Automatica 44 (2008) 2274–2287

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Recursive identification of switched ARX systems ☆

René Vidal

Center for Imaging Science, Department of Biomedical Engineering, Johns Hopkins University, 3400 N. Charles Street, Baltimore MD 21218, USA

ARTICLE INFO

Article history:
Received 19 July 2006
Received in revised form
13 August 2007
Accepted 24 January 2008
Available online 19 May 2008

Keywords:
Recursive identification
Switched ARX systems
Hybrid systems
Recursive least squares
Persistence of excitation

ABSTRACT

We consider the problem of recursively identifying the parameters of a deterministic discrete-time Switched Auto-Regressive eXogenous (SARX) model, under the assumption that the number of models, the model orders and the mode sequence are unknown. The key to our approach is to view the identification of multiple ARX models as the identification of a single, though more complex, lifted dynamical model built by applying a polynomial embedding to the input/output data. We show that the dynamics of this lifted model do not depend on the value of the discrete state or the switching mechanism, and are linear on the so-called *hybrid model parameters*. Therefore, one can identify the parameters of the lifted model using a standard recursive identifier applied to the embedded input/output data. The estimated hybrid model parameters are then used to build a polynomial whose derivatives at a regressor give an estimate of the parameters of the ARX model generating that regressor. The estimated ARX model parameters are shown to converge exponentially to their true values under a suitable persistence of excitation condition on a projection of the embedded input/output data. Such a condition is a natural generalization of the well known result for ARX models. Although our algorithm is designed for perfect input/output data, our experiments also evaluate its performance as a function of the level of noise for different choices of the number of models and model orders. We also present an application to temporal video segmentation.

1. Introduction

Hybrid systems are models of processes governed by differential or difference equations that exhibit both continuous and discontinuous behavior. In some instances, the discontinuous behavior arises genuinely from switching or interaction between separate physical systems, e.g., the switching among different gears on a car. More often, hybrid systems are a mere modeling abstraction: the physical process is intrinsically continuous, but its behavior is more easily described by certain discrete "events," "transitions," "phases," "modes," or "states."

Over the past decade or so, we have witnessed significant advances in the analysis, verification, computation, stability, and control of hybrid systems. These advances have been motivated by numerous potential applications in automatic highway systems, air traffic management systems, multi-agent control and coordination, control under quantization and communication constraints, and more recently, also in molecular biology, quantum computation, and computer vision. However, most of these theoretical developments assume that the parameters and the switching mechanism of the hybrid system are known. This is often not the case in practice. Therefore, we are faced with the task of simultaneously identifying the number of models, the model orders, the model parameters, and the mode sequence of a hybrid system from measurements of its input and output.

1.1. Prior work on hybrid system identification

Most existing hybrid system identification methods have been designed for the class of piecewise affine (PWA) (Münz & Krebs, 2005; Verdult & Verhaegen, 2004), or piecewise Auto-Regressive eXogenous (PWARX) systems (Bemporad, Garulli, Paoletti, & Vicino, 2005; Ferrari-Trecate, Muselli, Liberati, & Morari, 2003; Juloski, Weiland, & Heemels, 2005; Nakada, Takaba, & Katayama, 2005; Roll, Bemporad, & Ljung, 2004), i.e. models in which the regressor space is partitioned into polyhedra with affine or ARX submodels for each polyhedron. Due to the apparent coupling between the tasks of parameter identification and state estimation, most algorithms alternate between these two stages. Existing methods for PWARX systems include the following:

– The *optimization-based procedure* (Roll et al., 2004) finds the model parameters and the partition of the state space using mixed-integer linear and quadratic programming. The number of models and the model orders need to be known.

☆ This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Associate Editor Antonio Vicino under the direction of Editor Torsten Söderström.

E-mail address: rvidal@cis.jhu.edu.

0005-1098/$ – see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2008.01.025

– The *clustering-based procedure* (Ferrari-Trecate et al., 2003; Nakada et al., 2005) finds the partition of the state space using machine learning techniques such as K-means and support vector machines (Duda, Hart, & Stork, 2000), and the model parameters using linear least squares. The method can be extended to estimate the number of models (Ferrari-Trecate & Muselli, 2003), but the model orders need to be known.
– The *Bayesian procedure* (Juloski et al., 2005) iterates between assigning regressors to models and computing the model parameters using a probabilistic approach similar to Expectation Maximization (Duda et al., 2000). The number of models and the model orders need to be known.
– The *bounded-error procedure* (Bemporad et al., 2005) uses a greedy strategy to partition a set of infeasible inequalities into a minimum number of feasible subsystems, and then iterates between assigning regressors to models and computing the model parameters. This method can estimate the number of models, but needs to know the model orders.

For the class of Switched ARX (SARX) models, where the switching is arbitrary, the identification problem can be solved in closed form from the properties of the *hybrid decoupling polynomial* (Ma & Vidal, 2005; Vidal, 2004; Vidal, Soatto, Ma, & Sastry, 2003), an equation on the input/output data that does not depend on the mode sequence. This *algebraic procedure* estimates the model orders from a rank constraint on the input/output data, the number of models from the degree of the polynomial, and the model parameters from the derivatives of the polynomial.

Unfortunately, all the aforementioned hybrid system identification algorithms operate in a *batch mode*, i.e. the model parameters are identified *after* all the input/output data have been collected. This is a significant limitation for real time operation, because the complexity of batch methods depends on the number of data points.

Existing work (Skeppstedt, Ljung, & Millnert, 1992) combines a bank of standard recursive identifiers for ARX models with a classification procedure for identifying PWARX models. However, in this work the discrete state is assumed to be measured. Also, the dynamics of a standard recursive identifier applied to PWARX models are no longer linear, but rather hybrid, hence properties such as convergence of the parameter estimates are no longer straightforward to establish. To the best of our knowledge, other than Hashambhoy and Vidal (2005) and Vidal and Anderson (2004), there is no prior work addressing the *recursive* identification of hybrid dynamical models in the case in which the number of models, the model orders, the model parameters, and the mode sequence are all *unknown*.

1.2. Paper outline and contributions

This paper proposes an exponentially convergent recursive identification algorithm for deterministic SARX systems under the assumption that the number of models, the model orders, the model parameters, and the mode sequence are all *unknown*.

In Section 2, we review the classical output error recursive identification algorithm for ARX systems, both in the case of known and unknown orders. In each case, we derive a persistence of excitation condition on the input/output data that guarantees the exponential convergence of the recursive identifier.

In Section 3, we extend the output error identifier to SARX systems. Rather than defining a recursive identifier for each ARX model, we propose to identify the *hybrid model parameters* $\boldsymbol{h}$ of a lifted model. The estimates of $\boldsymbol{h}$ are then used to build a polynomial p whose derivatives at a regressor give an estimate of the parameters of the ARX model generating that regressor.

In Section 4, we derive a persistence of excitation condition on the embedded input/output data that guarantees the exponential convergence of the estimates of $\boldsymbol{h}$ when the number of models is *known* and the model orders are *known* and *equal*. Our condition is a natural generalization of the one for a single ARX model. However, it also imposes some restrictions on the mode sequence. For instance, it requires that each mode be visited a minimum number of times. Given an exponentially convergent estimate of $\boldsymbol{h}$, we prove that the estimates of the ARX model parameters also converge exponentially to their true values.

In Section 5, we show that when the number of models and the model orders are *unknown* and *possibly different*, the hybrid model parameters are no longer uniquely defined. Instead, they live on a manifold $\mathcal{H}$ of possible solutions. Although in principle this means that the recursive algorithm may not converge, we show that under a suitable persistence of excitation condition on a projection of the embedded regressors, the estimates of $\boldsymbol{h}$ do converge to a point in $\mathcal{H}$ that depends linearly on the initial condition. Furthermore, we show that after enforcing some of the entries of $\boldsymbol{h}$ to be zero, the estimates of the ARX model parameters converge exponentially to their true values.

In Section 6, we present an algorithm for the identification of SARX models. The algorithm runs two identifiers in parallel and increases their number of zeros until both identifiers converge to the same ARX parameters. We also present some variations to the algorithm that improve its performance with noisy data. Experiments for different choices of the number of models and the model orders are presented in Section 7 as well as an application to temporal video segmentation. Section 8 concludes the paper.

2. Recursive identification of ARX systems

Consider a discrete-time Auto-Regressive eXogenous (ARX) system whose dynamics are given by

$$y_t = \sum_{j=1}^{n_a} a_j y_{t-j} + \sum_{j=1}^{n_c} c_j u_{t-j}, \tag{1}$$

where $u_t \in \mathbb{R}$ is the *input*, $y_t \in \mathbb{R}$ is the *output*, n_a and n_c are the system orders, and $\{a_j\}_{j=1}^{n_a}$ and $\{c_j\}_{j=1}^{n_c}$ are the model parameters. Given input/output data $\{u_t, y_t\}_{t=0}^{\infty}$ generated by an ARX system such as (1), we wish to identify the model orders n_a, n_c, and the model parameters, $\{a_j\}_{j=1}^{n_a}$, $\{c_j\}_{j=1}^{n_c}$.

2.1. Identification of ARX systems of known orders

For the sake of simplicity, let us first assume that n_a and n_c are known. If we let $K = n_a + n_c + 1$,

$$\boldsymbol{b} \doteq [c_{n_c}, \ldots, c_1, a_{n_a}, \ldots, a_1, 1]^\top \in \mathbb{R}^K, \quad \text{and} \tag{2}$$

$$\boldsymbol{x}_t \doteq [u_{t-n_c}, \ldots, u_{t-1}, y_{t-n_a}, \ldots y_{t-1}, -y_t]^\top \in \mathbb{R}^K, \tag{3}$$

then we have that for all $t \geq \max\{n_a, n_c\}$ the *regressor* $\boldsymbol{x}_t$ lives in a hyperplane of $\mathbb{R}^K$

$$\boldsymbol{b}^\top \boldsymbol{x}_t = 0 \tag{4}$$

whose normal vector $\boldsymbol{b} \in \mathbb{R}^K$ is simply the vector of ARX model parameters (in homogeneous coordinates).

One may obtain an estimate $\hat{\boldsymbol{b}}_t$ of $\boldsymbol{b}$ from the data up to time t by following the gradient of the prediction error $(\boldsymbol{b}^\top \boldsymbol{x}_t)^2$. The *equation error identifier* (Ljung, 1987) estimates $\hat{\boldsymbol{b}}_t$ by following a normalized version of such gradient

$$\hat{\boldsymbol{b}}_{t+1} = \left(I_K - \frac{\mu \Pi_K \boldsymbol{x}_t \boldsymbol{x}_t^\top}{1 + \mu \|\Pi_K \boldsymbol{x}_t\|^2} \right) \hat{\boldsymbol{b}}_t, \tag{5}$$

where $\mu > 0$ is a parameter, $I_K \in \mathbb{R}^{K\times K}$ is the identity matrix, and $\Pi_K = \begin{bmatrix} I_{K-1} & \boldsymbol{0}_{K-1} \\ \boldsymbol{0}_{K-1}^\top & 0 \end{bmatrix} \in \mathbb{R}^{K\times K}$ with $\boldsymbol{0}_K \in \mathbb{R}^K$ the zero vector. For notational convenience, we will drop the subindex K in I_K and $\boldsymbol{0}_K$ whenever understood.