

Identification of piecewise constant parameters in nonlinear models

Original

Identification of piecewise constant parameters in nonlinear models / Calafiore, Giuseppe Carlo; Fracastoro, Giulia; Zino, Lorenzo. - In: EUROPEAN JOURNAL OF CONTROL. - ISSN 0947-3580. - 81:(2025), pp. 1-11.
[10.1016/j.ejcon.2024.101151]

Availability:

This version is available at: 11583/2994722 since: 2024-12-04T19:53:58Z

Publisher:

Elsevier

Published

DOI:10.1016/j.ejcon.2024.101151

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Identification of piecewise constant parameters in nonlinear models

Giuseppe Carlo Calafiore, Giulia Fracastoro, Lorenzo Zino*

Department of Electronics and Telecommunications, Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino, 10129, Italy

ARTICLE INFO

Recommended by T. Parisini

Keywords:

System identification
Nonlinear dynamics
Mixed-integer optimization
Convex programming

ABSTRACT

Many real-world dynamical systems are characterized by different temporal phases, with sudden changes in the values of the system's parameters in correspondence to variations from one phase to another. Identifying the system's parameters and these switching instants from potentially noisy measurements of the system's states is a relevant problem in several applications. We here propose a novel approach for estimating the time-varying parameters of a broad class of nonlinear dynamical systems from noisy state measurements. We formulate the problem as a mixed-integer quadratic program (MIQP) including a sparsity constraint to enforce the piecewise constant nature of the parameters. Then, we develop a convex relaxation of the problem in the form of a quadratic program (QP). The solution of the relaxed convex QP and/or the sub-optimal solutions of the MIQP returned by the MIQP solvers provide us with computationally-efficient approximations that can be used effectively in those large-dimensional cases in which the solution of the original MIQP is difficult to obtain. After validating our approach in a controlled experiment, we demonstrate its potential on two real-world case studies regarding marketing and epidemiological applications.

1. Introduction

Identifying the parameters of a dynamical system is a problem of paramount importance in many real-world applicative fields, spanning from engineering to physics, from economics and social science to epidemiology (Ljung, 2010). The field of system identification is mature and well-developed, especially in the context of linear systems, see, e.g., classical references such as Åström and Eykhoff (1971), Ljung (1998), Ljung, Hjalmarsson, and Ohlsson (2011), Söderström and Stoica (1989), which developed upon earlier approaches on least-squares methods for identifying the parameters of linear time-invariant systems (Strejc, 1980), and on design of techniques based on maximum-likelihood estimation for stochastic linear systems (Kashyap, 1970). In more recent times, optimization methods – in particular convex optimization – emerged as powerful tools for effective parameter identification, see, e.g., Vandenberghe (2012).

These classical approaches focus primarily on the scenario in which the dynamics involved are linear and the parameters do not change with time. However, many important applications do not fit in such context, and hence classical identification methods have limited applicability in these cases. In fact, the majority of real-world phenomena are characterized by inherent nonlinearities and non-stationary parameters' behavior. In this paper we focus in particular on the case where the parameters change abruptly, rather than smoothly with time. This happens often in systems that exhibit different temporal

phases, with sudden changes in the values of the parameters between phases, due to sudden changes in the environment. An example that, sadly, became very popular in the past few years comes from the mathematical modeling of epidemics, whereby the contagion mechanism is typically modeled by nonlinear dynamics (Paré, Beck, & Başar, 2020), and the implementation of nonpharmaceutical interventions, vaccination campaigns, as well as the appearance of new variants of the disease may yield abrupt changes in the value of the model parameters, see, e.g., Alisic, Paré, and Sandberg (2023), Calafiore, Novara, and Possieri (2020). This example illustrates the importance of developing algorithms that are able to identify piece-wise constant parameters for nonlinear dynamical systems from available noisy measurements, and thus to detect the key phases of the phenomenon under investigation.

Several methods have been developed for nonlinear system identification, including the use of maximum-likelihood methods (Voss, Timmer, & Kurths, 2004), particle swarm optimization (Schwaab, Biscaia, Monteiro, & Pinto, 2008), machine learning (Chiuso & Pillonetto, 2019), and neural networks (Forgione & Piga, 2021; Norgaard, Ravn, Poulsen, & Hansen, 2000); for more details, see, e.g., the surveys in Noël and Kerschen (2017), Schoukens and Ljung (2019). Also the problem of identifying time-varying parameters for linear dynamical systems has been addressed in several studies, which include approaches that are applicable to the piece-wise constant case. Efforts have been made to extend existing methods for constant parameters,

* Corresponding author.

E-mail addresses: giuseppe.calafiore@polito.it (G.C. Calafiore), giulia.fracastoro@polito.it (G. Fracastoro), lorenzo.zino@polito.it (L. Zino).

including the use of least-squares and wavelet functions (Goel & Bernstein, 2018; Li, Wei, & Billings, 2011; Paleologu, Benesty, & Ciochina, 2008; Tsatsanis & Giannakis, 1993), maximum-likelihood (Esfandiari & Narendra, 2022), generalization of the Extended Kalman filters (Wan & Van Der Merwe, 2000), and semidefinite programming methods (Liu & Vandenberghe, 2009). However, these methods do not enforce bounds on the number of changing points (switching instants) of the parameters, and typically yield a smooth continuously-changing estimation of the parameters. Other methods that enforce piecewise-constant variations often rely on the critical assumption that the switching instants of the parameters are known a priori (Esfandiari & Narendra, 2022); this assumption can be rather unrealistic and it is a key hypothesis that we lift in our approach. Approaches for linear regression with time-varying parameters include (Beck, 1983; Glushchenko & Lastochkin, 2023). Of particular interest in this context is the use of hybrid systems theory to address the problem of estimating piecewise-constant parameters of linear (or affine) dynamical systems, by framing it as a parameter identification problem for piecewise affine switched systems, assuming an underlying autoregressive model (Bemporad, Garulli, Paoletti, & Vicino, 2005; Roll, Bemporad, & Ljung, 2004).

Quite surprisingly, little material is available for the case of nonlinear dynamical systems with piecewise constant parameters and noisy measurements, which is relevant for many applications. Assuming underlying nonlinear autoregressive exogenous models, parameter identification algorithms for piecewise nonlinear switched systems have been developed by means of different approaches, including maximum likelihood frameworks, iterative filter-based techniques, and residual-based methods (Bako, Boukharouba, & Lecoeuche, 2010; Bianchi, Prandini, & Piroddi, 2020; Lauer, Bloch, & Vidal, 2011; Willsky, 2005). Most of the existing methods used to identify parameters in the presence of such abrupt changes assume an underlying Markov model that captures the changes in the parameters, and rely on the use of iterative filtering or bayesian inference (Fearnhead, 2006; Ljung & Gunnarsson, 1990). However, these methods based on hybrid systems are designed for systems that switch repeatedly between different modes according to purely stochastic rules, each one characterized by a set of parameters. Hence, these methods are unfit for the scenario of interest in this paper, in which parameters switch asynchronously and modes are in general not repeated and regulated by rules that are not purely stochastic, and for which one cannot assume absence of memory as in classical Markov chains. Moreover, these methods do not typically allow one to set a priori the total number of switches of the parameters. This limits the possibility of using them to identify and characterize a predetermined number of key switches, which could be useful in many real-world applications, including social and economic systems. Finally, ad-hoc methods have been developed to deal with specific real-world problems. For instance, in Calafiore and Fracastoro (2022) the authors deal with the problem of identifying epidemic parameters from publicly available data, and they propose the use of an ℓ_1 -type penalty function to enforce sparsity in the discrete derivatives of the time-varying parameters. However, a general and flexible formalization of an optimization framework to address this important problem is still missing.

In this paper, we aim at filling this gap by presenting a novel methodological framework for the identification of piecewise constant parameters from noisy observations for a broad class of nonlinear dynamical systems. Technically, our goal is to identify a sequence of parameters that: (i) minimizes the error between the predicted states and the observed states, (ii) controls the impact of noise in the measurements, while (iii) keeping small the number of times the parameters vary during the entire time-window, thus enforcing a piecewise-constant nature for such parameters. To this aim, we formulate the identification problem as an optimization problem with a quadratic objective function that encapsulates the prediction error and the impact of the noise, while the third goal is introduced as a sparsity constraint on the parameters' discrete-time derivatives. Using suitable algebraic

developments, we demonstrate that the latter can be conveniently re-cast as a linear constraint for a set of binary variables. Hence, the proposed parameter identification algorithm ultimately reduces to a mixed-integer quadratic program (MIQP), whose solution can be computed or approximated via state-of-the-art optimization solvers in small-to-medium size scenarios, see, e.g., Andersen, Roos, and Terlaky (2003). The computational complexity of the proposed method, however, remains NP-complete (Pia, Dey, & Molinaro, 2016), which may hinder its applicability in large-scale systems and/or when the parameters change several times in the observation time-window. To address these limitations, we further develop a convex relaxation of the proposed MIQP into a computationally efficient quadratic program (QP), which is then demonstrated to be a practically useful proxy of the original MIQP.

After a validation on a synthetic dataset, we test the proposed framework on two real-world case studies. First, we consider a marketing scenario in which the temporal evolution of the market shares of three competing brands is modeled using a generalized Lotka–Volterra model, based on Marasco, Picucci, and Romano (2016b), and we calibrate it on the Japanese beer market, using data from Marasco, Picucci, and Romano (2016a). This case study demonstrates the potentiality of our optimization-based approach, which is able to correctly identify the different phases of the market and the key switching instants discussed in the literature (Alexander, 2014). Moreover, the moderate size of the dataset allows us to perform a comparison between the MIQP and the relaxed QP. Such a comparison shows that sub-optimal solutions of the MIQP are good proxies of the exact solutions, but their computation becomes computationally challenging when the parameters have several changes over the time-window. In this scenarios, the QP is fast to be solved and its solution is a quite reliable proxy of the solution of the MIQP and can thus be used when solving the MIQP – or even computing good approximations of the solution – becomes practically infeasible. Second, we discuss a case study based on the spread of COVID-19 in Italy, using a susceptible-infected-removed-deceased (SIRD) model, similar to Calafiore et al. (2020). Here, we utilize a longer dataset spanning more than 2 years, which makes the solution of the MIQP computationally unfeasible (and also the computation of reliable sub-optimal solutions is challenging), highlighting the importance and effectiveness of the computationally-efficient QP relaxation. Our results allow us to gain interesting insight into the different phases that characterized the temporal evolution of the COVID-19 pandemic in Italy.

The rest of the paper is organized as follows. In Section 2, we present the setting and the considered model class. In Section 3, we introduce our optimization-based approach for identification, and in Section 4 we discuss in depth two application case studies. Section 5 concludes the paper.

2. Problem setup

2.1. Notation

We denote the set of real, nonnegative real, nonnegative integer, and strictly positive integer numbers as \mathbb{R} , \mathbb{R}_+ , \mathbb{N} , and \mathbb{N}_+ , respectively. Given a positive integer $n \in \mathbb{N}_+$, a (column) vector $x \in \mathbb{R}^n$ is denoted with bold lowercase font, with i th entry x_i . Given two positive integers $n, m \in \mathbb{N}_+$, a matrix $A \in \mathbb{R}^{n \times m}$ is denoted with bold capital font, with A_{ij} denoting the generic j th entry of the i th row and a_j denoting its j th column vector. Given a vector x or a matrix A , x^\top and A^\top are the transpose vector and matrix, respectively. Given a vector x , we denote by $\text{diag}(x)$ the square matrix with the elements of x on the diagonal and 0 on off-diagonal entries. The vector e_i denotes a vector of all zeros, except for a 1 in the i th entry. Given a matrix A , we denote by $\text{tr}(A)$ its trace. The matrix of all ones is denoted by $\mathbf{1}$, the matrix of all zeros is $\mathbf{0}$, and the identity matrix is I . When needed, dimensions are denoted as one positive integer index (for vectors) and two positive integer indices

(for matrices), which are the number of rows and columns, respectively. The set of n -dimensional symmetric positive semi-definite matrices is denoted by S^n . The sign operator $\text{sgn}(\cdot)$ and inequalities applied to a vector or a matrix are intended entrywise (i.e., $\mathbf{x} \geq 0 \iff x_i \geq 0$, for all entries i).

2.2. Model class

We consider a class of discrete-time nonlinear dynamical systems of the form

$$\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t); \boldsymbol{\theta}(t)), \quad (1)$$

where $\mathbf{x}(t) \in \mathbb{R}^n$ is the n -dimensional state variable of the system at time $t \in \mathbb{N}$, $\mathbf{u}(t) \in \mathbb{R}^p$ is the p -dimensional input vector, $\boldsymbol{\theta}(t) \in \mathbb{R}^m$ is a vector containing the m possibly time-varying parameters of the model, and $\mathbf{f} : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the state-transition map, on which we make the following assumption.

Assumption 1. The map $\mathbf{f}(\mathbf{x}, \mathbf{u}; \boldsymbol{\theta}) : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is differentiable with respect to \mathbf{x} and affine in the parameter vector $\boldsymbol{\theta}$, i.e., we can write

$$\mathbf{f}(\mathbf{x}, \mathbf{u}; \boldsymbol{\theta}) = \mathbf{f}_0(\mathbf{x}, \mathbf{u}) + \mathbf{F}(\mathbf{x}, \mathbf{u})\boldsymbol{\theta}, \quad (2)$$

with vector $\mathbf{f}_0(\mathbf{x}, \mathbf{u}) \in \mathbb{R}^n$ and matrix $\mathbf{F}(\mathbf{x}, \mathbf{u}) \in \mathbb{R}^{n \times m}$, for all $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{u} \in \mathbb{R}^p$.

At each time $t \in \mathbb{N}$, there is available an observation of the state of the system $\mathbf{x}(t)$, denoted as $\bar{\mathbf{x}}(t) \in \mathbb{R}^n$, which may be subject to noise and errors in the measurement. Specifically, we assume that the measurement is in the form

$$\bar{\mathbf{x}}(t) = \mathbf{x}(t) + \boldsymbol{\xi}(t), \quad (3)$$

where $\boldsymbol{\xi}(t) \in \mathbb{R}^n$ is a vector that accounts for errors in the measurement of $\mathbf{x}(t)$. We observe that we can write the evolution of the measured state $\bar{\mathbf{x}}(t)$ via the following recursion:

$$\begin{aligned} \bar{\mathbf{x}}(t+1) &= \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t); \boldsymbol{\theta}(t)) + \boldsymbol{\xi}(t+1) \\ &= \mathbf{f}(\bar{\mathbf{x}}(t) - \boldsymbol{\xi}(t), \mathbf{u}(t); \boldsymbol{\theta}(t)) + \boldsymbol{\xi}(t+1) \\ &= \mathbf{f}(\bar{\mathbf{x}}(t), \mathbf{u}(t); \boldsymbol{\theta}(t)) - \mathbf{J}_f(\bar{\mathbf{x}}(t))\boldsymbol{\xi}(t) + \boldsymbol{\xi}(t+1) + o(\|\boldsymbol{\xi}(t)\|_2), \end{aligned} \quad (4)$$

where $\mathbf{J}_f(\bar{\mathbf{x}}(t))$ denotes the Jacobian matrix of \mathbf{f} with respect to \mathbf{x} , evaluated in $\mathbf{x} = \bar{\mathbf{x}}(t)$. Hence, assuming that the error size $\|\boldsymbol{\xi}(t)\|_2$ is sufficiently small, we can approximate $\bar{\mathbf{x}}(t+1)$ by neglecting the higher-order terms, obtaining the approximate recursion:

$$\bar{\mathbf{x}}(t+1) \simeq \mathbf{f}(\bar{\mathbf{x}}(t), \mathbf{u}(t); \boldsymbol{\theta}(t)) - \mathbf{J}_f(\bar{\mathbf{x}}(t))\boldsymbol{\xi}(t) + \boldsymbol{\xi}(t+1). \quad (5)$$

As a consequence of **Assumption 1**, the Jacobian matrix $\mathbf{J}_f(\bar{\mathbf{x}}(t))$ takes the form

$$\mathbf{J}_f(\bar{\mathbf{x}}(t)) = \mathbf{J}_{f_0}(\bar{\mathbf{x}}(t)) + \sum_{i=1}^m \boldsymbol{\theta}_i \mathbf{J}_{f_i}(\bar{\mathbf{x}}(t)), \quad (6)$$

where \mathbf{J}_{f_0} is the Jacobian matrix of \mathbf{f}_0 and \mathbf{J}_{f_i} is the Jacobian matrix of the i th column of matrix $\mathbf{F}(\mathbf{x}, \mathbf{u})$, denoted as $\mathbf{f}_i(\mathbf{x}, \mathbf{u})$. Therefore, the approximate recursion in (5), neglecting the higher-order terms, becomes

$$\begin{aligned} \bar{\mathbf{x}}(t+1) &\simeq \mathbf{f}_0(\bar{\mathbf{x}}(t), \mathbf{u}(t)) + \mathbf{F}(\bar{\mathbf{x}}(t), \mathbf{u}(t))\boldsymbol{\theta}(t) - \mathbf{J}_{f_0}(\bar{\mathbf{x}}(t))\boldsymbol{\xi}(t) \\ &\quad - \sum_{i=1}^m \boldsymbol{\theta}_i \mathbf{J}_{f_i}(\bar{\mathbf{x}}(t))\boldsymbol{\xi}(t) + \boldsymbol{\xi}(t+1). \end{aligned} \quad (7)$$

Letting

$$\mathbf{G}(\boldsymbol{\theta}(t)) = \mathbf{G}(\boldsymbol{\theta}(t); \bar{\mathbf{x}}(t), \mathbf{u}(t)) := \mathbf{J}_{f_0}(\bar{\mathbf{x}}(t)) + \sum_{i=1}^m \boldsymbol{\theta}_i \mathbf{J}_{f_i}(\bar{\mathbf{x}}(t)), \quad (8)$$

we rewrite (7) more compactly as

$$\bar{\mathbf{x}}(t+1) \simeq \mathbf{f}_0(t) + \mathbf{F}(t)\boldsymbol{\theta}(t) - \mathbf{G}(\boldsymbol{\theta}(t))\boldsymbol{\xi}(t) + \boldsymbol{\xi}(t+1), \quad (9)$$

where $\mathbf{f}_0(t) := \mathbf{f}_0(\bar{\mathbf{x}}(t), \mathbf{u}(t))$ and $\mathbf{F}(t) := \mathbf{F}(\bar{\mathbf{x}}(t), \mathbf{u}(t))$.

2.3. Objectives

Our objective is to estimate the time-varying parameter vector $\boldsymbol{\theta}(t)$ from a given time sequence of noisy observations of the state variables $\bar{\mathbf{x}}(0), \dots, \bar{\mathbf{x}}(T-1)$, over a time-window of duration $T \in \mathbb{N}_+$. In particular, in our parameter identification, we aim at achieving the following three objectives:

1. Minimize the error between the predicted state of the system and the observations through the entire time-window;
2. Keep the impact of the measurement noise under control; and
3. Limit the number of times the parameters vary during the entire time-window.

3. A mixed-integer optimization approach

We next illustrate how to synthesize the identification objectives stated in Section 2.3 via a suitable mixed-integer optimization problem. Preliminarily, we make the following further assumptions on the noise sequence and on the parameter vector.

Assumption 2. The noise sequence $\boldsymbol{\xi}(t) \in \mathbb{R}^n$, $t = 0, 1, \dots$, is a sequence of independent random variables with mean equal to $\mathbf{0}$ and given sequence of covariance matrices $\mathbf{Q}(t) \in S^n$, $t = 0, 1, \dots$, known a priori.

Assumption 3. Bounds on the amplitude of $\boldsymbol{\theta}(t)$ are known a priori. That is, there exists a given constant $C > 0$ such that $\|\boldsymbol{\theta}(t)\|_\infty \leq C$, for all $t \in \mathbb{N}$.

We consider a given time horizon $T \in \mathbb{N}_+$, for which we have collected measured data $\bar{\mathbf{x}}(0), \dots, \bar{\mathbf{x}}(T-1)$. At each t , we recognize in (9) two independent error contributions to $\bar{\mathbf{x}}(t+1)$: the direct measurement error $\boldsymbol{\xi}(t+1)$, and the indirect term $\mathbf{G}(\boldsymbol{\theta}(t))\boldsymbol{\xi}(t)$ which quantifies the dynamic effect of the error in $\bar{\mathbf{x}}(t)$ on $\bar{\mathbf{x}}(t+1)$, up to the first order approximation. From **Assumption 2**, the covariance matrix of this error term is $\mathbf{G}(\boldsymbol{\theta}(t))\mathbf{Q}(t)\mathbf{G}(\boldsymbol{\theta}(t))^\top$.

As discussed in Section 2.3, the goal of this paper is to determine an estimate of the parameter vector $\boldsymbol{\theta}(t)$ from the noisy sequence of data $\bar{\mathbf{x}}(1), \dots, \bar{\mathbf{x}}(T)$ with the objective of minimizing the sum of two terms. The first objective is the total (weighted) squared prediction error

$$\mathcal{E} := \sum_{t=0}^{T-1} \|\mathbf{W}(t)(\bar{\mathbf{x}}(t+1) - \mathbf{f}_0(t) - \mathbf{F}(t)\boldsymbol{\theta}(t))\|_2^2, \quad (10)$$

where $\mathbf{W}(t) = \text{diag}(\mathbf{w}(t))$, with $\mathbf{w}(t) > 0$, $t = 0, \dots, T-1$ is a (possibly time-varying) diagonal weight matrix. The weights in $\mathbf{w}(t)$ re-scale the contribution of each component of the state vector \mathbf{x} to the total prediction error, giving flexibility to the identification approach for dealing with several issues, including (i) different orders of magnitude between the variables, (ii) different reliabilities of different state measurements, and (iii) focus of the identification process on fitting the parameters better on a limited time-interval.

The second objective is to control the total indirect measurement errors, by keeping small the sum of the traces of their covariance matrices:

$$\mathcal{V} := \sum_{t=0}^{T-1} \text{tr}(\mathbf{G}(\boldsymbol{\theta}(t))\mathbf{Q}(t)\mathbf{G}(\boldsymbol{\theta}(t))^\top) = \|\mathbf{G}(\boldsymbol{\theta}(t))\mathbf{Q}^{1/2}(t)\|_F^2. \quad (11)$$

Observe that the weighted prediction error in (10) is convex in $\boldsymbol{\theta}(t)$, and also the error cost in (11) is convex, since \mathbf{G} is affine in $\boldsymbol{\theta}(t)$. The parameter identification problem is next formalized as a minimization problem with respect to the parameter vector $\boldsymbol{\theta}(t)$, with a cost function equal to a positive linear combination of \mathcal{E} and \mathcal{V} , which is convex, and a constraint which encodes the third objective in our identification approach, that is the requirement that the time evolution of the parameter is piece-wise constant.

To this end, we let $\Delta\theta(t)$ denote the first-order difference operator, i.e., $\Delta\theta(t) = \theta(t) - \theta(t-1)$, for $t = 1, \dots, T-1$. Stacking all the T parameter vectors $\theta(0), \dots, \theta(T-1)$ in the Tm -dimensional vector $\Theta := [\theta(0)^\top, \dots, \theta(T-1)^\top]^\top$ and all the $T-1$ first-order differences $\Delta\theta(1), \dots, \Delta\theta(T-1)$ in the $(T-1)m$ -dimensional vector

$$\Delta\Theta := \begin{bmatrix} \Delta\theta(1) \\ \Delta\theta(2) \\ \vdots \\ \Delta\theta(T-1) \end{bmatrix} = \begin{bmatrix} \theta(1) - \theta(0) \\ \theta(2) - \theta(1) \\ \vdots \\ \theta(T-1) - \theta(T-2) \end{bmatrix} \quad (12)$$

we have that $\Delta\Theta = D\Theta$, where the matrix

$$D := \begin{bmatrix} -I_{m,m} & I_{m,m} & \mathbf{0}_{m,m} & \mathbf{0}_{m,m} & \cdots & \mathbf{0}_{m,m} \\ \mathbf{0}_{m,m} & -I_{m,m} & I_{m,m} & \mathbf{0}_{m,m} & \cdots & \mathbf{0}_{m,m} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{m,m} & \mathbf{0}_{m,m} & \cdots & \mathbf{0}_{m,m} & -I_{m,m} & I_{m,m} \end{bmatrix} \quad (13)$$

is a block-circulant matrix of dimension $(T-1)m \times Tm$. It is clear that the number of jumps, or switching instants, in the time evolution of the parameters in $\theta(t)$ is given by the number of nonzero entries in vector $\Delta\Theta$. The number of nonzero entries of a vector x is called the *cardinality*, denoted as $\|x\|_0$, and can be given a ‘‘variational’’ characterization, as specified in the following.

Lemma 1. *For any given vector $x \in \mathbb{R}^n$, the cardinality (i.e., the number of nonzero entries) of x , denoted by the 0-norm of x , $\|x\|_0$, can be characterized through the following identities:*

1. $\|x\|_0 = \min_{z \in \{0,1\}^n} \sum_{i=1}^n z_i$, subject to $x = \text{diag}(z)x$.
2. $\|x\|_0 = \min_{z \in \{0,1\}^n} \sum_{i=1}^n z_i$, subject to $x^\top x \leq x^\top \text{diag}(z)x$.
3. $\|x\|_0 = \min_{z \in \{-1,1\}^n} \|z\|_1$, subject to $\|x\|_1 \leq z^\top x$.
4. If $\|x\|_\infty \leq C$, it further holds that

$$\|x\|_0 = \min_{z \in \{0,1\}^n} \sum_{i=1}^n z_i, \text{ subject to } -Cz \leq x \leq Cz. \quad (14)$$

Proof. First, we observe that, for any $z \in \{0,1\}^n$ such that $x = \text{diag}(z)x$ it necessarily holds true that $x_i \neq 0 \implies z_i = 1$. Hence, we can immediately conclude that $\|x\|_0 \leq \sum_{i=1}^n z_i$. Then, we observe that the equality is attained for z^* defined entrywise as $z_i^* = 1$ if $x_i \neq 0$ and $z_i^* = 0$ otherwise. This yields the first identity. Identities 2. and 3. are obtained following a similar argument, with the unique difference that, in case 3., the equality is attained for $z^* = \text{sgn}(x)$. For 4, observe that for any $z \in \{0,1\}^n$ such that $-Cz \leq x \leq Cz$ it necessarily holds true that $x_i \neq 0 \implies z_i = 1$. Hence, we can immediately conclude that $\|x\|_0 \leq \sum_{i=1}^n z_i$. Then, we observe again that the equality is attained for z^* defined entry-wise as $z_i^* = 1$ if $x_i \neq 0$ and $z_i^* = 0$ otherwise, where the constraint $-Cz_i \leq x_i \leq Cz_i$ holds true, since $-C \leq x_i \leq C$, yielding the claim. \square

From item 4 of Lemma 1, we also obtain the following equivalent conditions

$$\|x\|_0 \leq k \iff \exists z \in \{0,1\}^n : \sum_{i=1}^n z_i \leq k, \quad -Cz \leq x \leq Cz. \quad (15)$$

Let now $e_i \in \mathbb{R}^m$ denote the i th standard basis vector of \mathbb{R}^m , and let us define the following block-circulant matrix:

$$E_i := \begin{bmatrix} e_i^\top & \mathbf{0}_m^\top & \mathbf{0}_m^\top & \cdots & \mathbf{0}_m^\top \\ \mathbf{0}_m^\top & e_i^\top & \mathbf{0}_m^\top & \cdots & \mathbf{0}_m^\top \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_m^\top & \mathbf{0}_m^\top & \mathbf{0}_m^\top & \cdots & e_i^\top \end{bmatrix}, \quad (16)$$

which is a block-diagonal matrix with $T-1$ rows and $(T-1)m$ columns. Then, we say that the i th parameter sequence $[\theta_i(0), \dots, \theta_i(T-1)]$ is k -piecewise constant if the vector $[\Delta\theta_i(1), \dots, \Delta\theta_i(T-1)]$ is k -sparse, i.e., if its cardinality is k . Using the matrix E_i , we have that the i th parameter sequence $[\theta_i(0), \dots, \theta_i(T-1)]$ is k -piecewise constant if and only if $\|E_i D \Theta\|_0 \leq k$.

Under Assumptions 1–3, we can finally formulate our parameter identification problem as an optimization problem where we aim to minimize a linear combination of the two terms \mathcal{E} and \mathcal{V} , weighted by a nonnegative coefficient $\gamma \geq 0$, while requiring the parameter sequences to be piecewise constant.

Problem 1. We formulate the following mixed-integer quadratic program (MIQP).

$$\begin{aligned} \min_{\Theta \in \mathbb{R}^{Tm}, Z \in \{0,1\}^{(T-1) \times m}} & \sum_{t=0}^{T-1} \|\mathbf{W}(t)(\bar{x}(t+1) - f_0(t) - F(t)\Theta(t))\|_2^2 \\ & + \gamma \sum_{t=0}^{T-1} \|\mathbf{G}(\theta(t))\mathbf{Q}^{1/2}(t)\|_F^2 \\ \text{s.t.} & \mathbf{1}^\top z_i \leq k_i, \quad i = 1, \dots, m, \\ & -Cz_i \leq E_i D \Theta \leq Cz_i, \quad i = 1, \dots, m, \\ & \|\Theta\|_\infty \leq C, \end{aligned} \quad (17)$$

where k_i , $i = 1, \dots, m$, are given upper bounds on the sparsity of $[\Delta\theta_i(1), \dots, \Delta\theta_i(T-1)]$, and $\gamma \geq 0$ is a tradeoff parameter controlling the penalty on the induced error term.

The cost function to be minimized in Problem 1 comes from the explicit expressions for \mathcal{E} and \mathcal{V} derived in (10) and (11). The constraint $\|\Theta\|_\infty \leq C$ comes from Assumption 3, while the other constraints in Problem 1 are derived from the k -sparsity constraints $\|E_i D \Theta\|_0 \leq k_i$, for all $i = 1, \dots, m$ by using (15). The $m(T-1)$ additional binary variables are collected in the matrix $Z \in \{0,1\}^{(T-1) \times m}$, with columns z_i , $i = 1, \dots, m$, thus obtaining the MIQP in (17).

3.1. Validation

Before further discussing our parameter identification method, presenting a convex relaxation, and illustrating their practical use through some case study, we report a validation on a simple example. Here, we assume that the ground truth is known a priori, and we show that the proposed method is able to correctly identify piecewise constant parameters of a nonlinear dynamical system from noisy measurements.

To this aim, we generate a time-series from a bi-dimensional generalized Lotka–Volterra model with no spontaneous growth (Brauer & Castillo-Chavez, 2012), which is characterized by the dynamical system

$$\begin{aligned} x_1(t+1) &= x_1(t) + \theta_1(t)x_1(t)x_2(t) \\ x_2(t+1) &= x_2(t) + \theta_2(t)x_1(t)x_2(t), \end{aligned} \quad (18)$$

with initial condition $x(0) = [0.2, 0.2]$ and time-varying parameters equal to

$$\theta_1(t) = \begin{cases} 0.2, & 0 \leq t \leq 11, \\ -0.1, & 12 \leq t \leq 19, \\ 0.05, & 20 \leq t \leq 29, \end{cases} \quad \theta_2(t) = \begin{cases} -0.12, & 0 \leq t \leq 6, \\ 0.07, & 7 \leq t \leq 21, \\ 0.15, & 22 \leq t \leq 29. \end{cases} \quad (19)$$

We can check that the map in (18) satisfies Assumption 1. In fact, we can write it as $f(x, \theta) = f_0(x) + F(x)\theta$, with

$$f_0(x) = x, \quad F(x) = \begin{bmatrix} x_1 x_2 & 0 \\ 0 & x_1 x_2 \end{bmatrix}, \quad (20)$$

where we have omitted the time-dependencies of the variables and of the parameters for brevity. We can write the Jacobian of the map in (18) as

$$J_f(x, \theta) = \begin{bmatrix} 1 + \theta_1 x_2 & \theta_1 x_1 \\ \theta_2 x_2 & 1 + \theta_2 x_1 \end{bmatrix} \quad (21)$$

and, following our approach, we define

$$\mathbf{G}(\theta) = \mathbf{I}_{2,2} + \begin{bmatrix} x_2 & x_1 \\ 0 & 0 \end{bmatrix} \theta_1 + \begin{bmatrix} 0 & 0 \\ x_2 & x_1 \end{bmatrix} \theta_2. \quad (22)$$

Using (18), we generate a sample path $x(0), \dots, x(T)$, and we generate the measurements $\bar{x}(0), \dots, \bar{x}(T)$, with $T = 30$, by adding to each entry $x(t)$ a noise with covariance matrix $\mathbf{Q}(t) = \sigma \mathbf{I}$, for all t . In our experiments, we test different values of σ , and we set the weight matrix equal to the identity, i.e., $\mathbf{W}(t) = \mathbf{I}_{2,2}$ for all $t = 0, \dots, T-1$, since the




Fig. 1. Validation on synthetic data with different levels of noise σ , with $k_1 = k_2 = 2$, $C = 0.25$, and $\gamma = 1$. The dashed curves illustrate the real profile of the two parameters.




Fig. 2. Parameter estimation replicating scenarios in Figs. 1(a) and 1(b), using an extended Kalman filter. The dashed curves illustrate the real profile of the two parameters.

two variables have the same order of magnitude. Then, we solved the MIQP on a 4.7 GHz 10-Core Intel Core i7 using Mosek (Andersen et al., 2003) implemented in MatLab R2022b with CVX (Grant & Boyd, 2008, 2014), which took just few seconds to reach a solution.

The results, reported in Fig. 1, validate our method. In fact, in the absence of noise (i.e., $\sigma = 0$), our method is able to identify the exact value of the parameters in (19); see Fig. 1(a). For small to moderate values of the noise (signal-to-noise ratios¹ of approximately 60 and 15 in Figs. 1(b) and 1(c), respectively), we observe that the MIQP is still very effective in detecting the correct switching instants of the parameters and provides very accurate estimates of their values, as one can see in Figs. 1(b) and 1(c), with a relative ℓ_2 -norm error averaged over 100 realization of the noise equal to 3.15% and 9.74%, respectively. Predictably, the accuracy decreases as the noise reaches an order of magnitude comparable with the measurements (signal-to-noise ratio of approximately 4), as in Fig. 1(d), where the mean relative ℓ_2 -norm error is larger than 40%.

Finally, we compare our novel technique with a classical method for estimating time-varying parameters. Specifically, we replicate the scenarios with no noise and small noise (i.e., $\sigma = 0$ and $\sigma = 0.005$), and we estimate the time-varying parameters using an extended Kalman filter (Wan & Van Der Merwe, 2000). The results, illustrated in Fig. 2, suggest that the constraint on the number of parameter switches, and the ensuing mixed-integer approach, is key to obtain piecewise-linear estimates for the parameters. In fact, even in the absence of noise, methods based on extended Kalman filters tend to enforce a smooth evolution of the parameters and hence generally fail in correctly detecting abrupt discontinuous changes.

3.2. Convex relaxation of Problem 1

While Problem 1 can be solved numerically in moderate-sized instances using state-of-the-art solvers for MIQP (Andersen et al., 2003), it is known that MIQPs are computationally challenging in general, due to the presence of binary variables which make them NP-Complete (Pia et al., 2016). It is therefore of interest to consider, along with the original MIQP formulation, also a tractable convex relaxation of Problem 1.

¹ The signal-to-noise ratio is the ratio between the mean square value of the state and the mean square value of the noise.

While different methods can be used to achieve this goal (Tillmann, Bienstock, Lodi, & Schwartz, 2024), one such relaxation is readily obtained from the characterization in (14) by relaxing the domain of $z \in \{0, 1\}^n$ to the (convex) domain $z \in [0, 1]^n$. The following lemma holds.

Lemma 2. For any $x \in \mathbb{R}^n$ with $\|x\|_\infty \leq C$ it holds that

$$\|x\|_0 \geq \frac{1}{C} \|x\|_1 = \min_{z \in [0, 1]^n} \sum_{i=1}^n z_i, \quad \text{subject to } -Cz \leq x \leq Cz. \quad (23)$$

Proof. Since enlarging the feasible set can only improve the optimization objective, we have from (14) that

$$\|x\|_0 \geq \min_{z \in [0, 1]^n} \sum_{i=1}^n z_i, \quad \text{subject to } -Cz \leq x \leq Cz. \quad (24)$$

Further, it is straightforward to check that the optimal z for the optimization problem in (24) is $z_i^* = |x_i|/C$, for all i , from which it follows that the optimal value of this problem is $\|x\|_1/C$. \square

From Lemma 2 we easily obtain the implications

$$\|x\|_0 \leq k \implies \frac{1}{C} \|x\|_1 \leq k \iff \exists z \in [0, 1]^n : \sum_{i=1}^n z_i \leq k, \quad -Cz \leq x \leq Cz, \quad (25)$$

which are used to obtain the following convex relaxed version of Problem 1.

Problem 2. A convex relaxation of the MIQP in 1 is given by the QP:

$$\begin{aligned} \min_{\theta \in \mathbb{R}^m, Z \in [0, 1]^{(T-1) \times m}} & \sum_{t=0}^{T-1} \|\mathbf{W}(t)(\bar{x}(t+1) - \mathbf{f}_0(t) - \mathbf{F}(t)\theta(t))\|_2^2 \\ & + \gamma \sum_{t=0}^{T-1} \|\mathbf{G}(\theta(t))\mathbf{Q}^{1/2}(t)\|_F^2 \\ \text{s.t.} & \mathbf{1}^\top z_i \leq k_i, \quad i = 1, \dots, m, \\ & -Cz_i \leq E_i \mathbf{D} \theta \leq Cz_i, \quad i = 1, \dots, m, \\ & \|\theta\|_\infty \leq C. \end{aligned} \quad (26)$$

The QP in Problem 2 can be easily solved in polynomial time using standard and efficient tools, see, e.g., Mosek (Andersen et al., 2003). However, it is worth noticing that this relaxation comes at a cost. In fact, the solution of the QP in Problem 2 may, in general, not satisfy the bound imposed on the number of switch instants k_i .




Fig. 3. Temporal evolution of the market shares of *Asahi* (red), *Kirin* (blue), and *Sapporo* (cyan) from 1963 to 2001. Data from [Marasco et al. \(2016a\)](#).

Remark 1 (*Use of the Estimated Model*). It is worth to briefly discuss about the possible uses of a model identified via the proposed approach. Once we solved [Problem 1](#), we find ourselves with the estimated time profiles of the model parameters $\theta(t)$ over the horizon $[0, T - 1]$. It is important to observe that these estimated parameters are determined so to “explain” the observations over the considered horizon, and shall not be used, in general, for extrapolating in time (i.e., for forward predicting) the system behavior. In fact, no hypotheses are made on the dynamic behavior of the parameter $\theta(t)$ itself, so that we do not know how the parameters will evolve from T onward. A forward prediction can only be made by making further assumptions on the *future* behavior of $\theta(t)$; e.g., one may assume that $\theta(t)$ will stay constant at its last estimated value $\theta(T - 1)$ for a certain number of forward periods. Besides this specific case, the main use of the estimated model is for analyzing the observed dynamic phenomenon inside the estimation horizon $[0, T - 1]$, and extracting useful information from the values of the parameters and their switching instants.

4. Case studies

We next present two case studies in which we use the proposed approach for calibrating nonlinear models based on real-world data. The first case study, discussed in [Section 4.1](#), presents the competition among different brands in the Japanese beer market and is focused on unveiling their interactions; the second one, reported in [Section 4.2](#), is an epidemiological case study, which focused on the estimation of the parameters of a COVID-19 SIRD model based on Italy epidemiological data covering the entire duration of the pandemic.

4.1. Japanese beer market

In the first case study, we consider the competition among the three main brands of beer in the Japanese market – namely, *Asahi*, *Kirin*, and *Sapporo* – for which data are available with a one-year temporal granularity from 1963 to 2000. The dataset, publicly available in [Marasco et al. \(2016a\)](#), comprises thus 38 data points consisting in the relative shares of the three brands in the market, as illustrated in [Fig. 3](#).

We focus our analysis on unveiling the relations among the three brands. To this aim, inspired by [Marasco et al. \(2016b\)](#), we represent the dynamics through a generalized Lotka–Volterra model, neglecting spontaneous growths and focusing on the interaction effects. This yields the following three-dimensional nonlinear dynamical system:

$$\begin{aligned} x_1(t+1) &= x_1(t) + \theta_1(t)x_1(t)x_2(t) + \theta_2(t)x_1(t)x_3(t) \\ x_2(t+1) &= x_2(t) + \theta_3(t)x_1(t)x_2(t) + \theta_4(t)x_2(t)x_3(t) \\ x_3(t+1) &= x_3(t) + \theta_5(t)x_1(t)x_3(t) + \theta_6(t)x_2(t)x_3(t), \end{aligned} \quad (27)$$

where $\mathbf{x}(t) = [x_1(t), x_2(t), x_3(t)]^\top$ are the shares of *Asahi*, *Kirin*, and *Sapporo* in the Japanese beer market at time t , respectively. The system is characterized by a six-dimensional parameter vector $\theta(t) = [\theta_1(t), \theta_2(t), \theta_3(t), \theta_4(t), \theta_5(t), \theta_6(t)]^\top$, where each entry captures the interaction effect between two brands. Specifically, $\theta_1(t)$ and $\theta_2(t)$ capture the effect on the temporal evolution of the shares of *Asahi* caused by

its interaction with *Kirin* and *Sapporo*, respectively. Similarly, $\theta_3(t)$ and $\theta_4(t)$ capture the effect on *Kirin* of its interaction with *Asahi* and *Sapporo*, respectively; and $\theta_5(t)$ and $\theta_6(t)$ the effect on *Sapporo* of its interaction with *Asahi* and *Kirin*, respectively. Similar to [\(18\)](#), we can check that the map in [\(27\)](#) satisfies [Assumption 1](#) and we can write the matrices F and G following similar arguments, obtaining

$$F(\mathbf{x}) = \begin{bmatrix} x_1x_2 & x_1x_3 & 0 & 0 & 0 & 0 \\ 0 & 0 & x_1x_2 & x_2x_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & x_1x_3 & x_2x_3 \end{bmatrix}, \quad (28)$$

and

$$\begin{aligned} G(\theta) &= I_{3,3} + \begin{bmatrix} x_2 & x_1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \theta_1 + \begin{bmatrix} x_3 & 0 & x_1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \theta_2 + \begin{bmatrix} 0 & 0 & 0 \\ x_2 & x_1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \theta_3 \\ &+ \begin{bmatrix} 0 & 0 & 0 \\ 0 & x_3 & x_2 \\ 0 & 0 & 0 \end{bmatrix} \theta_4 + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ x_3 & 0 & x_1 \end{bmatrix} \theta_5 + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & x_3 & x_2 \end{bmatrix} \theta_6. \end{aligned} \quad (29)$$

We performed the parameter identification using both the original MIQP in [Problem 1](#) and the QP in [Problem 2](#), in order to compare their outcome and the computational complexity. After re-formulating some constraints through the addition of slack variables, the MIQP involves 1875 variables (222 of which are binary) and 1122 scalar constraints. Again, we solved both problems using Mosek implemented in MatLab R2022b with CVX.

In our numerical experiments, we set the weight matrix equal to the identity, i.e., $\mathbf{W}(t) = I_{3,3}$ for all $t = 0, \dots, T - 1$, since the three variables have the same order of magnitude and comparable data quality. Moreover, since we had no information on the measurement error, but the three variables are bounded and slowly-changing over the entire time-window, we estimated the covariance matrix $\mathbf{Q}(t)$ using the data $\mathbf{x}(0), \dots, \mathbf{x}(T)$, by setting $\mathbf{Q}(t) = \text{diag}([\sigma_1, \sigma_2, \sigma_3])$ as a time-invariant diagonal matrix with i th diagonal entry equal to the sample variance of the i th variable, i.e., $\sigma_i := \frac{1}{T-1} \sum_{t=1}^T (x_i(t) - \frac{1}{T} \sum_{t=1}^T x_i(t))^2$. In this scenario, the objective function of our optimization problems reduces to

$$\sum_{t=0}^{T-1} \|\bar{\mathbf{x}}(t+1) - \mathbf{x}(t) - \mathbf{F}(\mathbf{x}(t))\theta(t)\|_2^2 + \gamma \|\text{diag}([\sigma_1, \sigma_2, \sigma_3])G(\theta(t))\|_F^2, \quad (30)$$

with $F(\mathbf{x})$ from [\(28\)](#) and $G(\theta)$ from [\(29\)](#).

First, we used the MIQP problem formulation to identify the parameters of [\(27\)](#), setting $k_i = 2$, for all $i = 1, \dots, 6$, $C = 0.25$, and $\gamma = 2$. The output of our parameter identification algorithms are reported in panels (a–c) of [Fig. 4](#). The results provide interesting insight into the system analyzed. In fact, they allow not only to shed light on the interactions between the three different brands, but also to identify the key switching instants in the Japanese beer market. In fact, following [Marasco et al. \(2016b\)](#), the signs of the parameters determine the (competitive or mutualistic) role between two brands: if both negative they are in pure competition; if both positive they are mutualistic; while if one is positive and the other negative, they have a prey–predator type of interaction. For instance, [Fig. 4\(c\)](#) shows two interesting turning points




Fig. 4. Comparison of the six time-varying parameters estimated using the MIQP (a–c) and using the QP relaxation (d–f) for the Japanese beer market case study, with $k_i = 2$, for all $i = 1, \dots, 6$, $C = 0.25$, and $\gamma = 2$.

in the interactions between *Kirin* and *Sapporo*. Until the second part of the 1970s, the two brands display a prey–predator type of interaction beneficial to *Kirin* (in fact, θ_4 is positive and θ_6 is negative). Then, the sign of such interaction is reverted in 1977 in favor of *Sapporo*. This coincides with the year in which *Sapporo* introduced its first draft beer, named “Black Label” (Alexander, 2014). Then, this period lasts until the end of the 1980s, when the interactions started becoming purely competitive. Indeed, it is known that this last period was characterized by an intense competition between the brands over dry beers, which goes under the name of “dry war” (Alexander, 2014). In fact, in 1987 Asahi launched *Super Dry* gaining momentum over the other brands, which started a strong competition, as we observed in panel (c). This is reflected in the parameters in Fig. 4. In fact, we observe an increase in both red curves at the end of the 1980s in panels (a) and (b), while the interaction between *Kirin* and *Sapporo* becomes purely competitive, as observed in the above. Hence, this simple case study elucidates how our algorithm can be used to understand the impact of historical events on the market and detect key phases and turning points.

Second, we performed the same parameter identification using the QP relaxation in Problem 2. The results are reported in panels (d–f) of Fig. 4. A comparison of these panels with the previous ones (a–c) suggests that the outcome of the QP relaxation was a good proxy for the MIQP in the case of interest. In fact, the solution of the QP approximates with sufficiently good accuracy not only the value of the majority of the parameters (the sign of the parameter, which determines whether the brands have a positive or negative interaction, coincides in more than 96% of the instances), but also their main switching instants (relative ℓ_2 -norm error of about 10%). Only minor differences are observed. In particular, since the solution of the QP is allowed to change value more than k_i times, some abrupt changes of parameters in the solutions of the MIQP are split into two smaller steps—see, for instance, the evolution of the parameter θ_4 (blue curve) in panels (c) and (f). Even for θ_4 , however, the signs coincide in all instants, except for one, in which they are both close to 0. It is worth noticing that, while the QP took less than one second to be solved, the computation of the exact solution of the MIQP required approximately 5.5 h.

Based on this observation, we performed an additional study to investigate the computational effort required to solve the MIQP. Specifically, we computed the time needed to compute sub-optimal solutions²

Table 1

Time needed (measured in seconds) to compute a optimal and sub-optimal solutions of the MIQP for the Japanese beer market case study, for increasing values of k and different relative optimality tolerance employed by the mixed-integer optimizer (5%, 1%, or 0.5%), compared to the time needed to solve the QP (last row).

	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
Optimal	43.67	19,645.12	>1 day	>1 day	>1 day	>1 day
0.5%	1.28	23.66	445.00	2,261.47	6,784.81	16,467.30
1%	0.23	1.33	2.73	11.75	22.58	41.72
5%	0.11	0.13	0.17	0.25	0.45	1.28
QP	0.81	0.78	0.95	0.82	1.09	1.04

of the MIQP for increasing values of k_i , by setting $k_i = k$, for all $i = 1, \dots, 6$, and testing values of k from 1 to 6. The results are reported in Table 1, and they also show that sub-optimal solutions of the MIQP with low relative optimality tolerance can be efficiently computed. Importantly, while sub-optimal solutions with larger tolerance show noticeable differences with respect to the exact solution of the MIQP reported in Fig. 4 (relative ℓ_2 -norm error greater than 25% when tolerance is 1%), the one obtained with 0.5% tolerance for $k = 2$ (computed in less than 24 s) is a good proxy of the optimal one (relative ℓ_2 -norm error smaller than 10%). To summarize, the analysis of this case study suggests that the MIQP approach developed in this paper could be useful to identify the parameters of low-dimensional systems, possibly using sub-optimal solutions. Furthermore, the solution of the convex relaxation in Problem 2 is in general a good proxy of MIQP solution and its use is suggested when dealing with higher dimensional systems, as we shall see in the next case study. In fact, our convex relaxation has performances that are comparable with a sub-optimal solutions of the MIQP with 0.5% tolerance (no more than 10% relative ℓ_2 -norm error), but its computational time does not grow with k_i . This suggests that, among the two approximations, the convex relaxation could be preferred when the parameters have many changing points, while for $k_i \leq 3$, sub-optimal solutions of the MIQP with at most 0.5% tolerance can be computed in a reasonable time and may have better performances.

4.2. COVID-19 in Italy

In the second case study, we calibrate an epidemic model using data on the COVID-19 pandemic in Italy, from its inception (February 24,

² A solution with 1% tolerance is obtained by allowing the objective function to be 1% larger than its optimal value.




Fig. 5. COVID-19 weekly figures from February 24, 2020 to July 12, 2022 in Italy, as reported in the official dataset available from [Dipartimento della Protezione Civile \(2020\)](#).




Fig. 6. Parameters estimated for the COVID-19 case study obtained as (a) a sub-optimal solution of the MIQP (40% tolerance) and (b) the solution of the QP relaxation for $k_1 = k_2 = k_3 = 3$ with $C = 0.5$ and $\gamma = 0.2$.

2020) till July 12, 2022. Specifically, we use publicly available data reported by the Civil Protection Department of the Italian Government on the number of reported cases, recoveries, and deaths ([Dipartimento della Protezione Civile, 2020](#)). We aggregated data on a weekly basis to account for the natural oscillations in the number of tests performed during the different days of the week. Hence, our dataset comprises 124 weekly data points for reported cases, total recovered people, and total deaths, as shown in [Fig. 5](#).

We formulate the epidemic model as a discrete-time SIRD model with time-varying parameters, similar to the one proposed in [Calafiore et al. \(2020\)](#), with time unit equal to one week. Such a temporal granularity allows for neglecting the delay between contagion and infectiousness, since the duration of the latency period is shorter than one week in the large majority of the cases ([Xin et al., 2021](#)). We observe that the SIRD model can be written in the form of (1), with affine mapping with respect to the parameters. Let us define the state variable vector $\mathbf{x}(t) := [x_1(t), x_2(t), x_3(t)]^T$, whose entries are respectively the total number of infected, recovered, and deceased individuals in the population at time t , with the number of susceptible individuals equal to $n - x_1(t) - x_2(t) - x_3(t)$, where n is the total population size. Let us further define the time-varying parameter vector $\boldsymbol{\theta}(t) = [\theta_1(t), \theta_2(t), \theta_3(t)]^T$ to contain the infection rate, recovery rate, and death rate at time t , respectively. Then, we can write the following dynamical system:

$$\begin{aligned} x_1(t+1) &= x_1(t) + \frac{1}{n}\theta_1(t)(n - x_1(t) - x_2(t) - x_3(t))x_1(t) - (\theta_2(t) + \theta_3(t))x_1(t) \\ x_2(t+1) &= x_2(t) + \theta_2(t)x_1(t) \\ x_3(t+1) &= x_3(t) + \theta_3(t)x_1(t), \end{aligned} \quad (31)$$

where we note that the system is autonomous, i.e., $\mathbf{u}(t) = \mathbf{0}$, for all $t \in \mathbb{N}$. Note that the second and third equations are linear and of straightforward interpretation: each week, a fraction $\theta_2(t)$ of the individuals who were infected the previous week have recovered and a fraction $\theta_3(t)$ of them have died. The first equation, instead, besides

having a negative contribution equal to the two terms described in the above, has a positive nonlinear term that accounts for new infections. In particular, the number of newly infected individuals is equal to the product between susceptible individuals, $n - x_1(t) - x_2(t) - x_3(t)$, the fraction of infected individuals in the population $\frac{1}{n}x_1(t)$, and the infection rate $\theta_1(t)$. This nonlinear term is typical from the epidemic literature, see, e.g., [Brauer and Castillo-Chavez \(2012\)](#) for more details.

We can check that the map in (31) satisfies [Assumption 1](#). In fact, we can write it as $\mathbf{f}(\mathbf{x}, \boldsymbol{\theta}) = \mathbf{f}_0(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}$, with

$$\mathbf{f}_0(\mathbf{x}) = \mathbf{x}, \quad \mathbf{F}(\mathbf{x}) = \begin{bmatrix} \frac{1}{n}x_1(n - x_1 - x_2 - x_3) & -x_1 & -x_1 \\ 0 & x_1 & 0 \\ 0 & 0 & x_1 \end{bmatrix}, \quad (32)$$

where we have omitted the time-dependencies of the variables and parameters for brevity. Note that we can write the Jacobian of the map in (32) as

$$\mathbf{J}_f(\mathbf{x}, \boldsymbol{\theta}) = \begin{bmatrix} 1 + \frac{1}{n}\theta_1(n - 2x_1 - x_2 - x_3) - \theta_2 - \theta_3 & 0 & 0 \\ \theta_2 & 1 & 0 \\ \theta_3 & 0 & 1 \end{bmatrix} \quad (33)$$

and

$$\mathbf{G}(\boldsymbol{\theta}) = \mathbf{I}_{3,3} + \begin{bmatrix} \frac{1}{n}(n - 2x_1 - x_2 - x_3) & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \boldsymbol{\theta}_1 + \begin{bmatrix} -1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \boldsymbol{\theta}_2 + \begin{bmatrix} -1 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \boldsymbol{\theta}_3. \quad (34)$$

Before formalizing the parameter identification problem, we observe from [Fig. 5](#) that the three components of the state variable vector $\mathbf{x}(t)$ have different orders of magnitude across them and large variations in time, due to different testing policies and accumulation of recovered and total deaths. This calls for the definition of a non-trivial sequence of weight matrices $\mathbf{W}(t)$ to re-scale the components of $\mathbf{x}(t)$ to comparable quantities, and thus avoid the objective function to be dominated by

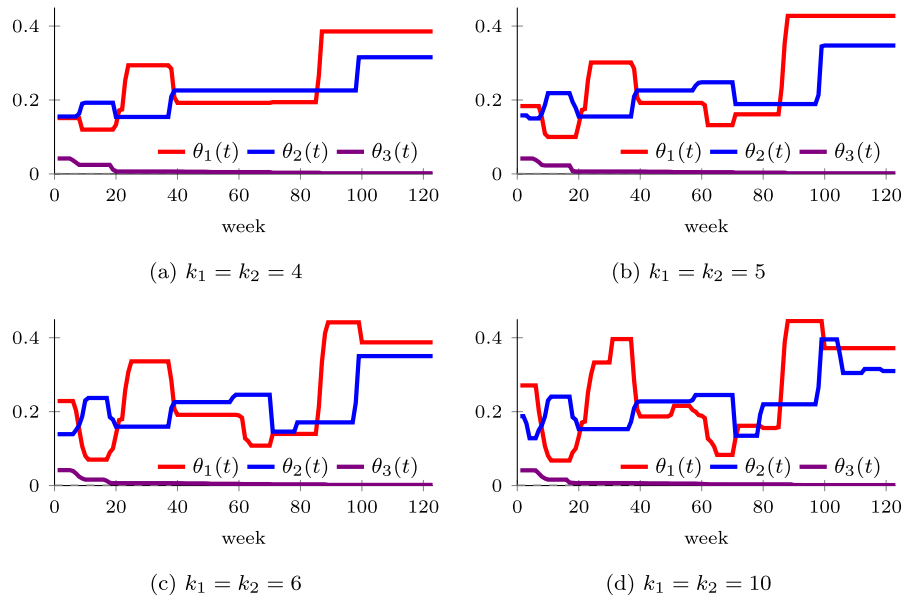


Fig. 7. Parameters estimated for the COVID-19 case study using the QP relaxation for different values of k_1 and k_2 . Common parameters are $k_3 = 3$, $C = 0.5$ and $\gamma = 0.2$.

just the few entries with larger order of magnitude. In particular, we set $\mathbf{W}(t) = \text{diag}(\bar{x}(t)^{-1})$, so that each the term in the first summand of the objective function measures the relative error of the predicted state with respect to the measured one. Then, similar to the previous case study, we set $\mathbf{Q}(t)$ as a time-invariant diagonal matrix with i th diagonal entry equal to the sample variance of the i th variable.

In the 124 weeks considered in the dataset, several events occurred, including two major lockdown periods (Spring 2020 and Winter–Spring 2021), the vaccination and booster campaigns (Spring–Summer 2021 and Fall 2021–Winter 2022, respectively), and several epidemic waves due to the diffusion of new variants. Hence, we expect the presence of several switching instants in the values of the piecewise constant parameters θ_1 , θ_2 , and θ_3 . We verify this conjecture by showing in Fig. 6 that the parameters estimated by setting too small values of k_i are not particularly informative on the different phases of the pandemic, since they cannot capture, e.g., the key changes that occurred in the infection rate when lockdown measures are imposed at mid March 2020. However, increasing k_i makes the solution of the MIQP hard to compute in practice. Indeed, computing a sub-optimal solution of the MIQP with $k_1 = k_2 = k_3 = 3$ at 20% sub-optimality tolerance takes almost 4 h, and for $k_1 = k_2 = k_3 = 6$ after 1 day of optimization the sub-optimality tolerance is still larger than 40%. This example highlights the importance of our computationally-efficient QP relaxation, which runs in few seconds even in this more complex case study, and provides a good proxy of the MIQP solution (see the results in Section 4.1, and compare the two panels in Fig. 6). Hence, in the following we focus on solving the QP, for which we use again Mosek implemented in MatLab R2022b with CVX. After the addition of suitable slack variables for obtaining a standard QP formulation, we get 4009 scalar variables and 1338 constraints.

We solved the QP in Problem 2 with different values of k_i and we compared the results obtained. After a preliminary analysis, we observe that the parameter $\theta_3(t)$ associated with the mortality of the disease tends to have few changes with respect the other parameters with a three-step decrease, corresponding with (i) the start of the decrease of the first wave (mid April 2020), (ii) the diffusion of antigenic tests, which increased the number of detected mild cases (end of summer 2020), and (iii) the peak of the second-dose vaccination campaign (summer 2021). Differently, the value of the parameter associated to the contagion and recovery rates show more switching instants. This may be due to the fact that these parameters are not only affected by

the characteristics of the disease and on the state of the vaccination campaign, but are also strongly impacted by human behavior. In view of these preliminary observations, we decided to fix the value of $k_3 = 3$, and vary the value of $k_1 = k_2$ in the range $k_1 = k_2 \in \{4, 5, 6, 10\}$.

The results are reported in Fig. 7. Unsurprisingly, we observe that, as we increase $k_1 = k_2$, the parameter identification method becomes able to detect an increasing number of key switching instants for the model parameters. More interestingly, we also observe that the effect of increasing k seems to be incremental: first, the algorithm detects the most crucial changes; then, as k grows large, smaller changes are added to those already detected. For $k_1 = k_2 \geq 5$, we observe that the algorithm identifies the key phases of the first pandemic wave, with a steep decrease of the infection rate and increase of the recovery rate in correspondence of the first lock down, which is then reverted at the end of the summer, slowly leading to the second wave in fall-winter 2020. Then, the infection rate $\theta_1(t)$ show a rapid decrease in correspondence to the start of the second lockdown, followed by another decrease, during the peak of the vaccination campaign; while the recovery rate shows an opposite, but less pronounced, behavior. Besides capturing parameter changes due to human behavior, our algorithm identifies also key natural events, such as the diffusion of new variants. For instance, the diffusion of the Omicron variant can be seen in the steep increase in θ_1 around week 90 (mid November 2021).

Finally, we would like to remark that, in this case study, we used a simple SIRD model to describe the evolution of COVID-19. Such a choice has been done for the sake of simplicity. More complex and realistic epidemic models that include, e.g., further health states (Giordano et al., 2020) or spatially-distributed populations (Rossa et al., 2020; Parino, Zino, Porfiri, & Rizzo, 2021), may be embedded within our framework, and our parameter identification method can be used to shed light on further important aspects, such as estimating the fraction of non-detected infected individuals or evaluating the impact of spatial heterogeneity.

5. Conclusion

In this paper we presented an optimization-based system identification approach that can be effectively used for estimating time-varying but piecewise constant parameters in a broad class of nonlinear dynamical models. The proposed method, which is based on casting the estimation problem as a MIQP, has several advantages, including the

possibility to set the number of changes in the model parameters and incorporating additional information available on the data quality, without relying on assumptions on the switching rule or on knowledge of the switching instants. The proposed optimization technique can be conveniently approximated via sub-optimal solutions of the MIQP as well as relaxed into a computationally-efficient QP, whose solution is shown to be in practice a good proxy of the original one. After validation on a synthetic case study, the practical use of our method and its approximations was finally demonstrated on two case studies.

The results presented in this paper, besides demonstrating the effectiveness of the proposed method, arise a few points for discussion, which pave the way for directions of future research. First, our algorithm requires the knowledge (or, at least, an estimate) of the covariance matrix of the measurement errors. In our case studies, we have assumed it to be constant and we have estimated it using the covariance of the measured data. However, more refined methods to estimate such a matrix should be designed in order to adopt our method in real-world scenarios, especially when measurement errors change over time. Second, our case studies suggest that the parameter k_i that determines the number of times the parameter varies during the entire time-window is crucial. In order to tune such a parameter, one should use the information available on the system, and combine it with heuristics. A simple heuristic could be to re-iterate the procedure, starting with k_i large, and then decrease it until the solution remains quite stable. Future work should focus on testing such heuristic and design more refined ones. Third, despite the empirical observation that the solution of the QP is often a sufficiently good proxy of the MIQP solution, more refined approximations methods should be designed, possibly leveraging finer regularization techniques for the cardinality bound (Tillmann et al., 2024), and rigorous theoretical guarantees on their performance must be derived. Fourth, an extensive comparison of the performance of different methods to estimate time-varying parameters with abrupt changes (e.g., comparing our method with algorithms based on recursive least-squares with forgetting factors (Goel & Bernstein, 2018; Paleologu et al., 2008) or switching systems (Bianchi et al., 2020)) is an important direction of future investigation in order to determine key advantages and disadvantages of each of these methods.

Despite its generality and flexibility, our parameter identification approach is not exempt from limitations. First, we assumed that the entire state of the system is observable, up to some level of uncertainty. In many realistic scenarios, however, only some of the variables of the systems can be observed (e.g., the number of detected infected individuals in an epidemiological application), while it is not possible to have information on other variables (e.g., the undetected infected individuals). Future extensions of the approach shall be developed to account for cases in which only part of the state is observable. Second, our method relies on a first-order expansion of the system's dynamics. While this is a standard practice and it is admissible in many real-world applications such as those considered in our case studies, further analysis is needed to deal with highly nonlinear dynamics, where high-order terms may not be neglected.

CRediT authorship contribution statement

Giuseppe Carlo Calafiore: Writing – original draft, Supervision, Resources, Project administration, Methodology, Investigation, Formal analysis, Conceptualization. **Giulia Fracastoro:** Writing – review & editing, Validation, Investigation, Formal analysis. **Lorenzo Zino:** Writing – original draft, Visualization, Validation, Software, Investigation, Formal analysis, Data curation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Alexander, J. W. (2014). *Brewed in Japan: the evolution of the Japanese beer industry*. Honolulu HI, US: University of Hawaii Press.
- Alicic, R., Paré, P. E., & Sandberg, H. (2023). Change time estimation uncertainty in nonlinear dynamical systems with applications to COVID-19. *International Journal of Robust and Nonlinear Control*, 33(9), 4732–4760. <http://dx.doi.org/10.1002/rnc.5974>.
- Andersen, E. D., Roos, C., & Terlaky, T. (2003). On implementing a primal–dual interior-point method for conic quadratic optimization. *Mathematical Programming*, 95(2), 249–277. <http://dx.doi.org/10.1007/s10107-002-0349-3>.
- Åström, K., & Eykhoff, P. (1971). System identification—a survey. *Automatica*, 7(2), 123–162. [http://dx.doi.org/10.1016/0005-1098\(71\)90059-8](http://dx.doi.org/10.1016/0005-1098(71)90059-8).
- Bako, L., Boukharouba, K., & Lecoeuche, S. (2010). An ℓ_0 – ℓ_1 norm based optimization procedure for the identification of switched nonlinear systems. In *Proceedings of the 49th IEEE Conference on Decision and Control* (pp. 4467–4472). <http://dx.doi.org/10.1109/CDC.2010.5717199>.
- Beck, N. (1983). Time-varying parameter regression models. *American Journal of Political Science*, 27(3), 557. <http://dx.doi.org/10.2307/2110985>.
- Bemporad, A., Garulli, A., Paoletti, S., & Vicino, A. (2005). A bounded-error approach to piecewise affine system identification. *IEEE Transactions on Automatic Control*, 50, 1567–1580. <http://dx.doi.org/10.1109/TAC.2005.856667>.
- Bianchi, F., Prandini, M., & Piroddi, L. (2020). A randomized two-stage iterative method for switched nonlinear systems identification. *Nonlinear Analysis. Hybrid Systems*, 35, Article 100818. <http://dx.doi.org/10.1016/j.nahs.2019.100818>.
- Brauer, F., & Castillo-Chavez, C. (2012). *Mathematical models in population biology and epidemiology* (2nd ed.). New York NY, USA: Springer, <http://dx.doi.org/10.1007/978-1-4614-1686-9>.
- Calafiore, G. C., & Fracastoro, G. (2022). Age structure in SIRD models for the COVID-19 pandemic—a case study on Italy data and effects on mortality. *PLoS ONE*, 17(2), 1–17. <http://dx.doi.org/10.1371/journal.pone.0264324>.
- Calafiore, G. C., Novara, C., & Possieri, C. (2020). A time-varying SIRD model for the COVID-19 contagion in Italy. *Annual Reviews in Control*, 50, 361–372. <http://dx.doi.org/10.1016/j.arcontrol.2020.10.005>.
- Chiuso, A., & Pillonetto, G. (2019). System identification: A machine learning perspective. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(1), 281–304. <http://dx.doi.org/10.1146/annurev-control-053018-023744>.
- Della Rossa, F., Salzano, D., Di Meglio, A., De Lellis, F., Coraggio, M., Calabrese, C., et al. (2020). A network model of Italy shows that intermittent regional strategies can alleviate the COVID-19 epidemic. *Nature Communications*, 11(1), 5106. <http://dx.doi.org/10.1038/s41467-020-18827-5>.
- Dipartimento della Protezione Civile (2020). COVID-19 Italia – Monitoraggio situazione. <https://github.com/pcm-dpc/COVID-19>.
- Esfandiari, K., & Narendra, K. S. (2022). Identification and control of linear systems with piece-wise constant parameters. In *2022 IEEE International Conference on Systems, Man, and Cybernetics* (pp. 3006–3011). <http://dx.doi.org/10.1109/SMC53654.2022.9945441>.
- Fearnhead, P. (2006). Exact and efficient bayesian inference for multiple changepoint problems. *Statistics and Computing*, 16(2), 203–213. <http://dx.doi.org/10.1007/s11222-006-8450-8>.
- Forgione, M., & Piga, D. (2021). Continuous-time system identification with neural networks: Model structures and fitting criteria. *European Journal of Control*, 59, 69–81. <http://dx.doi.org/10.1016/j.ejcon.2021.01.008>.
- Giordano, G., Blanchini, F., Bruno, R., Colaneri, P., Di Filippo, A., Di Matteo, A., et al. (2020). Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. *Nature Medicine*, 26(6), 855–860. <http://dx.doi.org/10.1038/s41591-020-0883-7>.
- Glushchenko, A., & Lastochkin, K. (2023). Unknown piecewise constant parameters identification with exponential rate of convergence. *International Journal of Adaptive Control and Signal Processing*, 37(1), 315–346. <http://dx.doi.org/10.1002/acs.3533>.
- Goel, A., & Bernstein, D. S. (2018). A targeted forgetting factor for recursive least squares. In *2018 IEEE Conference on Decision and Control* (pp. 3899–3903). IEEE, <http://dx.doi.org/10.1109/cdc.2018.8619181>.
- Grant, M., & Boyd, S. (2008). Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, & H. Kimura (Eds.), *Lecture notes in control and information sciences, Recent advances in learning and control* (pp. 95–110). Springer-Verlag Limited.
- Grant, M., & Boyd, S. (2014). CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>.
- Kashyap, R. (1970). Maximum likelihood identification of stochastic linear systems. *IEEE Transactions on Automatic Control*, 15(1), 25–34. <http://dx.doi.org/10.1109/TAC.1970.1099344>.
- Lauer, F., Bloch, G., & Vidal, R. (2011). A continuous optimization framework for hybrid system identification. *Automatica*, 47(3), 608–613. <http://dx.doi.org/10.1016/j.automatica.2011.01.020>.
- Li, Y., Wei, H.-L., & Billings, S. A. (2011). Identification of time-varying systems using multi-wavelet basis functions. *IEEE Transactions on Control Systems Technology*, 19(3), 656–663. <http://dx.doi.org/10.1109/TCST.2010.2052257>.

- Liu, Z., & Vandenberghe, L. (2009). Semidefinite programming methods for system realization and identification. In *Proceedings of the 48th IEEE Conference on Decision and Control held jointly with 2009 28th Chinese Control Conference* (pp. 4676–4681). <http://dx.doi.org/10.1109/CDC.2009.5400177>.
- Ljung, L. (1998). *System identification: theory for the user* (2nd ed.). London, UK: Pearson Education.
- Ljung, L. (2010). Perspectives on system identification. *Annual Reviews in Control*, 34(1), 1–12. <http://dx.doi.org/10.1016/j.arcontrol.2009.12.001>.
- Ljung, L., & Gunnarsson, S. (1990). Adaptation and tracking in system identification—a survey. *Automatica*, 26(1), 7–21. [http://dx.doi.org/10.1016/0005-1098\(90\)90154-a](http://dx.doi.org/10.1016/0005-1098(90)90154-a).
- Ljung, L., Hjalmarsson, H., & Ohlsson, H. (2011). Four encounters with system identification. *European Journal of Control*, 17(5), 449–471. <http://dx.doi.org/10.3166/ejc.17.449-471>.
- Marasco, A., Picucci, A., & Romano, A. (2016a). Determining firms' utility functions and competitive roles from data on market shares using Lotka–Volterra models. *Data in Brief*, 7, 709–713. <http://dx.doi.org/10.1016/j.dib.2016.03.020>.
- Marasco, A., Picucci, A., & Romano, A. (2016b). Market share dynamics using Lotka–Volterra models. *Technological Forecasting and Social Change*, 105, 49–62. <http://dx.doi.org/10.1016/j.techfore.2016.01.017>.
- Noël, J., & Kerschen, G. (2017). Nonlinear system identification in structural dynamics: 10 more years of progress. *Mechanical Systems and Signal Processing*, 83, 2–35. <http://dx.doi.org/10.1016/j.ymssp.2016.07.020>.
- Norgaard, M., Ravn, O., Poulsen, N., & Hansen, L. (2000). *Neural networks for modelling and control of dynamic systems*. London, UK: Springer.
- Paleologu, C., Benesty, J., & Ciochina, S. (2008). A robust variable forgetting factor recursive least-squares algorithm for system identification. *IEEE Signal Processing Letters*, 15, 597–600. <http://dx.doi.org/10.1109/LSP.2008.2001559>.
- Paré, P. E., Beck, C. L., & Başar, T. (2020). Modeling, estimation, and analysis of epidemics over networks: An overview. *Annual Reviews in Control*, 50, 345–360. <http://dx.doi.org/10.1016/j.arcontrol.2020.09.003>.
- Parino, F., Zino, L., Porfiri, M., & Rizzo, A. (2021). Modelling and predicting the effect of social distancing and travel restrictions on COVID-19 spreading. *Journal of the Royal Society Interface*, 18, Article 20200875. <http://dx.doi.org/10.1098/rsif.2020.0875>.
- Pia, A. D., Dey, S. S., & Molinaro, M. (2016). Mixed-integer quadratic programming is in NP. *Mathematical Programming*, 162(1–2), 225–240. <http://dx.doi.org/10.1007/s10107-016-1036-0>.
- Roll, J., Bemporad, A., & Ljung, L. (2004). Identification of piecewise affine systems via mixed-integer programming. *Automatica*, 40(1), 37–50. <http://dx.doi.org/10.1016/j.automatica.2003.08.006>.
- Schoukens, J., & Ljung, L. (2019). Nonlinear system identification: A user-oriented road map. *IEEE Control Systems Magazine*, 39(6), 28–99. <http://dx.doi.org/10.1109/MCS.2019.2938121>.
- Schwaab, M., Biscaia, J. L., Jr., Monteiro, J. L., & Pinto, J. C. (2008). Nonlinear parameter estimation through particle swarm optimization. *Chemical Engineering Science*, 63(6), 1542–1552. <http://dx.doi.org/10.1016/j.ces.2007.11.024>.
- Söderström, T., & Stoica, P. (1989). *System identification*. Hoboken NJ, US: Prentice Hall.
- Strejc, V. (1980). Least squares parameter estimation. *Automatica*, 16(5), 535–550. [http://dx.doi.org/10.1016/0005-1098\(80\)90077-1](http://dx.doi.org/10.1016/0005-1098(80)90077-1).
- Tillmann, A. M., Bienstock, D., Lodi, A., & Schwartz, A. (2024). Cardinality minimization, constraints, and regularization: A survey. *SIAM Review*, 66(3), 403–477. <http://dx.doi.org/10.1137/21M142770X>.
- Tsatsanis, M., & Giannakis, G. (1993). Time-varying system identification and model validation using wavelets. *IEEE Transactions on Signal Processing*, 41(12), 3512–3523. <http://dx.doi.org/10.1109/78.258089>.
- Vandenberghe, L. (2012). Convex optimization techniques in system identification. In *Proceedings of the 16th IFAC symposium on system identification, vol. 45* (pp. 71–76). <http://dx.doi.org/10.3182/20120711-3-BE-2027.00244>.
- Voss, H. U., Timmer, J., & Kurths, J. (2004). Nonlinear dynamical system identification from uncertain and indirect measurements. *International Journal of Bifurcation and Chaos*, 14(06), 1905–1933. <http://dx.doi.org/10.1142/s0218127404010345>.
- Wan, E., & Van Der Merwe, R. (2000). The unscented Kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Aystems for Signal Processing, Communications, and Control Symposium* (pp. 153–158). <http://dx.doi.org/10.1109/ASSPCC.2000.882463>.
- Willsky, A. S. (2005). *Detection of abrupt changes in dynamic systems* (pp. 27–49). Springer-Verlag. <http://dx.doi.org/10.1007/bfb0006388>.
- Xin, H., Li, Y., Wu, P., Li, Z., Lau, E. H. Y., Qin, Y., et al. (2021). Estimating the latent period of coronavirus disease 2019 (COVID-19). *Clinical Infectious Diseases*, 74(9), 1678–1681. <http://dx.doi.org/10.1093/cid/ciab746>.