Automatica 75 (2017) 1–10

Contents lists available at ScienceDirect

Automatica

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

Data-driven control of nonlinear systems: An on-line direct approach[☆]

Marko Tanaskovic [a,d], Lorenzo Fagiano [b,1], Carlo Novara [c], Manfred Morari [a]

[a] Automatic Control Laboratory, Swiss Federal Institute of Technology, Zurich, Switzerland
[b] ABB Switzerland Ltd., Corporate Research, Baden-Daettwil, Switzerland
[c] Dipartamento di Automatica e Informatica, Politechnico di Torino, Italy
[d] Singidunum University, Belgrade, Serbia

ARTICLE INFO

Article history:
Received 10 July 2015
Received in revised form
14 February 2016
Accepted 23 August 2016

Keywords:
Data-driven control
Dynamic inversion
Nonlinear control
Identification for control
Adaptive control

ABSTRACT

A data-driven method to design reference tracking controllers for nonlinear systems is presented. The technique does not derive explicitly a model of the system, rather it delivers directly a time-varying state-feedback controller by combining an on-line and an off-line scheme. Like in other on-line algorithms, the measurements collected in closed-loop operation are exploited to modify the controller in order to improve the tracking performance over time. At the same time, a predictable closed-loop behavior is guaranteed by making use of a batch of available data, which is a feature of off-line algorithms. The feedback controller is parameterized with kernel functions and the design approach exploits results in set membership identification and learning by projections. Under the assumptions of Lipschitz continuity and stabilizability of the system's dynamics, it is shown that if the initial batch of data is informative enough, then the resulting closed-loop system is guaranteed to be finite gain stable. In addition to the main theoretical properties of the approach, the design algorithm is demonstrated experimentally on a water tank system.

1. Introduction

Model-based control design approaches require the derivation of a mathematical model of the plant to be controlled, the identification of the model parameters and the design of a controller based on the derived model. This approach is widely used and gives good results in many applications. However, in several cases, building a detailed and accurate model of a nonlinear plant can be difficult, costly and time-consuming. In these situations, data-driven design techniques represent a possible alternative approach, since they do not require a detailed knowledge of the physics of the system and rely only on the available measured data and relatively little prior information (e.g. qualitative information on the relations between the involved variables, approximate knowledge of the system order, knowledge on the system's structure and on the most important internal states, etc.). In particular, in direct data-driven approaches the controller is designed directly from the measured data, eliminating completely the need for a model of the plant.

The existing direct data-driven approaches can be divided into on-line and off-line ones. In on-line schemes, the controller is modified with each new measurement obtained in closed-loop operation. Examples of on-line direct techniques are the perturbation stochastic approximation control (see Spall & Cristion, 1998), the model free adaptive control (see e.g. Hou & Jin, 2011, 2013a,b) and the unfalsified control (see Safonov & Tsao, 1997, Van Helvoort et al., 2007). The main advantage of on-line techniques is the ability to improve the control performance over time using the measured data. However, since the controller can change at any time, its behavior is often hard to predict. In addition, guaranteeing stability of these control schemes is quite challenging and requires restrictive assumptions on the controlled system.

In off-line procedures, the design is based on a batch of measurements, collected in preliminary experiments before the controller becomes operational, and no further modification is carried out during operation. Such techniques include the iterative

[☆] The material in this paper was partially presented at the 17th IFAC Symposium on System Identification (SYSID 2015), October 19–21, 2015, Beijing, China. This paper was recommended for publication in revised form by Associate Editor Juan I. Yuz under the direction of Editor Torsten Söderström.

E-mail addresses: markotanaskovic@yahoo.co.uk (M. Tanaskovic), lorenzo.fagiano@ch.abb.com (L. Fagiano), carlo.novara@polito.it (C. Novara), morari@control.ee.ethz.ch (M. Morari).

[1] L. Fagiano is now with the Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Italy.

http://dx.doi.org/10.1016/j.automatica.2016.09.032
0005-1098/

feedback tuning (see e.g. Hjalmarsson, 2002, Sjöberg et al., 2003), the correlation based tuning (see Miskovic, Karimi, Bonvin, & Gevers, 2007), the virtual reference feedback tuning (see Campi & Savaresi, 2006, Formentin et al., 2013) and the direct inversion based control (see e.g. Norgaard, Ravn, Poulsen, & Hansen, 2000 and the references therein). In most of these techniques, stability is not considered in the design phase and it is assessed by simulations or experimental verification before the controller becomes operational. Recently, an off-line direct technique that relies on nonlinear set-membership identification (see e.g. Milanese & Novara, 2011) has been proposed in Novara, Fagiano, and Milanese (2013). The approach guarantees theoretically finite-gain stability of the closed-loop system, as the number of data used for the design approaches infinity. The main disadvantage of off-line algorithms is that, unlike the on-line schemes, they do not exploit the additional measurements obtained during controller operation in order to improve the performance. On the other hand, the behavior of a controller designed off-line is usually more predictable.

In this paper, we propose a direct, data-driven design approach that combines the advantages of on-line and off-line techniques. The technique makes use of the theory of learning by projections (see Theodoridis, Slavakis, & Yamada, 2011 and the references therein) to update the controller on-line. At the same time, under the assumptions of Lipschitz continuity and stabilizability of the system's dynamics, it is shown that if the initial batch of data is informative enough, then the resulting closed-loop system is guaranteed to be finite gain stable. In particular, stability is achieved by enforcing a robust constraint on the control input; such a constraint is derived by means of set-membership identification. The mentioned theoretical results are obtained by considering the control design problem as a static inversion, where one aims to derive, from experimental data, an approximate inverse of the system's function. This is the same theoretical framework as in Novara, Fagiano et al. (2013), but the design approach and the results are completely different in order to enable on-line learning while retaining the stability guarantee. In addition, unlike the scheme in Novara, Fagiano et al. (2013), the proposed on-line design algorithm leads to closed-loop stability even when the number of initially available data is finite. After describing the approach and its properties, we present the experimental results obtained on a laboratory water tank system, where we compare our technique with a purely off-line direct design approach and a well tuned linear controller.

The paper is organized as follows. The control problem is defined in Section 2 and the design algorithm is presented in Section 3. The theoretical analysis of the proposed scheme is described in Section 4.1, while Section 4.2 discusses the tuning of the involved parameters. The experimental results are presented in Section 5 and conclusions are drawn in Section 6.

2. Problem statement

We consider a discrete-time nonlinear system with one input and n_x states, represented by the following state equation:

$$x_{t+1} = g(x_t, u_t) + e_{t+1}, \tag{1}$$

where $t \in \mathbb{Z}$ is the discrete time variable, $u_t \in \mathbb{R}$ is the control input, $x_t \in \mathbb{R}^{n_x}$ is the vector of measured states and $e_{t+1} \in \mathbb{R}^{n_x}$ is the vector of disturbance signals that accounts for both the measurement noise and process disturbances.

Assumption 1. The noise and disturbance term e_{t+1} is bounded in magnitude:

$$e_{t+1} \in B_\epsilon \doteq \{e_{t+1} : \|e_{t+1}\| \le \epsilon,\ \forall t \in \mathbb{Z}\}, \tag{2}$$

for some $\epsilon > 0$.

For a given compact domain Y and image set Z, let us denote the class of Lipschitz continuous functions over Y, with Lipschitz constant γ, with

$$\mathcal{F}(\gamma, Y) \doteq \Big\{f : Y \to Z : \|f(y^a) - f(y^b)\| \le \gamma \|y^a - y^b\|,\ \forall y^a,\ y^b \in Y\Big\}.$$

Remark 1. Throughout the paper, the notation $\|\cdot\|$ stands for a suitable vector norm chosen by the user (typically 2- or ∞-norm); the presented results hold for any norm.

We further consider that the reference trajectories of interest for the system to control are defined in a compact set $X \subset \mathbb{R}^{n_x}$, and that input constraints are present in the form of a compact interval $U \subset \mathbb{R}$. The system at hand is assumed to enjoy the following regularity property over these sets:

Assumption 2. For any $x \in X$, the function g is Lipschitz continuous with respect to u, i.e.

$$\forall x \in X, \quad g(x, \cdot) \in \mathcal{F}(\gamma_g, U). \tag{3}$$

The function g in (1) is unknown to the control designer, but a set $\mathcal{D}_N$ of past input and state measurements is available at time $t = 0$:

$$\mathcal{D}_N \doteq \{u_t, \omega_t\}_{t=-N}^{-1}, \tag{4}$$

where

$$\omega_t \doteq (x_t, x_{t+1}). \tag{5}$$

Assumption 3. The batch of data $\mathcal{D}_N$ is such that $u_t \in U$ and $w_t \in X \times X$, $\forall t = -N, \dots, -1$.

Remark 2. If the system is open-loop unstable, a pre-stabilizing controller (possibly a human operator, as in Fagiano & Novara, 2014) can be used to carry out the initial experiments to collect the data $\mathcal{D}_N$. Such data are usually collected and commonly used also in model-based approaches, to identify the parameters of the mathematical model of the system. Another scenario to which our approach applies is when a high-fidelity model of g is available, but it is too complex to carry out model-based control design. In this case, the data $\mathcal{D}_N$ can be also generated through simulations with such a model.

Remark 3. The problem settings introduced so far can be extended to the case of output-feedback control design, if the system state is not known or not fully measured. In this case, one can replace the state in (1) with a regressor or pseudo-state, composed by present and past values of the input and of the output. Then, under reasonable controllability/observability conditions, the dynamics can still be written in the form (1), and algorithms and results similar to those presented in the following can be derived. Another option is to employ a state observer (that can also be designed from data, see e.g. Novara, Ruiz, & Milanese, 2013), to obtain an estimate of the state, which would be affected by estimation errors that can be embedded in the term e_{t+1}. Similarly, when the system involves dynamics that are neglected, i.e. there are additional states with respect to those contained in x, our formulation and results are still valid as long as one assumes to embed the effects of such neglected dynamics into the additive term e_{t+1}.

In our theoretical derivations, we consider the notion of finite gain stability (see e.g. Khalil, 1996).