et al. (2022)). The control input and the uncertainty are assumed to be bounded and lie in the sets $\Omega = \{u \in \mathbb{R}^m || u_j| \leq \bar{u}_j, \forall j \in \{1,2,\ldots,m\}\}$ and $\Delta_\sigma = \{d_\sigma \in \mathbb{R}^n || d_{\sigma i}| \leq \bar{d}_{\sigma i}, \forall i \in \{1,2,\ldots,n\}\}$, respectively. The uncertainty upper bound $\bar{d}_\sigma = (\bar{d}_{\sigma 1},\cdots,\bar{d}_{\sigma n})$ is to be determined.

3.1. System Identification

Having defined the parameterized model of the system, we employ a system identification approach to update the system parameters. For each pair of samples obtained from the input and state of the system, i.e., (x^s, u^s), we first locate the element in the partition $\{\Upsilon_\sigma\}$ that contains the sampled state x^s. Then, we locally update the system coefficients of the particular piece from which the state is sampled. In Farsi and Liu (2020), the weights are updated according to

$$[\hat{W}_\sigma \quad \hat{W}_{1\sigma} \ \dots \ \hat{W}_{m\sigma}]_k = \arg\min_{\bar{W}} \quad \|\dot{X}_{k\sigma} - \bar{W}\Theta_{k\sigma}\|_2^2, \tag{4}$$

where k is the time step, and $\Theta_{k\sigma}$ includes a matrix of samples with

$$\Theta_k{}^s = [\Phi^T(x^s) \quad \Phi^T(x^s)u_1^s \quad \dots \quad \Phi^T(x^s)u_m^s]_k^T,$$

for the sth sample in the σth partition. Correspondingly, $\dot{X}_{k\sigma}$ contains the sampled state derivatives. While in principle any identification technique can be used, e.g., Yuan et al. (2019); Brunton et al. (2016), the linearity with respect to the coefficients allows us to employ least-squares techniques. In this paper, we implement the recursive least-squares (RLS) (Ljung and Söderström, 1983; Liu et al., 2016; Wu et al., 2015) technique that provides a more computationally efficient way to update the parameters. Accordingly, only one sample at each time is used to update the weights, instead of processing a history of samples.

3.2. Feedback Control

In Farsi and Liu (2020), a matrix differential equation is proposed using a quadratic parametrization in terms of the basis functions to obtain a feedback control. Here, we adopt a similar learning framework, but consider a family of n_σ differential equations, each of which corresponds to one particular mode of the system in the piecewise model. We integrate the following state-dependent Riccati differential equation in forward time:

$$\begin{aligned} -\dot{P}_\sigma =& \bar{Q} + P_\sigma \frac{\partial \Phi(x)}{\partial x} W_\sigma + W_\sigma^T \frac{\partial \Phi(x)}{\partial x}^T P_\sigma - \gamma P_\sigma \\ & - P_\sigma \frac{\partial \Phi(x)}{\partial x} \left(\sum_{j=1}^{m} W_{j\sigma} \Phi(x) r_j^{-1} \Phi(x)^T W_{j\sigma}^T \right) \frac{\partial \Phi(x)}{\partial x}^T P_\sigma. \end{aligned} \tag{5}$$

The solution to the differential equation (5) characterizes the value function defined by

$$V_\sigma = \Phi^T P_\sigma \Phi, \tag{6}$$

based on which we obtain a piecewise control

$$u_j = -r_j^{-1} \frac{\partial V_\sigma}{\partial x}^T g_j(x) = -\Phi(x)^T r_j^{-1} P_\sigma \frac{\partial \Phi(x)}{\partial x} W_{j\sigma} \Phi(x). \tag{7}$$

4. Analysis of Uncertainty Bounds

We use the uncertainty in the piecewise system (3) to capture approximation errors in identification. In this section, we analyze the worst-case bounds to provide guarantees for the proposed framework.

There exist two sources of uncertainty that affect the accuracy of the identified model. The first is the mismatch between the identified model and the observations made. The latter may also be affected by the measurement noise. The second is due to unsampled areas in the domain. We can estimate the uncertainty bound for any piece of the model by combining these two bounds. In what follows, we discuss the procedure of obtaining these bounds in more detail.

Assumption 1 *For any given (x^s, u^s), let $F_i(x^s, u^s)$ be the ith element of $F(x^s, u^s)$. We assume that $F_i(x^s, u^s)$ can be measured with some tolerance as $\tilde{F}_i(x^s, u^s)$, where $|\tilde{F}_i(x^s, u^s) - F_i(x^s, u^s)| \leq \varrho_e |\tilde{F}_i(x^s, u^s)|$ with $0 \leq \varrho_e < 1$ for all $i \in \{1, \cdots, n\}$.*

We make predictions $\hat{F}_i(x^s, u^s)$ of the state derivatives for any sample using the identified model. Hence, we can easily compute the distance between the prediction and the approximate evaluation of the system by using the samples collected for any piece. This gives the loss $|\hat{F}_i(x^s, u^s) - \tilde{F}_i(x^s, u^s)|$. The proof of the following result can be found in Appendix B.1 of Farsi et al. (2022).

Theorem 1 *Let Assumption 1 hold, and $S_{\Upsilon\sigma}$ denote the set of indices for sample pairs (x^s, u^s) such that $x^s \in \Upsilon_\sigma$. Then, an upper bound of the prediction error, regarding any sample (x^s, u^s) for $s \in \{1, \ldots, N_s\}$, is given by*

$$|\hat{F}_i(x^s, u^s) - F_i(x^s, u^s)| \leq \bar{d}_{e\sigma i} := \max_{s \in S_{\Upsilon\sigma}} \left(|\hat{F}_i(x^s, u^s) - \tilde{F}_i(x^s, u^s)| + \varrho_e |\tilde{F}_i(x^s, u^s)| \right),$$

where $\sigma \in \{1, \ldots, n_\sigma\}$, and $i \in \{1, \ldots, n\}$.

The samples may not be uniformly obtained from the domain. Depending on how smooth the dynamics are, there might be unpredictable behavior of the system in the gaps among the samples. Hence, the predictions made by the identified model may be misleading in the areas we have not visited yet. To take this into account, we assume a Lipschitz constant is given for the system. More specifically, we let $\varrho_x \in \mathbb{R}_+^n$ and $\varrho_u \in \mathbb{R}_+^n$ denote the Lipschitz constants of $F(x, u)$ with respect to x and u on $D \times \Omega$, respectively. We use this to bound the uncertainty for the unsampled areas.

The procedure starts with searching for the largest gaps in the state and control spaces that do not contain any samples as described in detail in Appendix B.2 of Farsi et al. (2022). Let (x^{s*}, u^{s*}) be the closest sample indexed in $S_{\Upsilon\sigma}$ to the center point $(c^*_{x\sigma}, c^*_{u\sigma})$ of the sample gap (as a Euclidean ball) with radii $(r^*_{x\sigma}, r^*_{u\sigma})$. We need to compute the worst case of the prediction error at the center point that is given by $|\hat{F}_i(c^*_{x\sigma}, c^*_{u\sigma}) - F_i(c^*_{x\sigma}, c^*_{u\sigma})|$, where $\hat{F}(\cdot, \cdot)$ denotes an evaluation of the identified model. However, according to Assumption 1, we do not have access to the original system to exactly evaluate $F(\cdot, \cdot)$. Therefore, we obtain the bound in terms of the approximate value instead.

Theorem 2 *Let Assumptions 1-4 hold and $(r^*_{x\sigma}, r^*_{u\sigma})$ be given by the solutions of (20) and (21) (details and the assumptions are given in Appendix B.2 of Farsi et al. (2022)). Then, an upper bound for the prediction error can be obtained regarding all unvisited points $x \in \Upsilon_\sigma$ and $u \in \Omega$ as below*

$$|F_i(x, u) - \hat{F}_i(x, u)| \leq \bar{d}_{\sigma i} = \varrho_{ui} r^*_{u\sigma} + \varrho_{xi} r^*_{x\sigma} + \bar{d}_{e\sigma i} + \hat{\varrho}_{ui} r^*_{u\sigma} + \hat{\varrho}_{xi} r^*_{x\sigma}. \tag{8}$$

The proof can be found in Appendix B.3 of Farsi et al. (2022).

5. Stability Verification for Piecewise-Affine Learning and Control

5.1. Piecewise Affine Models

A special case of system (3) can be obtained when we choose $\Phi(x) = \begin{bmatrix} 1 & x^T \end{bmatrix}$.

We consider system coefficients in the form of $W_\sigma = \begin{bmatrix} C_\sigma & A_\sigma \end{bmatrix}$ and $W_{j\sigma} = \begin{bmatrix} B_{j\sigma} & 0 \end{bmatrix}$. Clearly, A_σ, $B_{j\sigma}$, and C_σ can be used to rewrite the PWA system in the standard form

$$\dot{x} = A_\sigma x + \sum_{j=1}^{m} B_{j\sigma} u_j + C_\sigma + d_\sigma, \tag{9}$$

5.2. MIQP-based Stability Verification of PWA Systems

In this section, we adopt an MIQP-based verification technique based on the approach presented in Chen et al. (2020). In this framework, by considering a few steps ahead, we verify that the Lyapunov function is decreasing. However, it may not be necessarily monotonic, meaning that it may be increasing in some steps and then be decreasing greatly in some other steps to compensate. Regarding the fact that this approach is inherently a discrete technique, we need to consider a discretization of (9). By an Euler approximation, we have

$$x_{k+1} = \check{F}_d(x_k, u_k) = \check{A}_\sigma x_k + \sum_{j=1}^{m} \check{B}_{j\sigma} u_{jk} + \check{C}_\sigma + d_\sigma, \tag{10}$$

where $\check{A}_\sigma$, $\check{B}_{j\sigma}$, and $\check{C}_\sigma$ are the discrete system matrices of the same dimension as (9). Moreover, we re-adjust the uncertainty bound as $\bar{d}_\sigma := h\bar{d}_\sigma$, where h denotes the time step.

We refer the uncertain closed loop system with the control $u_{jk} = \omega_j(x_k)$ as

$$x_{k+1} = \check{F}_{d,cl}(x_k). \tag{11}$$

For this system, let the convex set $\bar{D} = \{x \in D | Z_{\bar{D}} x \leq z_{\bar{D}}\}$ be a user-defined region of interest (ROI), within which obtaining a region of attraction (ROA) is desirable.

5.2.1. Learning and Verification of a Lyapunov Function

Assuming $u_j = -r_j^{-1} B_{j\sigma}^T P_{3\sigma} x_k$, and defining $\check{A}_{cl,\sigma} = \check{A}_\sigma - \sum_{j=1}^{m} r_j^{-1} \check{B}_{j\sigma} B_{j\sigma}^T P_\sigma$, the discrete closed-loop system becomes $x_{k+1} = \check{A}_{cl,\sigma} x_k + \check{C}_\sigma + d_\sigma$.

Now, consider the Lyapunov function

$$V(x_k, \hat{P}) = \begin{bmatrix} x_k \\ x_{k+1} \end{bmatrix}^T \hat{P} \begin{bmatrix} x_k \\ x_{k+1} \end{bmatrix} \tag{12}$$

characterized by $\hat{P} \in \mathscr{F}$, where

$$\mathscr{F} = \{\hat{P} \in \mathbb{R}^{2n \times 2n} | 0 \leq \hat{P} \leq I, V(x_{k+1}, \hat{P}) - V(x_k, \hat{P}) < 0, \forall x_k \in \bar{D} \backslash \{0\}, d_\sigma \in \Delta_\sigma\}.$$

The structure of the Lyapunov function is suggested by Chen et al. (2020) that employs a piecewise quadratic function to parameterize the Lyapunov function. This approach combines the non-monotonic Lyapunov function Ahmadi and Parrilo (2008) and finite-step Lyapunov function Bobiti and Lazar (2016); Aeyels and Peuteman (1998) techniques to provide a guarantee by looking at the next few steps. It should be noted that the Lyapunov function may not be necessarily decreasing within any single step, while it must be decreasing within the finite steps taken into account.

The Learner: To realize a Lyapunov function, one needs a mechanism to look for the appropriate values of $\hat{P}$ within $\mathscr{F}$. For this purpose, we obtain an over-approximation of $\mathscr{F}$ by considering an only finite number of elements in $(\bar{D}, \Delta)$. Let us first define the increment on the Lyapunov function as

$$\Delta V(x, \hat{P}) = V(\check{F}_{d,cl}(x), \hat{P}) - V(x, \hat{P}) = \begin{bmatrix} \check{F}_{d,cl}(x) \\ \check{F}^{(2)}_{d,cl}(x) \end{bmatrix}^T \hat{P} \begin{bmatrix} \check{F}_{d,cl}(x) \\ \check{F}^{(2)}_{d,cl}(x) \end{bmatrix} - \begin{bmatrix} x \\ \check{F}_{d,cl}(x) \end{bmatrix}^T \hat{P} \begin{bmatrix} x \\ \check{F}_{cl}(x) \end{bmatrix},$$

where $\check{F}^{(2)}_{d,cl}(x) = \check{F}_{d,cl}(\check{F}_{d,cl}(x))$.

Furthermore, assume that the set of N_s number of samples is given as below

$$\mathscr{S} = \{(x, \check{F}_{d,cl}(x), \check{F}^{(2)}_{d,cl}(x))_1, \ldots, (x, \check{F}_{d,cl}(x), \check{F}^{(2)}_{d,cl}(x))_{N_s}\}.$$

Note that $\mathscr{S}$ implicitly includes samples of the disturbance input and the state.

Now, using $\mathscr{S}$ we obtain the over-approximation

$$\tilde{\mathscr{F}} = \{\hat{P} \in \mathbb{R}^{2n\times 2n} | 0 \leq \hat{P} \leq I, \Delta V(x, \hat{P}) \leq 0, \forall x \in \mathscr{S}, d_\sigma \in \Delta_\sigma\}.$$

To find an element in $\tilde{\mathscr{F}}$, there exist efficient iterative techniques that are well-known as cutting-plane approaches. See e.g. Atkinson and Vaidya (1995); Elzinga and Moore (1975); Boyd and Vandenberghe (2007). In Chen et al. (2020), the analytic center cutting-plane method (ACCPM) (Goffin and Vial, 1993; Nesterov, 1995; Boyd et al., 2004) is employed in an optimization problem:

$$\hat{P}^{(i)} = \arg\min_{\hat{P}} \quad -\sum_{x\in\mathscr{S}_i} \log(-\Delta V(x, \hat{P})) - \log\det(I - \hat{P}) - \log\det(\hat{P}) \tag{13}$$

where i is the iteration index. If feasible, the log-barrier function in the first term guarantees the solution within $\tilde{\mathscr{F}}$ for which the negativity of the Lyapunov difference holds. The other two terms ensure $0 \leq \hat{P}^{(i)} \leq I$. The solution gives a Lyapunov function V based on the set of the samples $\mathscr{S}_i$ in the ith stage. On the other hand, if a solution does not exist, the set $\mathscr{F}$ is concluded to be empty.

The Verifier: The Lyapunov function candidate suggested by (13) may not guarantee asymptotic stability for all $x \in \bar{D}$ and $d_\sigma \in \Delta_\sigma$ since only the sampled space was considered. Therefore, in the next step, we need to verify the Lyapunov function candidate for the uncertain system. To do so, a mixed-integer quadratic program is solved based on the convex hull formulation of the PWA:

$$\max_{\mathrm{x}^j, \mathrm{u}^j, \mathrm{d}^j, \mu^j} \quad \begin{bmatrix} \mathrm{x}^1 \\ \mathrm{x}^2 \end{bmatrix}^T \hat{P}^{(i)} \begin{bmatrix} \mathrm{x}^1 \\ \mathrm{x}^2 \end{bmatrix} - \begin{bmatrix} \mathrm{x}^0 \\ \mathrm{x}^1 \end{bmatrix}^T \hat{P}^{(i)} \begin{bmatrix} \mathrm{x}^0 \\ \mathrm{x}^1 \end{bmatrix} \tag{14}$$

subject to

$$Z_{\bar{D}}\mathrm{x}^0 \leq z_{\bar{D}}, \|\mathrm{x}^0\|_\infty \geq \epsilon \tag{15}$$

$$\mathrm{u}^j = \omega(\mathrm{x}^j) \tag{16}$$

$$Z_\sigma \mathrm{x}^j_\sigma \leq \mu^j_\sigma z_\sigma, Z_u \mathrm{u}_\sigma \leq \mu^j_\sigma z_u, |\mathrm{d}^j_{\sigma i}| \leq \mu^j_\sigma \bar{d}_{\sigma i}, \tag{17}$$

$$(1, \mathrm{x}^j, \mathrm{u}^j, \mathrm{d}^j, \mathrm{x}^{j+1}) = \sum_{\sigma=1}^{N_\sigma} (\mu^j_\sigma, \mathrm{x}^j_\sigma, \mathrm{u}^j_\sigma, \mathrm{d}^j_\sigma, A_\sigma \mathrm{x}^j_\sigma + B_\sigma \mathrm{u}^j_\sigma + \mu^j_\sigma c_\sigma + \mathrm{d}^j_\sigma) \tag{18}$$

$$\mu_\sigma \in \{0, 1\}, \forall \sigma \in \{1, \ldots, N_\sigma\}, i \in \{1, \ldots, n\}, j \in \{0, 1\}, \tag{19}$$

where a ball of radius ϵ around the origin is excluded from the set of states, and ϵ is chosen small enough in (15). This is due to Remark 5 in Appendix C.3 of Farsi et al. (2022) and the fact that the numerical value of the objective becomes considerably small when approaching the origin. This makes the negativity of the objective too hard to verify around the origin. For more details in the implementation of the algorithm, we refer the reader to Chen et al. (2020).

The system is given by (18) and (19). To define the piecewise system in a mixed-integer problem, similar to Chen et al. (2020), we use the convex-hull formulation of piecewise model that is presented in Marcucci and Tedrake (2019). However, to consider the uncertainty, we compose a slightly different system where we define extra variables to model the disturbance input.

Constraints (15), and (17) define the sets of the initial condition, the state, the control, and the disturbance inputs, respectively. Furthermore, the feedback control is implemented by (16).

To certify the closed-loop system as asymptotically stable, the optimal value returned by the MIQP (14) is required to be negative. Otherwise, the argument $(\mathrm{x}^{0*}, \mathrm{x}^{1*}, \mathrm{x}^{2*})$ of the optimal solution is added to the set of samples $\mathscr{S}$ as a counter-example.

5.3. Stability Analysis

Combining the uncertainty bounds in Section 4 and the Lyapunov-based verification results of this section, we are able to prove the following practical stability results of the closed-loop system.

Theorem 3 *Suppose that the MIQP (14) yields a negative optimal value. Let B_ϵ denote the set $\{x \in \mathbb{R}^n | \|x\|_\infty \leq \epsilon\}$, i.e., the ball of radius ϵ in infinity norm around the origin. Then the set B_ϵ is asymptotically stable for the closed-loop system (11). The largest sub-level set of V, i.e., $\{x \in \mathbb{R}^n \mid V(x) \leq c\}$ for some c, contained in $\bar{D}$ is a verified under-approximation of the real ROA.*

The proof can be found in Appendix C.3 of Farsi et al. (2022). Remark 5 in Appendix C.3 also discusses how to bridge the gap between convergence to B_ϵ and the convergence to the origin.

6. Numerical Results

To validate the proposed piecewise learning and verification technique we implemented the approach on the pendulum system and the dynamical vehicle system Pepy et al. (2006). Moreover, we compared the results with other techniques presented in the literature. To make a fair comparison, we have taken the parameters of the system from Chang et al. (2020). We performed all the simulations in Python 3.7 on a 2.6 GHz Intel Core i5 CPU.

6.1. Pendulum System

For the pendulum system, we discuss the simulation results in three sections. In the first section, we will explain the procedure of identifying the uncertain PWA model with a piecewise feedback control. In the second section, we verify the closed-loop uncertain system and obtain an ROA in $\bar{D}$. In the third section, we will present the comparison results.

6.1.1. IDENTIFY AND CONTROL

Control objective is to stabilize the pendulum at the top equilibrium point given by $x_{\text{eq}} = (0, 0)$. First, we start with learning a piecewise model together with the uncertainty bounds, and the feedback control. For this purpose, we sample the system, and update our model as discussed in section

3.1. We set the sampling time as $h = 5\text{ms}$. Accordingly, the value function and the control rule are updated online as in section 3.2. Then, to verify the value to be decreasing within each mode, it only remains to calculate the uncertainty bounds using the results obtained in section 4.

To make a visualization of the nonlinearity in the pendulum system possible, we portray the second dynamic assuming $u = 0$ in Fig. 1(g)subfigure, where the first dynamic is only linear. The procedure of learning is illustrated through several stages in Fig. 1. In the first column from the left, we illustrated the estimations only for the second dynamic with $u = 0$ to be comparable to Fig. 1(g)subfigure. Accordingly, it can be observed that the system identifier is able to closely approximate the nonlinearity with a piecewise model. More details on the uncertainty bounds obtained are provided in Appendix D of Farsi et al. (2022).

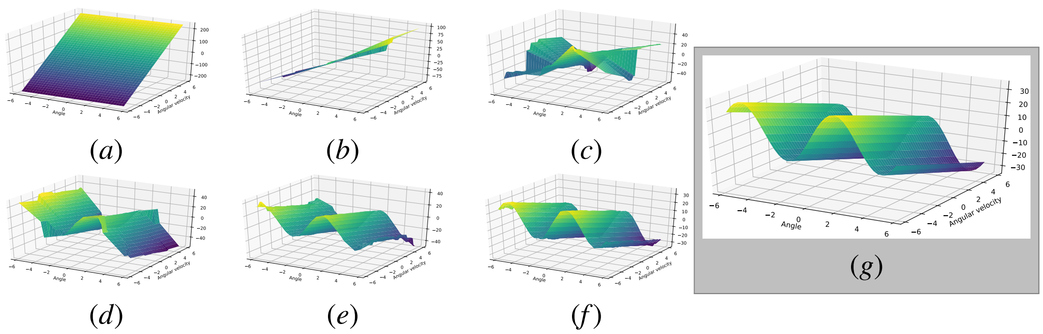

Figure 1: The procedure for learning the dynamics by the PWA model is illustrate step-by-step that shows the convergence of the identifier. Subfigure (g) illustrates the second dynamic of pendulum system assuming $u = 0$ that is $f_2(x_1, x_2)$. Subfigures (a)-(f) show the improvement of estimations of $f_2(x_1, x_2)$ given by (g), as the number of samples increases.

6.1.2. VERIFICATION

Having the system identified and the feedback control, we can apply the verification algorithm based on MIQP problem. As done in Chen et al. (2020), we implemented the learner in CVXpy Diamond and Boyd (2016) with MOSEK ApS (2020) solver, and the verifier in Gurobi 9.1.2 Gurobi (2020).

We choose $\bar{D}$ such that x_1 and $x_2 \in [-6, 6]$. To verify the system, we ran the algorithm and obtained a matrix $\hat{P}$ (who numerical values are given in Appendix D of Farsi et al. (2022)) in that characterizes the Lyapunov function as in (12).

The largest level set of the associated Lyapunov function in $\bar{D}$ is pictured in Fig. 2(a)subfigure as the the ROA of the closed-loop system. Moreover, we illustrate different trajectories of the controlled system that confirms the verified Lyapunov function by constructing an ROA around the origin.

6.1.3. COMPARISON RESULTS

To highlight the merits of the proposed piecewise learning approach, we compare the ROA obtained by different approaches in the literature. Chang et al. (2020) proposed a neural network (NN) Lyapunov function for stability verification. According to Chang et al. (2020), the comparison done on the pendulum system with linear quadratic regulator (LQR) and sums-of-squares (SOS) showed noticeable superiority of the NN-based Lyapunov approach. Following the comparison results from Chang et al. (2020), we compare the ROA obtained by our approach with NN, SOS, and LQR techniques in Fig. 2(b)subfigure. Clearly, the ROA obtained by the piecewise controller with the

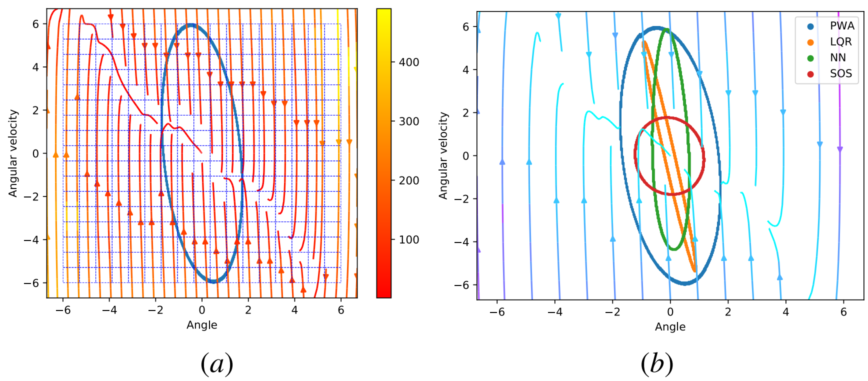

Figure 2: (a). The obtained ROA of the closed-loop PWA system is illustrated for x_1 and $x_2 \in [-6, 6]$. The uniform grids denote the modes of the PWA system. Multiple trajectories of the system are shown in a phase portrait where colormap represents the magnitude in the vector field. (b). The comparison results for ROA of the closed-loop system is illustrated for x_1 and $x_2 \in [-6, 6]$, together with the trajectories of the system. The comparison results for LQR, NN, and SOS are taken from Chang et al. (2020).

non-monotonic Lyapunov function is considerably larger than the ones obtained by NN, SOS, and LQR algorithms as shown in Chang et al. (2020).

6.2. Dynamic Vehicle System with Skidding

We have also implemented the proposed approach in a more complex dynamic vehicle system with skidding, which shows promising results. The results are omitted in the main paper due to the space limit. They can be found in Appendix D of Farsi et al. (2022).

7. Conclusion

For regulating nonlinear systems with uncertain dynamics, a piecewise nonlinear affine framework was proposed in which each piece is responsible for learning and controlling over a partition of the domain locally. Then, in a particular case of the proposed framework, we focused on learning in the form of the well-known PWA systems, for which we presented an optimization-based verification approach that takes into account the estimated uncertainty bounds. We used the pendulum system as a benchmark example for the numerical results. Accordingly, an ROA resulting from the level set of a learned Lyapunov function is obtained. Furthermore, the comparison with other control approaches in the literature illustrates a considerable improvement in the ROA using the proposed framework. As another example, we implemented the presented approach on a dynamical vehicle system with considerably higher number of partitions and dimensions. The results demonstrated that the approach can scale efficiently, hence, can be potentially implemented on more complex real-world problems in real-time.

Acknowledgments

This work is partially supported by Jiangsu Industrial Technology Research Institute (JITRI) through a JITRI-Waterloo project, the NSERC Canada Research Chairs (CRC) program, an NSERC Discovery Grant, and an Ontario Early Researcher Award (ERA).

References

Dirk Aeyels and Joan Peuteman. A new asymptotic stability criterion for nonlinear time-variant differential equations. *IEEE Transactions on automatic control*, 43(7):968–971, 1998.

Amir Ali Ahmadi and Pablo A Parrilo. Non-monotonic lyapunov functions for stability of discrete time nonlinear and switched systems. In *2008 47th IEEE conference on decision and control*, pages 614–621. IEEE, 2008.

Edoardo Amaldi, Stefano Coniglio, and Leonardo Taccari. Discrete optimization methods to fit piecewise affine models to data points. *Computers & Operations Research*, 75:214–230, 2016.

George Andrikopoulos, George Nikolakopoulos, Ioannis Arvanitakis, and Stamatis Manesis. Piecewise affine modeling and constrained optimal control for a pneumatic artificial muscle. *IEEE Transactions on Industrial Electronics*, 61(2):904–916, 2013.

MOSEK ApS. The mosek optimization toolbox for python manual, 2020.

David S Atkinson and Pravin M Vaidya. A cutting plane algorithm for convex programming that uses analytic centers. *Mathematical Programming*, 69(1):1–43, 1995.

SN Balakrishnan, Jie Ding, and Frank L Lewis. Issues on stability of adp feedback controllers for dynamical systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 38(4):913–917, 2008.

Mato Baotic. *Optimal control of piecewise affine systems: A multi-parametric approach*. PhD thesis, ETH Zurich, 2005.

Felix Berkenkamp, Riccardo Moriconi, Angela P Schoellig, and Andreas Krause. Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 4661–4666. IEEE, 2016.

Ruxandra Bobiti and Mircea Lazar. A sampling approach to finding lyapunov functions for nonlinear discrete-time systems. In *2016 European Control Conference (ECC)*, pages 561–566. IEEE, 2016.

Francesco Borrelli, Alberto Bemporad, Michael Fodor, and Davor Hrovat. An mpc/hybrid system approach to traction control. *IEEE Transactions on Control Systems Technology*, 14(3):541–552, 2006.

Stephen Boyd and Lieven Vandenberghe. Localization and cutting-plane methods. *From Stanford EE 364b lecture notes*, 2007.

Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

Valentina Breschi, Dario Piga, and Alberto Bemporad. Piecewise affine regression via recursive multiple least squares and multicategory discrimination. *Automatica*, 73:155–162, 2016.

Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.

Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural lyapunov control. *arXiv preprint arXiv:2005.00611*, 2020.

Shaoru Chen, Mahyar Fazlyab, Manfred Morari, George J. Pappas, and Victor M. Preciado. Learning lyapunov functions for piecewise affine systems with neural network controllers, 2020.

Frank J Christophersen, Mato Baotić, and Manfred Morari. Optimal control of piecewise affine systems: A dynamic programming approach. In *Control and Observer Design for Nonlinear Finite and Infinite Dimensional Systems*, pages 183–198. Springer, 2005.

Hongkai Dai, Benoit Landry, Lujie Yang, Marco Pavone, and Russ Tedrake. Lyapunov-stable neural-network control. *arXiv preprint arXiv:2109.14152*, 2021.

Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research*, 17(1):2909–2913, 2016.

Yingwei Du, Fangzhou Liu, Jianbin Qiu, and Martin Buss. Online identification of piecewise affine systems using integral concurrent learning. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 68(10):4324–4336, 2021.

Jack Elzinga and Thomas G Moore. A central cutting plane algorithm for the convex programming problem. *Mathematical Programming*, 8(1):134–145, 1975.

Milad Farsi and Jun Liu. Structured online learning-based control of continuous-time nonlinear systems. *IFAC-PapersOnLine*, 53(2):8142–8149, 2020.

Milad Farsi and Jun Liu. A structured online learning approach to nonlinear tracking with unknown dynamics. In *2021 American Control Conference (ACC)*, pages 2205–2211. IEEE, 2021.

Milad Farsi and Jun Liu. Structured online learning for low-level control of quadrotors. In *2022 American Control Conference (ACC)*. IEEE, 2022.

Milad Farsi, Yinan Li, Ye Yuan, and Jun Liu. A piecewise learning framework for control of unknown nonlinear systems with stability guarantees. 2022. doi: 10.48550/ARXIV.2204.02395. URL https://arxiv.org/abs/2204.02395.

Giancarlo Ferrari-Trecate, Marco Muselli, Diego Liberati, and Manfred Morari. A clustering technique for the identification of piecewise affine systems. *Automatica*, 39(2):205–217, 2003.

Claudio Gambella, Bissan Ghaddar, and Joe Naoum-Sawaya. Optimization problems for machine learning: A survey. *European Journal of Operational Research*, 290(3):807–828, 2021.

Andrea Garulli, Simone Paoletti, and Antonio Vicino. A survey on switched and piecewise affine system identification. *IFAC Proceedings Volumes*, 45(16):344–355, 2012.

Tobias Geyer, Georgios Papafotiou, and Manfred Morari. Hybrid model predictive control of the step-down dc–dc converter. *IEEE Transactions on Control Systems Technology*, 16(6):1112–1124, 2008.

Jean-Louis Goffin and Jean-Philippe Vial. On the computation of weighted analytic centers and dual ellipsoids with the projective algorithm. *Mathematical Programming*, 60(1):81–92, 1993.

Gurobi Optimizer Gurobi. Reference manual, gurobi optimization, 2020.

Rushikesh Kamalapurkar, Patrick Walters, Joel Rosenfeld, and Warren Dixon. Model-based reinforcement learning for approximate optimal control. In *Reinforcement Learning for Optimal Feedback Control*, pages 99–148. Springer, 2018.

Frank L Lewis and Draguna Vrabie. Reinforcement learning and adaptive dynamic programming for feedback control. *IEEE circuits and systems magazine*, 9(3):32–50, 2009.

XY Liu, Stefano Alfi, and Stefano Bruni. An efficient recursive least square-based condition monitoring approach for a rail vehicle suspension system. *Vehicle System Dynamics*, 54(6):814–830, 2016.

Lennart Ljung and Torsten Söderström. *Theory and practice of recursive identification*. MIT press, 1983.

Tobia Marcucci and Russ Tedrake. Mixed-integer formulations for optimal control of piecewise-affine systems. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, pages 230–239, 2019.

Tobia Marcucci, Robin Deits, Marco Gabiccini, Antonio Bicchi, and Russ Tedrake. Approximate hybrid model predictive control for multi-contact push recovery in complex environments. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 31–38. IEEE, 2017.

Yu Nesterov. Cutting plane algorithms from analytic centers: efficiency estimates. *Mathematical Programming*, 69(1):149–176, 1995.

Romain Pepy, Alain Lambert, and Hugues Mounier. Path planning using a dynamic vehicle model. In *2006 2nd International Conference on Information & Communication Technologies*, volume 1, pages 781–786. IEEE, 2006.

Steffen Rebennack and Vitaliy Krasko. Piecewise linear function fitting via mixed-integer linear programming. *INFORMS Journal on Computing*, 32(2):507–530, 2020.

Luis Rodrigues and Stephen Boyd. Piecewise-affine state feedback for piecewise-affine slab systems using convex optimization. *Systems & Control Letters*, 54(9):835–853, 2005.

Luis Rodrigues and Jonathan P How. Observer-based control of piecewise-affine systems. *International Journal of Control*, 76(5):459–477, 2003.

Nard Strijbosch, Isaac Spiegel, Kira Barton, and Tom Oomen. Monotonically convergent iterative learning control for piecewise affine systems. *IFAC-PapersOnLine*, 53(2):1474–1479, 2020.

Xiaoqiang Sun, Houzhong Zhang, Yingfeng Cai, Shaohua Wang, and Long Chen. Hybrid modeling and predictive control of intelligent vehicle longitudinal velocity considering nonlinear tire dynamics. *Nonlinear Dynamics*, 97(2):1051–1066, 2019.

Alejandro Toriello and Juan Pablo Vielma. Fitting piecewise linear continuous functions. *European Journal of Operational Research*, 219(1):86–95, 2012.

Cristina Vlad, Pedro Rodriguez-Ayerbe, Emmanuel Godoy, and Pierre Lefranc. Explicit model predictive control of buck converter. In *2012 15th International Power Electronics and Motion Control Conference (EPE/PEMC)*, pages DS1e–4. IEEE, 2012.

Fei-Yue Wang, Huaguang Zhang, and Derong Liu. Adaptive dynamic programming: An introduction. *IEEE computational intelligence magazine*, 4(2):39–47, 2009.

Lifu Wu, Xiaojun Qiu, Ian S Burnett, and Yecai Guo. A recursive least square algorithm for active control of mixed noise. *Journal of Sound and Vibration*, 339:1–10, 2015.

Ye Yuan, Xiuchuan Tang, Wei Zhou, Wei Pan, Xiuting Li, Hai-Tao Zhang, Han Ding, and Jorge Goncalves. Data driven discovery of cyber physical systems. *Nature communications*, 10(1):1–9, 2019.

Yuanyuan Zou and Shaoyuan Li. Robust model predictive control for piecewise affine systems. *Circuits, Systems & Signal Processing*, 26(3):393–406, 2007.