



HAL
open science

Global optimization for low-dimensional switching linear regression and bounded-error estimation

Fabien Lauer

► **To cite this version:**

Fabien Lauer. Global optimization for low-dimensional switching linear regression and bounded-error estimation. 2017. hal-01563788v1

HAL Id: hal-01563788

<https://hal.science/hal-01563788v1>

Preprint submitted on 18 Jul 2017 (v1), last revised 21 Nov 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Global optimization for low-dimensional switching linear regression and bounded-error estimation

Fabien Lauer

Université de Lorraine, LORIA, UMR 7503, F-54506 Vandœuvre-lès-Nancy, France
CNRS

Abstract

The paper provides global optimization algorithms for two particularly difficult nonconvex problems raised by hybrid system identification: switching linear regression and bounded-error estimation. While most works focus on local optimization heuristics without global optimality guarantees or with guarantees valid only under restrictive conditions, the proposed approach yields a solution with a certificate of global optimality. This approach relies on a branch-and-bound strategy for which we devise lower bounds that can be efficiently computed. In order to obtain scalable algorithms with respect to the number of data, we directly optimize the model parameters in a continuous optimization setting without involving integer variables. Numerical experiments show that the proposed algorithms offer a higher accuracy than convex relaxations with a reasonable computational burden for hybrid system identification. In addition, we discuss how bounded-error estimation is related to robust estimation in the presence of outliers and exact recovery under sparse noise, for which we also obtain promising numerical results.

1 Introduction

The paper tackles two problems that lie at the core of hybrid dynamical system identification, whose aim is to estimate, from input–output data, a model of a system switching at unknown instants between a number of linear subsystems. More precisely, we consider the minimization of the error of a switching linear model with a fixed number of modes and the iterative maximization of the number of data that can be approximated by a linear model with a bounded error. The latter problem, also known as bounded-error estimation, has an interest outside of hybrid systems as well and in particular for robust estimation in the presence of outliers.

The problems are understood as *global* minimization/maximization problems. However, due to their complexity, most of the literature, as reviewed in [17, 8], focuses on local optimization or heuristic approaches: for switching regression with a fixed number of modes in [21, 9, 13, 10, 14, 18] and for the bounded-error approach to switching regression in [5, 3, 16, 6]. Some of these methods can be proved to yield the global solution but only in specific conditions, such as in the absence of noise for [21] and under data-dependent conditions difficult to check in practice for [3]. Recent results showed that, though being NP-hard in general, some hybrid system identification problems, including the minimization of the error of a switching linear model, have a complexity no more than polynomial in the number of data for a fixed data dimension [11, 12]. However, in practice, the complexity of the corresponding polynomial algorithms remains too high except for small data sets in small dimensions.

Contribution Global optimization of such difficult problems in general is usually deemed impractical. Hence, we focus on instances where the data can be numerous but should live in a low-dimensional space, as is often the case in a system identification context (most examples in the literature on hybrid system identification have a dimension less than five). In this context, the paper proposes a branch-and-bound approach to the two problems above. Contrary to previous works, such an approach offers unconditional global optimality guarantees, while remaining computationally efficient with large data sets. Branch-and-bound is a standard approach to global

optimization, but it was only considered for hybrid system identification in [19], where an off-the-shelf solver is applied after a reformulation of the piecewise affine regression problem into a mixed-integer linear or quadratic program, with a number of binary variables proportional to the number of data. At the opposite, the proposed approach can handle larger data sets by developing dedicated optimization algorithms while focusing on the continuous variables of the problems, i.e., the model parameters, rather than the integer variables. Technically, the branch-and-bound approach relies on the derivation of a number of lower bounds on the different cost functions for parameters constrained to lie in a box (a hyperrectangle). In particular, efficiency is obtained thanks to two ingredients: i) simple lower bounds that can quickly discard boxes with very large costs, and ii) a constant-classification based criterion that allows us to more tightly lower bound the cost.

Paper organization Section 2 describes the general branch-and-bound approach adopted to tackle the problems of interest, which are formally described in dedicated sections: Sect. 3 for switching regression and Sect. 4 for the bounded-error approach. Then, Sect. 5 presents numerical results and Sect. 6 discusses open issues.

Notation Vectors are written in lowercase bold letters, while matrices are written in uppercase bold letters. For a vector \mathbf{u} , the k th entry is denoted by u_k , while for a vector \mathbf{u}_j , its k th entry is $u_{j,k}$. All inequalities between vectors, e.g., $\mathbf{u} \leq \mathbf{v}$, are meant entrywise. A box $B \subset \mathbb{R}^D$ is a hyperrectangular region of \mathbb{R}^D , i.e., $B = [\mathbf{u}, \mathbf{v}] = \prod_{k=1}^D [u_k, v_k]$ with $\mathbf{u} \in \mathbb{R}^D$, $\mathbf{v} \in \mathbb{R}^D$ such that $\mathbf{u} \leq \mathbf{v}$. The positive and negative parts of a scalar are denoted by $(\cdot)_+ = \max\{0, \cdot\}$ and $(\cdot)_- = \min\{0, \cdot\}$ and similar notations are used for the corresponding entrywise operations on vectors. Of course, $(\cdot)_+^2$ and $(\cdot)_-^2$ are understood as the squared positive and negative parts of a scalar, i.e., $(\cdot)_+^2 = ((\cdot)_+)^2$ and $(\cdot)_-^2 = ((\cdot)_-)^2$. The notation $|\cdot|$ denotes either the absolute value for real arguments or the cardinality for sets.

2 General approach

Consider the global minimization of some cost function $J(\mathbf{w})$ of a vector of parameters $\mathbf{w} \in \mathbb{R}^D$ over a box $B_{\text{init}} = [\mathbf{u}_{\text{init}}, \mathbf{v}_{\text{init}}] \subset \mathbb{R}^D$, where the different definitions of the cost function J for the problems of interest will be given in dedicated sections below. We attack these problems with a branch-and-bound approach, summarized in Algorithm 1, which takes a data set of regression vectors $\mathbf{x}_i \in \mathbb{R}^d$ and target outputs $y_i \in \mathbb{R}$ as inputs. In hybrid system identification, the regression vectors are typically built from lagged inputs and outputs of the system [17].

Algorithm 1 General branch-and-bound scheme.

Require: A data set $\{(\mathbf{x}_i, y_i)\}_{i=1}^N \subset \mathbb{R}^d \times \mathbb{R}$, initial box bounds $B_{\text{init}} = [\mathbf{u}_{\text{init}}, \mathbf{v}_{\text{init}}] \subset \mathbb{R}^D$ and $TOL > 0$. Optionally, an initial guess of $\mathbf{w} \in B_{\text{init}}$.

Initialize the global bounds $\underline{J} \leftarrow 0$, $\bar{J} \leftarrow +\infty$ or $\bar{J} \leftarrow J(\mathbf{w})$ if \mathbf{w} is provided, and the list of boxes $\mathcal{B} \leftarrow \{B_{\text{init}}\}$.

while $(\bar{J} - \underline{J})/\bar{J} > TOL$ **do**

Split the current box B into B^1 and B^2 such that $B = B^1 \cup B^2$.

Compute upper bounds $\bar{J}(B^1)$ and $\bar{J}(B^2)$.

 Update $\bar{J} \leftarrow \min\{\bar{J}, \bar{J}(B^1), \bar{J}(B^2)\}$ and the best solution \mathbf{w}^* .

Compute lower bounds $\underline{J}(B^1)$ and $\underline{J}(B^2)$.

 For $k = 1, 2$, append B^k to the list of active boxes \mathcal{B} if $\underline{J}(B^k) \leq \bar{J}$.

 Remove B from the list of active boxes: $\mathcal{B} \leftarrow \mathcal{B} \setminus \{B\}$.

 Select the next box $B \leftarrow \operatorname{argmin}_{B \in \mathcal{B}} \underline{J}(B)$ and set $\underline{J} \leftarrow \underline{J}(B)$.

end while

return \mathbf{w}^* and $\bar{J} = J(\mathbf{w}^*) \approx \min_{\mathbf{w} \in B_{\text{init}}} J(\mathbf{w})$.

The general branch-and-bound scheme relies on computing upper and lower bounds (\bar{J} and \underline{J} in Algorithm 1) on the global optimum $\min_{\mathbf{w} \in B_{\text{init}}} J(\mathbf{w})$. Then, regions B of the search space

in which the local lower bound $\underline{J}(B)$ is larger than the global upper bound \bar{J} can be discarded, reducing the volume left to explore until the relative optimality gap, $(\bar{J} - \underline{J})/\bar{J}$, decreases below a predefined tolerance TOL . Here, the considered regions are always boxes, i.e., hyperrectangles. Upper bounds $\bar{J}(B)$ can be easily computed by some local optimization or heuristic method for a problem of interest. Alternatively, $\bar{J}(B)$ can be computed merely as the cost function value at the box base point \mathbf{u} or at a random point inside the box, while local optimization is only used periodically. On the other hand, lower bounds $\underline{J}(B)$ require a careful derivation, the efficiency of the approach relying mostly on the tightness of these bounds.

Algorithm 1 retains only the solution yielding the best upper bound $\bar{J} = J(\mathbf{w}^*)$. Depending on the value of TOL , the algorithm can terminate while there are multiple remaining active boxes possibly containing equally good solutions within the tolerance. A possible modification would be to retain a list of solution candidates with cost function values close to the best one rather than a single solution. Since such a modification would be straightforward, in the following, we focus only on the version returning a single solution.

3 Switching linear regression

We consider the identification of a switching system with n modes generating a data set of N points $(\mathbf{x}_i, y_i) \in \mathbb{R}^d \times \mathbb{R}$, $i = 1, \dots, N$, with

$$y_i = \mathbf{w}_{q_i}^T \mathbf{x}_i + v_i, \quad (1)$$

where $q_i \in \mathcal{Q} = \{1, \dots, n\}$ is the index of the active mode for the i th point, $\{\mathbf{w}_j\}_{j=1}^n \subset \mathbb{R}^d$ is a collection of linear model parameter vectors and $v_i \in \mathbb{R}$ is a noise term. The aim here is to estimate, from the knowledge of $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ and n only, the concatenated parameter vector $\mathbf{w} = [\mathbf{w}_1^T, \dots, \mathbf{w}_n^T]^T \in \mathbb{R}^{nd}$. Throughout the paper, we assume a similar partitioning of all vectors from \mathbb{R}^{nd} , i.e., for $\mathbf{u} \in \mathbb{R}^{nd}$, \mathbf{u}_j refers to the j th subvector of dimension d in \mathbf{u} .

Least squares estimates¹ of \mathbf{w} and $\mathbf{q} = [q_1, \dots, q_N]^T$ are defined as the global solutions to

$$\begin{aligned} \min_{\mathbf{w} \in \mathbb{R}^{nd}, \mathbf{q} \in \mathcal{Q}^N} J_{\text{SWq}}(\mathbf{w}, \mathbf{q}), \quad (2) \\ \text{with } J_{\text{SWq}}(\mathbf{w}, \mathbf{q}) = \sum_{i=1}^N (y_i - \mathbf{w}_{q_i}^T \mathbf{x}_i)^2. \end{aligned}$$

Note that Problem (2) involves N integer variables in \mathbf{q} , which would imply a worst-case exponential complexity in the number of data for its direct global optimization. Other reformulations based on nN binary variables suffer from a similar limitation, which is why the following considers a continuous optimization point of view.

Using the classification rule²

$$q_i(\mathbf{w}) = \underset{j \in \mathcal{Q}}{\operatorname{argmin}} (y_i - \mathbf{w}_j^T \mathbf{x}_i)^2, \quad i = 1, \dots, N, \quad (3)$$

Problem (2) can be reformulated without integer variables as in [13], leading to

$$\begin{aligned} \min_{\mathbf{w} \in \mathbb{R}^{nd}} J_{\text{SW}}(\mathbf{w}), \quad (4) \\ \text{with } J_{\text{SW}}(\mathbf{w}) = \sum_{i=1}^N \min_{j \in \mathcal{Q}} (y_i - \mathbf{w}_j^T \mathbf{x}_i)^2, \end{aligned}$$

and the equivalence $J_{\text{SW}}(\mathbf{w}) = J_{\text{SWq}}(\mathbf{w}, \mathbf{q}(\mathbf{w}))$. Though equivalent, the formulation in (4) emphasizes the major role played by the continuous variables in \mathbf{w} , on which directly depends \mathbf{q} . Using

¹We restrict the presentation to the squared loss function $\ell(e) = e^2$, but similar results could be obtained for instance with the absolute loss $\ell(e) = |e|$.

²When the minimum is not unique in (3), ties are arbitrarily broken by returning the minimal index j of the minimum.

this fact, the dimension of the problem can be restricted to nd and becomes independent of N .³ However, the global optimization of Problem (4) remains nontrivial.

Note that for symmetry reasons, the cost function J_{SW} is invariant to permutations of the subvectors \mathbf{w}_j in \mathbf{w} , hence the minimizer is not unique. Such symmetries can be broken by arbitrarily imposing an ordering on the modes, for instance as

$$w_{j,1} \leq w_{j+1,1}, \quad j = 1, \dots, n-1, \quad (5)$$

where $w_{j,k}$ denotes the k th component of the j th parameter vector. Note that ties in the case $w_{j,1} = w_{j+1,1}$ can be broken by imposing similar constraints recursively on the remaining components. However, these additional constraints might be more difficult to take into account in the branch-and-bound approach, while they might also be of little use since the event corresponding to a tie in a global minimizer has zero measure with noisy data. Therefore, in the following we do not deal with such ties and focus on solving (4) under the constraints (5). More precisely, we consider the restriction of (4) subject to (5) where the domain \mathbb{R}^{nd} is replaced by a box $B_{\text{init}} \subset \mathbb{R}^{nd}$.

3.1 Branch-and-bound approach

Many heuristics have been proposed for switching linear regression (see, e.g., [8, 13, 10, 18]), and any of them can be used to compute upper bounds $\bar{J}(B)$. Here, we consider the simple and efficient k -LinReg method [10]. We choose to compute the initial guess of \mathbf{w} with it before starting the search, and additionally reuse it once in a while (e.g., every 100 iterations). Otherwise, at every iteration dealing with a box $B = [\mathbf{u}, \mathbf{v}]$, upper bounds are simply computed as $\bar{J}(B) = J_{\text{SW}}(\mathbf{u})$.

The symmetry-breaking constraints (5) can be simply imposed at the branching level by explicitly discarding regions of subboxes without feasible solutions. More precisely, we compute $B^1 = [\mathbf{u}^1, \mathbf{v}^1]$ and $B^2 = [\mathbf{u}^2, \mathbf{v}^2]$ from $B = [\mathbf{u}, \mathbf{v}]$ by first applying a standard split along the longest side of the box:

$$(j^*, k^*) = \underset{j \in \mathcal{Q}, k \in \{1, \dots, d\}}{\operatorname{argmax}} v_{j,k} - u_{j,k} \quad (6)$$

and

$$\mathbf{u}^1 = \mathbf{u}, \quad v_{j,k}^1 = \begin{cases} (u_{j,k} + v_{j,k})/2, & \text{if } (j, k) = (j^*, k^*), \\ v_{j,k}, & \text{otherwise,} \end{cases} \quad (7)$$

$$\mathbf{v}^2 = \mathbf{v}, \quad u_{j,k}^2 = \begin{cases} (u_{j,k} + v_{j,k})/2, & \text{if } (j, k) = (j^*, k^*), \\ u_{j,k}, & \text{otherwise.} \end{cases} \quad (8)$$

Then, in the case $k^* = 1$, which is the only one concerned by (5), we correct the box bounds recursively for $j = j^* - 1, \dots, 1$ with

$$v_{j,1}^1 = \min \{v_{j,1}^1, v_{j+1,1}^1\} \quad (9)$$

and, for $j = j^* + 1, \dots, n$, with

$$u_{j,1}^2 = \max \{u_{j,1}^2, u_{j-1,1}^2\}. \quad (10)$$

Figure 1 illustrates the splitting rule.

3.2 Lower bounds

Regarding the lower bounds, two different bounds are derived below, with increasing value (tightness) but also an increasing computational demand. In practice, we only compute the second lower bound if the first one is not large enough to discard the box (i.e., if it is smaller than the global upper bound \bar{J}).

We start with a preliminary result bounding the value of dot products involved in the cost function, which will be of interest throughout the paper.

³Note that the inner minimization over j in (4) merely amounts to taking the minimum value among n real numbers and should not be seen as an embedded optimization problem.

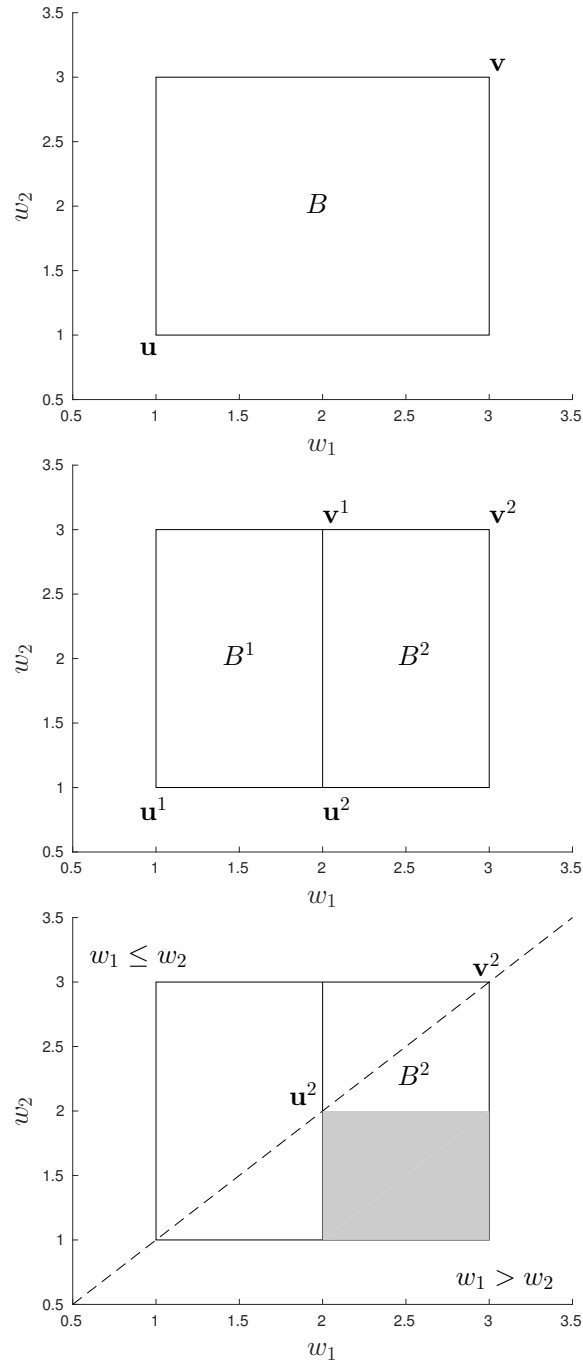


Figure 1: Illustration of the splitting procedure when $n = 2$ and $d = 1$. *Top*: a box $B = [\mathbf{u}, \mathbf{v}]$ is a rectangular region of the plane of axis (w_1, w_2) with bottom-left and top-right corners at \mathbf{u} and \mathbf{v} . *Middle*: B is split into $B^1 = [\mathbf{u}^1, \mathbf{v}^1]$ and $B^2 = [\mathbf{u}^2, \mathbf{v}^2]$ by application of (6)–(8). *Bottom*: B^2 is corrected as in (9)–(10) to remove the shaded area that does not contain any feasible solution according to (5), i.e., $w_1 > w_2$ for all \mathbf{w} in the shaded area.

Lemma 1. For any d -dimensional box $B_j = [\mathbf{u}_j, \mathbf{v}_j] \subset \mathbb{R}^d$, we have, for $i = 1, \dots, N$,

$$\begin{cases} \min_{\mathbf{w}_j \in B_j} \mathbf{w}_j^T \mathbf{x}_i = \mathbf{u}_j^T \mathbf{x}_i + L_i(B_j) \\ \max_{\mathbf{w}_j \in B_j} \mathbf{w}_j^T \mathbf{x}_i = \mathbf{u}_j^T \mathbf{x}_i + U_i(B_j), \end{cases} \quad (11)$$

where,

$$\begin{cases} L_i(B_j) = (\mathbf{v}_j - \mathbf{u}_j)^T (\mathbf{x}_i)_- \\ U_i(B_j) = (\mathbf{v}_j - \mathbf{u}_j)^T (\mathbf{x}_i)_+. \end{cases} \quad (12)$$

Proof. Any $\mathbf{w}_j \in [\mathbf{u}_j, \mathbf{v}_j]$ can be expressed as

$$\mathbf{w}_j = \mathbf{u}_j + \boldsymbol{\alpha} \odot (\mathbf{v}_j - \mathbf{u}_j), \quad (13)$$

where \odot denotes the entrywise product of vectors and $\boldsymbol{\alpha} \in [0, 1]^d$. Thus, for $i = 1, \dots, N$, we have

$$\begin{cases} \min_{\mathbf{w}_j \in [\mathbf{u}_j, \mathbf{v}_j]} \mathbf{w}_j^T \mathbf{x}_i = \mathbf{u}_j^T \mathbf{x}_i + \min_{\boldsymbol{\alpha} \in [0, 1]^d} (\boldsymbol{\alpha} \odot (\mathbf{v}_j - \mathbf{u}_j))^T \mathbf{x}_i \\ \max_{\mathbf{w}_j \in [\mathbf{u}_j, \mathbf{v}_j]} \mathbf{w}_j^T \mathbf{x}_i = \mathbf{u}_j^T \mathbf{x}_i + \max_{\boldsymbol{\alpha} \in [0, 1]^d} (\boldsymbol{\alpha} \odot (\mathbf{v}_j - \mathbf{u}_j))^T \mathbf{x}_i \end{cases} \quad (14)$$

with

$$\begin{aligned} \min_{\boldsymbol{\alpha} \in [0, 1]^d} (\boldsymbol{\alpha} \odot (\mathbf{v}_j - \mathbf{u}_j))^T \mathbf{x}_i &= \min_{\boldsymbol{\alpha} \in [0, 1]^d} \boldsymbol{\alpha}^T ((\mathbf{v}_j - \mathbf{u}_j) \odot \mathbf{x}_i) \\ &= \sum_{k=1}^d \min_{\alpha_k \in [0, 1]} \alpha_k (v_{j,k} - u_{j,k}) x_{i,k} \\ &= \mathbf{1}^T ((\mathbf{v}_j - \mathbf{u}_j) \odot \mathbf{x}_i)_- \\ &= L_i(B_j) \end{aligned} \quad (15)$$

and, similarly, $\max_{\boldsymbol{\alpha} \in [0, 1]^d} (\boldsymbol{\alpha} \odot (\mathbf{v}_j - \mathbf{u}_j))^T \mathbf{x}_i = U_i(B_j)$. \square

3.2.1 Lower bound based on pointwise minimum errors

The first lower bound is based on a pointwise decomposition of the optimization problem with respect to the index i of data points. In particular, we use the fact that, for any $i \in \{1, \dots, N\}$ and $j \in \mathcal{Q}$, the pointwise error of a parameter vector \mathbf{w}_j at a given point (\mathbf{x}_i, y_i) ,

$$e_i(\mathbf{w}_j) = y_i - \mathbf{w}_j^T \mathbf{x}_i, \quad (16)$$

can be made smaller in magnitude if we are not trying to simultaneously minimize the errors at other points. Therefore, the global cost $J_{\text{SW}}(\mathbf{w})$ must be at least as large as the sum of independently optimized pointwise errors. Formally, for any $i \in \{1, \dots, N\}$ and box $B_j = [\mathbf{u}_j, \mathbf{v}_j] \subset \mathbb{R}^d$, let

$$\begin{aligned} e_i^L(B_j) &= e_i(\mathbf{u}_j) - L_i(B_j) \\ e_i^U(B_j) &= e_i(\mathbf{u}_j) - U_i(B_j) \end{aligned} \quad (17)$$

with $L_i(B_j)$ and $U_i(B_j)$ as in (12). Then, we have the following lower bound.

Lemma 2. Given a box $B = B_1 \times \dots \times B_n$ with $B_j = [\mathbf{u}_j, \mathbf{v}_j] \subset \mathbb{R}^d$, $j = 1, \dots, n$, and the notations defined above,

$$J(B) = \sum_{i=1}^N \min_{j \in \mathcal{Q}} \left\{ (e_i^U(B_j))_+^2 + (e_i^L(B_j))_-^2 \right\} \quad (18)$$

is a lower bound on $\min_{\mathbf{w} \in B} J_{\text{SW}}(\mathbf{w})$.

Proof. For any $\mathbf{a} \in B$,

$$\min_{j \in \mathcal{Q}} e_i^2(\mathbf{a}_j) \geq \min_{\mathbf{w} \in B} \min_{j \in \mathcal{Q}} e_i^2(\mathbf{w}_j) = \min_{j \in \mathcal{Q}} \min_{\mathbf{w} \in B} e_i^2(\mathbf{w}_j) \quad (19)$$

and, by summing over i ,

$$J_{\text{SW}}(\mathbf{a}) \geq \sum_{i=1}^N \min_{j \in \mathcal{Q}} \min_{\mathbf{w} \in B} e_i^2(\mathbf{w}_j) = \sum_{i=1}^N \min_{j \in \mathcal{Q}} \min_{\mathbf{w}_j \in B_j} e_i^2(\mathbf{w}_j). \quad (20)$$

Since this holds for any $\mathbf{a} \in B$, it holds in particular for the one yielding the minimum of J_{SW} over B and we obtain

$$\min_{\mathbf{w} \in B} J_{\text{SW}}(\mathbf{w}) \geq \sum_{i=1}^N \min_{j \in \mathcal{Q}} \min_{\mathbf{w}_j \in B_j} e_i^2(\mathbf{w}_j). \quad (21)$$

On the other hand, Lemma 1 yields

$$\begin{cases} \min_{\mathbf{w}_j \in B_j} e_i(\mathbf{w}_j) = e_i(\mathbf{u}_j) - U_i(B_j) = e_i^U(B_j) \\ \max_{\mathbf{w}_j \in B_j} e_i(\mathbf{w}_j) = e_i(\mathbf{u}_j) - L_i(B_j) = e_i^L(B_j), \end{cases} \quad (22)$$

and thus

$$\min_{\mathbf{w}_j \in B_j} |e_i(\mathbf{w}_j)| = \begin{cases} e_i^U(B_j), & \text{if } e_i^U(B_j) > 0 \\ |e_i^L(B_j)|, & \text{if } e_i^L(B_j) < 0 \\ 0, & \text{otherwise,} \end{cases} \quad (23)$$

which can be rewritten as

$$\min_{\mathbf{w}_j \in B_j} e_i^2(\mathbf{w}_j) = (e_i^U(B_j))_+^2 + (e_i^L(B_j))_-^2. \quad (24)$$

Combining (21) and (24) completes the proof. \square

From a computing time perspective, once an upper bound has been computed as $\bar{J}(B) = J_{\text{SW}}(\mathbf{u})$, only the values of $L_i(B_j)$ and $U_i(B_j)$ are needed to compute $e_i^L(B_j)$ and $e_i^U(B_j)$ for the evaluation of the lower bound $\underline{J}(B)$ in Lemma 2; and, as will be seen below, these quantities will be used again.

3.2.2 Lower bound based on constant classifications

Let us denote by $I = \{1, \dots, N\}$ the set of all point indexes and, for all $j \in \mathcal{Q}$, the subset of indexes for which the classification remains constant and equal to j in a box $B = [\mathbf{u}, \mathbf{v}] \subset \mathbb{R}^{nd}$ by

$$I_j(B) = \{i \in I : \forall \mathbf{w} \in B, q_i(\mathbf{w}) = j\}, \quad (25)$$

where $q_i(\mathbf{w})$ is given by (3). We also define $I_0(B) = I \setminus \cup_{j=1}^n I_j(B)$ as the subset of the remaining indexes. Thus, any data index $i \in I$ is exactly in one and only one of the sets $I_j(B)$, $j = 0, \dots, n$. The following result shows how these index sets can be determined.

Lemma 3. *Given a box $B = B_1 \times \dots \times B_n$ with $B_j = [\mathbf{u}_j, \mathbf{v}_j] \subset \mathbb{R}^d$, $j = 1, \dots, n$, and the notations above, we have, for $j = 1, \dots, n$,*

$$I_j(B) = \left\{ i \in I : \max \{e_i^U(B_j)^2, e_i^L(B_j)^2\} < \min_{k < j} (e_i^U(B_k))_+^2 + (e_i^L(B_k))_-^2, \right. \\ \left. \max \{e_i^U(B_j)^2, e_i^L(B_j)^2\} \leq \min_{k > j} (e_i^U(B_k))_+^2 + (e_i^L(B_k))_-^2 \right\}.$$

Proof. For any $j \in \mathcal{Q}$, given the definition of q_i in (3) and the fact that ties are broken by setting $q_i(\mathbf{w})$ as the smallest mode index, the set $I_j(B)$ contains exactly the indexes $i \in I$ for which

$$\forall \mathbf{w} \in B, \forall k \in \mathcal{Q} \setminus \{j\}, \quad (y_i - \mathbf{w}_j^T \mathbf{x}_i)^2 \begin{cases} < (y_i - \mathbf{w}_k^T \mathbf{x}_i)^2, & \text{if } k < j \\ \leq (y_i - \mathbf{w}_k^T \mathbf{x}_i)^2, & \text{if } k > j. \end{cases}$$

Since the left-hand side and the right-hand side depend on different components, B_j and B_k , of B and since with the constraint $\mathbf{w} \in B$ the component $\mathbf{w}_j \in B_j$ can freely evolve independently of $\mathbf{w}_k \in B_k$, this condition is equivalent to

$$\forall \mathbf{w}_j \in B_j, \forall k \in \mathcal{Q} \setminus \{j\}, \forall \mathbf{w}_k \in B_k, \quad (y_i - \mathbf{w}_j^T \mathbf{x}_i)^2 \begin{cases} < (y_i - \mathbf{w}_k^T \mathbf{x}_i)^2, & \text{if } k < j \\ \leq (y_i - \mathbf{w}_k^T \mathbf{x}_i)^2, & \text{if } k > j, \end{cases}$$

which can be rewritten in the more compact form

$$\begin{cases} \max_{\mathbf{w}_j \in B_j} e_i^2(\mathbf{w}_j) < \min_{k < j} \min_{\mathbf{w}_k \in B_k} e_i^2(\mathbf{w}_k) \\ \max_{\mathbf{w}_j \in B_j} e_i^2(\mathbf{w}_j) \leq \min_{k > j} \min_{\mathbf{w}_k \in B_k} e_i^2(\mathbf{w}_k). \end{cases} \quad (26)$$

Using (24) with k instead of j in the right-hand sides and noting that (22) yields the left-hand sides as

$$\max_{\mathbf{w}_j \in B_j} e_i^2(\mathbf{w}_j) = \max \{e_i^U(B_j)^2, e_i^L(B_j)^2\} \quad (27)$$

completes the proof. \square

Once the sets $I_j(B)$, $j = 1, \dots, n$, have been determined, the following gives an improved lower bound by constraining the error over points with index in a set $I_j(B)$ to be computed with respect to a single linear model.

Lemma 4. *Given a box $B = [\mathbf{u}, \mathbf{v}] \subset \mathbb{R}^{nd}$ and the notations above, for any $\mathcal{J} \subseteq \mathcal{Q}$,*

$$\begin{aligned} \underline{J}(B) &= \sum_{i \in I_0(B)} \min_{j \in \mathcal{Q}} \left\{ (e_i^U(B_j))_+^2 + (e_i^L(B_j))_-^2 \right\} + \sum_{j \in \mathcal{J}} \min_{\mathbf{w}_j \in B_j} \sum_{i \in I_j(B)} (y_i - \mathbf{w}_j^T \mathbf{x}_i)^2 \\ &\quad + \sum_{j \in \mathcal{Q} \setminus \mathcal{J}} \sum_{i \in I_j(B)} (e_i^U(B_j))_+^2 + (e_i^L(B_j))_-^2 \end{aligned} \quad (28)$$

is a lower bound on $\min_{\mathbf{w} \in B} J_{\text{SW}}(\mathbf{w})$.

Proof. Since each data point index $i \in I$ is exactly in one and only one of the sets $I_j(B)$, $j = 0, \dots, n$, for all $\mathbf{w} \in B$, the cost function in (4) can be rewritten as

$$\begin{aligned} J_{\text{SW}}(\mathbf{w}) &= \sum_{j=0}^n \sum_{i \in I_j(B)} \min_{j \in \mathcal{Q}} (y_i - \mathbf{w}_j^T \mathbf{x}_i)^2 \\ &= \sum_{i \in I_0(B)} \min_{j \in \mathcal{Q}} e_i^2(\mathbf{w}_j) + \sum_{j=1}^n \sum_{i \in I_j(B)} e_i^2(\mathbf{w}_j), \end{aligned} \quad (29)$$

where the second line is due to the definition of the $I_j(B)$'s in (25). Then, the sum over $I_0(B)$ can be lower bounded as in (21), while each of the other terms only involves the j th subvector \mathbf{w}_j :

$$\min_{\mathbf{w} \in B} J_{\text{SW}}(\mathbf{w}) \geq \sum_{i \in I_0(B)} \min_{j \in \mathcal{Q}} \min_{\mathbf{w}_j \in B_j} e_i^2(\mathbf{w}_j) + \sum_{j=1}^n \min_{\mathbf{w}_j \in B_j} \sum_{i \in I_j(B)} e_i^2(\mathbf{w}_j). \quad (30)$$

Substituting (24) in the above for all $i \in I_0(B)$ proves the Lemma for the case $\mathcal{J} = \mathcal{Q}$. To complete the proof for any $\mathcal{J} \subset \mathcal{Q}$, we write

$$\sum_{j=1}^n \min_{\mathbf{w}_j \in B_j} \sum_{i \in I_j(B)} e_i^2(\mathbf{w}_j) \geq \sum_{j \in \mathcal{J}} \min_{\mathbf{w}_j \in B_j} \sum_{i \in I_j(B)} e_i^2(\mathbf{w}_j) + \sum_{j \in \mathcal{Q} \setminus \mathcal{J}} \sum_{i \in I_j(B)} \min_{\mathbf{w}_j \in B_j} e_i^2(\mathbf{w}_j) \quad (31)$$

and lower bound the second term by invoking again (24) for all $i \in I_j(B)$ and $j \in \mathcal{Q} \setminus \mathcal{J}$. \square

Lemma 4 can lower bound the error over the points with index in $I_j(B)$ in two different manners: the second term in (28) corresponds to a box-constrained least squares error, while the third term

is the sum of pointwise minimum errors for the j th model. Thus, the degree of freedom left by the choice of \mathcal{J} in Lemma 4 can be used to trade off computing time for accuracy: solving a constrained least squares problem yields a larger lower bound but is more demanding than summing pointwise minimum errors (the latter being computed with very little effort given that all quantities involved have already been computed when determining $I_j(B)$). Thus, Lemma 4 can be used to provide a sequence of lower bounds with increasing complexity: starting with $\mathcal{J} = \emptyset$, we increment the cardinality of \mathcal{J} and stop as soon as the lower bound reaches the upper bound, at which point the box can be discarded from the search.

3.3 Convergence

The convergence of the branch-and-bound algorithm follows from the tightness of the lower bound in Lemma 4. It is formally stated below under two assumptions.

Assumption 1. *The global optimum of Problem (4) is strictly positive:*

$$J^* = \min_{\mathbf{w} \in B_{\text{init}}} J_{\text{SW}}(\mathbf{w}) > 0. \quad (32)$$

Assumption 2. *Upper bounds in a box $B = [\mathbf{u}, \mathbf{v}]$ are computed as $\bar{J}(B) = J_{\text{SW}}(\mathbf{u})$ or such that $\bar{J}(B) \leq J_{\text{SW}}(\mathbf{u})$.*

Assumption 1 merely requires that the data cannot be exactly fitted by n linear models, which is most often the case with noisy measurements. Otherwise, the convergence of Algorithm 1 can be proved similarly for noiseless data by changing the stopping criterion on the relative accuracy, $(\bar{J} - \underline{J})/\bar{J} \leq \text{TOL}$, for one on the absolute accuracy: $\bar{J} - \underline{J} \leq \text{TOL}$. Assumption 2 simply requires that the upper bounds are at least as accurate as the straightforward computation of the cost function value at the box base point \mathbf{u} . Though this depends on the precise choice of heuristic for computing the upper bounds, Assumption 2 can always be made to hold simply by setting $\bar{J}(B) \leftarrow \min\{\bar{J}(B), J_{\text{SW}}(\mathbf{u})\}$.

Theorem 1. *Under Assumptions 1–2, Algorithm 1 as described above with lower bounds $\underline{J}(B)$ computed as in Lemma 4 with $\mathcal{J} = \mathcal{Q}$ converges in a finite number of iterations for any $\text{TOL} > 0$.*

Proof. Let t denote the iteration counter in Algorithm 1. The splitting rule detailed in Sect. 3.1 guarantees that the side lengths of the boxes decrease towards 0, i.e., $\mathbf{v} - \mathbf{u} \xrightarrow{t \rightarrow \infty} \mathbf{0}$ and $B \xrightarrow{t \rightarrow \infty} \{\mathbf{u}\}$, for all active boxes B with lower bound $\underline{J}(B)$ smaller than the upper bound \bar{J} . Then, by their definitions in (12), this implies the convergence of $L_i(B_j)$ and $U_i(B_j)$ towards zero for all data index $i \in I$ and mode $j \in \mathcal{Q}$. Recalling (17), this leads to $e_i^L(B_j) \xrightarrow{t \rightarrow \infty} e_i(\mathbf{u}_j)$ and $e_i^U(B_j) \xrightarrow{t \rightarrow \infty} e_i(\mathbf{u}_j)$. Hence, by the definition of $\underline{J}(B)$ in Lemma 4, we obtain for any active B that

$$\begin{aligned} \underline{J}(B) &\xrightarrow{t \rightarrow \infty} \sum_{i \in I_0(B)} \min_{j \in \mathcal{Q}} e_i^2(\mathbf{u}_j) + \sum_{j \in \mathcal{Q}} \min_{\mathbf{w}_j \in B_j} \sum_{i \in I_j(B)} e_i^2(\mathbf{w}_j) \\ &\xrightarrow{t \rightarrow \infty} \sum_{i \in I_0(B)} \min_{j \in \mathcal{Q}} e_i^2(\mathbf{u}_j) + \sum_{j \in \mathcal{Q}} \sum_{i \in I_j(B)} e_i^2(\mathbf{u}_j). \end{aligned} \quad (33)$$

By the definition of the index sets $I_j(B)$, $j = 0, \dots, n$, and since each data index $i \in I$ is exactly in one and only one of these sets, the right-hand side above equals $\sum_{i \in I} \min_{j \in \mathcal{Q}} e_i^2(\mathbf{u}_j)$ and we have

$$\underline{J}(B) \xrightarrow{t \rightarrow \infty} J_{\text{SW}}(\mathbf{u}). \quad (34)$$

Thus, for all $\epsilon > 0$, there is a finite iteration number T such that for all subsequent iterations $t \geq T$ and all boxes B in the list of active boxes \mathcal{B} ,

$$\epsilon \geq J_{\text{SW}}(\mathbf{u}) - \underline{J}(B) \geq \bar{J}(B) - \underline{J}(B), \quad (35)$$

where the second inequality is due to Assumption 2. In particular, consider some $\epsilon \in (0, \text{TOL} \cdot J^*)$ (such an ϵ exists by Assumption 1) and let B be the current box at iteration T . Recall that

Algorithm 1 is such that, at this iteration, $\bar{J} \leq \bar{J}(B)$ and $\underline{J} = \underline{J}(B)$. Then, the fact that \bar{J} is an upper bound on J^* ensures that the stopping criterion is met:

$$\frac{\bar{J} - \underline{J}}{\bar{J}} \leq \frac{\bar{J}(B) - \underline{J}(B)}{\bar{J}} \leq \frac{\epsilon}{\bar{J}} \leq \frac{TOL \cdot J^*}{\bar{J}} \leq TOL, \quad (36)$$

and the algorithm stops at iteration T . \square

Remark 1. *The worst-case time complexity of the branch-and-bound approach described above remains exponential in the number of variables, nd , as is naturally expected from the NP-hardness of Problem (4) [12]. However, in practice, the average time complexity can be much lower thanks to the rapid contraction of the search space offered by the bounding scheme. In the product nd , the main parameter influencing the computing time is the number of modes, n . Indeed, increasing n does not only increase the number of splits required to obtain small boxes with small values of $L_i(B_j)$ and $U_i(B_j)$ as for d , but also affects the quality of the lower bounds involving the operation $\min_{j \in \mathcal{Q}}$. Therefore, for large n , the iterative approach described in the next section can be more suitable.*

4 Bounded-error estimation

Consider now the problem of estimating both the minimal number n and the collection of n parameter vectors $\{\mathbf{w}_j\}_{j=1}^n$ such that, given a threshold ϵ on the error,

$$\min_{j \in \{1, \dots, n\}} |y_i - \mathbf{w}_j^T \mathbf{x}_i| \leq \epsilon, \quad i = 1, \dots, N. \quad (37)$$

This setting has a long history in the hybrid system identification literature [5, 3]. In this framework, the number of models is not fixed a priori but estimated in order to satisfy the bound on the error. It is shown in [2] that this problem is NP-hard and standard approaches work in a greedy manner by estimating one model at each iteration until the bounded-error constraint is satisfied for all points. After estimating the j th model, the data correctly approximated in the sense of (37) are removed from the data set before proceeding with the next iteration. In [5], the j th model is defined as the one correctly approximating the maximum number of points and is estimated by finding the feasible subsystem of inequalities (37) of maximum cardinality (the MAX FS problem). However, finding such a subsystem is itself shown to be NP-hard in [1] and [5] has to rely on a suboptimal heuristic.

Focusing on the noiseless case, [3] poses this problem as a sparse optimization one: the model yielding a perfect fit of the maximum number of points is the one that yields the sparsest error vector $\mathbf{e}_j = \mathbf{y}_j - \mathbf{X}_j \mathbf{w}_j$ (with \mathbf{X}_j and \mathbf{y}_j containing the remaining data at the j th iteration). Though sparse optimization problems are NP-hard in general [15], the advantage of this point of view is that guarantees of convergence of heuristics (such as the ℓ_1 -norm based convex relaxation) to the exact solution can be obtained [7]. However, in the noisy case, \mathbf{e}_j is not sparse anymore but “compressible” and the guarantees become weaker. Then, the method of [3] basically behaves as a robust estimator based on the ℓ_1 -loss when considering points from the non-dominant mode as outliers [4].

The two points of view from [5] and [3] can be united by considering that a vector is sparse when a large fraction of its entries are not larger than a given threshold ϵ . This leads to the minimization of the loss function $\ell_{0,\epsilon}(e) = \mathbf{1}_{|e| > \epsilon}$, plotted on the left of Fig. 2, instead of the ℓ_0 -pseudo norm of \mathbf{e}_j . This minimization is directly equivalent to the MAX FS problem of [5], while the setting of [3] is recovered with $\epsilon = 0$. By generalizing these ideas, we also consider the minimization of saturated loss functions,

$$\forall p \in \{1, 2\}, \epsilon > 0, \quad \ell_{p,\epsilon}(e) = (\min\{|e|, \epsilon\})^p, \quad (38)$$

which yield a nonzero loss for $0 < e < \epsilon$ (see the middle and right plots of Fig. 2). Indeed, depending on the noise model, a better alternative might be to not only obtain error vectors with many small entries, but also to minimize a standard loss over these small entries, such as a squared error or absolute deviation.

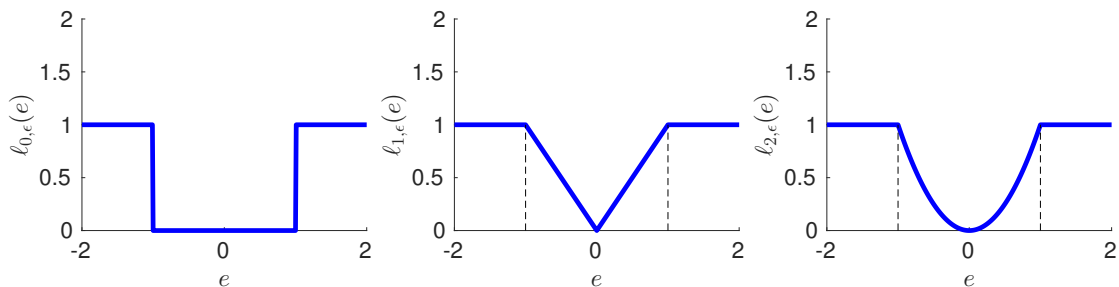


Figure 2: Saturated loss functions $\ell_{p,\epsilon}(e)$ plotted for $\epsilon = 1$ with, from left to right, $p = 0$, $p = 1$ and $p = 2$.

Overall, the considered bounded-error approach relies on iteratively solving for $j = 1, 2, \dots$ the nonconvex optimization problem

$$\min_{\mathbf{w}_j \in B \subset \mathbb{R}^d} J_{\text{BE}}^p(\mathbf{w}_j) = \sum_{i \in I_j} \min \{ |y_i - \mathbf{w}_j^T \mathbf{x}_i|^p, \epsilon^p \} \quad (39)$$

or

$$\min_{\mathbf{w}_j \in B \subset \mathbb{R}^d} J_{\text{BE}}^0(\mathbf{w}_j) = \sum_{i \in I_j} \mathbf{1}_{|y_i - \mathbf{w}_j^T \mathbf{x}_i| > \epsilon}, \quad (40)$$

where $I_1 = \{1, \dots, N\}$ and $I_j = \{i \in I_{j-1} : |y_i - \mathbf{w}_j^T \mathbf{x}_i| > \epsilon\}$, until the set of indexes I_j of points left in the data set and unassigned to a mode becomes empty.

Note that these problems are also of interest for the robust estimation of a single (non-hybrid) linear model in the presence of outliers, as will be considered in the examples of Sect. 5.3–5.4 (see also [4] for an analysis of the sparse optimization method of [3] in this context).

4.1 Branch-and-bound approach

We now focus on the minimization of J_{BE}^p in the cases $p = 2$ and $p = 0$. In devising a branch-and-bound algorithm to solve (39) and (40), we will use techniques developed in Section 3 for switching regression. Indeed, the bounded-error cost functions are closely related to J_{SW} and their minimization can be seen as a switching regression problem with two modes, one of which having a model with constant error equal to ϵ^2 for J_{BE}^2 or 1 for J_{BE}^0 .

Upper bounds on the bounded-error cost functions are simply computed as the cost function value at the base point \mathbf{u} of the current box $B = [\mathbf{u}, \mathbf{v}] \subset \mathbb{R}^d$. In the case $p = 2$, we also periodically use the heuristic described below with an initialization at the center of the box.

Since in bounded-error estimation, each optimization problem, either (39) or (40), aims at the estimation of a single parameter vector, there is no symmetry that needs to be broken with additional constraints as we did for switching regression with (5). Therefore, the splitting rule remains simple: we split in the middle of the longest side of the box with (6)–(8).

4.2 Heuristic for bounded-error estimation

We introduce a new heuristic inspired by the k -LinReg algorithm [10] to minimize the bounded-error cost with the $\ell_{2,\epsilon}$ loss. Starting at iteration $t = 0$ from a parameter vector \mathbf{w}^0 and a data set $\{(\mathbf{x}_i, y_i)\}_{i \in I_j}$, the algorithm alternates between a classification step and a least squares regression as follows.

1. Identify the set of points satisfying the bounded-error criterion:

$$I_B(t) = \{i \in I_j : (y_i - \mathbf{x}_i^T \mathbf{w}_j^t)^2 \leq \epsilon^2\}. \quad (41)$$

2. Update the parameter vector with

$$\mathbf{w}_j^{t+1} = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \sum_{i \in I_B(t)} (y_i - \mathbf{w}^T \mathbf{x}_i)^2. \quad (42)$$

3. Set $t \leftarrow t + 1$ and repeat from step (1) until convergence.

The following shows that this algorithm is a descent method for problem (39) with $p = 2$.

Proposition 1. *The algorithm above monotonically decreases the bounded-error cost function with $\ell_{2,\epsilon}$ loss, J_{BE}^2 .*

Proof. Partitioning the set of all indexes into three disjoint subsets as

$$I_j = I_B(t) \cup (I_B(t+1) \setminus I_B(t)) \cup (I_j \setminus (I_B(t) \cup I_B(t+1))), \quad (43)$$

we write the difference between the cost function values at the new estimate, \mathbf{w}_j^{t+1} , and at the old one, \mathbf{w}_j^t , as

$$\begin{aligned} J_{\text{BE}}^2(\mathbf{w}_j^{t+1}) - J_{\text{BE}}^2(\mathbf{w}_j^t) &= \sum_{i \in I_B(t)} \ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^{t+1}) - \ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^t) \\ &+ \sum_{i \in I_B(t+1) \setminus I_B(t)} \ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^{t+1}) - \ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^t) \\ &+ \sum_{I_j \setminus (I_B(t) \cup I_B(t+1))} \ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^{t+1}) - \ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^t). \end{aligned}$$

The first sum is negative since the new estimate \mathbf{w}_j^{t+1} is precisely defined in (42) as the minimizer of the loss over $I_B(t)$. The second sum is also negative since, for any $i \in I_B(t+1) \setminus I_B(t)$, the point (\mathbf{x}_i, y_i) is within the bounded-error tolerance for the new estimate but saturates the loss for the old one:

$$\ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^{t+1}) \leq \epsilon^2 = \ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^t). \quad (44)$$

Finally, the third sum is zero because for $i \notin I_B(t) \cup I_B(t+1)$, the loss is saturated for both estimates:

$$\ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^{t+1}) = \ell_{2,\epsilon}(y_i - \mathbf{x}_i^T \mathbf{w}_j^t) = \epsilon^2. \quad (45)$$

Therefore, $J_{\text{BE}}^2(\mathbf{w}_j^{t+1}) \leq J_{\text{BE}}^2(\mathbf{w}_j^t)$, and the cost function value can only decrease. \square

To initialize the first upper bound in the global optimization Algorithm 1, we use 100 runs of the algorithm above with random initializations.

4.3 Lower bounds

Following the path of Sect. 3, we derive a number of lower bounds of increasing tightness.

Lemma 5. *Given a box $B = [\mathbf{u}, \mathbf{v}] \subset \mathbb{R}^d$, with the notations of Sect. 3 and in particular (17),*

$$\underline{J}(B) = \sum_{i \in I_j} \min \left\{ (e_i^U(B))_+^2 + (e_i^L(B))_-^2, \epsilon^2 \right\} \quad (46)$$

is a lower bound on $\min_{\mathbf{w}_j \in B} J_{\text{BE}}^2(\mathbf{w}_j)$.

Proof. We proceed as in the beginning of the proof of Lemma 2: since for any $\mathbf{a} \in B$,

$$\min \{e_i^2(\mathbf{a}), \epsilon^2\} \geq \min_{\mathbf{w}_j \in B} \min \{e_i^2(\mathbf{w}_j), \epsilon^2\} = \min \left\{ \min_{\mathbf{w}_j \in B} e_i^2(\mathbf{w}_j), \epsilon^2 \right\}, \quad (47)$$

we have

$$\min_{\mathbf{w}_j \in B} \sum_{i \in I_j} \min \{e_i^2(\mathbf{w}_j), \epsilon^2\} \geq \sum_{i \in I_j} \min \left\{ \min_{\mathbf{w}_j \in B} e_i^2(\mathbf{w}_j), \epsilon^2 \right\}. \quad (48)$$

Then, noticing that the sum in the left-hand side is $J_{\text{BE}}^2(\mathbf{w}_j)$ and introducing (24) in the right-hand side completes the proof. \square

Let us define the index sets

$$I_1(B) = \left\{ i \in I_j : \max_{\mathbf{w} \in [\mathbf{u}, \mathbf{v}]} e_i^2(\mathbf{w}) \leq \epsilon^2 \right\}, \quad (49)$$

$$I_2(B) = \left\{ i \in I_j : \min_{\mathbf{w} \in [\mathbf{u}, \mathbf{v}]} e_i^2(\mathbf{w}) > \epsilon^2 \right\} \quad (50)$$

and

$$I_0(B) = I_j \setminus (I_1(B) \cup I_2(B)). \quad (51)$$

These sets can be easily determined by using (27) for $I_1(B)$ and (24) for $I_2(B)$. Then, we obtain the following lower bound.

Lemma 6. *Given a box $B = [\mathbf{u}, \mathbf{v}] \subset \mathbb{R}^d$ and the notations above,*

$$\underline{J}(B) = \sum_{i \in I_0(B)} \min \left\{ (e_i^U(B))_+^2 + (e_i^L(B))_-^2, \epsilon^2 \right\} + \min_{\mathbf{w}_j \in B} \sum_{i \in I_1(B)} e_i^2(\mathbf{w}_j) + |I_2(B)|\epsilon^2 \quad (52)$$

is a lower bound on $\min_{\mathbf{w}_j \in B} J_{\text{BE}}^2(\mathbf{w}_j)$.

Proof. Since each data index $i \in I_j$ is exactly in one and only one of the sets $I_0(B)$, $I_1(B)$ and $I_2(B)$, the cost function can be decomposed as

$$J_{\text{BE}}^2(\mathbf{w}_j) = \sum_{k=0}^2 \sum_{i \in I_k(B)} \min \{ e_i^2(\mathbf{w}_j), \epsilon \}. \quad (53)$$

By definition of the index sets, this yields, for any $\mathbf{w}_j \in B$,

$$J_{\text{BE}}^2(\mathbf{w}_j) = \sum_{i \in I_0(B)} \min \{ e_i^2(\mathbf{w}_j), \epsilon \} + \sum_{i \in I_1(B)} e_i^2(\mathbf{w}_j) + |I_2(B)|\epsilon^2. \quad (54)$$

We can use the fact that Lemma 5 holds similarly for a sum over any index set instead of I_j , and in particular for $I_0(B)$, to lower bound the first term. Then, the result follows since the minimum of the sum of three terms is larger than or equal to the sum of the minimum of each term. \square

As for the switching regression case of Section 3, the only demanding task in the lower bound of Lemma 6 is to solve a box-constrained least squares problem over the points with index in $I_1(B)$.

Using similar index sets, we also obtain a version of Lemma 6 for the $\ell_{0,\epsilon}$ loss function.

Lemma 7. *Given a box $B = [\mathbf{u}, \mathbf{v}] \subset \mathbb{R}^d$ and the notations above, $\underline{J}(B) = |I_2(B)|$ is a lower bound on $\min_{\mathbf{w}_j \in B} J_{\text{BE}}^0(\mathbf{w}_j)$.*

4.4 Convergence

As for the switching regression case studied in Sect. 3.3, convergence is obtained, for both the $\ell_{2,\epsilon}$ and the $\ell_{0,\epsilon}$ loss functions, from the tightness of the bounds, under the following assumptions.

Assumption 3. *For $p = 2$ (respectively, $p = 0$), the global optimum of Problem (39) (resp. Problem (40)) is strictly positive:*

$$J^* = \min_{\mathbf{w}_j \in B_{\text{init}}} J_{\text{BE}}^p(\mathbf{w}_j) > 0. \quad (55)$$

Assumption 4. *For any $p \in \{0, 2\}$, upper bounds in a box $B = [\mathbf{u}, \mathbf{v}]$ are computed as $\bar{J}(B) = J_{\text{BE}}^p(\mathbf{u})$ or such that $\bar{J}(B) \leq J_{\text{BE}}^p(\mathbf{u})$.*

Theorem 2. *Under Assumptions 3–4, the branch-and-bound algorithm described above to minimize J_{BE}^p with $p \in \{0, 2\}$ and lower bounds $\underline{J}(B)$ computed as in Lemma 6 for $p = 2$ or Lemma 7 for $p = 0$ converges in a finite number of iterations for any $TOL > 0$.*

Proof. The proof is similar to the one of Theorem 1. Its adaptation requires only to show that the lower bounds converge towards the cost function value at the box base point \mathbf{u} as the iteration counter t tends to the infinity: $\underline{J}(B) \xrightarrow{t \rightarrow \infty} J_{\text{BE}}^p(\mathbf{u})$. As in Theorem 1, this is a consequence of the splitting rule which guarantees that the remaining boxes shrink towards a single point: $B \xrightarrow{t \rightarrow \infty} \{\mathbf{u}\}$. Then, for $p = 0$, this directly leads to

$$\underline{J}(B) = |I_2(B)| \xrightarrow{t \rightarrow \infty} |\{i \in I_j : e_i^2(\mathbf{u}) > \epsilon^2\}| = J_{\text{BE}}^0(\mathbf{u}). \quad (56)$$

For $p = 2$, $B \xrightarrow{t \rightarrow \infty} \{\mathbf{u}\}$ yields, for all data index $i \in I_j$, the convergence of $L_i(B)$ and $U_i(B)$ towards zero and thus $e_i^L(B) \xrightarrow{t \rightarrow \infty} e_i(\mathbf{u})$ and $e_i^U(B) \xrightarrow{t \rightarrow \infty} e_i(\mathbf{u})$, which imply

$$\begin{aligned} \underline{J}(B) &\xrightarrow{t \rightarrow \infty} \sum_{i \in I_0(B)} \min \{e_i^2(\mathbf{u}), \epsilon^2\} + \min_{\mathbf{w}_j \in B} \sum_{i \in I_1(B)} e_i^2(\mathbf{w}_j) + |I_2(B)|\epsilon^2 \\ &\xrightarrow{t \rightarrow \infty} \sum_{i \in I_0(B)} \min \{e_i^2(\mathbf{u}), \epsilon^2\} + \sum_{i \in I_1(B)} e_i^2(\mathbf{u}) + |I_2(B)|\epsilon^2 \\ &\xrightarrow{t \rightarrow \infty} J_{\text{BE}}^2(\mathbf{u}). \end{aligned} \quad (57)$$

Then, we conclude as in Theorem 1 using Assumptions 3–4. \square

5 Numerical experiments

Four sets of experiments are performed to validate the proposed algorithms: for switching regression in Sect. 5.1, for bounded-error identification in Sect. 5.2, for robust estimation in Sect. 5.3 and for exact recovery under sparse noise in Sect. 5.4. In all experiments, the initial box bounds on all variables are set to $[-10, 10]$. The tolerance on the relative optimality gap is set to $TOL = 0.001$. The accuracy is measured in terms of the normalized parametric mean squared error, $\text{NMSE} = \sum_{j=1}^n \|\boldsymbol{\theta}_j - \mathbf{w}_j\|_2^2 / \|\boldsymbol{\theta}_j\|_2^2$, where $\boldsymbol{\theta}_j$ is the j th parameter vector of the true system, and the classification error rate (CE) equal to the fraction of data points for which the mode q_i is incorrectly estimated. The computing times refer to Matlab implementations running on a laptop with a 3Ghz i7-dual core processor. To set the accuracy reference, we use an *oracle* based on independent least squares estimations with knowledge of the true classification of the data points (the mode q_i or the inlier/outlier categorization).

5.1 Switching linear regression

We first evaluate the global optimization approach to switching regression proposed in Sect. 3 with respect to its average computing time as a function of the number of modes n , the dimension d and the number of data N . For each problem size, we report the average and standard deviation of the computing time over 10 trials, in which the regression vectors \mathbf{x}_i and the true parameter vectors $\boldsymbol{\theta}_j$ are randomly drawn from a uniform distribution in $[-5, 5]^d$. For regression vectors, we slightly alterate this in order to avoid the presence of vectors close to the origin (which yield data points that are consistent with all linear models and do not bring information for the estimation). The outputs are then generated with $y_i = \boldsymbol{\theta}_{q_i}^T \mathbf{x}_i + v_i$, where the mode q_i is uniformly drawn in \mathcal{Q} and v_i is a centered Gaussian noise of standard deviation $\sigma_v = 0.1$.

The results in Table 1 show that switching regression problems with up to 10 000 points in dimension 5 can be solved in about one minute on a standard laptop. But, as expected, the computing time quickly increases with the dimension and the number of modes. Yet, these results support the claim that the complexity of the proposed approach remains reasonably low with respect to the number of data. In particular, Table 1 suggests an even less than linear complexity in N , indicating that the number of data does not critically influences the number of iterations and mostly affects the linear algebra and convex optimization operations. This would explain why the observed complexity in N is sub-linear instead of linear: when N increases, such operations benefit from standard parallel processing features that do not offer significant speed-ups for small

Table 1: Average and standard deviation of the computing time, the NMSE and the classification error rate (CE) for the global optimization of a switching linear model with n modes for different dimensions d and number of data N .

n	d	N	Time (s)	NMSE $\times 10^4$	CE (%)	
2	2	500	0.1 ± 0.1	0.252 ± 0.419	1.2 ± 0.4	
		1 000	0.2 ± 0.1	0.092 ± 0.079	1.2 ± 0.2	
		10 000	0.6 ± 0.4	0.018 ± 0.022	1.6 ± 0.7	
	3	500	0.8 ± 1.3	0.101 ± 0.058	0.9 ± 0.5	
		1 000	0.6 ± 0.3	0.335 ± 0.862	0.8 ± 0.5	
		10 000	2.3 ± 1.5	0.009 ± 0.008	1.0 ± 0.4	
	4	500	4.1 ± 2.9	0.128 ± 0.063	1.0 ± 0.6	
		1 000	5.8 ± 6.9	0.064 ± 0.065	0.8 ± 0.4	
		10 000	11.9 ± 11.9	0.005 ± 0.002	0.7 ± 0.1	
	5	500	24.0 ± 20.7	0.145 ± 0.107	0.7 ± 0.5	
		1 000	35.3 ± 29.9	0.078 ± 0.009	0.6 ± 0.3	
		10 000	66.7 ± 20.1	0.009 ± 0.005	0.6 ± 0.1	
	3	2	500	1.3 ± 0.9	0.899 ± 0.824	2.3 ± 0.6
			1 000	1.8 ± 1.4	0.184 ± 0.126	2.3 ± 0.8
			10 000	3.8 ± 2.2	0.036 ± 0.045	2.2 ± 0.3
3		500	22.8 ± 23.2	0.289 ± 0.177	1.8 ± 0.6	
		1 000	50.7 ± 45.3	0.143 ± 0.088	1.9 ± 0.8	
		10 000	72.4 ± 38.8	0.026 ± 0.017	1.8 ± 0.4	
4		500	783 ± 626	0.267 ± 0.131	1.5 ± 0.2	
		1 000	1404 ± 977	0.098 ± 0.015	1.4 ± 0.5	
		10 000	2061 ± 1239	0.015 ± 0.005	1.5 ± 0.1	

N . Note that given the NP-hardness of the problem, the high complexity with respect to n and d appears hardly overcomable by any global optimization approach.

Regarding the accuracy of the approach, both the parametric (NMSE) and classification (CE) errors reported in Table 1 are very low and comparable to the ones obtained by the oracle (not reproduced here). In particular, the few remaining classification errors are inherently due to the fact that noisy data points generated with one mode can be better approximated by another one.

Switched system identification example Consider the benchmark example from [20], where the aim is to identify the dynamical system arbitrarily switching between $n = 2$ modes as

$$y_i = \begin{cases} -0.9y_{i-1} + u_i + v_i, & \text{if } q_i = 1 \\ 0.7y_{i-1} - u_i + v_i, & \text{if } q_i = 2 \end{cases} \quad (58)$$

from $N = 1000$ data points with $\mathbf{x}_i = [y_{i-1}, u_i]^T$ and v_i a centered Gaussian noise of standard deviation $\sigma_v = 0.2$. Over 100 trials with random input (u_i) and noise sequences, the proposed algorithm obtains an average NMSE of 1.1547×10^{-4} very close to the one of the oracle equal to 1.1381×10^{-4} and about 4 times smaller than the 4.8356×10^{-4} reported in [20]. In this experiment, the average computing time was only 0.43 ± 0.05 seconds and the dynamical nature of the data did not seem to alter the efficiency of the algorithm.

5.2 Bounded-error identification of switched dynamical systems

Consider now the example in [3] where the aim is to estimate the system arbitrarily switching between $n = 3$ modes as

$$y_i = \begin{cases} -0.4y_{i-1} + 0.25y_{i-2} - 0.15u_i + 0.08u_{i-1} + v_i, & \text{if } q_i = 1 \\ 1.55y_{i-1} - 0.58y_{i-2} - 2.1u_i + 0.96u_{i-1} + v_i, & \text{if } q_i = 2 \\ y_{i-1} - 0.24y_{i-2} - 0.65u_i + 0.3u_{i-1} + v_i, & \text{if } q_i = 3 \end{cases} \quad (59)$$

Table 2: NMSE and computing time (seconds) for the identification of the system (59) with different loss functions.

n	N	$\ell_{0,\epsilon}$ -loss		$\ell_{2,\epsilon}$ -loss		ℓ_1 -loss [3]	
		NMSE	Time	NMSE	Time	NMSE	Time
3	300	0.0030	4.5	0.0018	67	0.0502	0.2
	3000	0.00046	35.0	0.00003	123	0.00003	3.4
5	300	0.0069	14.3	0.0052	417.2	2.8128	0.2
	3000	0.0029	95.3	0.00025	926.3	2.0917	5.6

from $N = 300$ data points with $\mathbf{x}_i = [y_{i-1}, y_{i-2}, u_i, u_{i-1}]^T$ and a centered Gaussian noise v_i of standard deviation σ_v such that the signal-to-noise ratio is about 30 dB. Applying the iterative bounded-error approach depicted in Sect. 4 with the $\ell_{0,\epsilon}$ and $\ell_{2,\epsilon}$ losses for $\epsilon = 1.5\sigma_v$, we obtain rather accurate estimates of the parameters, as reported in Table 2. In this setting, the overall computing time for the multiple global optimizations of (40) or (39) remains reasonable: a few seconds for the $\ell_{0,\epsilon}$ -loss and close to one minute for the $\ell_{2,\epsilon}$ -loss.

In addition, the low complexity with respect to the number of data of the proposed approach is here also verified: in a similar experiment with ten times more data ($N = 3000$), the computing time was only multiplied by two for the $\ell_{2,\epsilon}$ loss and grew linearly for the $\ell_{0,\epsilon}$ loss.

Of course, on this example taken from [3], the sparse optimization method developed in [3] and based on ℓ_1 -minimization can also achieve a good accuracy in a shorter amount of time. However, extending the system (59) with 2 more modes while maintaining a uniform distribution for q_i yields a much more difficult problem for this method which then fails to return relevant estimates. Indeed, in this case, the dominant mode generates only about 20% of the data points and the optimal error vector becomes far from sparse/compressible, hence breaking the desired behavior of convex relaxations. In such a case, it is critical to consider the global optimization of the nonconvex problem. Applying the proposed approach, we obtain accurate estimates of the five parameter vectors in reasonable time (see the two last rows of Table 2).

5.3 Bounded-error estimation in the presence of outliers

We now evaluate the bounded-error approach for robust estimation in the presence of outliers. In this case, the problems (39) or (40) are solved only once to estimate a single model from the maximal number of points that can be considered as inliers. We compare the proposed algorithms with the standard ℓ_1 -minimization in a setting similar to the one in [4]: an increasing fraction r of 500 data points in dimension 4 are corrupted by outliers ξ_i drawn from a Gaussian distribution with mean 100 and variance 1000^2 : $y_i = \boldsymbol{\theta}^T \mathbf{x}_i + v_i + \xi_i$. The results are reported in Fig. 3 (top) for $\epsilon = 1.5\sigma_v$ and $\sigma_v = 0.1$.

The proposed global minimizations of the $\ell_{0,\epsilon}$ and $\ell_{2,\epsilon}$ losses perform similarly well at rejecting outliers and yielding accurate estimates with an error close to the one of the oracle. On the other hand, the error of the ℓ_1 -minimization quickly increases with more than 60% of outliers. Regarding the computing time, these results are obtained in a few seconds for the $\ell_{0,\epsilon}$ loss and about 30 seconds for the $\ell_{2,\epsilon}$. This is slower than the ℓ_1 -minimization heuristic (which takes less than a second) but still reasonable given the gain in accuracy.

In addition, the ℓ_1 -based method is known to break down when estimating affine models with data corrupted by more than 50% of positive gross errors ($\xi_i > 0$) [4]. Indeed, the bottom plot of Fig. 3 shows that, when \mathbf{x}_i includes a constant component in order to implement an affine model and ξ_i is replaced by $|\xi_i|$, the error of this approach quickly increases with the number of outliers. At the opposite, the proposed global optimizations are not affected by such adversarial conditions and the error remains comparable to the one of the oracle.

5.4 Exact recovery with sparse and gross measurement errors

Most of the analysis of convex relaxations in [4] actually deals with exact recovery in the case where the data is noiseless ($v_i = 0$) except for sparse and gross errors. The proposed bounded-error algorithms can also be applied in this setting, e.g., with $\epsilon = 10^{-6}$, and evaluated on the basis

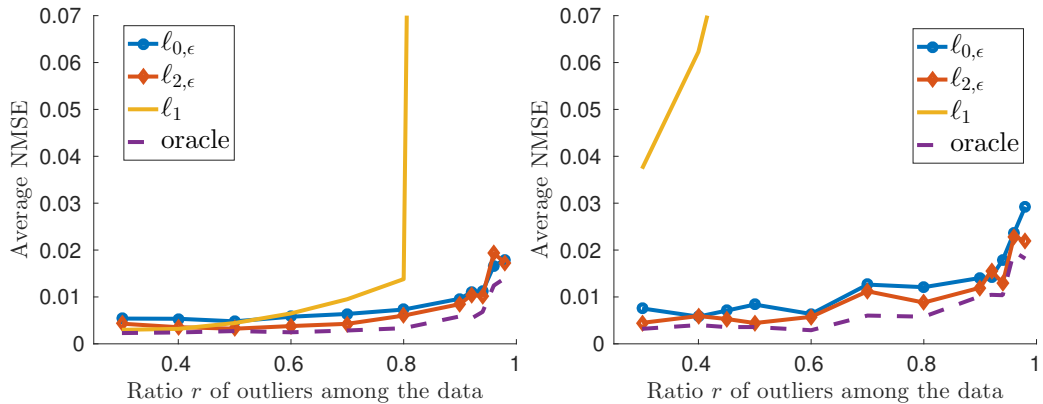


Figure 3: Average NMSE over 10 trials vs the ratio of outliers for a linear model (left) and an affine model (right).

Table 3: Exact recovery with sparse and gross measurement errors occurring only at rN points. The percentages are the maximal values of r for which the methods achieve exact recovery over all trials.

Loss function	$\ell_{0,\epsilon}$	$\ell_{2,\epsilon}$	ℓ_1
Linear case	98%	98%	80%
Affine case with positive outliers	98%	98%	43%

of the probability of exact recovery, estimated as the fraction of successful trials⁴. We study this probability as a function of the ratio $r \in [0.3, 0.99]$ of outliers in the data for experiments in a setting similar to section 5.3 except that $\sigma_v = 0$.⁵ In particular, we report in Table 3 the maximal ratio r leading to exact recovery over all trials, i.e., for which the probability estimate is one. This experiment is conducted in two scenarios: with a linear model and outliers ξ_i or with an affine model (\mathbf{x}_i includes a constant component) and positive outliers $|\xi_i|$.

For both scenarios, the results show that the global optimization of nonconvex cost functions provides a larger range of cases for exact recovery than the standard convex relaxation. In addition, this range is unaffected by the linear or affine nature of the model and the constant sign of the outliers, whereas these factors can seriously alter the performance of the ℓ_1 -minimization. In all these trials, the proposed algorithms performed global optimization in less than one second for $r \leq 90\%$ and less than 10 seconds for $r \leq 98\%$. For $r = 99\%$, we stopped the algorithm after 2 minutes without obtaining the solution.

6 Conclusions

The paper introduced new branch-and-bound algorithms to solve several nonconvex optimization problems relevant to hybrid system identification and robust estimation. Compared with most of the literature, a specificity of the proposed approach is that it offers global optimality guarantees that are independent of the data. In addition, by focusing on the continuous variables the approach remains scalable with respect to the number of data. Indeed, switching regression problems with thousands of points could be solved in seconds with global optimality certificates for the first time.

However, the worst-case complexity remains exponential in the dimension of the data and the number of modes. Given the NP-hardness of the considered problems, such a worst-case complexity can be expected for any global optimization approach. In this regard, an interesting open issue concerns the characterization of the (typically lower) theoretical average time complexity, for instance under a given probability distribution of the data or of the noise.

⁴In practice, exact recovery is said to occur when $\|\mathbf{w} - \boldsymbol{\theta}\|_2 < 10^{-6}$.

⁵Note that with $N = 500$, a ratio $r = 0.99$ of outliers leaves 5 uncorrupted data points, which should be enough to exactly recover $\boldsymbol{\theta} \in \mathbb{R}^d$ with $d = 4$.

References

- [1] E. Amaldi and V. Kann. The complexity and approximability of finding maximum feasible subsystems of linear relations. *Theoretical Computer Science*, 147(1-2):181–210, 1995.
- [2] E. Amaldi and M. Mattavelli. The MIN PFS problem and piecewise linear model estimation. *Discrete Applied Mathematics*, 118:115–143, 2002.
- [3] L. Bako. Identification of switched linear systems via sparse optimization. *Automatica*, 47(4):668–677, 2011.
- [4] L. Bako and H. Ohlsson. Analysis of a nonsmooth optimization approach to robust estimation. *Automatica*, 66:132–145, 2016.
- [5] A. Bemporad, A. Garulli, S. Paoletti, and A. Vicino. A bounded-error approach to piecewise affine system identification. *IEEE Transactions on Automatic Control*, 50(10):1567–1580, 2005.
- [6] G. Diehm, S. Maier, M. Flad, and S. Hohmann. An identification method for individual driver steering behaviour modelled by switched affine systems. In *Proc. of the 52nd IEEE Int. Conf. on Decision and Control (CDC)*, pages 3547–3553, 2013.
- [7] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Springer, 2013.
- [8] A. Garulli, S. Paoletti, and A. Vicino. A survey on switched and piecewise affine system identification. In *Proc. of the 16th IFAC Symp. on System Identification (SYSID)*, pages 344–355, 2012.
- [9] A. L. Juloski, S. Weiland, and W. Heemels. A Bayesian approach to identification of hybrid systems. *IEEE Transactions on Automatic Control*, 50(10):1520–1533, 2005.
- [10] F. Lauer. Estimating the probability of success of a simple algorithm for switched linear regression. *Nonlinear Analysis: Hybrid Systems*, 8:31–47, 2013. Supplementary material available at <https://members.loria.fr/FLauer/files/klinreg/>.
- [11] F. Lauer. On the complexity of piecewise affine system identification. *Automatica*, 62:148–153, 2015.
- [12] F. Lauer. On the complexity of switching linear regression. *Automatica*, 74:80–83, 2016.
- [13] F. Lauer, G. Bloch, and R. Vidal. A continuous optimization framework for hybrid system identification. *Automatica*, 47(3):608–613, 2011.
- [14] V. L. Le, F. Lauer, and G. Bloch. Selective ℓ_1 minimization for sparse recovery. *IEEE Transactions on Automatic Control*, 59(11):3008–3013, 2014.
- [15] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995.
- [16] N. Ozay, M. Sznajder, C.M. Lagoa, and O. Camps. A sparsification approach to set membership identification of switched affine systems. *IEEE Transactions on Automatic Control*, 57(3):634–648, 2012.
- [17] S. Paoletti, A. L. Juloski, G. Ferrari-Trecate, and R. Vidal. Identification of hybrid systems: a tutorial. *European Journal of Control*, 13(2-3):242–262, 2007.
- [18] T. Pham Dinh, H.A. Le Thi, H.M. Le, and F. Lauer. A difference of convex functions algorithm for switched linear regression. *IEEE Transactions on Automatic Control*, 59(8):2277–2282, 2014.
- [19] J. Roll, A. Bemporad, and L. Ljung. Identification of piecewise affine systems via mixed-integer programming. *Automatica*, 40(1):37–50, 2004.

- [20] R. Vidal. Recursive identification of switched ARX systems. *Automatica*, 44(9):2274–2287, 2008.
- [21] R. Vidal, S. Soatto, Y. Ma, and S. Sastry. An algebraic geometric approach to the identification of a class of linear hybrid systems. In *Proc. of the 42nd IEEE Conf. on Decision and Control (CDC), Maui, Hawaiï, USA*, pages 167–172, 2003.