Contents lists available at ScienceDirect

Automatica

Brief paper

A fast algorithm to assess local structural identifiability☆

Johannes D. Stigter, Jaap Molenaar

Mathematical and Statistical Methods Group, Wageningen University, Droevendaalsesteeg 1, 6708 PB Wageningen, The Netherlands

ARTICLE INFO

ABSTRACT

The paper presents a novel method for assessing the local structural identifiability question for a general non-linear state-space model. The method is a combination of (i) the application of a singular value decomposition to a parametric output sensitivity matrix that is created by simply integrating the model once and, (ii) a symbolic computation for a *reduced model* that is guided by the SVD results and allows a confirmation of the conclusions regarding identifiability obtained in the first step. In case there is a lack of identifiability, the symbolic computation quickly results in determination of the exact structure of the nullspace and a suitable re-parametrisation. The method is discussed in detail and applied to three case studies, of which the last two are considerably large, containing 22 and 43 parameters to be identified, respectively.

1. Introduction

The local structural identifiability problem for dynamic state-space models is a well-known and well-studied topic in the system identification literature. For a given model the problem amounts to finding unique values (at least locally) for a set of model parameters obtained from calibrating the model against a (hypothetical) data-set that is assumed to be measured without any measurement noise. The type of model we are interested in is the (general) continuous time, non-linear state-space model

$$\frac{dx(t)}{dt} = f(x(t), u(t), \theta) \qquad (1)$$

$$y(t) = h(x(t), \theta) \qquad (2)$$

where f denotes the dynamic model structure, h the output or observation function, $x(t)$ is the state vector ($\dim(x) = n$), $u(t)$ the vector of input signals ($\dim(u) = r$), $y(t)$ the vector of output signals or sensor readings ($\dim(y) = m$), and θ a p-dimensional vector of model parameters. We recall that if θ^* denotes the true parameter values, then a system is said to be locally structural identifiable if, for any admissible input $u(t)$ and any parameter vector $\theta \in \Theta$ within an open neighbourhood of θ^*, equivalence of the

output signals $y(u, \theta)$ and $y(u, \theta^*)$ implies $\theta = \theta^*$. For a detailed discussion on local and global structural identifiability we refer to Walter and Pronzato (1997).

Before a calibration of a given model and associated experimental set-up can be performed, it is of course essential to know in advance whether values for the considered set of model parameters can be uniquely determined. The least we need to establish is a confirmation of local structural identifiability of the model for the true values of the parameters $\theta = \theta^*$. In addition, in case the local structural identifiability question is answered in the affirmative, one ideally would like to quantify the uncertainty in the parameter estimate and to analyse possible correlations between the parameters that hamper accurate parameter estimation if *real* measurements are used that include measurement noise. A well-known example of a strong correlation between two parameters is discussed by Holmberg, who shows that in the case of a batch-reactor model describing the growth dynamics of a biomass, the growth parameters μ and K_M are heavily correlated. This correlation leads to large uncertainty bounds on the parameter estimates (Holmberg & Ranta, 1982). In summary, addressing the local structural identifiability problem and analysing (parametric) uncertainty propagation for a given model is an important question in any modelling exercise where real data are used for model calibration.

Of course, an answer to the local structural identifiability problem very much depends on the available input/output signals. Some model parameters may be better identifiable with a sensor that measures the associated state variable, e.g., a temperature sensor for determination of a heat-exchange constant in a heat-diffusion model. Other parameters may not be identifiable, simply because there is a lack of informative data from the available

sensor. It is well known, for example, that oxygen uptake rate data from a so-called rapid oxygen demand toxicity device (RODTOX) are not sufficient for estimation of a few essential bio-kinetic parameters that characterise the growth of a population of micro-organisms feeding on a substrate in a batch reactor (Dochain, Vanrolleghem, & Van Daele, 1995).

Because of its fundamental nature the local structural identifiability problem has a long history. Many different approaches to the problem exist including (i) Pohjanpalo's power series expansion (Pohjanpalo, 1978), (ii) application of Fliess' series expansion (Tunali & Tarn, 1987; Walter & Pronzato, 1996) (iii) application of non-linear control theory and, more specifically, non-linear observability (Hermann & Krener, 1977; Isidori, 1989; Kwatny & Blankenship, 2000), (iii) the similarity transformation approach (Denis-Vidal & Joly-Blanchard, 2004; Denis-Vidal, Joly-Blanchard, & Noiret, 2001; Tunali & Tarn, 1987), (iv) differential algebra methods (Anguelova, Karlsson, & Jirstrand, 2012; Denis-Vidal et al., 2001; Ljung & Glad, 1994; Saccomani, 2004; Saccomani, Audoly, & D'Angió, 2003; Sedoglavic, 2002), and (v) a sensitivity based analysis (Brun & Reichert, 2001; Quaiser & Mönnigmann, 2009). A summary of the various approaches can be found in a recent review paper by Miao and co-workers (Miao, Xia, Perelson, & Wu, 2011). Of course, each of these approaches has its own advantages/disadvantages.

The current literature also shows that a persistent bottle-neck in the analysis of local structural identifiability lies in the fact that for *large models* the analysis is highly complex and impractical because of the very large symbolic expressions that arise when using computer algebra software to address the problem. For example, already when using a Taylor series expansion for a model with only 4 states, 5 unknown parameters, and only one output (sensor), the symbolic expressions can become so long that the results are not tractable any more (Chappell, Godfrey, & Vajda, 1990; Walter & Pronzato, 1990). This computational complexity has already been noted in the early nineties by Chappell et al. and, recently, was re-evaluated and re-confirmed by Chis and co-workers in Chis, Banga, and Balsa-Canto (2011) for several non-linear state-space models. The results show that the available methods for addressing the local structural identifiability problem are very limited in their applicability when dealing with a large non-linear state-space model with many parameters and a very limited sensor availability. In such cases, the number of derivatives of the output signal $y(t)$, i.e. $\dot{y}(t)$, $\ddot{y}(t)$, ..., that are needed in a structural identifiability analysis is practically infeasible.

Our primary goal is to present a hybrid approach to assess local structural identifiability of the non-linear state-space model (1)–(2). The approach attempts to bridge the gap between what is often called theoretical and practical identifiability, i.e. the purely theoretical yes/no question of local structural identifiability (Walter & Pronzato, 1997), and, on the other hand, the propagation of parametric uncertainty in non-linear state space models that is studied in more practically oriented identification studies (for example Brun & Reichert, 2001). The method presented here combines the strong points of both approaches. The numerical calculations in our method are complemented with a tractable symbolic computation that allows a firm conclusion on local structural identifiability. In addition, since large models are central to the discussion, our approach seeks for an efficient analysis of identifiability that can be calculated rapidly for relatively large models with many states and parameters. The algorithm suggested in this paper can be applied, for example, in the emerging field of systems biology, where complicated molecular interactions are studied as part of the dynamic behaviour of mRNA and protein concentrations. In the example section we will demonstrate this for two such models, i.e., a circadian clock model by Goldbeter, and the so-called NFκB model by Lipniacki and co-workers that describes a two-feedback-loop of the nuclear factor κB signalling pathway. First, however,
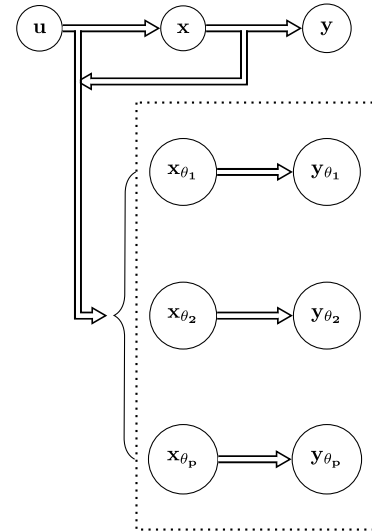

Fig. 1. Diagram of the augmented system, including state vector, parametric state- and output-sensitivities.

we will make the discussion more precise by defining several concepts, such as parametric state- and output sensitivity, Fisher information, etc.

2. Parametric state and output sensitivities—information flow in the model

2.1. Introduction

Parametric output sensitivities can be easily derived for a given (continuous-time) state space model (1)–(2). The associated parametric state/output sensitivity system follows straightforwardly from a differentiation of Eqs. (1)–(2) w.r.t. θ and interchanging the order of differentiation. Denoting with x_θ and y_θ the partial derivatives of the state and output vector with respect to θ, respectively, we get

$$\frac{dx_\theta(t)}{dt} = \frac{\partial f}{\partial x} x_\theta(t) + \frac{\partial f}{\partial \theta} \tag{3}$$

$$y_\theta(t) = \frac{\partial h}{\partial x} x_\theta(t) + \frac{\partial h}{\partial \theta}. \tag{4}$$

Note that the sensitivities x_θ constitute a time dependent ($n \times p$) *matrix* where each column represents the sensitivity of the various states to *one* parameter θ_i taken from the vector θ. Another important observation is that Eqs. (3)–(4) form a *linear time-varying system* for $x_\theta(t)$ with time-dependent coefficients and input signals $\frac{\partial f}{\partial \theta}$ and $\frac{\partial h}{\partial \theta}$. Eqs. (1)–(4) are to be solved simultaneously.

To better understand the information flow (the sensitivity dynamics) in this set of equations, we refer to Fig. 1. Clearly, there is no coupling between sensitivities of different parameters $\{\theta_i, i = 1, \ldots, p\}$, whereas there is coupling between the original system model and *each individual* parametric state and/or output sensitivity. Put differently, we can think of the augmented system (1)–(4) as p sub-systems that run in parallel and are driven by the original system. For each of the p output sensitivities $\{y_{\theta_i}, i = 1, \ldots, p\}$ we observe that its dynamic behaviour is determined by the original state vector $x(t)$, input vector $u(t)$, state sensitivity vector $x_{\theta_i}(t)$, and, possibly, $\frac{\partial h}{\partial \theta}$.[1]

---

[1] In Fig. 1 we have assumed for simplicity that the output function h does not depend on θ. Hence, no arrows have been drawn from the output node $y(t)$ to the sensitivity nodes $y_{\theta_i}(t)$. Of course, dependencies of h w.r.t. θ can easily be included in the analysis.