Learning-Based Adaptive Control

An Extremum Seeking Approach – Theory and Applications

Mouhacine Benosman

BH

LEARNING-BASED ADAPTIVE CONTROL

An Extremum Seeking Approach - Theory and Applications

MOUHACINE BENOSMAN

ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Butterworth-Heinemann is an imprint of Elsevier

Butterworth-Heinemann is an imprint of Elsevier
The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, UK
50 Hampshire Street, 5th Floor, Cambridge, MA 02139, USA

Notices

Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

Library of Congress Cataloging-in-Publication Data
A catalog record for this book is available from the Library of Congress

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library

ISBN: 978-0-12-803136-0

For information on all Butterworth Heinemann publications
visit our website at https://www.elsevier.com/

Publisher: Joe Hayton
Acquisition Editor: Sonnini R. Yura
Editorial Project Manager: Ana Claudia Abad Garcia
Production Project Manager: Kiruthika Govindaraju
Cover Designer: Victoria Pearson Esser

Typeset by SPi Global, India

LEARNING-BASED ADAPTIVE CONTROL

PREFACE

Paris 18th, 1914: The day was, by all accounts, a sunny and pleasant one, the blue sky a perfect backdrop for spectacle. A large crowd had gathered along the banks of the Seine, near the Argenteuil bridge in the city's northwestern fringes, to witness the Concours de la Sécurité en Aéroplane, an aviation competition organized to show off the latest advances in flight safety. Nearly sixty planes and pilots took part, demonstrating an impressive assortment of techniques and equipment. Last on day's program, flying a Curtiss C-2 biplane, was a handsome American pilot named Lawrence Sperry. Sitting beside him in the C-2's open cockpit was his French mechanic, Emil Cachin. As Sperry flew past the ranks of spectators and approached the judges' stand, he let go of the plane's controls and raised his hands. The crowd roared. The plane was flying itself!

Carr (2014, Chapter 3)

This was, a hundred years ago, one of the first demonstrations of adaptation and control implemented in hardware at that time. Since then, adaptive control has been one of the main problems studied in control theory. The problem is well understood, yet it has a very active research frontier (cf. Chapter 2). This monograph focuses on a specific subclass of adaptive control, namely learning-based adaptive control.

As we will see in Chapter 2, adaptive control can be divided into three main subclasses: the classical model-based adaptive control, which mainly uses physics-based models of the controlled system; the model-free adaptive control, which is solely based on the interaction of the controller with the system; and learning-based adaptive control, which uses both model-based and model-free techniques to design flexible yet fast and stable (i.e., safe) adaptive controllers. The basic idea of learning-based modular adaptive control that we introduce in this monograph is depicted in Fig. A. We see that there are two main blocks: the model-based block and the model-free block. The model-based part is concerned with ensuring some type of stability during the learning process. The model-free (i.e., learning) part is concerned with improving the performance of the controllers by tuning online some parameters of the model-based controllers. Due to the modular design, the two blocks can be connected safely, that is, without jeopardizing the stability (in the sense of boundedness) of the whole system.

We argue that one of the main advantages of this type of adaptive controllers, compared to other approaches of adaptation, is the fact that they ensure stability of the system, yet they take advantage of the flexibility of model-free learning algorithms. Model-based adaptive controllers can be

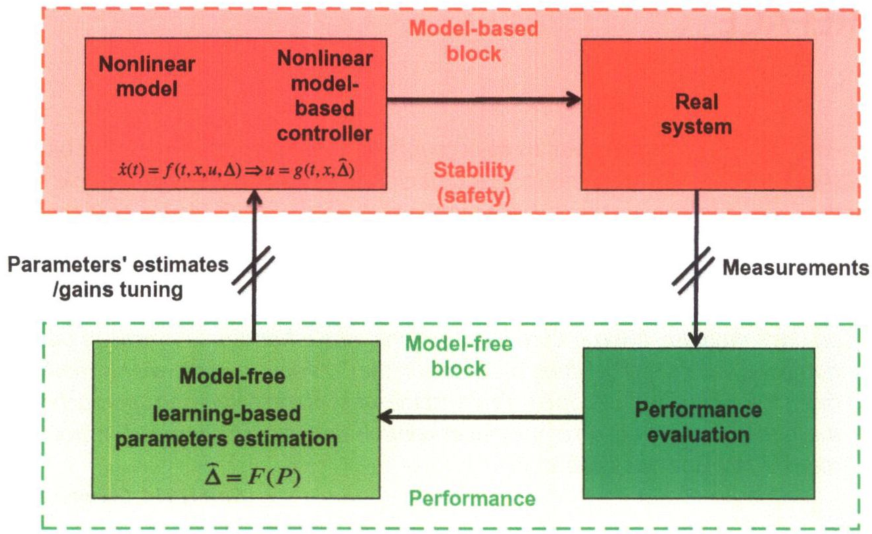

Fig. A Block diagram of the modular learning-based adaptive control.

very efficient and stable. However, they impose many constraints on the model, as well as on the uncertainties' structure (e.g., linear vs nonlinear structures, etc.). Model-free adaptive controllers, on the other hand, allow a lot of flexibility in terms of model structures because they do not rely on any model. However, they lack some of the stability guarantees which characterize model-based adaptive controllers. Furthermore, model-free adaptive algorithms have to learn the best control action (or policy) over a large domain of control options because they do not take advantage of any physical knowledge of the system; that is, they do not use any model of the system. Learning-based adaptive controllers strike a balance; they have some of the stability guarantees due to their model-based part, but also are expected to converge faster to the optimal performance, compared to their model-free counterparts. This is due to the fact that they use some initial knowledge and modeling of the system, albeit uncertain or incomplete.

To make the book easier to read, without the need to refer to other sources, we will recall in Chapter 1 the main definitions and tools used in control theory. This includes the classical definitions of vector spaces, Hilbert spaces, invariant sets, and so on. Other useful stability concepts like Lyapunov, Lagrange stability, and input-to-state stability (ISS) will also be recalled. Finally, some important notions of passivity and nonminimum phase will be presented as well.

In Chapter 2, we will present a general survey of the adaptive control field. We will classify some of the main relevant results from a few subfields of adaptive control theory. The main goal of this chapter is to situate the results of this monograph in the global picture of adaptive control, so that the reader can better understand where our results stand, and how they differ from other works.

The remaining chapters are more technical because they are about more specific results of learning-based adaptation, which we have been working on for the past 5 years.

Starting with Chapter 3, we focus on a very specific problem in learning-based adaptation, namely the problem of iterative feedback tuning (IFT). The main goal of IFT is to automate the tuning of feedback gains for linear or nonlinear feedback controllers. We will first give a brief overview of some IFT research results, and then introduce our work in this field. More specifically, we will focus on extremum seeking-based nonlinear IFT. Throughout this book, we will often focus on nonlinear models (maybe with the exception of Chapter 6) because we believe that, with some simplifications, the nonlinear results can easily be applied to linear models. We will not, however, explicitly derive such simplifications; we leave it to the interested reader to apply the techniques presented here to linear plants.

Chapter 4 presents the general formulation of our modular extremum seeking-based adaptive control for nonlinear models. We will start this chapter with the case of general nonlinear models, without imposing any structural constraints on the model's equations or on the uncertainties (besides some basic smoothness constraints). For this rather general class of models we will argue that, under the assumption of input-to-state stabilizability (by feedback), we can design modular learning-based indirect adaptive controllers, where a model-free learning algorithm is used to estimate online the model's parametric uncertainties. We then focus on a more specific class of nonlinear systems, namely nonlinear systems affine in the control vector. For this class of nonlinear systems we present a constructive control design, which ensures the ISS, and then complement it with an extremum seeking model-free learning algorithm to estimate the model's uncertainties.

In Chapter 5, we will study the problem of real-time nonlinear model identification. What we mean by real time is that we want to identify some parameters of the system online while the system is performing its nominal tasks, without the need to stop or change the system's tasks solely for the purpose of identification. Indeed, real-time identification, if achieved properly, can be of great interest in industrial applications, where interrupting a system's task can lead to big financial losses. If we can identify the system's parameters and keep updating them in real time, we can track their drift, for instance due to the aging of the system or due to a change in the nominal task (a manipulator arm moving different parts, with different masses, etc.) in real time, and then update the model accordingly.

We will study the problem of extremum seeking-based parametric model identification for both finite dimension ordinary differential equation models and infinite dimension partial differential equations (PDEs). We will also study in Chapter 5 a related problem, namely reduced order stabilization for PDEs. In this problem we will use model-free extremum seekers to auto-tune stabilizing terms, known as closure models, which are used to stabilize reduced order models obtained by projecting the PDEs onto a finite dimensional space.

Finally, as a by-product of the approach advocated in this book, we will study in Chapter 6 the specific problem of model predictive control (MPC) for linear models with parametric uncertainties. This case can be seen as a special case of the general results presented in Chapter 4, where the controllers ensuring ISS, are in the form of a model predictive controller. We will use recent results in the field to design an MPC with ISS properties (which is a rather standard result, due to the numerous recent papers on the ISS-MPC topic), and then *carefully and properly* complement it with a model-free extremum seeker to iteratively learn the model's uncertainties and improve the overall performance of the MPC.

"Conclusions and Further Notes" chapter summarizes the presented results. We close with a few thoughts about possible extensions of the results of this book, and mention some open problems which we think are important to investigate in future adaptive control research.

M. Benosman
Cambridge, MA, United States
March 2016

REFERENCES

N. Carr, The Glass Cage: How Our Computers Are Changing Us, W.W. Norton & Company Ltd., New York, 2014.

ACKNOWLEDGMENTS

I would like to thank everyone who has supported me along the way, including my family, collaborators, and coworkers—in particular, Anthony Vetro, Petros Boufonous, Arvind Raghunathan, Daniel Nikovski, Abraham Goldsmith, and Piyush Grover. I would also like to thank Matthew Brand and Alan Sullivan for proofreading the opening and closing chapters of this book, and for their friendship and support. My gratitude also goes to all the students with whom I have worked over the past 5 years; some parts of the research presented in this book are based on collaborations with Gokhan Atinc, Anantharaman Subbaraman, Meng Xia, and Boris Kramer. A special thanks to my friends, Jennifer Milner and Marjorie Berkowitz, for their encouragement. Finally, I want to thank my editors for their patience throughout the course of this project.

As a special note, I would also like to dedicate this modest work to my late colleague, Mr John C. Barnwell III, a genius hardware expert, and most important, a genuine person, who left us prematurely, while I was finishing this book.

Dove si urla, non c'è vera conoscenza. (Where there is shouting, there is no true knowledge.)

Leonardo da Vinci

Mon verre n'est pas grand mais je bois dans mon verre. (My glass is not large, but I drink from my glass.)

Alfred Louis Charles de Musset

CONTENTS

CHAPTER 1
Some Mathematical Tools

We will report here most of the mathematical tools which will be used throughout this book. The goal is to give the reader the main tools to be able to understand the remaining chapters of this monograph. The concepts presented here might seem too general because the chapter includes all the mathematical definitions which will be used at some point in the book. However, later on, to make it more specific, in each of the chapters we will start with a brief recall of the main mathematical tools which will be needed in each specific chapter.
We start with some useful definitions and properties related to vectors, matrices, and functions, see Golub and Van Loan (1996).

1.1 NORMS DEFINITIONS AND PROPERTIES

Let us start by recalling the important definition of vector spaces, which are often used in control theory.

Definition 1.1. *A vector space is a set V on which two operations $+$ and $\cdot$ are defined, called vector addition and scalar multiplication.*

The operation $+$ (vector addition) must satisfy the following conditions:
- *Closure: if u and v are any vectors in V, then the sum $u + v$ belongs to V.*
- *Commutative law: for all vectors u and v in V, $u + v = v + u$.*
- *Associative law: for all vectors u, v, and w in V, $u + (v + w) = (u + v) + w$.*
- *Additive identity: the set V contains an additive identity element, denoted by 0, such that for any vector v in V, $0 + v = v$ and $v + 0 = v$.*
- *Additive inverses: for each vector v in V, the equations $v + x = 0$ and $x + v = 0$ have a solution x in V, called an additive inverse of v, and denoted by $-v$.*

The operation $\cdot$ (scalar multiplication) is defined between real numbers (or scalars) and vectors, and must satisfy the following conditions:
- *Closure: if v is any vector in V, and c is any real number, then the product $c \cdot v$ belongs to V.*
- *Distributive law: for all real numbers c and all vectors u, v in V, $c \cdot (u + v) = c \cdot u + c \cdot v$.*

Learning-Based Adaptive Control
http://dx.doi.org/10.1016/B978-0-12-803136-0.00001-4

- *Distributive law: for all real numbers c, d and all vectors v in V, $(c + d) \cdot v = c \cdot v + d \cdot v$.*
- *Associative law: for all real numbers c, d and all vectors v in V, $c \cdot (d \cdot v) = (cd) \cdot v$.*
- *Unitary law: for all vectors v in V, $1 \cdot v = v$.*

A very well-known example of vector spaces is the space $\mathbb{R}^n$, $n \in \mathbb{Z}^+$. It will often be used in this book to define our mathematical models.

Definition 1.2. *A Hilbert space is a vector space $\mathcal{H}$ associated with an inner product $\langle ., . \rangle$ such that the norm defined by $|f|^2 = \langle f, f \rangle, \quad \forall f \in \mathcal{H}$, makes $\mathcal{H}$ a complete metric space.*

A well-known example of finite-dimension Hilbert spaces is the $\mathbb{R}^n$ associated with the inner product $\langle u, v \rangle = u^T v$, that is, the vector dot product.

Next, we recall the definition of norms of a vector x in a vector space $\mathbb{V}$. The norm of x can be seen as the extension to vector objects of the classical absolute value for a scalar element $|a|$, $a \in \mathbb{R}$. A more rigorous definition is given next.

Definition 1.3. *The norm of a vector $x \in \mathbb{V}$ is a real valued function $\|.\| : \mathbb{V} \to \mathbb{R}$, s.t.,*

1. $\|x\| \geq 0$, *with* $\|x\| = 0$ *iff* $x = 0$
2. $\|ax\| = |a| \|x\|, \quad \forall a \in \mathbb{R}$
3. $\|x + y\| \leq \|x\| + \|y\|, \quad \forall y \in \mathbb{V}$

Some examples of vector norms are as follows:

- *The infinity norm:* $\|x\|_\infty = \max_i |x_i|$
- *The one norm:* $\|x\|_1 = \sum_i |x_i|$
- *The Euclidean (or two) norm:* $\|x\|_2 = \sqrt{\sum_i |x_i|^2}$
- *The p-norm:* $\|x\|_p = \left(\sum_i |x_i|^p\right)^{1/p}, \quad \forall p \in \mathbb{Z}^+$

We recall that these norms are equivalent in $\mathbb{R}^n$, that is, $\forall x \in \mathbb{R}^n$; these norms satisfy the inequalities

$$\begin{aligned} \|x\|_\infty &\leq \|x\|_2 \leq \sqrt{n}\|x\|_\infty, \\ \|x\|_\infty &\leq \|x\|_1 \leq n\|x\|_\infty. \end{aligned} \tag{1.1}$$

The Euclidean norm satisfies the following (Cauchy-Schwarz) inequality

$$|x^T y| \leq \|x\|_2 \|y\|_2, \quad \forall x, y \in \mathbb{R}^n. \tag{1.2}$$

We define now the induced matrix norms associated with vector norms. We recall that a matrix $A \in \mathbb{R}^{m \times n}$ consists of m rows, and n columns of real elements. It can be defined as the linear operator $A(.) : \mathbb{R}^n \to \mathbb{R}^m$, s.t.,

$$\tilde{x} = Ax, \ \tilde{x} \in \mathbb{R}^m, \ x \in \mathbb{R}^n. \tag{1.3}$$

Definition 1.4. *For a given vector norm $\|x\|$, $\forall x \in \mathbb{R}^n$, we define the induced matrix norm $\|A\|$, $\forall A \in \mathbb{R}^{m\times n}$, as*

$$\|A\| = \sup_{x\neq 0} \frac{\|Ax\|}{\|x\|} = \sup_{\|x\|=1} \|Ax\|, \quad \forall x \in \mathbb{R}^n. \tag{1.4}$$

Some properties of induced matrix norms are

1. $\|Ax\| \leq \|A\|\|x\|, \quad \forall x \in \mathbb{R}^n$
2. $\|A + B\| \leq \|A\| + \|B\|$
3. $\|AB\| \leq \|A\| + \|B\|$

Examples of induced norms for $A \in \mathbb{R}^{m\times n}$ are

- *The infinity induced matrix norm:* $\|A\|_\infty = \max_i \sum_j |a_{ij}|$
- *The one induced matrix norm:* $\|A\|_1 = \max_j \sum_i |a_{ij}|$
- *The two induced matrix norm:* $\|A\|_2 = \sqrt{\lambda_{\max}(A^T A)}$, where $\lambda_{\max}(A)$ is the maximum eigenvalue of A

Another frequently used matrix norm is the Frobenius norm, defined as $\|A\|_F = \sqrt{\sum_i \sum_j |a_{ij}|^2}$. It is worth noting that the Frobenius norm is not induced by any vector p-norm.

Some other useful definitions and properties related to matrices are recalled below.

Definition 1.5. *A real symmetric matrix $A \in \mathbb{R}^{n\times n}$ is positive (negative) semidefinite if $x^T Ax \geq 0$ ($x^T Ax \leq 0$), for all $x \neq 0$.*

Definition 1.6. *A real symmetric matrix $A \in \mathbb{R}^{n\times n}$ is positive (negative) definite if $x^T Ax > 0$ ($x^T Ax < 0$), for all $x \neq 0$.*

Definition 1.7. *The leading principles of a real symmetric matrix $A \in \mathbb{R}^{n\times n}$ are defined as the submatrices* $A_i = \begin{bmatrix} a_{11} & \dots & a_{1i} \\ \vdots & & \vdots \\ a_{i1} & \dots & a_{ii} \end{bmatrix}$, $i = 2, \dots, n$.

The positive (negative) definiteness of a matrix can be tested by using the following properties:

1. A real symmetric matrix $A \in \mathbb{R}^{n\times n}$ is positive (negative) definite iff $\lambda(A)_i > 0$ $(\lambda(A)_i < 0)$, $\quad \forall i = 1, \dots, n$.
2. A real symmetric matrix $A \in \mathbb{R}^{n\times n}$ is positive (negative) definite iff all the determinants of its leading principles $\det(A_i)$ are positive (negative).

Other useful properties are as follows:

1. If $A > 0$, then its inverse A^{-1} exists and is positive definite.
2. If $A_1 > 0$, and $A_2 > 0$, then $\alpha A_1 + \beta A_2 > 0$ for all $\alpha > 0$, $\beta > 0$.

3. If $A \in \mathbb{R}^{n\times n}$ is positive definite, and $C \in \mathbb{R}^{m\times n}$ is of rank m, then $B = CAC^T \in \mathbb{R}^{m\times m}$ is positive definite.
4. The Rayleigh-Ritz inequality: for any real symmetric matrix $A \in \mathbb{R}^{n\times n}$, $\lambda(A)_{\min} x^T x \leq x^T A x \leq \lambda(A)_{\max} x^T x, \ \forall x \in \mathbb{R}^n$.

The norms and properties defined so far are for constant objects, that is, constant vectors and matrices. It is possible to extend some of these definitions to the case of time-varying objects, that is, time-varying vector functions, and matrices. We recall some of these definitions next.

Consider a vector function $x(t): \ \mathbb{R}^+ \to \mathbb{R}^n$, with $\mathbb{R}^+ = [0, \ \infty[$, its p-norm is defined as

$$\|x(t)\|_p = \left(\int_0^\infty \|x(\tau)\|^p d\tau \right)^{1/p}, \tag{1.5}$$

where $p \in [1, \infty[$, and the norm inside the integral can be any vector norm in $\mathbb{R}^n$. The function is said to be in $\mathcal{L}_p$ if its p-norm is finite. Furthermore, the infinity-norm in this case is defined as

$$\|x(t)\|_\infty = \sup_{t \in \mathbb{R}^+} \|x(t)\|. \tag{1.6}$$

The function is said to be in $\mathcal{L}_\infty$ if its infinity-norm is finite.

Some properties of these function norms are given below:

1. Hölder's inequality: for $p, \ q \in [1, \infty[$, such that $\frac{1}{p} + \frac{1}{q} = 1$, if $f \in \mathcal{L}_p$, $g \in \mathcal{L}_q$, then $\|fg\|_1 \leq \|f\|_p \|g\|_q$.
2. Schwartz inequality: for $f, g \in \mathcal{L}_2$, $\|fg\|_1 \leq \|f\|_2 \|g\|_2$.
3. Minkowski inequality: for $f, g \in \mathcal{L}_p, \ p \in [1, \infty[$, $\|f+g\|_p \leq \|f\|_p + \|g\|_p$.

Next, to slowly move toward the definitions related to dynamical systems and their stability, we introduce some basic functions' definitions and properties, see Perko (1996) and Khalil (1996).

1.2 VECTOR FUNCTIONS AND THEIR PROPERTIES

We will recall here some basic definitions of functions' properties which are often used in dynamical systems theory.

Definition 1.8. *A function $f: \ [0, \ \infty[\to \mathbb{R}$ is continuous on $[0, \ \infty[$, if for any $\epsilon > 0$ there exists a $\delta(t, \epsilon)$, such that for any $t, \ \tilde{t} \in [0, \ \infty[$, with $|t - \tilde{t}| < \delta(t, \epsilon)$ we have $|f(t) - f(\tilde{t})| < \epsilon$.*

Definition 1.9. *A function $f: \ [0, \ \infty[\to \mathbb{R}$ is uniformly continuous on $[0, \ \infty[$, if for any $\epsilon > 0$ there exists a $\delta(\epsilon)$, such that for any $t, \ \tilde{t} \in [0, \ \infty[$, with $|t - \tilde{t}| < \delta(\epsilon)$ we have $|f(t) - f(\tilde{t})| < \epsilon$.*