Decentralized Reinforcement Learning Applied To Le Robots

This project investigated learning systems consisting of multiple interacting controllers, or agents: each of which employed a modern reinforcement learning method. The objective was to study the utility of reinforcement learning as an approach to complex decentralized control problems. The major accomplishment was a detailed study of multi-agent reinforcement learning applied to a large-scale decentralized stochastic control problem. This study included a very successful demonstration that a multi-agent reinforcement learning system using neural networks could learn high-performance dispatching of multiple elevator cars in a simulated multi-story building. This problem is representative of very large-scale dynamic optimization problems of practical importance that are intractable for standard methods. The performance achieved by the distributed elevator controller surpassed that of the best of the elevator control algorithms accessible in the literature, showing that reinforcement learning can be a useful approach to difficult decentralized control problems. Additional empirical results demonstrated the performance of reinforcement learning-systems in the setting of nonzero-sum games, with mixed results. Some progress was also made in improving theoretical understanding of multi-agent reinforcement learning.

The purpose of this book is to develop in greater depth some of the methods from the author's Reinforcement Learning and Optimal Control recently published textbook (Athena Scientific, 2019). In particular, we present new research, relating to systems involving multiple agents, partitioned architectures, and distributed asynchronous computation. We pay special attention to the contexts of dynamic programming/policy iteration and control theory/model predictive control. We also discuss in some detail the application of the methodology to challenging discrete/combinatorial optimization problems, such as routing, scheduling, assignment, and mixed integer programming, including the use of neural network approximations within these contexts. The book focuses on the fundamental idea of policy iteration, i.e., start from some policy, and successively generate one or more improved policies. If just one improved policy is generated, this is called rollout, which, based on broad and consistent computational experience, appears to be one of the most versatile and reliable of all reinforcement learning methods. In this book, rollout algorithms are developed for both discrete deterministic and stochastic DP problems, and the development of distributed implementations in both multiagent and multiprocessor settings, aiming to take advantage of parallelism. Approximate policy iteration is more ambitious than rollout, but it is a strictly off-line method, and it is generally far more computationally intensive. This motivates the use of parallel and distributed computation. One of the purposes of the monograph is to discuss distributed (possibly asynchronous) methods that relate to rollout and policy iteration, both in the context of an exact and an approximate implementation involving neural networks or other approximation architectures. Much of the new research is inspired by the remarkable AlphaZero chess program, where policy iteration, value and policy networks, approximate lookahead minimization, and parallel computation all play an important role.

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.

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 book presents the proceedings of the International Conference on Aerospace System Science and Engineering (ICASSE 2019), held in Toronto, Canada, on July 30–August 1, 2019, and jointly organized by the University of Toronto Institute for Aerospace Studies (UTIAS) and the Shanghai Jiao Tong University School of Aeronautics and Astronautics. ICASSE 2019 provided a forum that brought together experts on aeronautics and astronautics to share new ideas and findings. These proceedings present high-quality contributions in the areas of aerospace system science and engineering, including topics such as trans-space vehicle system design and integration, air vehicle systems, space vehicle systems, near-space vehicle systems, aerospace robotics and unmanned systems, communication, navigation and surveillance, aerodynamics and aircraft design, dynamics and control, aerospace propulsion, avionics systems, optoelectronic systems, and air traffic management.

The three volume proceedings LNAI 11051 – 11053 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2018, held in Dublin, Ireland, in September 2018. The total of 131 regular papers presented in part I and part II was carefully reviewed and selected from 535 submissions; there are 52 papers in the applied data science, nectar and demo track. The contributions were organized in topical sections named as follows: Part I: adversarial learning; anomaly and outlier detection; applications; classification; clustering and unsupervised learning; deep learningensemble methods; and evaluation. Part II: graphs; kernel methods; learning paradigms; matrix and tensor analysis; online and active learning; pattern and sequence mining; probabilistic models and statistical methods; recommender systems; and transfer learning. Part III: ADS data science applications; ADS e-commerce; ADS engineering and design; ADS financial and security; ADS health; ADS sensing and positioning; nectar track; and demo track.

These proceedings of the SAI Intelligent Systems Conference 2016 (IntelliSys 2016) offer a remarkable collection of chapters on a wide range of topics in intelligent systems, artificial intelligence and their applications to the real world. Authors hailing from 56 countries on 5 continents submitted 404 papers to the conference, attesting to the global importance of the conference's themes. After being reviewed, 222 papers were accepted for presentation, and 168 were ultimately selected for these proceedings. Each has been reviewed on the basis of its originality, novelty and rigorousness. The papers not only present state-of-the-art methods and valuable experience from researchers in the related research areas; they also outline the field's future development.

This book constitutes the revised post-conference proceedings of the 16th European Conference on Multi-Agent Systems, EUMAS 2018, held at Bergen, Norway, in December 2018. The 18 full papers presented in this volume were carefully reviewed and selected from a total of 34 submissions. The papers report on both early and mature research and cover a wide range of topics in the field of multi-agent systems.

Reinforcement learning (RL) is a powerful machine learning paradigm that studies the interaction between a single agent with an unknown environment. A plethora of applications fit into the RL framework, however, in many cases of interest, a team of agents will need to interact with the environment and with each other to achieve a common goal. This is the object study of collaborative multi-agent RL (MARL). Several challenges arise when considering collaborative MARL. One of these challenges is decentralization. In many cases, due to design constraints, it is undesirable or inconvenient to constantly relay data between agents and a centralized location. Therefore, fully distributed solutions become preferable. The first part of this dissertation addresses the challenge of designing fully decentralized MARL algorithms. We consider two problems: policy evaluation and policy optimization. In the policy evaluation problem, the objective is to estimate the performance of a target team policy in a particular environment. This problem has been studied before for the case with streaming data, however, in most implementations the target policy is evaluated using a finite data set. For this case, existing algorithms guarantee convergence at a sub-linear rate. In this dissertation we introduce Fast Diffusion for Policy Evaluation (FDPE), an algorithm that converges at linear rate for the finite data set case. We then consider the policy optimization problem, where the objective is for all agents to learn an optimal team policy. This problem has also been studied recently, however, existing solutions are data inefficient and converge to Nash equilibria (whose performance can be catastrophically bad) as opposed to team optimal policies. For this case we introduce the Diffusion for Team Policy Optimization (DTPO) algorithm. DTPO is more data efficient than previous algorithms and does not converge to Nash equilibria. For both of these cases, we provide experimental studies that show the effectiveness of the proposed methods. Another challenge that arises in collaborative MARL, which is orthogonal to the decentralization problem, is that of scalability. The parameters that need to be estimated when full team policies are learned, grow exponentially with the number of agents. Hence, algorithms that learn joint team policies quickly become intractable. A solution to this problem is for each agent to learn an individual policy, such that the resulting joint team policy is optimal. This problem has been the object of much research lately. However, most solution methods are data inefficient and often make unrealistic assumptions that greatly limit the applicability of these approaches. To address this problem we introduce Logical Team Q-learning (LTQL), an algorithm that learns factored policies in a data efficient manner and is applicable to any cooperative MARL environment. We show that LTQL outperforms previous methods in a challenging predator-prey task. Another challenge is that of efficient exploration. This is a problem both in the single-agent and multi-agent settings, although in MARL it becomes more severe due to the larger state-action space. The challenge of deriving policies that are efficient at exploring the state space has been addressed in many recent works. However, most of these approaches rely on heuristics, and more importantly, they consider the problem of exploring the state space separately from that of learning an optimal policy (even though they are related, since the state-space is explored to collect data to learn an optimal policy). To address this challenge, we introduce the Information Seeking Learner (ISL), an algorithm that displays state of the art performance in difficult exploration benchmarks. The fundamental value of our work on exploration is that we take a fundamentally different approach from previous works. As opposed to earlier methods we consider the problem of exploring the state space and learning an optimal policy jointly. The main insight of our approach is that in RL, obtaining point estimates of the quantities of interest is not sufficient and confidence bound estimates are also necessary.

Multi-Agent Reinforcement Learning and Adaptive Neural Networks

Reinforcement learning (RL) will deliver one of the biggest breakthroughs in AI over the next decade, enabling algorithms to learn from their environment to achieve arbitrary goals. This exciting development avoids constraints found in traditional machine learning (ML) algorithms. This practical book shows data science and AI professionals how to learn by reinforcementand enable a machine to learn by itself. Author Phil Winder of Winder Research covers everything from basic building blocks to state-of-the-art practices. You'll explore the current state of RL, focus on industrial applications, learnnumerous algorithms, and benefit from dedicated chapters on deploying RL solutions to production. This is no cookbook; doesn't shy away from math and expects familiarity with ML. Learn what RL is and how the algorithms help solve problems Become grounded in RL fundamentals including Markov decision processes, dynamic programming, and temporal difference learning Dive deep into a range of value and policy gradient methods Apply advanced RL solutions such as meta learning, hierarchical learning, multi-agent, and imitation learning

Understand cutting-edge deep RL algorithms including Rainbow, PPO, TD3, SAC, and more Get practical examples through the accompanying website

This book introduces multiagent planning under uncertainty as formalized by decentralized partially observable Markov decision processes (Dec-POMDPs). The intended audience is researchers and graduate students working in the fields of artificial intelligence related to sequential decision making: reinforcement learning, decision-theoretic planning for single agents, classical multiagent planning, decentralized control, and operations research.

This volume contains the technical papers presented in the four high-quality workshops associated with the European Conference on Service-Oriented and Cloud Computing, ESOCC 2014, held in Manchester, UK, in September 2014: 4th International Workshop on Adaptive Services for the Future Internet, WAS4FI 2014, 2nd International Workshop on Cloud for IoT, CLIoT 2014, 2nd International Workshop on Cloud Service Brokerage, CSB 2014, and Seamless Adaptive Multi-cloud Management of Service-based Applications, SeaCloudS Workshop. The 19 revised full papers and 3 short papers were carefully reviewed and selected from 39 submissions. They focus on specific topics in service-oriented and cloud computing domains as cloud computing, service buses, Web services, service-oriented architectures, event-driven architectures, enterprise architectures, business process management, software selection and adaptation.

Deep reinforcement learning is the combination of reinforcement learning (RL) and deep learning. This field of research has recently been able to solve a wide range of complex decision-making tasks that were previously out of reach for a machine. Deep RL opens up many new applications in domains such as healthcare, robotics, smart grids, finance, and many more. This book provides the reader with a starting point for understanding the topic. Although written at a research level it provides a comprehensive and accessible introduction to deep reinforcement learning models, algorithms and techniques. Particular focus is on the aspects related to generalization and how deep RL can be used for practical applications. Written by recognized experts, this book is an important introduction to Deep Reinforcement Learning for practitioners, researchers and students alike.

Reinforcement learning encompasses both a science of adaptive behavior of rational beings in uncertain environments and a computational methodology for finding optimal behaviors for challenging problems in control, optimization and adaptive behavior of intelligent agents. As a field, reinforcement learning has progressed tremendously in the past decade. The main goal of this book is to present an up-to-date series of survey articles on the main contemporary sub-fields of reinforcement learning. This includes surveys on partially observable environments, hierarchical task decompositions, relational knowledge representation and predictive state representations. Furthermore, topics such as transfer, evolutionary methods and continuous spaces in reinforcement learning are surveyed. In addition, several chapters review reinforcement learning methods in robotics, in games, and in computational neuroscience. In total seventeen different subfields are presented by mostly young experts in those areas, and together they truly represent a state-of-the-art of current reinforcement learning research. Marco Wiering works at the artificial intelligence department of the University of Groningen in the Netherlands. He has published extensively on various reinforcement learning topics. Martijn van Otterlo works in the cognitive artificial intelligence group at the Radboud University Nijmegen in The Netherlands. He has mainly focused on expressive knowledge representation in reinforcement learning settings.

Introducing a new approach to multiagent reinforcement learning and distributed artificial intelligence, this guide shows how classical game theory can be used to compose basic learning units. This approach to creating agents has the advantage of leading to powerful, yet intuitively simple, algorithms that can be analyzed. The setup is demonstrated here in a number of different settings, with a detailed analysis of agent learning behaviors provided for each. A review of required background materials from game theory and reinforcement learning is also provided, along with an overview of related multiagent learning methods.

Deep Learning and back-propagation have been successfully used to perform centralized training with communication protocols among multiple agents in a cooperative Multi-Agent Deep Reinforcement Learning (MARL) environment. In this work, I present techniques for centralized training of MARL agents in large scale environments and compare my work against current state of the art techniques. This work uses model-free Deep Q-Network (DQN) as the baseline model and allows inter agent communication for cooperative policy learning. I present two novel, scalable and centralized MARL training techniques (MA-MeSN, MA-BoN), which are developed under the principle that the behavior policy and message/communication policies have different optimization criteria. Thus, this work presents models which separate the message learning module from the behavior policy learning module. As shown in the experiments, the separation of these modules helps in faster convergence in complex domains like autonomous driving simulators and achieves better results than the current techniques in literature. Subsequently, this work presents two novel techniques for achieving decentralized execution for the communication based cooperative policy. The first technique uses behavior cloning as a method of cloning an expert cooperative policy to a decentralized agent without message sharing. In the second method, the behavior policy is coupled with a memory module which is local to each model. This memory model is used by the independent agents to mimic the communication policies of other agents and thus generate an independent behavior policy. This decentralized approach has minimal effect on degradation of the overall cumulative reward achieved by the centralized policy. Using a fully decentralized approach allows us to address the challenges of noise and communication bottlenecks in real-time communication channels. In this work, I theoretically and empirically compare the centralized and decentralized training algorithms to current research in the field of MARL. As part of this thesis, I also developed a large scale multi-agent testing environment. It is a new OpenAI-Gym environment which can be used for large scale multi-agent research as it simulates multiple autonomous cars driving cooperatively on a highway in the presence of a bad actor. I compare the performance of the centralized algorithms to existing state-of-the-art algorithms, for ex, DIAL and IMS which are based on cumulative reward achieved per episode and other metrics. MA-MeSN and MA-BoN achieve a cumulative reward of at-least 263% higher than the reward achieved by the DIAL and IMS. I also present an ablation study of the scalability of MA-BoN and show that MA-MeSN and MA-BoN algorithms only exhibit a linear increase in inference time and number of trainable parameters compared to quadratic increase for DIAL.

Decentralized Reinforcement Learning for the Online Optimization of Distributed Systems.

Move beyond the foundations of machine learning and game theory in cyber security to the latest research in this cutting-edge field In Game Theory and Machine Learning for Cyber Security, a team of expert security researchers delivers a collection of central research contributions from both machine learning and game theory applicable to cybersecurity. The distinguished editors have included resources that address open research questions in game theory and machine learning applied to cyber security systems and examine the strengths and limitations of current game theoretic models for cyber security. Readers will explore the vulnerabilities of traditional machine learning algorithms and how they can be mitigated in an adversarial machine learning approach. The book offers a comprehensive suite of solutions to a broad range of technical issues in applying game theory and machine learning to solve cyber security challenges. Beginning with an introduction to foundational concepts in game theory, machine learning, cyber security, and cyber deception, the editors provide readers with resources that discuss the latest in hypergames, behavioral game theory, adversarial machine learning, generative adversarial networks, and multi-agent reinforcement learning. Readers will also enjoy: A thorough introduction to game theory for cyber deception, including scalable algorithms for identifying stealthy attackers in a game theoretic framework, honeypot allocation over attack graphs, and behavioral games for cyber deception An exploration of game theory for cyber security, including actionable game-theoretic adversarial intervention detection against persistent and advanced threats Practical discussions of adversarial machine learning for cyber security, including adversarial machine learning in 5G security and machine learning-driven fault injection in cyber-physical systems In-depth examinations of generative models for cyber security Perfect for researchers, students, and experts in the fields of computer science and engineering, Game Theory and Machine Learning for Cyber Security is also an indispensable resource for industry professionals, military personnel, researchers, faculty, and students with an interest in cyber security.

This book constitutes the thoroughly refereed proceedings of the First International Conference on Machine Learning for Networking, MLN 2018, held in Paris, France, in November 2018. The 22 revised full papers included in the volume were carefully reviewed and selected from 48 submissions. They present new trends in the following topics: Deep and reinforcement learning; Pattern recognition and classification for networks; Machine learning for network slicing optimization, 5G system, user behavior prediction, multimedia, IoT, security and protection; Optimization and new innovative machine learning methods; Performance analysis of machine learning algorithms; Experimental evaluations of machine learning; Data mining in heterogeneous networks; Distributed and decentralized machine learning algorithms; Intelligent cloud-support communications, resource allocation, energy-aware/green communications, software defined networks, cooperative networks, positioning and navigation systems, wireless communications, wireless sensor networks, underwater sensor networks.

This book constitutes the thoroughly refereed proceedings of the Second International Conference on Machine Learning for Networking, MLN 2019, held in Paris, France, in December 2019. The 26 revised full papers included in the volume were carefully reviewed and selected from 75 submissions. They present and discuss new trends in deep and reinforcement learning, pattern recognition and classification for networks, machine learning for network slicing optimization, 5G system, user behavior prediction, multimedia, IoT, security and protection, optimization and new innovative machine learning methods, performance analysis of machine learning algorithms, experimental evaluations of machine learning, data mining in heterogeneous networks, distributed and decentralized machine learning algorithms, intelligent cloud-support communications, ressource allocation, energy-aware communications, software de ned networks, cooperative networks, positioning and navigation systems, wireless communications, wireless sensor networks, underwater sensor networks.

This three-volume set LNAI 8188, 8189 and 8190 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2013, held in Prague, Czech Republic, in September 2013. The 111 revised research papers presented together with 5 invited talks were carefully reviewed and selected from 447 submissions. The papers are organized in topical sections on reinforcement learning; Markov decision processes; active learning and optimization; learning from sequences; time series and spatio-temporal data; data streams; graphs and networks; social network analysis; natural language processing and information extraction; ranking and recommender systems; matrix and tensor analysis; structured output prediction, multi-label and multi-task learning; transfer learning; bayesian learning; graphical models; nearest-neighbor methods; ensembles; statistical learning; semi-supervised learning; unsupervised learning; subgroup discovery, outlier detection and anomaly detection; privacy and security; evaluation; applications; and medical applications.

This book summarizes several industry happenings in the blockchain world. While the blockchain and cryptocurrency industry has grown leaps and bounds, within the past several years including newer forms of products that use decentralized finance, non-fungible tokens, smart contracts, etc. Overall, the decentralized finance (Defi) markets have grown to include protocols such as MakerDAO, Uniswap, wrapped tokens such as synthetics, etc. However, Non Fungible tokens create a large marketplace, giving digital art a unique and powerful outlet.

Overall non-fungible tokens can be traded for other assets such as Ethereum, Bitcoin, etc. on public decentralized and centralized marketplaces. Overall, smart contracts provide a unique opportunity to enable different types of marketplaces that include financial assets and non-fungible tokens that include digital art and smart contract-based financial instruments.

Many complex systems found in nature can be viewed as function optimizers. In particular, they can be viewed as such optimizers of functions in extremely high dimensional spaces. Given the difficulty of performing such high-dimensional op timization with modern computers, there has been a lot of exploration of computa tional algorithms that try to emulate those naturally-occurring function optimizers. Examples include simulated annealing (SA [15,18]), genetic algorithms (GAs) and evolutionary computation [2,3,9,11,20-22,24,28]. The ultimate goal of this work is an algorithm that can, for any provided high-dimensional function, come close to extremizing that function. Particularly desirable would be such an algorithm that works in an adaptive and robust manner, without any explicit knowledge of the form of the function being optimized. In particular, such an algorithm could be used for distributed adaptive control---one of the most important tasks engineers will face in the future, when the systems they design will be massively distributed and horribly messy congeries ofcomputational systems.

The 15 chapters in this book explore the theoretical as well as a number of technical research outcomes on all aspects of UAVs. UAVs has widely differing applications such as disaster management, structural inspection, goods delivery, transportation, localization, mapping, pollution and radiation monitoring, search and rescue, farming, etc. The advantages of using UAVs are countless and have led the way for the full integration of UAVs, as intelligent objects into the IoT system. The book covers cover such subjects as: Efficient energy management systems in UAV based IoT networks IoE enabled UAVs Mind-controlled UAV using Brain-Computer Interface (BCI) The importance of AI in realizing autonomous and intelligent flying IoT Blockchain-based solutions for various security issues in UAV-enabled IoT The challenges and threats of UAVs such as hijacking, privacy, cyber-security, and physical safety.

In this carefully edited book some selected results of theoretical and applied research in the field of broadly perceived intelligent systems are presented. The problems vary from industrial to web and problem independent applications. All this is united under the slogan: "Intelligent systems conquer the world”. The book brings together innovation projects with analytical research, invention, retrieval and processing of knowledge and logical applications in technology. This book is aiming to a wide circle of readers and particularly to the young generation of IT/ICT experts who will build the next generations of intelligent systems.

Brains rule the world, and brain-like computation is increasingly used in computers and electronic devices. Brain-like computation is about processing and interpreting data or directly putting forward and performing actions. Learning is a very important aspect. This book is on reinforcement learning which involves performing actions to achieve a goal. The first 11 chapters of this book describe and extend the scope of reinforcement learning. The remaining 11 chapters show that there is already wide usage in numerous fields. Reinforcement learning can tackle control tasks that are too complex for traditional, hand-designed, non-learning controllers. As learning computers can deal with technical complexities, the tasks of human operators remain to specify goals on increasingly higher levels. This book shows that reinforcement learning is a very dynamic area in terms of theory and applications and it shall stimulate and encourage new research in this field.

Multi-agent systems consist of multiple agents that interact and coordinate with each other to work towards to certain goal. Multi-agent systems naturally arise in a variety of domains such as robotics, telecommunications, and economics. The dynamic and complex nature of these systems entails the agents to learn the optimal solutions on their own instead of following a pre-programmed strategy. Reinforcement learning provides a framework in which agents learn optimal behavior based on the response obtained from the environment. In this thesis, we propose various novel de- centralized, learning automaton based algorithms which can be employed by a group of interacting learning automata. We propose a completely decentralized version of the estimator algorithm. As compared to the completely centralized versions proposed before, this completely decentralized version proves to be a great improvement in terms of space complexity and convergence speed. The decentralized learning algorithm was applied; for the first time; to the domains of distributed object tracking and distributed watershed management. The results obtained by these experiments show the usefulness of the decentralized estimator algorithms to solve complex optimization problems. Taking inspiration from the completely decentralized learning algorithm, we propose the novel concept of partial decentralization. The partial decentralization bridges the gap between the completely decentralized and completely centralized algorithms and thus forms a comprehensive and continuous spectrum of multi-agent algorithms for the learning automata. To demonstrate the applicability of the partial decentralization, we employ a partially decentralized team of learning automata to control multi-agent Markov chains. More flexibility, expressiveness and flavor can be added to the partially decentralized framework by allowing different decentralized modules to engage in different types of games. We propose the novel framework of heterogeneous games of learning automata which allows the learning automata to engage in disparate games under the same formalism. We propose an algorithm to control the dynamic zero-sum games using heterogeneous games of learning automata.

Multiagent systems is an expanding field that blends classical fields like game theory and decentralized control with modern fields like computer science and machine learning. This monograph provides a concise introduction to the subject, covering the theoretical foundations as well as more recent developments in a coherent and readable manner. The text is centered on the concept of an agent as decision maker. Chapter 1 is a short introduction to the field of multiagent systems. Chapter 2 covers the basic theory of singleagent decision making under uncertainty. Chapter 3 is a brief introduction to game theory, explaining classical concepts like Nash equilibrium. Chapter 4 deals with the fundamental problem of coordinating a team of collaborative agents. Chapter 5 studies the problem of multiagent reasoning and decision making under partial observability. Chapter 6 focuses on the design of protocols that are stable against manipulations by self-interested agents. Chapter 7 provides a short introduction to the rapidly expanding field of multiagent reinforcement learning. The material can be used for teaching a half-semester course on multiagent systems covering, roughly, one chapter per lecture.

The book begins with a chapter on traditional methods ofsupervised learning, covering recursive least squares learning,mean square error methods, and stochastic approximation. Chapter 2covers single agent reinforcement learning. Topics include learningvalue functions, Markov games, and TD learning with eligibilitytraces. Chapter 3 discusses two player games including two playermatrix games with both pure and mixed strategies. Numerousalgorithms and examples are presented. Chapter 4 covers learning inmulti-player games, stochastic games, and Markov games, focusing onlearning multi-player grid games—two player grid games,Q-learning, and Nash Q-learning. Chapter 5 discusses differentialgames, including multi player differential games, actor critiquestructure, adaptive fuzzy control and fuzzy interference systems,the evader pursuit game, and the defending a territory games.Chapter 6 discusses new ideas on learning within robotic swarms andthe innovative idea of the evolution of personality traits. • Framework for understanding a variety of methods andapproaches in multi-agent machine learning. • Discusses methods of reinforcement learning such as anumber of forms of multi-agent Q-learning • Applicable to research professors and graduatestudents studying electrical and computer engineering, computerscience, and mechanical and aerospace

engineering
This book constitutes the refereed procedings of the 6th International Conference on Autonomic and Trusted Computing, ATC 2009, held in Brisbane, Australia, in July 2009, co-located with UIC 2009, the 6th International Conference on Ubiquitous Intelligence and Computing. The 17 revised full papers presented together with one invited paper and one keynote talk were carefully reviewed and selected from 52 submissions. The regular papers are organized in topical sections on organic and autonomic computing, trusted computing, wireless sensor networks, and trust.
This two-volume set constitutes the refereed proceedings of the workshops which complemented the 19th Joint European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD, held in Würzburg, Germany, in September 2019. The 70 full papers and 46 short papers presented in the two-volume set were carefully reviewed and selected from 200 submissions. The two volumes (CCIS 1167 and CCIS 1168) present the papers that have been accepted for the following workshops: Workshop on Automating Data Science, ADS 2019; Workshop on Advances in Interpretable Machine Learning and Artificial Intelligence and eXplainable Knowledge Discovery in Data Mining, AIMLAI-XKDD 2019; Workshop on Decentralized Machine Learning at the Edge, DMLE 2019; Workshop on Advances in Managing and Mining Large Evolving Graphs, LEG 2019; Workshop on Data and Machine Learning Advances with Multiple Views; Workshop on New Trends in Representation Learning with Knowledge Graphs; Workshop on Data Science for Social Good, SoGood 2019; Workshop on Knowledge Discovery and User Modelling for Smart Cities, UMCIT 2019; Workshop on Data Integration and Applications Workshop, DINA 2019; Workshop on Machine Learning for Cybersecurity, MLCS 2019; Workshop on Sports Analytics: Machine Learning and Data Mining for Sports Analytics, MLSA 2019; Workshop on Categorising Different Types of Online Harassment Languages in Social Media; Workshop on IoT Stream for Data Driven Predictive Maintenance, IoTStream 2019; Workshop on Machine Learning and Music, MML 2019; Workshop on Large-Scale Biomedical Semantic Indexing and Question Answering, BioASQ 2019.
How can we build machines that collaborate and learn more seamlessly with humans, and with each other? How do we create fairer societies? How do we minimize the impact of information manipulation campaigns, and fight back? How do we build machine learning algorithms that are more sample efficient when learning from each other's sparse data, and under time constraints? At the root of these questions is a simple one: how do agents, human or machines, learn from each other, and can we improve it and apply it to new domains? The cognitive and social sciences have provided innumerable insights into how people learn from data using both passive observation and experimental intervention. Similarly, the statistics and machine learning communities have formalized learning as a rigorous and testable computational process. There is a growing movement to apply insights from the cognitive and social sciences to improving machine learning, as well as opportunities to use machine learning as a sandbox to test, simulate and expand ideas from the cognitive and social sciences. A less researched and fertile part of this intersection is the modeling of social learning: past work has been more focused on how agents can learn from the 'environment', and there is less work that borrows from both communities to look into how agents learn from each other. This thesis presents novel contributions into the nature and usefulness of social learning as an inductive bias for reinforced learning. I start by presenting the results from two large-scale online human experiments: first, I observe Dunbar cognitive limits that shape and limit social learning in two different social trading platforms, with the additional contribution that synthetic financial bots that transcend human limitations can obtain higher profits even when using naive trading strategies. Second, I devise a novel online experiment to observe how people, at the individual level, update their belief of future financial asset prices (e.g. S&P 500 and Oil prices) from social information. I model such social

learning using Bayesian models of cognition, and observe that people make strong distributional assumptions on the social data they observe (e.g. assuming that the likelihood data is unimodal). I were fortunate to collect one round of predictions during the Brexit market instability, and find that social learning leads to higher performance than when learning from the underlying price history (the environment) during such volatile times. Having observed the cognitive limits and biases people exhibit when learning from other agents, I present an motivational example of the strength of inductive biases in reinforcement learning: I implement a learning model with a relational inductive bias that pre-processes the environment state into a set of relationships between entities in the world. I observe strong improvements in performance and sample efficiency, and even observe the learned relationships to be strongly interpretable. Finally, given that most modern deep reinforcement learning algorithms are distributed (in that they have separate learning agents), I investigate the hypothesis that viewing deep reinforcement learning as a social learning distributed search problem could lead to strong improvements. I do so by creating a fully decentralized, sparsely-communicating and scalable learning algorithm, and observe strong learning improvements with lower communication bandwidth usage (between learning agents) when using communication topologies that naturally evolved due to social learning in humans. Additionally, I provide a theoretical upper bound (that agrees with our empirical results) regarding which communication topologies lead to the largest learning performance improvement. Given a future increasingly filled with decentralized autonomous machine learning systems that interact with humans, there is an increasing need to understand social learning to build resilient, scalable and effective learning systems, and this thesis provides insights into how to build such systems.

The LNCS series reports state-of-the-art results in computer science research, development, and education, at a high level and in both printed and electronic form. Enjoying tight cooperation with the R&D community, with numerous individuals, as well as with prestigious organizations and societies, LNCS has grown into the most comprehensive computer science research forum available. The scope of LNCS, including its subseries LNAI and LNBI, spans the whole range of computer science and information technology including interdisciplinary topics in a variety of application fields. In parallel to the printed book, each new volume is published electronically in LNCS Online.

The book addresses some of the most recent issues, with the theoretical and methodological aspects, of evolutionary multi-objective optimization problems and the various design challenges using different hybrid intelligent approaches. Multi-objective optimization has been available for about two decades, and its application in real-world problems is continuously increasing. Furthermore, many applications function more effectively using a hybrid systems approach. The book presents hybrid techniques based on Artificial Neural Network, Fuzzy Sets, Automata Theory, other metaheuristic or classical algorithms, etc. The book examines various examples of algorithms in different real-world application domains as graph growing problem, speech synthesis, traveling salesman problem, scheduling problems, antenna design, genes design, modeling of chemical and biochemical processes etc.