Asymptotic Tracking By A Reinforcement Learning Based

Mathematics—Advances in Research and Application: 2012 Edition is a ScholarlyEditions™ eBook that delivers timely, authoritative, and comprehensive information about Mathematics. The editors have built Mathematics—Advances in Research and Application: 2012 Edition on the vast information databases of ScholarlyNews.™ You can expect the information about Mathematics in this eBook to be deeper than what you can access anywhere else, as well as consistently reliable, authoritative, informed, and relevant. The content of Mathematics—Advances in Research and Application: 2012 Edition has been produced by the world's leading scientists, engineers, analysts, research institutions, and companies. All of the content is from peer-reviewed sources, and all of it is written, assembled, and edited by the editors at ScholarlyEditions™ and available exclusively from us. You now have a source you can cite with authority, confidence, and credibility. More information is available at http://www.ScholarlyEditions.com/.
Morphing Aerospace Vehicles and Structures provides a highly timely presentation of the state-of-the-art, future directions and technical requirements of morphing aircraft. Divided into three sections it addresses morphing aircraft, bio-inspiration, and smart structures with specific focus on the flight control, aerodynamics, bio-mechanics, materials, and structures of these vehicles as well as power requirements and the use of advanced piezo materials and smart actuators. The tutorial approach adopted by the contributors, including underlying concepts and mathematical formulations, unifies the methodologies and tools required to provide practicing engineers and applied researchers with the insight to synthesize morphing air vehicles and morphing structures, as well as offering direction for future research.
Reinforcement Learning for Optimal Feedback Control develops model-based and data-driven reinforcement learning methods for solving optimal control problems in nonlinear deterministic dynamical systems. In order to achieve learning under uncertainty, data-driven methods for identifying system models in real-time are also developed. The book illustrates the advantages gained from the use of a model and the use of previous experience in the form of recorded data through simulations and experiments. The book's focus on deterministic systems allows for an in-depth Lyapunov-based analysis of the performance of the methods described during the learning phase and during execution. To yield an approximate optimal controller, the authors focus on theories and methods that fall under the umbrella of actor–critic methods for machine learning. They concentrate on establishing stability during the learning phase and the execution phase, and adaptive model-based and data-driven reinforcement learning, to assist readers in the learning process, which typically relies on instantaneous input-output measurements. This monograph provides academic researchers with backgrounds in diverse disciplines from aerospace engineering to computer science, who are interested in optimal reinforcement learning functional analysis and functional approximation theory, with a good introduction to the use of model-based methods. The thorough treatment of an advanced treatment to control will also interest practitioners working in the chemical-process and power-supply industry.
Networked Control Systems (NCSs) are spatially distributed systems for which the communication between sensors, actuators and controllers is realized by a shared (wired or wireless) communication network. NCSs offer several advantages, such as reduced installation and maintenance costs, as well as greater flexibility, over conventional control systems in which parts of control loops exchange information via dedicated point-to-point connections. The principal goal of this book is to present a coherent and versatile framework applicable to various settings investigated by the authors over the last several years. This framework is applicable to nonlinear time-varying dynamic plants and controllers with delayed dynamics; a large class of static, dynamic, probabilistic and priority-oriented scheduling protocols; delayed, noisy, lossy and intermittent information exchange; decentralized control problems of heterogeneous agents with time-varying directed (not necessarily balanced) communication topologies; state- and output-feedback; off-line and on-line intermittent feedback; optimal intermittent feedback through Approximate Dynamic Programming (ADP) and Reinforcement Learning (RL); and control systems with exogenous disturbances and modeling uncertainties.
A comprehensive exploration of the control schemes of human-robot interactions In Human-Robot Interaction Control Using Reinforcement Learning, an expert team of authors delivers a concise overview of human-robot interaction control schemes and insightful presentations of novel, model-free and reinforcement learning controllers. The book begins with a brief introduction to state-of-the-art human-robot interaction control and reinforcement learning before moving on to describe the typical environment model. The authors also describe some of the most famous identification techniques for parameter estimation. Human-Robot Interaction Control Using Reinforcement Learning offers rigorous mathematical treatments and demonstrations that facilitate the understanding of control schemes and algorithms. It also describes stability and convergence analysis of human-robot interaction control and reinforcement learning based control. The authors also discuss advanced and cutting-edge topics, like inverse and velocity kinematics solutions, H2 neural control, and likely upcoming developments in the field of robotics. Readers will also enjoy: A thorough introduction to model-based human-robot interaction control Comprehensive explorations of model-free human-robot interaction control and human-in-the-loop control using Euler angles Practical discussions of reinforcement learning for robot position and force control, as well as continuous time reinforcement learning for robot force control In-depth examinations of robot control in worst-case uncertainty using reinforcement learning and the control of redundant robots using multi-agent reinforcement learning Perfect for senior undergraduate and graduate students, academic researchers, and industrial practitioners studying and working in the fields of robotics, learning control systems, neural networks, and computational intelligence, Human-Robot Interaction Control Using Reinforcement Learning is also an indispensable resource for students and professionals studying reinforcement learning.
The Republic of Korea's industrial policy has directed that nation's economy through nearly three decades of spectacular growth. But the authors of this paper maintain that this policy is showing signs of being outmoded. The time has come, the authors argue, for the Korean government to stop managing the economy's structural development and to redefine the responsibilities of business and government. Under this proposed compact, the allocation of resources would shift from the government to the private industrial and financial sectors. The transformation of the government bureaucracy from an ad hoc policy role to one of a transparent and predictable regulator is a key to the success of this undertaking. These new directions would present the government with enormous challenges. Greater competitive discipline and regulatory oversight would be required. While dealing with the complexities of the transition, the government would have to maintain macroeconomic stability and the momentum of savings and investment. For comparison, the study examines the industrial economies of France, Germany, Japan, and the United States, which underwent similar shifts.
There are many methods of stable controller design for nonlinear systems. In seeking to go beyond the minimum requirement of stability,

Adaptive Dynamic Programming in Discrete Time approaches the challenging topic of optimal control for nonlinear systems using the tools of adaptive dynamic programming (ADP). The range of systems treated is extensive; affine, switched, singularly perturbed and time-delay nonlinear systems are discussed as are the uses of neural networks and techniques of value and policy iteration. The text features three main aspects of ADP in which the methods proposed for stabilization and for tracking and games benefit from the incorporation of optimal control methods: • infinite-horizon control for which the difficulty of solving partial differential Hamilton–Jacobi–Bellman equations directly is overcome, and proof provided that the iterative value function updating sequence converges to the infimum of all the value functions obtained by admissible control law sequences; • finite-horizon control, implemented in discrete-time nonlinear systems showing the reader how to obtain suboptimal control solutions within a fixed number of control steps and with results more easily applied in real systems than those usually gained from infinite-horizon control; • nonlinear games for which a pair of mixed optimal policies are derived for solving games both when the saddle point does not exist, and, when it does, avoiding the existence conditions of the saddle point. Non-zero-sum games are studied in the context of a single network scheme in which policies are obtained guaranteeing system stability and minimizing the individual performance function yielding a Nash equilibrium. In order to make the coverage suitable for the student as well as for the expert reader, Adaptive Dynamic Programming in Discrete Time: • establishes the fundamental theory involved clearly with each chapter devoted to a clearly identifiable control paradigm; • demonstrates convergence proofs of the ADP algorithms to deepen understanding of the derivation of stability and convergence with the iterative computational methods used; and • shows how ADP methods can be put to use both in simulation and in real applications. This text will be of considerable interest to researchers interested in optimal control and its applications in operations research, applied mathematics computational intelligence and engineering. Graduate students working in control and operations research will also find the ideas presented here to be a source of powerful methods for furthering their study.

This text focuses on the scientific study of animal intelligence. It celebrates comparative cognition's first quarter century, with a collection of chapters, covering the realm of the scientific study of animal intelligence.

Psychology of Learning and Motivation

A comprehensive look at state-of-the-art ADP theory and real-world applications This book fills a gap in the literature by providing a theoretical framework for integrating techniques from adaptive dynamic programming (ADP) and modern nonlinear control to address data-driven optimal control design challenges arising from both parametric and dynamic uncertainties. Traditional model-based approaches leave much to be desired when addressing the challenges posed by the ever-increasing complexity of real-world engineering systems. An alternative which has received much interest in recent years are biologically-inspired approaches, primarily RADP. Despite their growing popularity worldwide, until now books on ADP have focused nearly exclusively on analysis and design, with scant consideration given to how it can be applied to address robustness issues, a new challenge arising from dynamic uncertainties encountered in common engineering problems. Robust Adaptive Dynamic Programming zeros in on the practical concerns of engineers. The authors develop RADP theory from linear systems to partially-linear, large-scale, and completely nonlinear systems. They provide in-depth coverage of state-of-the-art applications in power systems, supplemented with numerous real-world examples implemented in MATLAB. They also explore fascinating reverse engineering topics, such how ADP theory can be applied to the study of the human brain and cognition. In addition, the book: Covers the latest developments in RADP theory and applications for solving a range of systems' complexity problems Explores multiple real-world implementations in power systems with illustrative examples backed up by reusable MATLAB code and Simulink block sets Provides an overview of nonlinear control, machine learning, and dynamic control Features discussions of novel applications for RADP theory, including an entire chapter on how it can be used as a computational mechanism of human movement control Robust Adaptive Dynamic Programming is both a valuable working resource and an intriguing exploration of contemporary ADP theory and applications for practicing engineers and advanced students in systems theory, control engineering, computer science, and applied mathematics.

Reinforcement learning is a learning paradigm concerned with learning to control a system so as to maximize a numerical performance measure that expresses a long-term objective. What distinguishes reinforcement learning from supervised learning is that only partial feedback is given to the learner about the learner's predictions. Further, the predictions may have long term effects through influencing the future state of the controlled system. Thus, time plays a special role. The goal in reinforcement learning is to develop efficient learning algorithms, as well as to understand the algorithms' merits and limitations. Reinforcement learning is of great interest because of the large number of practical applications that it can be used to address, ranging from problems in artificial intelligence to operations research or control engineering. In this book, we focus on those algorithms of reinforcement learning that build on the powerful theory of dynamic programming.We give a fairly comprehensive catalog of learning problems, describe the core ideas, note a large number of state of the art algorithms, followed by the discussion of their theoretical properties and limitations.

This book presents a class of novel optimal control methods and games schemes based on adaptive dynamic programming techniques. For systems with one control input, the ADP-based optimal control is designed for different objectives, while for systems with multi-players, the optimal control inputs are proposed based on games. In order to verify the effectiveness of the proposed methods, the book analyzes the properties of the adaptive dynamic programming methods, including convergence of the iterative value functions and the stability of the system under the iterative control laws. Further, to substantiate the mathematical analysis, it presents various application examples, which provide reference to real-world practices.

The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

Reinforcement learning is a learning paradigm concerned with learning to control a system so as to maximize a numerical performance measure that expresses a long-term objective. What distinguishes reinforcement learning from supervised learning is that only partial feedback is given to the learner about the learner's predictions. Further, the predictions may have long term effects through influencing the future state of the controlled system. Thus, time plays a special role. The goal in reinforcement learning is to develop efficient learning algorithms, as well as to understand the algorithms' merits and limitations. Reinforcement learning is of great interest because of the large number of practical applications that it can be used to address, ranging from problems in artificial intelligence to operations research or control engineering. In this book, we focus on those algorithms of reinforcement learning that build on the powerful theory of dynamic programming. We give a fairly comprehensive catalog of learning problems, describe the core ideas, note a large number of state of the art

algorithms, followed by the discussion of their theoretical properties and limitations. Table of Contents: Markov Decision Processes / Value Prediction Problems / Control / For Further Exploration

Comparative Cognition celebrates comparative cognitions first quarter century with a state-of-the-art collection of chapters covering the broad realm of the scientific study of animal intelligence. It will be an invaluable resource for students and professional researchers in all areas of psychology and neuroscience.

In reinforcement learning (RL) problems, learning agents sequentially execute actions with the goal of maximizing a reward signal. The RL framework has gained popularity with the development of algorithms capable of mastering increasingly complex problems, but learning difficult tasks is often slow or infeasible when RL agents begin with no prior knowledge. The key insight behind "transfer learning" is that generalization may occur not only within tasks, but also across tasks. While transfer has been studied in the psychological literature for many years, the RL community has only recently begun to investigate the benefits of transferring knowledge. This book provides an introduction to the RL transfer problem and discusses methods which demonstrate the promise of this exciting area of research. The key contributions of this book are: Definition of the transfer problem in RL domains Background on RL, sufficient to allow a wide audience to understand discussed transfer concepts Taxonomy for transfer methods in RL Survey of existing approaches In-depth presentation of selected transfer methods Discussion of key open questions By way of the research presented in this book, the author has established himself as the pre-eminent worldwide expert on transfer learning in sequential decision making tasks. A particular strength of the research is its very thorough and methodical empirical evaluation, which Matthew presents, motivates, and analyzes clearly in prose throughout the book. Whether this is your initial introduction to the concept of transfer learning, or whether you are a practitioner in the field looking for nuanced details, I trust that you will find this book to be an enjoyable and enlightening read. Peter Stone, Associate Professor of Computer Science Inthesummerof2008,reinforcementlearningresearchersfromaroundtheworld gathered in the north of France for a week of talks and discussions on reinfor- ment learning, on how it could be made more e?cient, applied to a broader range of applications, and utilized at more abstract and symbolic levels. As a participant in this 8th European Workshop on Reinforcement Learning, I was struck by both the quality and quantity of the presentations. There were four full days of short talks, over 50 in all, far more than there have been at any p- vious meeting on reinforcement learning in Europe, or indeed, anywhere else in the world. There was an air of excitement as substantial progress was reported in many areas including Computer Go, robotics, and ?tted methods. Overall, the work reported seemed to me to be an excellent, broad, and representative sample of cutting-edge reinforcement learning research. Some of the best of it is collected and published in this volume. The workshopandthe paperscollectedhere provideevidence thatthe ?eldof reinforcement learning remains vigorous and varied. It is appropriate to re?ect on some of the reasons for this. One is that the ?eld remains focused on a pr- lem — sequential decision making — without prejudice as to solution methods. Another is the existence of a common terminology and body of theory.

This combined survey of operant and classical conditioning provides professional and academic readers with an up-to-date, inclusive account of a core field of psychology research, with in-depth coverage of the basic theory, its applications, and current topics including behavioral economics. Provides comprehensive coverage of operant and classical conditioning, relevant fundamental theory, and applications including the latest techniques Features chapters by leading researchers, professionals, and academicians Reviews a range of core literature on conditioning Covers cutting-edge topics such as behavioral economics

A comprehensive and rigorous introduction for graduate students and researchers, with applications in sequential decision-making problems.

The explosive growth in computational power over the past several decades offers new tools and opportunities for economists. This handbook volume surveys recent research on Agent-based Computational Economics (ACE), the computational study of economic processes modeled as dynamic systems of interacting agents. Empirical referents for "agents" in ACE models can range from individuals or social groups with learning capabilities to physical world features with no cognitive function. Topics covered include: learning; empirical validation; network economics; social dynamics; financial markets; innovation and technological change; organizations; market design; automated markets and trading agents; political economy; social-ecological systems; computational laboratory development; and general methodological issues. *Every volume contains contributions from leading researchers *Each Handbook presents an accurate, self-contained survey of a particular topic *The series provides comprehensive and accessible surveys

Reinforcement learning (RL) and adaptive dynamic programming (ADP) has been one of the most critical research fields in science and engineering for modern complex systems. This book describes the latest RL and ADP techniques for decision and control in human engineered systems, covering both single player decision and control and multi-player games. Edited by the pioneers of RL and ADP research, the book brings together ideas and methods from many fields and provides an important and timely guidance on controlling a wide variety of systems, such as robots, industrial processes, and economic decision-making.

The four volume set LNCS 9489, LNCS 9490, LNCS 9491, and LNCS 9492 constitutes the proceedings of the 22nd International Conference on Neural Information Processing, ICONIP 2015, held in Istanbul, Turkey, in November 2015. The 231 full papers presented were carefully reviewed and selected from 375 submissions. The 4 volumes represent topical sections containing articles on Learning Algorithms and Classification Systems; Artificial Intelligence and Neural Networks: Theory, Design, and Applications; Image and Signal Processing; and Intelligent Social Networks.

A comprehensive guide to the friction, contact and impact on robot control and force feedback mechanism Dynamics and Control of Robotic Manipulators with Contact and Friction offers an authoritative guide to the basic principles of robot dynamics and control with a focus on contact and friction. The authors discuss problems in interaction between human and real or virtual robot where dynamics with friction and contact are relevant. The book fills a void in the literature with a need for a text that considers the contact and friction generated in robot joints during their movements. Designed as a practical resource, the text provides the information needed for task planning in view of contact, impact and friction for the designer of a robot control system for high accuracy and long durability. The authors include a review of the most up-to-date advancements in robot dynamics and control. It contains a comprehensive resource to the effective design and fabrication of robot systems and components for engineering and scientific purposes. This important guide: Offers a comprehensive reference with systematic treatment and a unified framework Includes simulation and experiments used in dynamics and control of robot considering contact, impact and friction Discusses the most current tribology methodology used to treat the multiple–scale effects Contains valuable descriptions of experiments and software used Presents illustrative accounts on the methods employed to handle friction in the closed loop, including the principles, implementation, application scope, merits and demerits Offers a cohesive treatment that covers tribology and multi-scales, multi-physics and nonlinear stochastic dynamics control Written for graduate students of robotics, mechatronics, mechanical engineering, tracking control and practicing professionals and industrial researchers, Dynamics and Control of Robotic Manipulators with Contact and Friction offers a review to effective design and fabrication of stable and durable robot system and components.

Centered around major topic areas of both theoretical and practical importance, the World Congress on Neural Networks provides its registrants -- from a diverse background encompassing industry, academia, and government -- with the latest research and

applications in the neural network field.