

Technical report 25-007

Learning-Based Model Predictive Control for Piecewise Affine Systems with Feasibility Guarantees*

S. Mallick, A. Dabiri, and B. De Schutter

To cite this work, please refer to the published version:

S. Mallick, A. Dabiri, and B. De Schutter, “Learning-based model predictive control for piecewise affine systems with feasibility guarantees,” *Proceedings of the 2025 European Control Conference*, Thessaloniki, Greece, pp. 345–350, June 2025. doi:[10.23919/ECC65951.2025.11187256](https://doi.org/10.23919/ECC65951.2025.11187256)

Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
phone: +31-15-278.24.73 (secretary)
URL: <https://www.dcsc.tudelft.nl>

* This report can also be downloaded via <https://dpub.eu/25-007>

Learning-Based Model Predictive Control for Piecewise Affine Systems with Feasibility Guarantees

Samuel Mallick, Azita Dabiri, Bart De Schutter

Abstract—Online model predictive control (MPC) for piecewise affine (PWA) systems requires the online solution to an optimization problem that implicitly optimizes over the switching sequence of PWA regions, for which the computational burden can be prohibitive. Alternatively, the computation can be moved offline using explicit MPC; however, the online memory requirements and the offline computation can then become excessive. In this work we propose a solution in between online and explicit MPC, addressing the above issues by partially dividing the computation between online and offline. To solve the underlying MPC problem, a policy, learned offline, specifies the sequence of PWA regions that the dynamics must follow, thus reducing the complexity of the remaining optimization problem that solves over only the continuous states and control inputs. We provide a condition, verifiable during learning, that guarantees feasibility of the learned policy’s output, such that an optimal continuous control input can always be found online. Furthermore, a method for iteratively generating training data offline allows the feasible policy to be learned efficiently, reducing the offline computational burden. A numerical experiment demonstrates the effectiveness of the method compared to both online and explicit MPC.

I. INTRODUCTION

Traditional online model predictive control (MPC) for piecewise affine (PWA) systems requires the online solution to an optimization problem that optimizes over, not only the continuous states and control actions, but also implicitly over the sequence of PWA regions. Through the introduction of auxiliary variables this problem can be reformulated as a mixed-integer linear/quadratic program (MILP/MIQP) [1]. Unfortunately, MILP/MIQPs are NP-hard problems, and in the context of MPC the computational burden can grow exponentially with the size of the system and the length of the prediction horizon [2], often limiting a real-time implementation.

Explicit MPC methodology has been proposed to address this issue, shifting the computational burden offline [3]. Here, a multiparametric MILP/MIQP (mp-MILP/MIQP) is solved offline, characterizing the optimal control input as a function of the system state. This solution takes the form of a partition of the state space, with a PWA control law associated to each region of the partition, such that the online computation reduces to a lookup table and a function evaluation. However, the computational complexity of the

mp-MILP/MIQP is such that this approach is in general limited to small-scale systems. Furthermore, the solution’s partition can be exceedingly complex, requiring significant online memory space to store the regions.

As an alternative to offline MPC, machine learning methods have been employed to address the computational burden of mixed-integer-based MPC. In [4] a learned policy provides warm-start guesses for an MIQP solver, in order to decrease the required solution time; however, in the worst case the computational burden remains exponential in the problem size. In [5], [6], [7], [8] a policy is learned that selects and fixes the configuration of integer variables prior to solving the MPC optimization problem, which is then convex, and far less demanding to solve. However, these approaches do not guarantee that the remaining optimization problem is feasible after the integer variables are chosen.

In light of the above issues, this work presents a learning-based MPC controller for PWA systems in which the feasibility of the learned policy’s output is guaranteed. We propose an approach that follows the principle of [5], [6], [7], [8] in decoupling the selection of discrete components, i.e., the sequence of PWA regions, and the continuous components, i.e., the state and control input trajectories, into a learned policy and a linear MPC optimization problem, respectively. Leveraging the properties of PWA dynamics, we formulate a class of classifiers whose output can be verified for feasibility during learning. Furthermore, we introduce an iterative procedure for generating training data such that the feasible policy is learned efficiently.

The contributions of this work are as follows. As only a linear MPC problem is solved online, relative to the original online MPC problem the proposed approach has a much lower online computational burden. Furthermore, the online memory requirement for the learned policy is significantly less than that of offline MPC, as learning the sequence of PWA regions is less complex than learning the optimal control signal. Finally, in contrast to existing approaches, the learned policy provably selects a feasible output.

The remainder of this paper is organized as follows. In Section II the problem setting is defined. In Section III useful structural properties of the MPC optimization problem are provided. Section IV introduces our proposed controller, which is then demonstrated in a numerical example in Section V. Finally, Section VI concludes the paper.

A. Definitions

We define a *polytope* as the intersection of a finite number of (open or closed) halfspaces, which is then convex. A

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 101018826 - ERC Advanced Grant CLariNet).

The authors are with the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands, {s.h.mallick, a.dabiri, b.deschutter}@tudelft.nl

collection of sets P_1, \dots, P_l forms a *partition* of the set P if $\bigcup_{i=1}^l P_i = P$, $P_i \neq \emptyset$ for all i , and $P_i \cap P_j = \emptyset$ for all $i \neq j$. If each set P_i is a polytope they form a *convex partition* of P . For a set \mathcal{X} we denote its convex hull as $\text{Conv}(\mathcal{X})$, and its closure as $\bar{\mathcal{X}}$. The set of vertices of a polytope P is denoted as $\text{Vert}(P)$. Finally, indexing via subscripts represents closed-loop time steps, e.g., x_k , while brackets represent time steps within an MPC prediction horizon, e.g., $x(k)$.

II. PROBLEM SETTING

Consider the discrete-time PWA system with state $x \in \mathcal{X} \subseteq \mathbb{R}^n$, input $u \in \mathcal{U} \subseteq \mathbb{R}^m$, and affine dynamics over a convex partition $\{P_i\}_{i=1}^l$ of \mathcal{X} with l polytopes indexed by $i \in \mathcal{L} = \{1, \dots, l\}$, referred to as regions,

$$x(k+1) = A_i x(k) + B u(k) + c_i, \quad x(k) \in P_i. \quad (1)$$

The states and inputs are constrained to the sets \mathcal{X} and \mathcal{U} , respectively.

Assumption 1. The dynamics are assumed to be continuous, i.e., $A_i x + B u + c_i = A_j x + B u + c_j$ for all $x \in \bar{P}_i \cap \bar{P}_j$, such that the dynamics can be considered over the closures $\{\bar{P}_i\}_{i \in \mathcal{L}}$ without the need for a set-valued function. The sets \mathcal{X} and \mathcal{U} are assumed to be compact polytopes with the origin in their interior. There exists a terminal set $\mathcal{X}_f \subseteq \mathcal{X}$ that is a polytope, assumed to contain the origin in its interior, and to be positive invariant under some linear control laws, i.e., there exists $K_i \in \mathbb{R}^{m \times n}$, for all i such that $P_i \cap \mathcal{X}_f \neq \emptyset$, such that $(A_i + B K_i)x \in \mathcal{X}_f$ for all $x \in P_i \cap \mathcal{X}_f$.

We consider an MPC controller for the system (1). Define $\delta = (\delta(0), \dots, \delta(N)) \in \mathbb{Z}_{[0,l]}^{N+1}$ as a switching sequence that specifies the PWA regions over a prediction horizon of N . The control input at state x is then computed via the following optimization problem:

$$J(x) = \min_{\mathbf{x}, \mathbf{u}, \delta} J(\mathbf{x}, \mathbf{u}) \quad (2a)$$

$$\text{s.t. } x(0) = x, \quad (2b)$$

$$x(k+1) = A_{\delta(k)} x(k) + B u(k) + c_{\delta(k)} \quad (2c)$$

$$\text{for } k = 0, \dots, N-1,$$

$$x(k) \in \bar{P}_{\delta(k)} \text{ for } k = 0, \dots, N, \quad (2d)$$

$$(x(k), u(k)) \in \mathcal{X} \times \mathcal{U} \text{ for } k = 0, \dots, N-1, \quad (2e)$$

$$x(N) \in \mathcal{X}_f, \quad (2f)$$

where $\mathbf{x} = (x^\top(0), \dots, x^\top(N))^\top$ and $\mathbf{u} = (u^\top(0), \dots, u^\top(N-1))^\top$. In (2d) the closure $\bar{P}_{\delta(k)}$ can be used as the dynamics are continuous.

If no solution exists for problem (2), by convention $J(x) = \infty$. Define the set of feasible states $\mathcal{X}_0 = \{x | J(x) < \infty\}$. Once (2) is solved online numerically, the first of the optimal control inputs $u^*(0)$ is applied to the system, with the problem resolved again at each time step in a receding horizon fashion. However, due to optimizing over the switching sequence δ , problem (2) is non-convex, and may be computationally intensive to solve. We hence introduce

an alternative MPC controller as a function of both x and δ

$$J(x, \delta) = \min_{\mathbf{x}, \mathbf{u}} J(\mathbf{x}, \mathbf{u}), \quad \text{s.t. (2b) - (2f)}, \quad (3)$$

where again $J(x, \delta) = \infty$ if no solution exists. With δ prespecified, the dynamics (2c) are linear time-varying and, for convex J , problem (3) is a convex problem that can be solved efficiently online. Define the set of feasible states for a given δ , with $\delta(0) = i$, as $\mathcal{X}_{0,\delta} = \{x | J(x, \delta) < \infty\} \subseteq \bar{P}_i \cap \mathcal{X}_0$. As (3) includes polytopic constraints and linear dynamics, $\mathcal{X}_{0,\delta}$ is a polytope [9]. In the following we address how δ is prespecified, and how it is ensured that $J(x, \delta) < \infty$ if $J(x) < \infty$.

III. STRUCTURE OF $J(x, \delta)$

In this section we analyze the structure of the optimal switching sequence $\delta^*(x) = \arg \min_{\delta} J(x, \delta)$, and provide a result on the feasibility of choices for δ , i.e., when $J(x, \delta) < \infty$, that is useful in the sequel. Note that $\delta^*(x)$ is not necessarily unique due to the closure \bar{P}_δ in (2d), e.g., when $x \in \bar{P}_i \cap \bar{P}_j$ then there are at least two valid $\delta^*(x)$ values, with i and j as the first elements, respectively.

Consider cost functions of the form

$$J(\mathbf{x}, \mathbf{u}) = \sum_{k=0}^{N-1} \left(\|Qx(k)\|_{p_1} + \|Ru(k)\|_{p_2} \right) + \|Px(N)\|_{p_3}, \quad (4)$$

where $p_1, p_2, p_3 \in \{1, \infty\}$, and Q , R , and P are of rank n, m , and n , respectively, such that (2) can be reformulated as an MILP and (3) as a linear program (LP).

A. Structure of $\delta^*(x)$

In [10], the structure of the optimal control law $u^*(0)$ for (2) is proven. Here, we use a similar line of reasoning to show the structure of $\delta^*(x)$. A graphical illustration of the following Proposition is given in Figure 1.

Proposition 1. Consider the problem $\delta^*(x) \in \arg \min_{\delta} J(x, \delta)$ for $x \in \mathcal{X}_0$. Define $\mathcal{L}_i = \left\{ \delta \in \mathbb{Z}_{[0,l]}^{N+1} \mid \delta(0) = i \right\}$. Then

- 1) $\delta^*(x) \in \mathcal{L}_i$ for all $x \in \bar{P}_i$.
- 2) For $i \in \mathcal{L}$ we can define the collection $\mathcal{R}_{i,1}, \dots, \mathcal{R}_{i,R_i}$ as a convex partition of $P_i \cap \mathcal{X}_0$, where $\delta^*(x) = \delta_{i,j}^*$ for all $x \in \bar{\mathcal{R}}_{i,j}$.
- 3) The collection $\mathcal{R}_{1,1}, \dots, \mathcal{R}_{1,R_1}, \dots, \mathcal{R}_{l,1}, \dots, \mathcal{R}_{l,R_l}$ is a convex partition of \mathcal{X}_0 .

Proof. 1) For $x \in \bar{P}_i$ there exists a $\delta^*(x)$ for which the first element must be i by constraints (2b) and (2d) for $k = 0$. Hence, $\delta^*(x) \in \mathcal{L}_i$ for all $x \in \bar{P}_i$. 2) As states can be feasible for multiple choices of δ , the sets $\mathcal{X}_{0,\delta}$ can overlap, e.g., Figure 1b. An optimal $\delta^*(x)$ for a state $x \in \bar{P}_i$ where multiple switching sequences $\delta_a, \delta_b, \dots \in \mathcal{L}_i$ are feasible is found by comparing the values $J(x, \delta_a), J(x, \delta_b), \dots$, and taking the minimum. Consider the simple case of two such sets $\mathcal{X}_{0,\delta_a} \cup \mathcal{X}_{0,\delta_b} \neq \emptyset$. Then $\delta^*(x) = \delta_a$ for x in

$$\{x | \mathcal{X}_{0,\delta_a} \cup \mathcal{X}_{0,\delta_b}, J(x, \delta_a) - J(x, \delta_b) \leq 0\} \cap (\mathcal{X}_{0,\delta_a} \setminus \mathcal{X}_{0,\delta_b}), \quad (5)$$

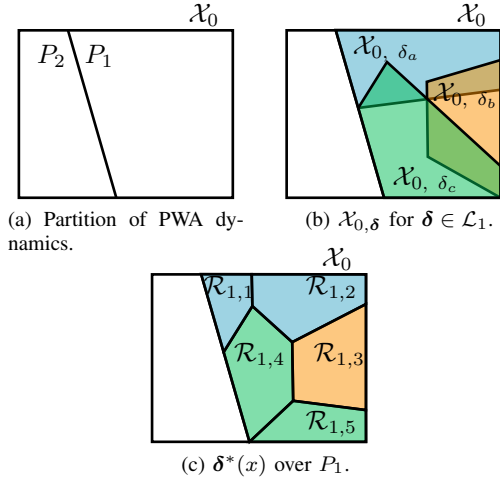


Fig. 1: Illustration of the proof of Proposition 1. $\mathcal{R}_{1,1}$ and $\mathcal{R}_{1,2}$ are associated with $\delta^*(x) = \delta_a$, $\mathcal{R}_{1,3}$ is associated with $\delta^*(x) = \delta_b$, and $\mathcal{R}_{1,4}$ and $\mathcal{R}_{1,5}$ are associated with $\delta^*(x) = \delta_c$.

which, as $J(x, \delta_a)$ and $J(x, \delta_b)$ are piecewise-linear functions of x [9], can be represented as the union of a finite number of polytopes. Generalizing this to many overlapping sets, the region defining where a given sequence is optimal can be decomposed into a finite number of polytopes, i.e., the regions $\mathcal{R}_{i,j}$. with the definition of \mathcal{L}_i we have $\bigcup_{\delta \in \mathcal{L}_i} \mathcal{X}_{0,\delta} = \bar{P}_i \cap \mathcal{X}_0$, e.g., Figure 1c. It follows that the collection $\mathcal{R}_{i,1}, \dots, \mathcal{R}_{i,\bar{R}_i}$ is a convex partition of $P_i \cap \mathcal{X}_0$. 3) Finally, as $\bigcup_{i \in \mathcal{L}} (P_i \cap \mathcal{X}_0) = \mathcal{X}_0$, the collection $\mathcal{R}_{1,1}, \dots, \mathcal{R}_{1,\bar{R}_1}, \dots, \mathcal{R}_{l,1}, \dots, \mathcal{R}_{l,\bar{R}_l}$ is a convex partition of \mathcal{X}_0 . \square

This result informs us that $\delta^*(x)$ can be characterized by a convex partition, where each polytope in the partition is associated with a value for $\delta^*(x)$. Furthermore, as \mathcal{X}_0 can be expressed as a union of polytopes, $\mathcal{X} \setminus \mathcal{X}_0$ can also be expressed as a union of polytopes, such that by learning a convex partition over \mathcal{X} we can represent both \mathcal{X}_0 and $\delta^*(x)$ over \mathcal{X}_0 .

B. Feasibility of $J(x, \delta)$

Proposition 2. For any δ and the set of states $\tilde{\mathcal{X}} \subseteq \mathcal{X}_0$ such that $J(x, \delta) < \infty$ for all $x \in \tilde{\mathcal{X}}$, we have $J(x, \delta) < \infty$ for all $x \in \text{Conv}(\tilde{\mathcal{X}})$.

Proof. Assume there exists $x \in \text{Conv}(\tilde{\mathcal{X}})$ such that $J(x, \delta) = \infty$. Then the set $\mathcal{X}_{0,\delta}$ is non-convex. This is a contradiction; hence, $J(x, \delta) < \infty$ for all $x \in \text{Conv}(\tilde{\mathcal{X}})$. \square

Proposition 2 provides a tool for checking the feasibility of a given δ over an infinite number of states by checking feasibility over a finite number of states, i.e., given a polytope $\mathcal{R} \subseteq \mathcal{X}_0$, the feasibility of δ for the entire polytope can be confirmed by confirming that $J(x, \delta) < \infty$ for all $x \in \text{Vert}(\mathcal{R})$.

IV. LEARNING-BASED MPC CONTROLLER

In this section we formulate our learning-based MPC solution. We propose to learn a policy $\pi_\theta(x) = \delta$ offline in a supervised-learning manner, where θ is a vector of parameters to be learned, and pairs of training data $(x, \delta^*(x))$ are generated by solving (2). Then, for deployment online, only the LP (3) must be solved for $J(x, \pi_\theta(x))$.

A. Classifier

Proposition 1 suggests that π_θ should be a classifier that can effectively approximate a convex partition, while Proposition 2 provides a tool to verify feasibility, provided π_θ partitions its outputs into polytopes. In light of this, we formulate the following structure that defines a class of suitable classifiers: $\pi_\theta(x) = \phi_\theta(x)(f_{\theta,i}(x))$ for $x \in P_i$, with

$$\begin{aligned} f_{\theta,i} : P_i &\rightarrow \{\mathcal{R}_{i,1}, \dots, \mathcal{R}_{i,\bar{R}_i}\} \\ \phi_{\theta,i} : \{\mathcal{R}_{i,1}, \dots, \mathcal{R}_{i,\bar{R}_i}\} &\rightarrow \mathcal{L}_i \cup \{-1\}, \end{aligned} \quad (6)$$

where each collection $\mathcal{R}_{i,1}, \dots, \mathcal{R}_{i,\bar{R}_i}$ is a convex partition of P_i . The function $\phi_{\theta,i}$ maps the polytopes to either a switching sequence δ , or to -1 representing infeasibility, i.e., $x \notin \mathcal{X}_0$. For convenience, denote the codomains of $f_{\theta,i}$ and $\phi_{\theta,i}$ as $\mathcal{F}_{\theta,i}$ and Φ_i , respectively. Note that the number of polytopes \bar{R}_i depends on the training data.

Many classifiers can satisfy this structure, e.g., decision trees, oblique decision trees, neural networks with linear activation functions, and support vector machines [11]. In [12] a classifier tailored to learning functions that are PWA across polytopes is presented, alongside a thorough review of alternative approaches.

B. Iterative Training Procedure

Learning π_θ in a supervised manner requires the solution to problem (2) in order to generate training data. As solving (2) is computationally demanding, it is desirable to train π_θ with as little data as possible. Furthermore, given the structure of $\delta^*(x)$ outlined in Section III, it is clear that some states are more informative than others. In particular, states near the boundaries of the polytopes forming the convex partition are more useful than those in the interiors.

Addressing both these points, Algorithm 1 details the offline training procedure. In step 3, $\text{train}(\mathcal{T})$ is pseudo-code for a supervised learning algorithm that trains the classifier given a set of labeled data $\mathcal{T} = \{(x^{(i)}, \delta^*(x^{(i)}))\}_i^{N_{\text{data}}}$, and is specific to the form of π_θ . Also specific to the form of π_θ is the presence and degree of classifier misclassification, i.e., how often π_θ chooses the incorrect switching sequence for a state in the training data \mathcal{T} , which can aid in finding a feasible π_θ with a simple partition, i.e., with less polytopes, but can increase suboptimality. Beginning from an initial training set, π_θ is trained iteratively, where at each iteration, all vertices of all learned polytopes are checked for feasibility. Only vertices that are infeasible are added to the training set, naturally concentrating the training data in the areas of high information gain, i.e., at the boundaries of the polytopes of the partition.

Algorithm 1 Train π_θ (offline)

```
1: Inputs: Initial training set  $\mathcal{T} \subset \mathcal{X} \times \bigcup_{i \in \mathcal{L}} \Phi_i$ 
2: while true do
3:    $\theta \leftarrow \text{train}(\mathcal{T})$ 
4:   feasible  $\leftarrow 1$ 
5:    $\mathcal{Z} \leftarrow \{\}$ 
6:   for  $i \in \mathcal{L}$  do
7:     for  $\mathcal{R} \in \mathcal{F}_{\theta,i}$  do
8:       for  $x \in \text{Vert}(\mathcal{R})$  do
9:         if  $(\phi_{\theta,i}(\mathcal{R}) = -1 \text{ and } J_i(x) < \infty)$  or
            $(\phi_{\theta,i}(\mathcal{R}) \neq -1 \text{ and } J(x, \pi_\theta(x)) = \infty)$  then
10:            $\mathcal{Z} \leftarrow \mathcal{Z} \cup \{x\}$ 
11:           feasible  $\leftarrow 0$ 
12:         end if
13:       end for
14:     end for
15:   end for
16:   if feasible = 1 then
17:     Return:  $\pi_\theta$ 
18:   else
19:     for  $x \in \mathcal{Z}$  do
20:       Solve (2) for  $J(x)$  and  $\delta^*$ 
21:       if  $J(x) < \infty$  then
22:          $\mathcal{T} \leftarrow \mathcal{T} \cup \{(x, \delta^*)\}$ 
23:       else
24:          $\mathcal{T} \leftarrow \mathcal{T} \cup \{(x, -1)\}$ 
25:       end if
26:     end for
27:   end if
28: end while
```

Remark 1. Clearly the number of iterations and the amount of data required for convergence of Algorithm 1 depends on the underlying structure of $\delta^*(x)$, which in general grows in complexity with N and l . An analysis of the scalability of Algorithm 1 with N and l is left for future work.

C. Approximating \mathcal{X}_0

In order to approximate the feasible region \mathcal{X}_0 , the boundary of \mathcal{X}_0 is handled as a special case. As the PWA dynamics are continuous, several switching sequences can be feasible for a given state, e.g., $x \in \bar{P}_i \cap \bar{P}_j$, such that the shared vertices of adjacent polytopes can be feasible for the switching sequence of each polytope. However, it is not the case that the shared vertices of polytopes can be used to verify feasibility and infeasibility, as this would imply both $J(x) < \infty$ and $J(x) = \infty$ at the boundary of \mathcal{X}_0 . In light of this, we introduce the tightened version of (2):

$$\begin{aligned} J_l(x) &= \min_{\mathbf{x}, \mathbf{u}, \delta} J(\mathbf{x}, \mathbf{u}) \\ \text{s.t.} \quad & (2b) - (2d) \\ & (x(k), u(k)) \in (\mathcal{X} \ominus \mathcal{O}) \times \mathcal{U} \\ & \text{for } k = 0, \dots, N-1, \\ & x(N) \in \mathcal{X}_f \ominus \mathcal{O}, \end{aligned}$$

where \mathcal{O} is an arbitrarily small set containing the origin. Again, if no solution exists for (7) then $J_l(x) = \infty$. The feasible set for (7) is contained within \mathcal{X}_0 , such that on the boundary of \mathcal{X}_0 we have $J(x) < \infty$ and $J_l(x) = \infty$. In Algorithm 1, when verifying that the vertices of a polytope labeled with -1 are infeasible, in step 9, we then use $J_l(x)$. The result is that an under-approximation of \mathcal{X}_0 is learned. However the approximation error can be made arbitrarily small by reducing \mathcal{O} . Define this feasible set, representing the valid region for π_θ , as $\mathcal{X}_{0,\theta} = \{x | \pi_\theta(x) \neq -1\}$.

Proposition 3. For π_θ trained with Algorithm 1, $J(x, \pi_\theta(x)) < 0$ for all $x \in \mathcal{X}_{0,\theta}$.

Proof. For π_θ trained with Algorithm 1 it holds, for all $i \in \mathcal{L}$ and for all $\mathcal{R} \in \mathcal{F}_{\theta,i}$ such that $\phi(\mathcal{R}) \neq -1$, that $J(x, \pi_\theta(x)) < \infty$ for all $x \in \text{Vert}(\mathcal{R})$ and hence, by Proposition (2), for all $x \in \bar{\mathcal{R}}$. As, by the definition (6), the collection of sets \mathcal{R} for all $\phi_{\theta,i}(\mathcal{R}) \neq -1$ is a convex partition of $P_i \cap \mathcal{X}_{0,\theta}$, it holds that $J(x, \pi_\theta(x)) < 0$ for all $x \in \mathcal{X}_{0,\theta}$. \square

Remark 2. In Algorithm 1, an MILP is solved only when generating training data at the infeasible vertices, i.e., step 20. Checking for $J_l(x) < \infty$ in step 9 involves checking for the existence of a feasible solution only, which is much easier than solving the problem to optimality.

D. Control Law

With π_θ we now define the controller in Algorithm 2. The check in step 2 is required as $\mathcal{X}_{0,\theta}$ can be an under-approximation of \mathcal{X}_0 , such that a feasible trajectory could in theory pass out of $\mathcal{X}_{0,\theta}$ before entering \mathcal{X}_f . We now prove recursive feasibility of the controller.

Proposition 4. For $x_0 \in \mathcal{X}_{0,\theta}$, for the closed-loop system and with δ_k chosen as by Algorithm 2, we have $J(x_k, \delta_k) < \infty$ for all $k \geq 0$.

Proof. Assume that at time step k we have $J(x_k, \delta_k) < \infty$. Applying $u_k^*(0)$, obtained from (3), will propagate the system to $x_{k+1} = x_k^*(1)$. If $\pi_\theta(x_{k+1}) = -1$, then the shifted solution $\bar{u}_k = (u_k^*(1), \dots, u_k^*(N-1), K_j x_k^*(N))$ is feasible for problem (3) with the shifted sequence $\delta_{k+1} = (\delta_k(1), \dots, \delta_k(N), i)$ and with $(A_j + BK_j)x_k^*(N) \in P_i$ and $x_k^*(N) \in P_j$, as \mathcal{X}_f is forward invariant under the linear controllers. Hence, we have $J(x_{k+1}, \delta_{k+1}) < \infty$. If $\pi_\theta(x_{k+1}) \neq -1$, then $\delta_{k+1} = \pi_\theta(x_{k+1})$ and $J(x_{k+1}, \delta_{k+1}) < \infty$ by Proposition 3. Therefore, if $J(x_k, \delta_k) < \infty$ at time step k then we have $J(x_{k+1}, \delta_{k+1}) < \infty$ at time step $k+1$. For $x_0 \in \mathcal{X}_{0,\theta}$ we have $J(x_0, \delta_0) < \infty$ again by Proposition 3, and by induction $J(x_k, \delta_k) < \infty$ for all $k \geq 0$. \square

V. NUMERICAL EXAMPLE

Consider a representative numerical PWA system (1) with $l = 2$,

$$A_1 = \begin{bmatrix} 1 & 0.2 \\ 0 & 1 \end{bmatrix} \quad A_2 = \begin{bmatrix} 0.5 & 0.2 \\ 0 & 1 \end{bmatrix} \quad c_1 = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad c_2 = \begin{bmatrix} 0.5 \\ 0 \end{bmatrix}, \quad (7)$$

Algorithm 2 Controller at time step k (online)

-
- 1: **Inputs:** Current state x_k , previous switching sequence δ_{k-1} (except if $k = 0$)
 - 2: **if** $k = 0$ or $\pi_\theta(x_k) \neq -1$ **then**
 - 3: $\delta_k \leftarrow \pi_\theta(x_k)$
 - 4: **else**
 - 5: $\delta_k \leftarrow (\delta_{k-1}(1), \dots, \delta_{k-1}(N), i)$ with $(A_j + BK_j)x_{k-1}^*(N) \in P_i$ and $x_{k-1}^*(N) \in P_j$
 - 6: **end if**
 - 7: Solve $J(x_k, \delta_k)$ and apply $u^*(0)$
-

$B = [0.1 \ 1]^\top$, and $P_{1(2)} = \{x \mid [1 \ 0]x \leq (>)1\}$. The input and state constraints are $\mathcal{U} = \{u \mid |u| \leq 3\}$ and $\mathcal{X} = \{x \mid Dx \leq E\}$, respectively, with

$$D = \begin{bmatrix} -1 & -3 & 0.2 & -1 & 1 & 0 \\ 1 & -1 & 1 & 0 & 0 & -1 \end{bmatrix}^\top \quad (8)$$

$$E = [15 \ 25 \ 9 \ 6 \ 8 \ 10]^\top.$$

Consider the MPC controller (2) with cost matrices $Q = P = \mathbf{I}_2$, $R = 1$, and using the 1-norm, i.e., $p_1 = p_2 = p_3 = 1$. The terminal set is computed as the maximal constraint admissible set for the system $x(k+1) = (A_1 + BK)x$, restricted to P_1 , where K is the optimal LQR gain controller for (A_1, B, Q, R) . The example is simulated on an 11th Gen Intel laptop with four i7 cores, 3.00GHz clock speed, and 16Gb of RAM. For fairness, all MILP and LP optimization problems are solved with Gurobi [13] without warm-starting, while explicit MPC controllers are generated with the MPT3 toolbox [14]. Source code is available at <https://github.com/SamuelMallick/supervised-learning-pwa-mpc>. For π_θ we use an ensemble of the classifier presented in [12], with one for each P_i , with all tunable hyperparameters for the classifier available in the source code.

Algorithm 1 is used to train π_θ using an initial training set $\mathcal{T} = \{(x^{(i)}, \delta^*(x^{(i)}))\}_{i=1}^{45} \cup \{(x^{(j)}, \delta^*(x^{(j)}))\}_{j=1}^{45}$ with $x^{(i)}$ and $x^{(j)}$ sampled uniformly from P_1 and P_2 , respectively. The tightened MPC problem (7) is implemented with $\mathcal{O} = \{x \mid \|x\|_\infty < 0.1\}$. Figure 2 demonstrates Algorithm 1 at iterations 0, 8 and 42, for $N = 12$. It can be seen that across the iterations, training data is added around areas of importance, allowing a feasible π_θ to be learned efficiently. At the first iteration training data is sampled randomly, and the partition is poor, with many infeasible vertices. By the eighth iteration, following Algorithm 1, more points have been sampled at the infeasible vertices of the previous iterations, and the partition is significantly improved. Finally, at iteration 42 the partition is validated to be feasible for all $x \in \mathcal{X}_{0,\theta}$, and the training is terminated.

Table I shows the number of regions in the learned convex partition of \mathcal{X} and the number of iterations required to guarantee feasibility, as the horizon N increases from 5 to 12. Additionally, the number of training data generated is reported, normalized by that for $N = 5$, indicating how

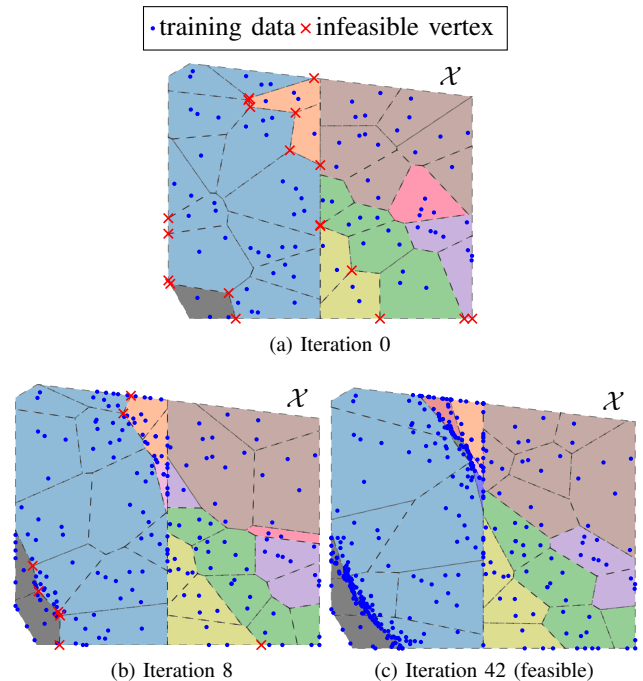


Fig. 2: Training for $N = 12$. Each color represents a different δ , while the dark gray region represents infeasibility.

TABLE I: Training π_θ : Number of training data normalized by the value for $N = 5$, with true values in brackets.

N	# regions	# iterations	# training data
5	29	34	1 (377)
6	19	46	1.170 (441)
7	19	56	1.191 (449)
8	26	26	0.801 (302)
9	27	33	0.889 (335)
10	24	52	1.016 (383)
11	25	42	0.939 (354)
12	30	42	1.053 (397)

the offline computation scales with N . For comparison, an explicit MPC controller [3] is computed for each N . This involves solving an mp-MILP, resulting in a partition of \mathcal{X} that defines the optimal control law. Table II shows the number of regions required by the explicit MPC approach, and the computation time required to compute the controller, again normalized by the time for $N = 5$. For π_θ , both the number of regions and offline computation are relatively uncorrelated with N , while for explicit MPC the number of regions and the offline computation time scale poorly with N .

To investigate the performance of π_θ , the open-loop cost $J(x, \pi_\theta(x))$ is compared against the optimal $J(x)$ with $N = 12$. Figure 3 shows the percentage suboptimality: $\Delta J = 100 \cdot (J(x, \pi_\theta(x)) - J(x)) / J(x)$, for 100709 states sampled densely from $\mathcal{X}_{0,\theta}$. Suboptimality is only introduced around the boundaries of the convex partition in Figure 2c, demonstrating that π_θ learns a convex partition close to the

TABLE II: Explicit MPC: Computation time normalized by the value for $N = 5$, with true values in brackets.

N	# regions	comp. time (s)	
5	388	1	(25.189)
6	469	5.908	(148.807)
7	520	14.781	(372.299)
8	549	40.761	(1.027·10 ³)
9	574	150.494	(3.791·10 ³)
10	600	464.185	(1.169·10 ⁴)
11	741	1.307·10 ³	(3.292·10 ⁴)
12	2541	6.965·10 ³	(1.754·10 ⁵)




Fig. 3: Open-loop suboptimality ΔJ for $N = 12$.

optimal one. Furthermore, for all states $J(x, \pi_\theta(x)) < \infty$, demonstrating Proposition 3.

To investigate the closed-loop performance, Algorithm 2 is compared against the optimal online MPC controller¹, where (2) is solved, for $N = 12$. Simulations are run for 1000 initial states, sampled uniformly from $\mathcal{X}_{0,\theta}$, until $\|x\|_2 < 0.01$, with the cumulative closed-loop costs compared. Table III gives the distribution of the percentage suboptimality, as well as the computation time for each approach. It can be seen that, as the proposed approach optimizes for the continuous variables online, and as $\delta = \pi_\theta(x)$ is suboptimal over only a subset of the state space (see Figure 3), the closed-loop suboptimality is limited. Furthermore, the computation time is significantly improved by Algorithm 2. Finally, for all time steps of all simulations, $J(x_k, \delta_k) < \infty$, demonstrating Proposition 4.

¹The performance of the optimal online MPC controller is the same as that of the explicit MPC controller; however, as the memory requirements for explicit MPC can be unreasonable, we use an online approach and explore its online computational burden.

TABLE III: Closed-loop comparison $N = 12$.

	mean	median	min	max	std
$\sum_k \Delta J$	0.371	0	0	35.970	1.808
Time (2) (s)	0.0323	0.0152	0.00361	0.315	0.0362
Time (3) (s)	0.00143	0.00125	0.000605	0.0517	0.00231

VI. CONCLUSIONS

In this work we have proposed a learning-based MPC controller for PWA systems where the PWA switching sequences are selected by a learned policy, such that only a linear optimization problem must be solved online. In contrast to existing works in this direction, feasibility of the controller is guaranteed. To this end a class of suitable classifiers has been formulated, and an efficient training strategy has been presented. A numerical example has demonstrated the benefits of the proposed approach in terms of offline computation and online memory with respect to offline MPC, and online computation with respect to online MPC. Future work will look at bounding the suboptimality of the proposed approach, as well as further experimentation on larger systems with more regions in the PWA dynamics.

REFERENCES

- [1] A. Bemporad and M. Morari, "Control of systems integrating logic, dynamics, and constraints," *Automatica*, vol. 35, no. 3, pp. 407–427, 1999.
- [2] L. A. Wolsey and G. L. Nemhauser, *Integer and Combinatorial Optimization*. John Wiley & Sons, 2014.
- [3] R. Oberdieck and E. N. Pistikopoulos, "Explicit hybrid model-predictive control: The exact solution," *Automatica*, vol. 58, pp. 152–159, 2015.
- [4] D. Masti and A. Bemporad, "Learning binary warm starts for multiparametric mixed-integer quadratic programming," in *2019 18th European Control Conference (ECC)*, Naples, Italy, 2019, pp. 1494–1499.
- [5] D. Masti, T. Pippia, A. Bemporad, and B. De Schutter, "Learning approximate semi-explicit hybrid MPC with an application to microgrids," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 5207–5212, 2020.
- [6] A. Cauligi, P. Culbertson, E. Schmerling, M. Schwager, B. Stellato, and M. Pavone, "CoCo: Online mixed-integer control via supervised learning," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1447–1454, 2021.
- [7] A. Cauligi, A. Chakrabarty, S. D. Cairano, and R. Quirynen, "PRISM: Recurrent neural networks and presolve methods for fast mixed-integer optimal control," *Proceedings of Machine Learning Research*, vol. 168, pp. 1–13, 2022.
- [8] C. F. O. da Silva, A. Dabiri, and B. De Schutter, "Integrating reinforcement learning and model predictive control with applications to microgrids," *arXiv preprint arXiv:2409.11267*, 2024.
- [9] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2016.
- [10] F. Borrelli, M. Baotić, A. Bemporad, and M. Morari, "Dynamic programming for constrained optimal control of discrete-time linear hybrid systems," *Automatica*, vol. 41, no. 10, pp. 1709–1721, 2005.
- [11] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*. Springer, 2006, vol. 4, no. 4.
- [12] A. Bemporad, "A piecewise linear regression and classification algorithm with application to learning and model predictive control of hybrid systems," *IEEE Transactions on Automatic Control*, vol. 68, no. 6, pp. 3194–3209, 2023.
- [13] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2024. [Online]. Available: <https://www.gurobi.com>
- [14] M. Herceg, M. Kvasnica, C. Jones, and M. Morari, "Multi-Parametric Toolbox 3.0," in *Proc. of the European Control Conference*, Zürich, Switzerland, July 2013, pp. 502–510, <http://control.ee.ethz.ch/~mpt>.