

Piecewise regression via mixed-integer programming for MPC

Dieter Teichrib

DIETER.TEICHTRIB@TU-DORTMUND.DE

Moritz Schulze Darup

MORITZ.SCHULZEDARUP@TU-DORTMUND.DE

Control and Cyberphysical Systems Group, TU Dortmund University, Germany

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

Abstract

Piecewise regression is a versatile approach used in various disciplines to approximate complex functions from limited, potentially noisy data points. In control, piecewise regression is, e.g., used to approximate the optimal control law of model predictive control (MPC), the optimal value function, or unknown system dynamics. Neural networks are a common choice to solve the piecewise regression problem. However, due to their nonlinear structure, training is often based on gradient-based methods, which may fail to find a global optimum or even a solution that leads to a small approximation error. To overcome this problem and to find a global optimal solution, methods based on mixed-integer programming (MIP) can be used. However, the known MIP-based methods are either limited to a special class of functions, e.g., convex piecewise affine functions, or they lead to complex approximations in terms of the number of regions of the piecewise defined function. Both complicate a usage in the framework of control. We propose a new MIP-based method that is not restricted to a particular class of piecewise defined functions and leads to functions that are fast to evaluate and can be used within an optimization problem, making them well suited for use in control.

Keywords: piecewise regression, mixed-integer programming, model predictive control, neural networks

1. Introduction

Piecewise regression is used in many disciplines to fit a continuous piecewise defined function to an unknown function based on a finite number of possibly noisy samples. In the context of model predictive control (MPC) for linear systems, piecewise regression is used, e.g., to approximate the piecewise affine (PWA) optimal control law (Chen et al., 2018; Karg and Lucia, 2020) or the piecewise quadratic optimal value function (OVF) (Teichrib and Schulze Darup, 2021). Both functions are continuous and defined on polyhedral regions, which partition the state space (c.f. Bemporad et al. (2002)). In addition, piecewise regression may be used to find a surrogate model for an unknown nonlinear system that is used as a prediction model in nonlinear MPC. Neural networks (NN) are very popular in piecewise regression as they allow to approximate a large class of functions with arbitrary accuracy (Hornik et al., 1989). However, due to the nonlinear structure of NN, computing NN parameters, which minimize the approximation error, requires solving a nonlinear optimization problem (OP). This OP is typically solved approximately by gradient-based methods, such as stochastic gradient descent, RMSProb, or Adam (Kingma and Ba, 2017). However, these methods are often sensitive to the choice of hyperparameters, and they may even fail to converge for some choices of the learning rate. Moreover, even if the hyperparameters are chosen properly, general convergence guarantees are lacking for the popular RMSProb and Adam algorithms. Fur-

thermore, in case of approximated system dynamics or OVF, trained NN may subsequently be used as a prediction model or to derive optimal control actions, respectively. For both applications, the NN will be included in an OP (either as constraints or costs), and an efficient inclusion is hard for some types of NN. We will consider maxout NN with a special topology that allows to solve the problem of piecewise regression globally and that is suitable for a subsequent usage in an OP.

An alternative for solving the problem of piecewise regression is the direct use of PWA functions. When using PWA functions for regression, either the regions are fixed and a linear regression is performed for each region, or the computation is split into two parts. In the first part, the data is divided into clusters that define the polyhedral regions of the PWA function. In the second part, the parameters of the function are computed by performing a linear regression for every region (see, e.g., [Magnani and Boyd \(2009\)](#)). Often, these algorithms alternate between clustering and regression until a termination criterion is met ([Bemporad, 2023](#)). The main drawback of these methods is that the regression problem is only solved approximately and the approximation quality depends on the initial clustering of the data points. To overcome this problem and to find a global optimal solution, clustering and regression should be done simultaneously. Most of the approaches where clustering and regression are formulated in a combined OP are based on mixed-integer programming (MIP). One of the first MIP-based approaches was introduced in [Roll et al. \(2004\)](#), where so-called hinging hyperplanes ([Breiman, 1993](#)) are fitted via MIP to the data. However, only PWA functions that possess the consistent variation property ([Chua and Deng, 1988](#), Def. 5) can be exactly represented by hinging hyperplanes. Since the consistent variation property typically does not hold in MPC, hinging hyperplanes are not the ideal choice. Another MIP-based method was introduced in [Bertsimas and Shioda \(2007\)](#). They use binary variables to assign points to regions, we will use similar ideas in our method. In addition, we will also use clustering as in [Bertsimas and Shioda \(2007, Sec. 3.2.\)](#) to cluster points and then assign these clusters to regions of the function. The main weakness of this method is that the function is discontinuous and it is not possible to enforce continuity for problems with $n > 1$ (c.f. [Bertsimas and Shioda \(2007, Sec. 5.3.\)](#)). Moreover, the regions of the function are not necessarily polytopes. Hence, further computations are required to find polytopes that contain the majority of the points associated with an individual region. Further MIP-based methods were recently introduced in [Rebennack and Krasko \(2020\)](#); [Warwicker and Rebennack \(2023\)](#). However, these methods can only be applied to univariate or bivariate functions. The MIP-based method from [Toriello and Vielma \(2012, Sec. 3.2.\)](#) solves some of the aforementioned problems. This method allows the fitting of a convex continuous PWA function to data via a mixed-integer quadratic program (MIQP). The main idea is to represent the convex function as the maximum of its affine segments and use binary variables to determine which affine segment provides the largest value for a given data point. The advantage of this method is that the polyhedral regions of the function are implicitly given by the structure of the function and do not need to be computed separately. However, the method is not extendable to non-convex functions. In [Siahkamari et al. \(2020\)](#), an alternative method is given, which allows to combine the benefits of [Bertsimas and Shioda \(2007\)](#) and [Toriello and Vielma \(2012\)](#), i.e., it is possible to fit non-convex continuous functions that are defined on polyhedral regions. In [Siahkamari et al. \(2020\)](#), the fitted PWA function has one region per data point, this allows to solve the piecewise regression problem by solving a quadratic program (QP) instead of an MIQP as in [Roll et al. \(2004\)](#); [Bertsimas and Shioda \(2007\)](#); [Toriello and Vielma \(2012\)](#). Unfortunately, the coupling between the number of regions and the number of data points leads to rather complex functions for large data sets, which is disadvantageous, especially when the function is used in an OP.

Contributions. The main contribution of this paper is to provide a new MIP-based method to globally solve the piecewise regression problem, which combines the advantages of the known approaches. More specifically, our new method

1. allows to freely choose the number of segments as in [Bertsimas and Shioda \(2007\)](#) without the need to compute the regions of the function explicitly,
2. has a simple representation of the fitted function as in [Toriello and Vielma \(2012\)](#) without being restricted to convex functions,
3. allows to fit a continuous function as in [Siahamari et al. \(2020\)](#) without coupling the number of regions and the number of data points,
4. results in a model that is suitable for the use in the constraints and the cost function of an OP as in [Bemporad \(2023\)](#).

We also provide two approaches for reducing the complexity of the proposed method to make it applicable to larger data sets.

2. Problem formulation and preliminaries

We consider the problem of continuous piecewise regression, i.e., our aim is to fit a continuous piecewise defined function $\Phi(\mathbf{x}; \boldsymbol{\theta}) : \mathbb{R}^n \rightarrow \mathbb{R}$ to samples of an unknown function $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$, which typically involves solving the nonlinear OP

$$\min_{\boldsymbol{\theta}} \sum_{k=1}^{N_D} \ell(y^{(k)} - \Phi(\mathbf{x}^{(k)}; \boldsymbol{\theta})), \quad (1)$$

for a given data set $\mathcal{D} := \{(\mathbf{x}^{(i)}, y^{(i)}) \in \mathcal{R}^{n+1} \mid y^{(i)} = f(\mathbf{x}^{(i)}), i \in \{1, \dots, N_D\}\}$, where $\boldsymbol{\theta}$ are the parameters of the function and ℓ is a convex loss function. In the following, we briefly summarize the methods from [Toriello and Vielma \(2012\)](#) and [Siahamari et al. \(2020\)](#), as our method is most closely related to these and contains both as special cases. In [Toriello and Vielma \(2012\)](#), the problem of piecewise affine regression is solved by the MIQP

$$\min_{\mathbf{V}, \alpha^{(i)}, \boldsymbol{\delta}^{(i)}} \sum_{i=1}^{N_D} (y^{(i)} - \alpha^{(i)})^2 \quad (2a)$$

$$\text{s.t.} \quad \mathbf{V} \begin{pmatrix} \mathbf{x}^{(i)} \\ 1 \end{pmatrix} \leq \mathbf{1}\alpha^{(i)}, \quad \mathbf{V} \begin{pmatrix} \mathbf{x}^{(i)} \\ 1 \end{pmatrix} + M(\mathbf{1} - \boldsymbol{\delta}^{(i)}) \geq \mathbf{1}\alpha^{(i)}, \quad \forall i \in \{1, \dots, N_D\}, \quad (2b)$$

$$\sum_{k=1}^{p_1} \boldsymbol{\delta}_k^{(i)} = \mathbf{1}, \quad \boldsymbol{\delta}^{(i)} \in \{0, 1\}^{p_1}, \quad \forall i \in \{1, \dots, N_D\}. \quad (2c)$$

with $\mathbf{V} \in \mathbb{R}^{p_1 \times (n+1)}$ and the subscript j refers to the j -th row of a matrix or column vector. The $N_D p_1$ binary variables select which of the p_1 affine segments $\mathbf{V}_j(\mathbf{x}^{(i)\top} \mathbf{1})^\top$ with $j \in \{1, \dots, p_1\}$ provides the largest value for a given data point. The constraints (2b) ensure that $\max\{\mathbf{V}_j \mathbf{g}(\mathbf{x}^{(i)})\} = \alpha^{(i)}$ holds, i.e., for a sufficiently large $M \in \mathbb{R}_{>0}$, the MIQP (2) solves the nonlinear OP (1) with $\ell(\cdot) := \|\cdot\|_2^2$, $\boldsymbol{\theta} := \mathbf{V}$, and $\Phi(\mathbf{x}; \boldsymbol{\theta}) = \max\left\{\mathbf{V} \begin{pmatrix} \mathbf{x} \\ 1 \end{pmatrix}\right\} := \max_{1 \leq i \leq p_1} \left\{\mathbf{V}_i \begin{pmatrix} \mathbf{x} \\ 1 \end{pmatrix}\right\}$. Thus, a convex PWA function is fitted to the data set. In contrast to [Bertsimas and Shioda \(2007\)](#), this method allows to make predictions by evaluating the function $\Phi(\mathbf{x})$ without explicitly computing the regions of the function. Moreover, the function $\Phi(\mathbf{x})$ is continuous by construction. However, the disadvantage is

that the fitted function is always convex and PWA. In our method, we also use the binary variables to select the largest segments, but we will overcome the limitation to convex functions by extending the MIQP (2).

A method that combines the ability to fit non-convex functions with the continuity property of (2) is presented in [Siahkamari et al. \(2020\)](#), where the problem of piecewise affine regression is solved by solving a QP (see [Siahkamari et al. \(2020, Eq. \(6\)\)](#)) instead of an MIQP. The constraints in the QP ensure that the function

$$\Phi(\mathbf{x}) = \max_{1 \leq i \leq N_D} \{\mathbf{a}_i^\top (\mathbf{x} - \mathbf{x}^{(i)}) + \hat{y}_i + z_i\} - \max_{1 \leq i \leq N_D} \{\mathbf{b}_i^\top (\mathbf{x} - \mathbf{x}^{(i)}) + z_i\} \quad (3)$$

is such that we have $\Phi(\mathbf{x}^{(i)}) = \hat{y}_i$ and the cost function is the sum of the residuals $(y^{(i)} - \hat{y}_i)^2$, i.e., the function (3) is fitted to the data. Since every PWA can be decomposed into the difference of two convex PWA functions and (3) is the difference of two convex PWA functions, a general PWA function is fitted to the data points in [Siahkamari et al. \(2020\)](#). Unfortunately, unlike [Bertsimas and Shioda \(2007\)](#) or [Toriello and Vielma \(2012\)](#), this method does not allow to choose the number of regions of the PWA function freely. As apparent from (3), the number of regions is always equal to the number of data points, resulting in a rather complex function for large data sets.

3. A new MIP-based method for piecewise regression

In this section, we present our new method for piecewise regression via MIP, which combines advantageous properties of the known approaches while avoiding some of their limitations. Our starting point is the continuous function $\Phi(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$ parameterized by

$$\Phi(\mathbf{x}) := \max\{\mathbf{V}\mathbf{g}(\mathbf{x})\} - \max\{\mathbf{W}\mathbf{h}(\mathbf{x})\} \quad (4)$$

with $\mathbf{V} \in \mathbb{R}^{p_1 \times r_1}$ and $\mathbf{W} \in \mathbb{R}^{p_2 \times r_2}$ as well as $\mathbf{g}(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^{r_1}$ and $\mathbf{h}(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^{r_2}$. We denote $\mathbf{V}_j\mathbf{g}(\mathbf{x})$ with $j \in \{1, \dots, p_1\}$ and $\mathbf{W}_j\mathbf{h}(\mathbf{x})$ with $j \in \{1, \dots, p_2\}$ as segments of the max-functions. In addition, we refer to a segment i as being active if it provides the largest value for a given data point \mathbf{x} , i.e., if $\mathbf{V}_i\mathbf{g}(\mathbf{x}) \geq \mathbf{V}_j\mathbf{g}(\mathbf{x})$ holds for all $j \in \{1, \dots, p_1\}$. This form of parameterization is chosen because every PWA function can be parameterized in this way if $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ contain affine and constant functions, i.e., for

$$\mathbf{g}(\mathbf{x}) = \mathbf{h}(\mathbf{x}) = \begin{pmatrix} \mathbf{x} \\ 1 \end{pmatrix}. \quad (5)$$

Moreover, in [Teichrib and Schulze Darup \(2022\)](#), it is shown that every PWQ with $n = 1$ can be represented in the form (4) if $\mathbf{g}(\mathbf{x})$ further contains all quadratic monomials of \mathbf{x} . Alternatively, the function (4) can be interpreted as a shallow maxout NN ([Goodfellow et al., 2013](#)) with two neurons of rank p_1 and p_2 , respectively. For the training of such NN, the nonlinear OP (1) with $\Phi(\mathbf{x})$ as in (4) is often solved with gradient-based methods. However, these gradient-based methods may get stuck in poor local optima. This problem can be avoided by reformulating the nonlinear OP as the MIQP

$$\min_{\mathbf{V}, \mathbf{W}, \alpha^{(i)}, \beta^{(i)}, \gamma^{(i)}, \delta^{(i)}} \sum_{i=1}^{N_D} (y^{(i)} - (\alpha^{(i)} - \beta^{(i)}))^2 \quad (6a)$$

$$\text{s.t.} \quad \mathbf{V}\mathbf{g}(\mathbf{x}^{(i)}) \leq \mathbf{1}\alpha^{(i)}, \quad \mathbf{V}\mathbf{g}(\mathbf{x}^{(i)}) + M(\mathbf{1} - \delta^{(i)}) \geq \mathbf{1}\alpha^{(i)}, \quad (6b)$$

$$\sum_{k=1}^{p_1} \delta_k^{(i)} = 1, \quad \delta^{(i)} \in \{0, 1\}^{p_1}, \quad (6c)$$

$$\mathbf{W}\mathbf{h}(\mathbf{x}^{(i)}) \leq \mathbf{1}\beta^{(i)}, \quad \mathbf{W}\mathbf{h}(\mathbf{x}^{(i)}) + M(\mathbf{1} - \gamma^{(i)}) \geq \mathbf{1}\beta^{(i)}, \quad (6d)$$

$$\sum_{k=1}^{p_2} \gamma_k^{(i)} = 1, \quad \gamma^{(i)} \in \{0, 1\}^{p_2}, \quad (6e)$$

for all $i \in \{1, \dots, N_D\}$. The binary variable vectors $\delta^{(i)}$ and $\gamma^{(i)}$ in (6b) and (6d) select which segments of the max-functions are active for the point $\mathbf{x}^{(i)}$. The constant $M \in \mathbb{R}_{>0}$, often referred to as big- M , is typically very large. We assume without loss of generality that $p_j \leq N_D$ for $j \in \{1, 2\}$, because otherwise there are segments which are inactive for all $\mathbf{x}^{(i)}$ with $i \in \{1, \dots, N_D\}$. Note that for the special case $p_2 = 0$ as well as $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ as in (5) the MIQP (6) becomes equivalent to (2). In the following theorem, we will show that the MIQP (6) can be used to solve the OP (1).

Theorem 1 *Let $\Phi(\mathbf{x})$ be as in (4), $\ell(\cdot) := \|\cdot\|_2^2$, and $\theta := \{\mathbf{V}, \mathbf{W}\}$. Then there exists an $M \in \mathbb{R}_{>0}$ for which the MIQP (6) is equivalent to the nonlinear OP (1).*

Proof The constraints (6c) are such that for every $i \in \{1, \dots, N_D\}$ there exist exactly one k_i with $\delta_{k_i}^{(i)} = 1$ and $\delta_k^{(i)} = 0$ for all $k \in \{1, \dots, p_1\} \setminus k_i$. Thus, the constraints (6b) and (6c) can be equivalently written as

$$\mathbf{V}_{k_i}\mathbf{g}(\mathbf{x}^{(i)}) = \alpha^{(i)}, \quad \mathbf{V}_k\mathbf{g}(\mathbf{x}^{(i)}) \leq \mathbf{V}_{k_i}\mathbf{g}(\mathbf{x}^{(i)}), \quad \forall i \in \{1, \dots, N_D\}, \quad \forall k \in \{1, \dots, p_1\} \quad (7a)$$

$$\mathbf{V}_k\mathbf{g}(\mathbf{x}^{(i)}) + M \geq \mathbf{V}_{k_i}\mathbf{g}(\mathbf{x}^{(i)}), \quad \forall i \in \{1, \dots, N_D\}, \quad \forall k \in \{1, \dots, p_1\}. \quad (7b)$$

As a consequence, the constraints (7a) are equivalent to $\alpha^{(i)} = \max\{\mathbf{V}\mathbf{g}(\mathbf{x}^{(i)})\}$. By using similar arguments, we can show that similar conditions hold for $\beta^{(i)}$, which leads to the OP

$$\min_{\mathbf{V}, \mathbf{W}} \sum_{i=1}^{N_D} (y^{(i)} - (\max\{\mathbf{V}\mathbf{g}(\mathbf{x}^{(i)})\} - \max\{\mathbf{W}\mathbf{h}(\mathbf{x}^{(i)})\}))^2 \quad (8a)$$

$$\text{s.t.} \quad \max\{\mathbf{V}\mathbf{g}(\mathbf{x}^{(i)})\} = \alpha^{(i)}, \quad \max\{\mathbf{W}\mathbf{h}(\mathbf{x}^{(i)})\} = \beta^{(i)}, \quad \forall i \in \{1, \dots, N_D\}, \quad (8b)$$

$$\mathbf{V}\mathbf{g}(\mathbf{x}^{(i)}) + \mathbf{1}M \geq \mathbf{1}\alpha^{(i)}, \quad \mathbf{W}\mathbf{h}(\mathbf{x}^{(i)}) + \mathbf{1}M \geq \mathbf{1}\beta^{(i)}, \quad \forall i \in \{1, \dots, N_D\}. \quad (8c)$$

This OP is always feasible since the trivial solution ($\mathbf{V} = \mathbf{0}$, $\mathbf{W} = \mathbf{0}$, and $\alpha^{(i)} = \beta^{(i)} = 0$ for all $i \in \{1, \dots, N_D\}$) is in the feasible set. Furthermore, since $\mathbf{1}\alpha^{(i)} - \mathbf{V}\mathbf{g}(\mathbf{x}^{(i)})$ and $\mathbf{1}\beta^{(i)} - \mathbf{W}\mathbf{h}(\mathbf{x}^{(i)})$ are finite, we can always find a finite $M \in \mathbb{R}_{>0}$ such that the constraints (8c) never hold with equality. For such an M , the constraints (8b) and (8c) can be neglected. Consequently, the OP (8) and thus (6) are equivalent to the nonlinear OP (1) with $\Phi(\mathbf{x})$ as in (4), $\ell(\cdot) := \|\cdot\|_2^2$, and $\theta := \{\mathbf{V}, \mathbf{W}\}$. ■

According to Theorem 1, the MIQP (6) allows to compute a global optimal solution for a special instance of the nonlinear OP (1). Remarkably, the MIQP (6) can be used to find an approximation of the form (4) for every continuous $\mathbf{g}_i(\mathbf{x})$ and $\mathbf{h}_j(\mathbf{x})$, e.g., by choosing $\mathbf{g}(\mathbf{x})$ quadratic and

$\mathbf{h}(\mathbf{x})$ affine, the function represented by (4) is a PWQ function. Even more complex functions are possible. For example, Figure 1 shows a piecewise defined function with sinusoidal segments, approximated by (4) with $\mathbf{g}(x) = (1 \ x \ \sin(0.2x))^\top$ and $\mathbf{h}(x) = (1 \ x)^\top$. In all figures the functions $\max\{\mathbf{V}\mathbf{g}(\mathbf{x})\}$ and $\max\{\mathbf{W}\mathbf{h}(\mathbf{x})\}$ are represented in blue and green, respectively. The approximation of $f(x)$ by $\Phi(x)$ is illustrated in black, where the data points $(x^{(i)}, f(x^{(i)}))$ used in the MIQP (6) are the black crosses.

Regarding the M from the MIQP (6), it should be noted that methods from the literature (see, e.g., Bonami et al. (2015)) to compute the smallest possible value of M cannot be applied here, because the weights \mathbf{V} and \mathbf{W} are optimization variables and thus not known in advance. This also means that as apparent from the proof of Theorem 1, unlike usual in the big- M formulation, the M does not influence the feasibility of the problem. However, the value of M influences the approximation error if it is not sufficiently large. In fact, a small M leads to smaller values in the entries of the weights \mathbf{V} and \mathbf{W} . Thus, M implicitly regularizes the weights. Assuming the PWA case, we have the following result.

Theorem 2 *Let $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ be as in (5). Then the norm of the matrices $\mathbf{V}_{k,1:n}$ and $\mathbf{W}_{l,1:n}$ of the OP (6) are bounded by the constants v_k and w_l , respectively, i.e., $\|\mathbf{V}_{k,1:n}\|_2 \leq v_k, \forall k \in \{1, \dots, p_1\}$ and $\|\mathbf{W}_{l,1:n}\|_2 \leq w_l, \forall l \in \{1, \dots, p_2\}$.*

Proof Assume that $\delta_{k_i}^{(i)} = 1$ holds for $\mathbf{x}^{(i)}$, then we have according to (7b)

$$\mathbf{V}_{k_i,1:n}\mathbf{x}^{(i)} + \mathbf{V}_{k_i,n+1} \leq \mathbf{V}_{k,1:n}\mathbf{x}^{(i)} + \mathbf{V}_{k,n+1} + M, \quad \forall i \in \{1, \dots, N_D\}, \forall k \in \{1, \dots, p_1\}. \quad (9)$$

We now construct vectors $\tilde{\mathbf{x}}^{(1)}, \dots, \tilde{\mathbf{x}}^{(n)}$ such that they are an orthonormal basis containing the vector $\tilde{\mathbf{x}}^{(\kappa)} := \frac{\mathbf{V}_{k_i,1:n}^\top}{\|\mathbf{V}_{k_i,1:n}\|_2}$. By representing $\mathbf{x}^{(i)}$ in the orthonormal basis we can rewrite (9) as follows

$$\begin{aligned} & \mathbf{V}_{k_i,1:n} \sum_{l=1}^n \tilde{\mathbf{x}}^{(l)} \tilde{\mathbf{x}}^{(l)\top} \mathbf{x}^{(i)} + \mathbf{V}_{k_i,n+1} \leq \mathbf{V}_{k,1:n} \mathbf{x}^{(i)} + \mathbf{V}_{k,n+1} + M \\ \Leftrightarrow & \|\mathbf{V}_{k_i,1:n}\|_2 \|\tilde{\mathbf{x}}^{(\kappa)}\|_2 \tilde{\mathbf{x}}^{(\kappa)\top} \mathbf{x}^{(i)} + \mathbf{V}_{k_i,1:n} \sum_{l=1, l \neq \kappa}^n \tilde{\mathbf{x}}^{(l)} \tilde{\mathbf{x}}^{(l)\top} \mathbf{x}^{(i)} + \mathbf{V}_{k_i,n+1} \leq \mathbf{V}_{k,1:n} \mathbf{x}^{(i)} + \mathbf{V}_{k,n+1} + M \\ \Leftrightarrow & \|\mathbf{V}_{k_i,1:n}\|_2 \leq \frac{1}{\tilde{\mathbf{x}}^{(\kappa)\top} \mathbf{x}^{(i)}} \left(M + \mathbf{V}_{k,n+1} - \mathbf{V}_{k_i,n+1} + (\mathbf{V}_{k,1:n} - \mathbf{V}_{k_i,1:n} \sum_{l=1, l \neq \kappa}^n \tilde{\mathbf{x}}^{(l)} \tilde{\mathbf{x}}^{(l)\top}) \mathbf{x}^{(i)} \right) =: v_k \end{aligned}$$

for all $i \in \{1, \dots, N_D\}$ and for all $k \in \{1, \dots, p_1\}$. The upper bound w_l for $\|\mathbf{W}_{l,1:n}\|_2$ can be derived in a similar way, which completes the proof. \blacksquare

Theorem 2 shows that it is not necessary to explicitly include a regularization in the cost function as in Siahkamari et al. (2020) to ensure small weights since M implicitly regularizes the weights. Figure 2 shows the effect of this implicit regularization. For the example, the MIQP (6) is solved with $p_1 = p_2 = 2$ and (5), for different values of M . In Figure 2 (a), the constant is not sufficiently large and thus an exact fit of the data is impossible. For $M = 50$ and $M = 100$, an exact fit is possible. However, $M = 50$ leads to a fit with smaller weights.

Another advantage of the proposed method is that since for $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ as in (5), (4) is a maxout NN with affine pre- and post-activation, we can directly use Teichrib and Schulze Darup (2023, Lem. 2) to show that (4) can be modelled as an MI feasibility problem and is thus well-suited for the use inside an OP.




Figure 1: $\Phi(x)$ with $p_i = 1$ in (a), $p_i = 2$ in (b), and $p_i = 3$ in (c) for $i \in \{1, 2\}$. Figure 2: PWA $\Phi(x)$ with $M = 10$ in (a), $M = 50$ in (b), and $M = 100$ in (c).

Corollary 3 Let $g(x)$ and $h(x)$ be as in (5). Then, the function (4) can be modelled as an MI feasibility problem.

Proof According to Teichrib and Schulze Darup (2023, Lem. 2), an MI feasibility problem with the constraints Teichrib and Schulze Darup (2023, Eq. (18)) can be used to replace the nonlinear relation (4). ■

This corollary allows the replacement of the nonlinear function (4) inside an OP by MI linear constraints. Thus, a nonlinear OCP, where (4) is used either as a prediction model or to approximate the cost function, can be reformulated as MIP. This enables the use of standard MI solvers to find a global optimum of the OP.

4. Complexity reduction

In this section, we describe two approaches to reduce the complexity of the MIQP (6). The approaches mainly focus on reducing the number of binary variables. Before describing our approaches, we first show that the MIQP (6) can be solved by a QP without using any binary variables if we are not interested in a small number of segments p_1 and p_2 in (4). Consider the special case where we have exactly as many segments as data points, i.e. $p_1 = p_2 = N_D$. In this case, the QP

$$\min_{\mathbf{V}, \mathbf{W}, \alpha^{(i)}, \beta^{(i)}} \sum_{i=1}^{N_D} (y^{(i)} - (\alpha^{(i)} - \beta^{(i)}))^2 \quad (10a)$$

$$\text{s.t. } \mathbf{V}g(\mathbf{x}^{(i)}) \leq \mathbf{1}\alpha^{(i)}, \quad \mathbf{W}h(\mathbf{x}^{(i)}) \leq \mathbf{1}\beta^{(i)}, \quad \forall i \in \{1, \dots, N_D\}, \quad (10b)$$

$$\mathbf{V}_\kappa g(\mathbf{x}^{(\kappa)}) \geq \alpha^{(\kappa)}, \quad \mathbf{W}_\kappa h(\mathbf{x}^{(\kappa)}) \geq \beta^{(\kappa)}, \quad \forall \kappa \in \{1, \dots, N_D\} \quad (10c)$$

can be used to compute optimal \mathbf{V} , \mathbf{W} , $\alpha^{(i)}$, and $\beta^{(i)}$ of the MIQP (6). Since there is one segment for each data point, no binary variables are needed to decide which segment is active, thus (6b)-(6e) can be replaced by (10b) and (10c). However, this leads to the fact that the number of data points determines the number of segments. As a result, we have $2N_D(N_D + 1)$ inequality constraints and $2N_D + N_D(r_1 + r_2)$ optimization variables in (10). The QP (10) can be reformulated to obtain the QP from Siahkamari et al. (2020). Thus, our method generalizes known results for piecewise linear regression and includes the OP from Siahkamari et al. (2020) as a special case. The elimination of the binary variables in the OP (10) is only possible for $p_1 = p_2 = N_D$, which leads to a complex function $\Phi(x)$ with a high number of segments, for large data sets. Therefore, in the remainder of this section, we will present two approaches for reducing the number of binary variables in the

MIQP (6) without increasing the complexity of the fitted function $\Phi(\mathbf{x})$, in terms of the number of segments.

4.1. Symmetry breaking

Due to symmetries in the MIQP (6), there are many equivalent solutions. These symmetries result from the fact that the values of the max-functions in (4) are independent of the order of the p_1 and p_2 , respectively, segments. In the following proposition, we will present a method to eliminate some of these redundant solutions.

Proposition 4 *There exists an optimal solution of the MIQP (6) with*

$$\delta_{2:p_1}^{(1)} = \mathbf{0}, \delta_{3:p_1}^{(2)} = \mathbf{0}, \dots, \delta_{p_1}^{(p_1-1)} = 0 \quad \text{and} \quad \gamma_{2:p_2}^{(1)} = \mathbf{0}, \gamma_{3:p_2}^{(2)} = \mathbf{0}, \dots, \gamma_{p_2}^{(p_2-1)} = 0. \quad (11)$$

Proof Every permutation of the k -indices of an optimal solution of the MIQP (6) is also an optimal solution since the maximum is invariant under a permutation of the segments. Assume $\delta^{(i)}$ and $\gamma^{(i)}$ are the binary variables of an optimal solution of the MIQP (6). Due to the discussed invariance of the maximum, we can assume that the optimal solution is such that the first segment is active for the first data point, i.e., $\delta_1^{(1)} = 1$ and $\gamma_1^{(1)} = 1$. In combination with (6c) and (6e) this leads to $\delta_{2:p_1}^{(1)} = \mathbf{0}$ and $\gamma_{2:p_2}^{(1)} = \mathbf{0}$. For the second data point, there are two possibilities. Either the same segment is active as for the first data point, or a different segment is active, which we can assume to be the second segment, i.e., $\delta_1^{(2)} + \delta_2^{(2)} = 1$ and $\gamma_1^{(2)} + \gamma_2^{(2)} = 1$ or equivalent $\delta_{3:p_1}^{(2)} = \mathbf{0}$ and $\gamma_{3:p_2}^{(2)} = \mathbf{0}$. We can proceed in a similar way with the remaining data points until the $(p_1 - 1)$ -th and $(p_2 - 1)$ -th, respectively, data point, which leads to the remaining constraints in (11) and completes the proof. \blacksquare

According to Proposition 4 we can reduce the number of binary variables by including the constraints (11) in the MIQP (6). Symmetry breaking is also achieved in Toriello and Vielma (2012, Prop. 3.2) by enforcing an order to the parameter of the fitted function through additional constraints. However, unlike our method, this does not reduce the number of binary variables.

4.2. Preclustering

Since the number of data points determines the number of binary variables in (6), reducing the number of data points is a very effective way to reduce the computation time needed to solve the MIQP (6). Unfortunately, a certain number of data points is often required to obtain an approximation that generalizes well to unseen data. This means that it is impossible to simply reduce the size of the data set. However, we can precompute a clustering of the data set and modify the OP so that instead of assigning individual data points to segments of the max-function, whole clusters of multiple data points are assigned. Assume the data set is partitioned into K disjoint clusters of the form $\mathcal{C}_j := \{(\mathbf{x}^{(i)}, y^{(i)}) \in \mathcal{D} \mid i \in \mathcal{I}_j\}$ with $\mathcal{I}_j \subseteq \{1, \dots, N_D\}$, $\mathcal{C}_i \cap \mathcal{C}_j = \emptyset$ for all $i \neq j$, and $\cup_{j=1}^K \mathcal{C}_j = \mathcal{D}$. Then, we can formulate a modified version of MIQP (6) with

$$\mathbf{V}\mathbf{g}(\mathbf{x}^{(i)}) \leq \mathbf{1}\alpha^{(i)}, \quad \mathbf{V}\mathbf{g}(\mathbf{x}^{(i)}) + M(\mathbf{1} - \delta^{(l)}) \geq \mathbf{1}\alpha^{(i)}, \quad \forall i \in \{1, \dots, N_D\}, \quad (12a)$$

$$\mathbf{W}\mathbf{h}(\mathbf{x}^{(i)}) \leq \mathbf{1}\beta^{(i)}, \quad \mathbf{W}\mathbf{h}(\mathbf{x}^{(i)}) + M(\mathbf{1} - \gamma^{(l)}) \geq \mathbf{1}\beta^{(i)}, \quad \forall i \in \{1, \dots, N_D\}, \quad (12b)$$

instead of (6b) and (6d), where l is such that we have $(\mathbf{x}^{(i)}, y^{(i)}) \in \mathcal{C}_l$. The binary variable vectors $\delta^{(l)} \in \{0, 1\}^{p_1}$ and $\gamma^{(l)} \in \{0, 1\}^{p_2}$ select which segment of the max-functions is active for the

whole cluster \mathcal{C}_i and not just for a single data point as in the initial MIQP (6). Any algorithm that computes a clustering with the desired properties can be used here, e.g. k-means++ from Arthur and Vassilvitskii (2007). Note that although we perform a clustering before solving the optimization problem, the modified MIQP (6) with (12) instead of (6b) and (6d) does not compute a linear regression for each cluster individually as in Magnani and Boyd (2009), since there is not a distinct segment for each cluster. As long as $p_i < K$ with $i \in \{1, 2\}$, multiple clusters are assigned to the same segment. With the preclustering, we can reduce the number of binary variables from $N_D(p_1 + p_2)$ to $K(p_1 + p_2)$. For $K = N_D$, i.e., each cluster contains exactly one data point, the constraints (6b), (6d) and (12a), (12b) are equivalent. Furthermore, analogous to (10) for $p_1 = p_2 = K$ the MIQP with preclustering can be reformulated as a QP by replacing (10c) in the QP (10) with

$$\mathbf{V}_\kappa \mathbf{g}(\mathbf{x}^{(i)}) \geq \alpha^{(i)}, \quad \mathbf{W}_\kappa \mathbf{h}(\mathbf{x}^{(i)}) \geq \beta^{(i)}, \quad \forall i \in \{1, \dots, N_D\}, \quad (13)$$

where κ in (13) is such that we have $(\mathbf{x}^{(i)}, y^{(i)}) \in \mathcal{C}_\kappa$. The resulting QP has $2N_D(K + 1)$ inequality constraints and $2N_D + K(r_1 + r_2)$ optimization variables.

5. Numerical examples

In this section, we use the MIQP (6) with different $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ to approximate a nonlinear MPC and compare the results with an approximation by an NN.

5.1. Nonlinear MPC

Consider the nonlinear OCP

$$\begin{aligned} V(\mathbf{x}) := \min_{\substack{\mathbf{x}(0), \dots, \mathbf{x}(N) \\ u(0), \dots, u(N-1)}} \mathbf{x}(N)^\top \begin{pmatrix} 1.3 & 1.9 \\ 1.9 & 3.0 \end{pmatrix} \mathbf{x}(N) + \sum_{k=0}^{N-1} \|\mathbf{x}(k)\|_2^2 + \|u(k)\|_2^2 & (14) \\ \text{s.t. } \mathbf{x}_1(k+1) = \mathbf{x}_1(k) + T_s \mathbf{x}_2(k), & \forall k \in \{0, \dots, N-1\}, \\ \mathbf{x}_2(k+1) = \mathbf{x}_2(k) - T_s \frac{\rho_0}{m} e^{-\mathbf{x}_1(k)} \mathbf{x}_1(k) - T_s \frac{h_d}{m} \mathbf{x}_2(k) + T_s \frac{u(k)}{m}, & \forall k \in \{0, \dots, N-1\}, \\ -1.3 \leq \mathbf{x}_2(k) \leq 1.3, \quad -4 \leq u(k) \leq 4, & \forall k \in \{0, \dots, N-1\}, \\ \mathbf{x}(0) = \mathbf{x}, \quad \mathbf{x}(N)^\top \begin{pmatrix} 1.3 & 1.9 \\ 1.9 & 3.0 \end{pmatrix} \mathbf{x}(N) \leq 0.001, & \end{aligned}$$

from Magni et al. (2003) with parameters $N = 15$, $T_s = 0.4$, $m = 1$, $\rho_0 = 0.33$ and $h_d = 1.1$. The OCP (14) is solved for 1350 different values of \mathbf{x} on a regular grid of size 45×30 with \mathbf{x}_1 from -2 to 5 and \mathbf{x}_2 from -1.3 to 1.3 to generate a data set of the form $(\mathbf{x}^{(i)}, y^{(i)} = u^*(\mathbf{x}^{(i)}))$, where $u^*(\mathbf{x}^{(i)})$ is the optimal $u(0)$ for the state $\mathbf{x} = \mathbf{x}^{(i)}$. The nonlinear OCP (14) is solved using the software framework CasADi from Andersson et al. (2019). Based on the data set, two different maxout NN of the form (4) with $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ as in (5) are trained using the Adam algorithm (Kingma and Ba, 2017). The values of p_1 and p_2 of the two NN are chosen as $p_1 = p_2 = 2$ and $p_1 = p_2 = 50$ (row 1 and 2 of Table 1). Such NN can be interpreted as maxout NN with affine pre- and post-activation and 2 neurons of rank 2 and 50, respectively, in one hidden layer. We also used the MIQP (6) with $p_1 = p_2 = 2$, different choices of $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ and the symmetry breaking constraints (11) to compute global optimal parameters \mathbf{V} and \mathbf{W} of the function $\Phi(\mathbf{x})$. We start with $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ as in (5), which results in a function that is equal to the first maxout NN, and

add monomials up to the order of three to the functions (row 3 and 4 of Table 1). The results are summarized in Table 1. The last two columns show the quality of the approximation in terms of the mean squared error (MSE) and the maximum absolute value of the error on a test data set. The second row indicates the method used to compute the parameters. As apparent from Table 1, the

Table 1: Different approximations of the nonlinear optimal control law $u^*(\mathbf{x})$.

| No. | Method | p_1, p_2 | $\mathbf{g}(\mathbf{x}), \mathbf{h}(\mathbf{x})$ | MSE | e_{\max} |
|-----|----------|------------|---|--------|------------|
| 1. | NN | 2 | $(1 \ \mathbf{x}^\top)^\top$ | 0.0532 | 1.6054 |
| 2. | NN | 50 | $(1 \ \mathbf{x}^\top)^\top$ | 0.0252 | 1.4971 |
| 3. | MIQP (6) | 2 | $(1 \ \mathbf{x}^\top)^\top$ | 0.0503 | 1.0797 |
| 4. | MIQP (6) | 2 | $(1 \ \mathbf{x}^\top \ \mathbf{x}_1^2 \ \mathbf{x}_1\mathbf{x}_2 \ \mathbf{x}_2^2 \ \mathbf{x}_1^3 \ \mathbf{x}_1^2\mathbf{x}_2 \ \mathbf{x}_1\mathbf{x}_2^2 \ \mathbf{x}_2^3)^\top$ | 0.0240 | 1.0687 |

approximation provided by the MIQP (6) with cubic terms in $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ gives the lowest error, even though it has only two segments. By increasing the number of segments p_1 and p_2 , the MSE of the first NN approaches the MSE of the 4th function. However, the maximum error barely changes, which may be a problem since it is shown in Fabiani and Goulart (2022) that the maximum error is crucial for the stability and performance of an approximated controller. Another advantage of the proposed method is that adding additional terms to the $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ functions can improve the approximation quality without increasing the number of segments. This is beneficial for a fast evaluation of the approximation $\Phi(\mathbf{x})$. Figure 3 shows how the function 4 from Table 1 leads to a better result by approximating the curvature of the optimal control law.




Figure 3: (a) shows the optimal control law with the sampled data points. (b) and (c) are approximations of the optimal control law by the functions 1 and 4, respectively, from Table 1.

6. Conclusion

We proposed a new MIP-based method for piecewise regression, which allows to fit an arbitrary piecewise defined function to data points by using the difference of two max-functions (4). The presented method is very flexible as it is not restricted to a special class of functions and thus allows to trade off the complexity of the fitted function against the approximation error. In addition, we showed in a numerical example that the method can be used to find a simple approximation, in terms of the number of regions, of the optimal control law of nonlinear MPC.

References

- Joel A. E. Andersson, Joris Gillis, Greg Horn, James B. Rawlings, and Moritz Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019.
- David Arthur and Sergei Vassilvitskii. K-means++: The advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, page 1027–1035, USA, 2007. Society for Industrial and Applied Mathematics.
- Alberto Bemporad, Manfred Morari, Vivek Dua, and Efstratios N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, 2002.
- Alberto Bemporad. 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, 2023.
- Dimitris Bertsimas and Romy Shioda. Classification and regression via integer optimization. *Operations Research*, 55(2):252–271, 2007.
- Pierre Bonami, Andrea Lodi, Andrea Tramontani, and Sven Wiese. On mathematical programming with indicator constraints. *Math. Program.*, 151(1):191–223, 2015.
- Leo Breiman. Hinging hyperplanes for regression, classification, and function approximation. *IEEE Transactions on Information Theory*, 39(3):999–1013, 1993.
- Steven Chen, Kelsey Saulnier, Nikolay Atanasov, Daniel D. Lee, Vijay Kumar, George J. Pappas, and Manfred Morari. Approximating explicit model predictive control using constrained neural networks. In *Proceedings of the 2018 American Control Conference*, pages 1520–1527, 2018.
- Leon O. Chua and An-Chang Deng. Canonical piecewise-linear representation. *IEEE Transactions on Circuits and Systems*, 35(1):101–111, 1988.
- Filippo Fabiani and Paul J. Goulart. Reliably-stabilizing piecewise-affine neural network controllers. *IEEE Transactions on Automatic Control*, 68(9):5201–5215, 2023.
- Ian Goodfellow, David Warde-Farley, Mehdi Mirza, Aaron Courville, and Yoshua Bengio. Max-out networks. In *Proceedings of the 30th International Conference on Machine Learning*, 28(3):1319–1327, 2013.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- Benjamin Karg and Sergio Lucia. Efficient representation and approximation of model predictive control laws via deep learning. *IEEE Transactions on Cybernetics*, 50(9):3866–3878, 2020.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv:1412.6980 [cs.LG], 2017.
- Alessandro Magnani and Stephen P. Boyd. Convex piecewise-linear fitting. *Optimization and Engineering*, 10(1):1–17, 2009.

- Lalo Magni, Giuseppe De Nicolao, Riccardo Scattolini, and Frank Allgöwer. Robust model predictive control for nonlinear discrete-time systems. *International Journal of Robust and Nonlinear Control*, 13(3-4):229–246, 2003.
- Steffen Rebennack and Vitaliy Krasko. Piecewise linear function fitting via mixed-integer linear programming. *INFORMS Journal on Computing*, 32(2):507–530, 2020.
- Jacob Roll, Alberto Bemporad, and Lennart Ljung. Identification of piecewise affine systems via mixed-integer programming. *Automatica*, 40(1):37–50, 2004.
- Ali Siahkamari, Aditya Gangrade, Brian Kulis, and Venkatesh Saligrama. Piecewise linear regression via a difference of convex functions. In *Proceedings of the 37th International Conference on Machine Learning*, 119:8895–8904, 2020.
- Dieter Teichrib and Moritz Schulze Darup. Tailored neural networks for learning optimal value functions in mpc. In *Proceedings of the 2021 IEEE Conference on Decision and Control*, pages 5281–5287, 2021.
- Dieter Teichrib and Moritz Schulze Darup. Tailored max-out networks for learning convex PWQ functions. In *Proceedings of the 2022 European Control Conference*, pages 2272–2278, 2022.
- Dieter Teichrib and Moritz Schulze Darup. Error bounds for maxout neural network approximations of model predictive control. *IFAC-PapersOnLine*, 56(2):10113:10119, 2023.
- Alejandro Toriello and Juan Pablo Vielma. Fitting piecewise linear continuous functions. *European Journal of Operational Research*, 219(1):86–95, 2012.
- John A. Warwicker and Steffen Rebennack. A unified framework for bivariate clustering and regression problems via mixed-integer linear programming. *Discrete Applied Mathematics*, 336: 15–36, 2023.