

Artificial Intelligence And Le Robots Case Studies Of Successful Robot Systems

Answer Set Programming (ASP) is a declarative problem solving approach, initially tailored to modeling problems in the area of Knowledge Representation and Reasoning (KRR). More recently, its attractive combination of a rich yet simple modeling language with high-performance solving capacities has sparked interest in many other areas even beyond KRR. This book presents a practical introduction to ASP, aiming at using ASP languages and systems for solving application problems. Starting from the essential formal foundations, it introduces ASP's solving technology, modeling language and methodology, while illustrating the overall solving process by practical examples

Automated trading in electronic markets is one of the most common and consequential applications of autonomous software agents. Design of effective trading strategies requires thorough understanding of how market mechanisms operate, and appreciation of strategic issues that commonly manifest in trading scenarios. Drawing on research in auction theory and artificial intelligence, this book presents core principles of strategic reasoning that apply to market situations. The author illustrates trading strategy choices through examples of concrete market environments, such as eBay, as well as abstract market models defined by configurations of auctions and traders. Techniques for addressing these choices constitute essential building blocks for the design of trading strategies for rich market applications. The lecture assumes no prior background in game theory or auction theory, or artificial intelligence. Table of Contents: Introduction / Example: Bidding on eBay / Auction Fundamentals / Continuous Double Auctions / Interdependent Markets / Conclusion

Adversarial Machine Learning Morgan & Claypool Publishers

A guide to understanding the inner workings and outer limits of technology and why we should never assume that computers always get it right. In *Artificial Unintelligence*, Meredith Broussard argues that our collective enthusiasm for applying computer technology to every aspect of life has resulted in a tremendous amount of poorly designed systems. We are so eager to do everything digitally—hiring, driving, paying bills, even choosing romantic partners—that we have stopped demanding that our technology actually work. Broussard, a software developer and journalist, reminds us that there are fundamental limits to what we can (and should) do with technology. With this book, she offers a guide to understanding the inner workings and outer limits of technology—and issues a warning that we should never assume that computers always get things right. Making a case against technochauvinism—the belief that technology is always the solution—Broussard argues that it's just not true that social problems would inevitably retreat before a digitally enabled Utopia. To prove her point, she undertakes a series of adventures in computer programming. She goes for an alarming ride in a driverless car, concluding “the cyborg future is not coming any time soon”; uses artificial intelligence to investigate why students can't pass standardized tests; deploys machine learning to predict which passengers survived the Titanic disaster; and attempts to repair the U.S. campaign finance system by building AI software. If we understand the limits of what we can do with technology, Broussard tells us, we can make better choices about what we should do with it to make the world better for everyone.

Urban mobility is not only one of the pillars of modern economic systems, but also a key issue in the quest for equality of opportunity, once it can improve access to other services. Currently, however, there are a number of negative issues related to traffic, especially in mega-cities, such as economical issues (cost of opportunity caused by delays), environmental (externalities related to emissions of pollutants), and social (traffic accidents). Solutions to these issues are more and more closely tied to information and communication technology. Indeed, a search in the technical literature (using the keyword “urban traffic” to filter out articles on data network traffic) retrieved the following number of articles (as of December 3, 2013): 9,443 (ACM Digital Library), 26,054 (Scopus), and 1,730,000 (Google Scholar).

Moreover, articles listed in the ACM query relate to conferences as diverse as MobiCom, CHI, PADS, and AAMAS. This means that there is a big and diverse community of computer scientists and computer engineers who tackle research that is connected to the development of intelligent traffic and transportation systems. It is also possible to see that this community is growing, and that research projects are getting more and more interdisciplinary. To foster the cooperation among the involved communities, this book aims at giving a broad introduction into the basic but relevant concepts related to transportation systems, targeting researchers and practitioners from computer science and information technology. In addition, the second part of the book gives a panorama of some of the most exciting and newest technologies, originating in computer science and computer engineering, that are now being employed in projects related to car-to-car communication, interconnected vehicles, car navigation, platooning, crowd sensing and sensor networks, among others. This material will also be of interest to engineers and researchers from the traffic and transportation community.

Leveraging Artificial Intelligence in Global Epidemics provides readers with a detailed technical description of the role Artificial Intelligence plays in various stages of a disease outbreak, using COVID-19 as a case study. In the fight against epidemics, medical staff are on the front line; but behind the lines the battle is fought by researchers, and data scientists. Artificial Intelligence has been helping researchers with computer modeling and simulation for predictions about disease progression, the overall economic situation, tax incomes and population development. In the same manner, AI can prepare researchers for any emergency situation by backing the medical science. Artificial Intelligence plays a key and cutting-edge role in the preparedness for and dealing with the outbreak of global epidemics. It can help researchers analyze global data about known viruses to predict the patterns of the next pandemic and the impacts it will have. Not only prediction, AI plays an increasingly important role in assessing readiness, early detection, identification of patients, generating recommendations, situation awareness and more. It is up to the right input and the innovative ways by humans to leverage what AI can do. As COVID-19 has grabbed the world and its economy today, an analysis of the COVID-19 outbreak and the global responses and analytics will pay a long way in preparing humanity for such future situations. Provides readers with understanding of how Artificial Intelligence can be applied to the prediction, forecasting, detection, and testing of global epidemics, using COVID-19 and other recent epidemics such as Ebola, Corona viruses, Zika, influenza, Dengue, Chikungaya, and malaria as case studies Includes background material regarding readiness for coping with epidemics, including Machine Learning models for prediction of epidemic outbreaks based on existing data Includes technical coverage of key topics such as generating recommendations to combat outbreaks, genome sequencing, AI-assisted testing, AI-assisted contact tracing, situation awareness and combating disinformation, and the role of Artificial Intelligence and Machine Learning in drug discovery, vaccine development, and drug re-purposing This interdisciplinary and international handbook captures and shapes much needed reflection on normative frameworks for the production, application, and use of artificial intelligence in all spheres of individual, commercial, social, and public life.

Many machine learning algorithms require real-valued feature vectors of data instances as inputs. By projecting data into vector spaces, representation learning techniques have achieved promising performance in many areas such as computer vision and natural language processing. There is also a need to learn representations for discrete relational data, namely networks or graphs. Network Embedding (NE) aims at learning vector representations for each node or vertex in a network to encode the topologic structure. Due to its convincing performance and efficiency, NE has been widely applied in many network applications such as node classification and link prediction. This book provides a comprehensive introduction to the basic concepts, models, and applications of network representation learning (NRL). The book starts with an introduction to the background and rising of network embeddings as a general overview for readers. Then it introduces the development of NE techniques by

presenting several representative methods on general graphs, as well as a unified NE framework based on matrix factorization. Afterward, it presents the variants of NE with additional information: NE for graphs with node attributes/contents/labels; and the variants with different characteristics: NE for community-structured/large-scale/heterogeneous graphs. Further, the book introduces different applications of NE such as recommendation and information diffusion prediction. Finally, the book concludes the methods and applications and looks forward to the future directions.

Most subfields of computer science have an interface layer via which applications communicate with the infrastructure, and this is key to their success (e.g., the Internet in networking, the relational model in databases, etc.). So far this interface layer has been missing in AI. First-order logic and probabilistic graphical models each have some of the necessary features, but a viable interface layer requires combining both. Markov logic is a powerful new language that accomplishes this by attaching weights to first-order formulas and treating them as templates for features of Markov random fields. Most statistical models in wide use are special cases of Markov logic, and first-order logic is its infinite-weight limit. Inference algorithms for Markov logic combine ideas from satisfiability, Markov chain Monte Carlo, belief propagation, and resolution. Learning algorithms make use of conditional likelihood, convex optimization, and inductive logic programming. Markov logic has been successfully applied to problems in information extraction and integration, natural language processing, robot mapping, social networks, computational biology, and others, and is the basis of the open-source Alchemy system.

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

The International Conference on Intelligent Computing (ICIC) was formed to provide an annual forum dedicated to the emerging and challenging topics in artificial intelligence, machine learning, bioinformatics, and computational biology, etc. It aims to bring together researchers and practitioners from both academia and industry to share ideas, problems, and solutions related to the multifaceted aspects of intelligent computing. ICIC 2009, held in Ulsan, Korea, September 16-19, 2009, constituted the 5th International Conference on Intelligent Computing. It built upon the success of ICIC 2008, ICIC 2007, ICIC 2006, and ICIC 2005 held in Shanghai, Qingdao, Kunming, and Hefei, China, 2008, 2007, 2006, and 2005, respectively. This year, the conference concentrated mainly on the theories and methodologies as well as the emerging applications of intelligent computing. Its aim was to unify the picture of contemporary intelligent computing techniques as an integral concept that highlights the trends in advanced computational intelligence and bridges theoretical research with applications. Therefore, the theme for this conference was "Emerging Intelligent Computing Technology and Applications." Papers focusing on this theme were solicited, addressing theories, methodologies, and applications in science and technology.

Planning is the branch of Artificial Intelligence (AI) that seeks to automate reasoning about plans, most importantly the reasoning that goes into formulating a plan to achieve a given goal in a given situation. AI planning is model-based: a planning system takes as input a description (or model) of the initial situation, the actions available to change it, and the goal condition to output a plan composed of those actions that will accomplish the goal when executed from the initial situation. The Planning Domain Definition Language (PDDL) is a formal knowledge representation language designed to express planning models. Developed by the planning research community as a means of facilitating systems comparison, it has become a de-facto standard input language of many planning systems, although it is not the only modelling language for planning. Several variants of PDDL have emerged that capture planning problems of different natures and complexities, with a focus on deterministic problems. The purpose of this book is two-fold. First, we present a unified and current account of PDDL, covering the subsets of PDDL that express discrete, numeric, temporal, and hybrid planning. Second, we want to introduce readers to the art of modelling planning problems in this language, through educational examples that demonstrate how PDDL is used to model realistic planning problems. The book is intended for advanced students and researchers in AI who want to dive into the mechanics of AI planning, as well as those who want to be able to use AI planning systems without an in-depth explanation of the algorithms and implementation techniques they use.

The hidden costs of artificial intelligence, from natural resources and labor to privacy and freedom What happens when artificial intelligence saturates political life and depletes the planet? How is AI shaping our understanding of ourselves and our societies? In this book Kate Crawford reveals how this planetary network is fueling a shift toward undemocratic governance and increased inequality. Drawing on more than a decade of research, award-winning science, and technology, Crawford reveals how AI is a technology of extraction: from the energy and minerals needed to build and sustain its infrastructure, to the exploited workers behind "automated" services, to the data AI collects from us. Rather than taking a narrow focus on code and algorithms, Crawford offers us a political and a material perspective on what it takes to make artificial intelligence and where it goes wrong. While technical systems present a veneer of objectivity, they are always systems of power. This is an urgent account of what is at stake as technology companies use artificial intelligence to reshape the world.

The increasing abundance of large high-quality datasets, combined with significant technical advances over the last several decades have made machine learning into a major tool employed across a broad array of tasks including vision, language, finance, and security. However, success has been accompanied with important new challenges: many applications of machine learning are adversarial in nature. Some are adversarial because they are safety critical, such as autonomous driving. An adversary in these applications can be a malicious party aimed at causing congestion or accidents, or may even model unusual situations that expose vulnerabilities in the prediction engine. Other applications are adversarial because their task and/or the data they use are. For example, an important class of problems in security involves detection, such as malware, spam, and intrusion detection. The use of machine learning for detecting malicious entities creates an incentive among adversaries to evade detection by changing their behavior or the content of malicious objects they develop. The field of adversarial machine learning has emerged to study vulnerabilities of machine learning approaches in adversarial settings and to develop techniques to make learning robust to adversarial manipulation. This book provides a technical overview of this field. After reviewing machine learning concepts and approaches, as well as common use cases of these in adversarial settings, we present a general categorization of attacks on machine learning. We then address two major categories of attacks and associated defenses: decision-time attacks, in which an adversary changes the nature of instances seen by a learned model at the time of prediction in order to cause errors, and poisoning or training time attacks, in which the actual training dataset is maliciously modified. In our final chapter devoted to technical content, we discuss recent techniques for attacks on deep learning, as well as approaches for improving robustness of deep neural networks. We conclude with a discussion of several important issues in the area of adversarial learning that in our view warrant further research. Given the increasing interest in the area of adversarial machine learning, we hope this book provides readers with the tools necessary to successfully engage in research and practice of machine learning in adversarial settings.

Logic Programming is a style of programming in which programs take the form of sets of sentences in the language of Symbolic Logic. Over the years, there has been growing interest in Logic Programming due to applications in deductive databases, automated worksheets, Enterprise Management (business rules), Computational Law, and General Game Playing. This book introduces Logic Programming theory, current technology, and popular applications. In this volume, we take an innovative, model-theoretic approach to logic programming. We begin with the fundamental notion of datasets, i.e., sets of ground atoms. Given this fundamental notion, we introduce views, i.e., virtual relations; and we define classical logic programs as sets of view definitions, written using traditional Prolog-like notation but with semantics given in terms of datasets rather than implementation. We then introduce actions, i.e., additions and deletions of ground atoms; and we define dynamic logic programs as sets of action definitions. In addition to the printed book, there is an online version of the text with an interpreter and a compiler for the language used in the text and an integrated development environment for use in developing and deploying practical logic programs. "This is a book for the 21st century: presenting an elegant and innovative perspective on logic programming. Unlike other texts, it takes datasets as a fundamental notion, thereby bridging the gap between programming languages and knowledge representation languages; and it treats updates on an equal footing with datasets, leading to a sound and practical treatment of action and change." - Bob Kowalski, Professor Emeritus, Imperial College London "In a world where Deep Learning and Python are the talk of the day, this book is a remarkable development. It introduces the reader to the fundamentals of traditional Logic Programming and makes clear the benefits of using the technology to create runnable specifications for complex systems." - Son Cao Tran, Professor in Computer Science, New Mexico State University "Excellent introduction to the fundamentals of Logic Programming. The book is well-written and well-structured. Concepts are explained clearly and the gradually increasing complexity of exercises makes it so that one can understand easy notions quickly before moving on to more difficult ideas." - George Younger, student, Stanford University

Machine Learning for Healthcare: Handling and Managing Data provides in-depth information about handling and managing healthcare data through machine learning methods. This book expresses the long-standing challenges in healthcare informatics and provides rational explanations of how to deal with them. Machine Learning for Healthcare: Handling and Managing Data provides techniques on how to apply machine learning within your organization and evaluate the efficacy, suitability, and efficiency of machine learning applications. These are illustrated in a case study which examines how chronic disease is being redefined through patient-led data learning and the Internet of Things. This text offers a guided tour of machine learning algorithms, architecture design, and applications of learning in healthcare. Readers will discover the ethical implications of machine learning in healthcare and the future of machine learning in population and patient health optimization. This book can also help assist in the creation of a machine learning model, performance evaluation, and the operationalization of its outcomes within organizations. It may appeal to computer science/information technology professionals and researchers working in the area of machine learning, and is especially applicable to the healthcare sector. The features of this book include: A unique and complete focus on applications of machine learning in the healthcare sector. An examination of how data analysis can be done using healthcare data and bioinformatics. An investigation of how healthcare companies can leverage the tapestry of big data to discover new business values. An exploration of the concepts of machine learning, along with recent research developments in healthcare sectors.

As science continues to advance, researchers are continually gaining new insights into the way living beings behave and function, and into the composition of the smallest molecules. Most of these biological processes have been imitated by many scientific disciplines with the purpose of trying to solve different problems, one of which is artificial intelligence. Advancing Artificial Intelligence through Biological Process Applications presents recent advances in the study of certain biological processes related to information processing that are applied to artificial intelligence. Describing the benefits of recently discovered and existing techniques to adaptive artificial intelligence and biology, this book will be a highly valued addition to libraries in the neuroscience, molecular biology, and behavioral science spheres.

Solving challenging computational problems involving time has been a critical component in the development of artificial intelligence systems almost since the inception of the field. This book provides a concise introduction to the core computational elements of temporal reasoning for use in AI systems for planning and scheduling, as well as systems that extract temporal information from data. It presents a survey of temporal frameworks based on constraints, both qualitative and quantitative, as well as of major temporal consistency techniques. The book also introduces the reader to more recent extensions to the core model that allow AI systems to explicitly represent temporal preferences and temporal uncertainty. This book is intended for students and researchers interested in constraint-based temporal reasoning. It provides a self-contained guide to the different representations of time, as well as examples of recent applications of time in AI systems.

Life in the Caspian Republic has taught Agent Nikolai South two rules. Trust No One. And work just hard enough not to make enemies. Here, in the last sanctuary for the dying embers of the human race in a world run by artificial intelligence, if you stray from the path – your life is forfeit. But when a Party propagandist is killed – and is discovered as a “machine” – he’s given a new mission: chaperone the widow, Lily, who has arrived to claim her husband’s remains. But when South sees that she, the first “machine” ever allowed into the country, bears an uncanny resemblance to his late wife, he’s thrown into a maelstrom of betrayal, murder, and conspiracy that may bring down the Republic for good. WHEN THE SPARROW FALLS illuminates authoritarianism, complicity, and identity in the digital age, in a page turning, darkly-funny, frightening and touching story that recalls Philip K. Dick, John le Carré and Kurt Vonnegut in equal measure. At the Publisher's request, this title is being sold without Digital Rights Management Software (DRM) applied.

Human decision-making often transcends our formal models of "rationality." Designing intelligent agents that interact proficiently with people necessitates the modeling of human behavior and the prediction of their decisions. In this book, we explore the task of automatically predicting human decision-making and its use in designing intelligent human-aware automated computer systems of varying natures—from purely conflicting interaction settings (e.g., security and games) to fully cooperative interaction settings (e.g., autonomous driving and personal robotic assistants). We explore the techniques, algorithms, and empirical methodologies for meeting the challenges that arise from the above tasks and illustrate major benefits from the use of these computational solutions in real-world application domains such as security, negotiations, argumentative interactions, voting systems, autonomous driving, and games. The book presents both the traditional and classical methods as well as the most recent and cutting edge advances, providing the reader with a panorama of the challenges and solutions in predicting human decision-making.

Semi-supervised learning is a learning paradigm concerned with the study of how computers and natural systems such as humans learn in the presence of both labeled and unlabeled data. Traditionally, learning has been studied either in the unsupervised paradigm (e.g., clustering, outlier detection) where all the data are unlabeled, or in the supervised paradigm (e.g., classification, regression) where all the data are labeled. The goal of semi-supervised learning is to understand how combining labeled and unlabeled data may change the learning behavior, and design algorithms that take advantage of such a combination. Semi-supervised learning is of great interest in machine learning and data mining because it can use readily available unlabeled data to improve supervised learning tasks when the labeled data are scarce or expensive. Semi-supervised learning also shows potential as a quantitative tool to understand human category learning, where most of the input is self-evidently unlabeled. In this introductory book, we present some popular semi-supervised learning models, including self-training, mixture models, co-training and multiview learning, graph-based methods, and semi-supervised support vector machines. For each model, we discuss its basic mathematical formulation. The success of semi-supervised learning depends critically on some underlying assumptions. We emphasize the assumptions made by each model and give counterexamples when appropriate to demonstrate the limitations of the different models. In addition, we discuss semi-supervised learning for cognitive psychology. Finally, we give a computational learning theoretic perspective on semi-supervised learning, and we conclude the book with a brief discussion of open questions in the field. Table of Contents: Introduction to Statistical Machine Learning / Overview of Semi-Supervised Learning / Mixture Models and EM / Co-Training / Graph-Based Semi-Supervised Learning / Semi-Supervised Support Vector Machines / Human Semi-Supervised Learning / Theory and Outlook

This bestselling book gives business leaders and executives a foundational education on how to leverage artificial intelligence and machine learning solutions to deliver ROI for your business.

This book constitutes the refereed proceedings of the 17th Conference on Artificial Intelligence in Medicine, AIME 2019, held in Poznan, Poland, in June 2019. The 22 revised full and 31 short papers presented were carefully reviewed and selected from 134 submissions. The papers are organized in the following topical sections: deep learning; simulation; knowledge representation; probabilistic models; behavior monitoring; clustering, natural language processing, and decision support; feature selection; image processing; general machine learning; and unsupervised learning.

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

This book provides a comprehensive and self-contained introduction to Federated Learning, ranging from the basic knowledge and theories to various key applications, and the privacy and incentive factors are the focus of the whole book. This book is timely needed since Federated Learning is getting popular after the release of the General Data Protection Regulation (GDPR). As Federated Learning aims to enable a machine model to be collaboratively trained without each party exposing private data to others. This setting adheres to regulatory requirements of data privacy protection such as GDPR. This book contains three main parts. First, it introduces different privacy-preserving methods for protecting a Federated Learning model against different types of attacks such as Data Leakage and/or Data Poisoning. Second, the book presents incentive mechanisms which aim to encourage individuals to participate in the Federated Learning ecosystems. Last but not the least, this book also describes how Federated Learning can be applied in industry and business to address data silo and privacy-preserving problems. The book is intended for readers from both academia and industries, who would like to learn federated learning from scratch, practice its implementation, and apply it in their own business. Readers are expected to have some basic understanding of linear algebra, calculus, and neural network. Additionally, domain knowledge in FinTech and marketing are preferred.

Artificial systems that think and behave intelligently are one of the most exciting and challenging goals of Artificial Intelligence. Action Programming is the art and science of devising high-level control strategies for autonomous systems which employ a mental model of their environment and which reason about their actions as a means to achieve their goals. Applications of this programming paradigm

include autonomous software agents, mobile robots with high-level reasoning capabilities, and General Game Playing. These lecture notes give an in-depth introduction to the current state-of-the-art in action programming. The main topics are knowledge representation for actions, procedural action programming, planning, agent logic programs, and reactive, behavior-based agents. The only prerequisite for understanding the material in these lecture notes is some general programming experience and basic knowledge of classical first-order logic.

Cognition and artificial intelligence are entering a new era in which the aspects of symbolic manipulation and of connectionism begin to come together. This leads to a dialog of truly interdisciplinary character. The book covers aspects of fuzzy logic, case based reasoning, learning as well as meaning, language, and consciousness. The authors of this topical volume have their background in logic, computer science, physics and mathematics, philosophy, psychology and neurobiology.

The widespread adoption of AI and machine learning is revolutionizing many industries today. Once these technologies are combined with the programmatic availability of historical and real-time financial data, the financial industry will also change fundamentally. With this practical book, you'll learn how to use AI and machine learning to discover statistical inefficiencies in financial markets and exploit them through algorithmic trading. Author Yves Hilpisch shows practitioners, students, and academics in both finance and data science practical ways to apply machine learning and deep learning algorithms to finance. Thanks to lots of self-contained Python examples, you'll be able to replicate all results and figures presented in the book. In five parts, this guide helps you: Learn central notions and algorithms from AI, including recent breakthroughs on the way to artificial general intelligence (AGI) and superintelligence (SI) Understand why data-driven finance, AI, and machine learning will have a lasting impact on financial theory and practice Apply neural networks and reinforcement learning to discover statistical inefficiencies in financial markets Identify and exploit economic inefficiencies through backtesting and algorithmic trading--the automated execution of trading strategies Understand how AI will influence the competitive dynamics in the financial industry and what the potential emergence of a financial singularity might bring about

General game players are computer systems able to play strategy games based solely on formal game descriptions supplied at "runtime" (in other words, they don't know the rules until the game starts).

Unlike specialized game players, such as Deep Blue, general game players cannot rely on algorithms designed in advance for specific games; they must discover such algorithms themselves. General game playing expertise depends on intelligence on the part of the game player and not just intelligence of the programmer of the game player. GGP is an interesting application in its own right. It is intellectually engaging and more than a little fun. But it is much more than that. It provides a theoretical framework for modeling discrete dynamic systems and defining rationality in a way that takes into account problem representation and complexities like incompleteness of information and resource bounds. It has practical applications in areas where these features are important, e.g., in business and law. More fundamentally, it raises questions about the nature of intelligence and serves as a laboratory in which to evaluate competing approaches to artificial intelligence. This book is an elementary introduction to General Game Playing (GGP). (1) It presents the theory of General Game Playing and leading GGP technologies. (2) It shows how to create GGP programs capable of competing against other programs and humans. (3) It offers a glimpse of some of the real-world applications of General Game Playing. Table of Contents: Preface / Introduction / Game Description / Game Management / Game Playing / Small Single-Player Games / Small Multiple-Player Games / Heuristic Search / Probabilistic Search / Propositional Nets / General Game Playing With Propnets / Factoring / Discovery of Heuristics / Logic / Analyzing Games with Logic / Solving Single-Player Games with Logic / Discovering Heuristics with Logic / Games with Incomplete Information / Games with Historical Constraints / Incomplete Game Descriptions / Advanced General Game Playing / Authors' Biographies

The first edition of this popular textbook, *Contemporary Artificial Intelligence*, provided an accessible and student friendly introduction to AI. This fully revised and expanded update, *Artificial Intelligence: With an Introduction to Machine Learning, Second Edition*, retains the same accessibility and problem-solving approach, while providing new material and methods. The book is divided into five sections that focus on the most useful techniques that have emerged from AI. The first section of the book covers logic-based methods, while the second section focuses on probability-based methods. Emergent intelligence is featured in the third section and explores evolutionary computation and methods based on swarm intelligence. The newest section comes next and provides a detailed overview of neural networks and deep learning. The final section of the book focuses on natural language understanding. Suitable for undergraduate and beginning graduate students, this class-tested textbook provides students and other readers with key AI methods and algorithms for solving challenging problems involving systems that behave intelligently in specialized domains such as medical and software diagnostics, financial decision making, speech and text recognition, genetic analysis, and more.

This book constitutes the refereed proceedings of the 12th Biennial Conference of the Canadian Society for Computational Studies of Intelligence, AI'98, held in Vancouver, BC, Canada in June 1998. The 28 revised full papers presented together with 10 extended abstracts were carefully reviewed and selected from a total of more than twice as many submissions. The book is divided in topical sections on planning, constraints, search and databases; applications; genetic algorithms; learning and natural language; reasoning; uncertainty; and learning.

Lifelong Machine Learning, Second Edition is an introduction to an advanced machine learning paradigm that continuously learns by accumulating past knowledge that it then uses in future learning and problem solving. In contrast, the current dominant machine learning paradigm learns in isolation: given a training dataset, it runs a machine learning algorithm on the dataset to produce a model that is then used in its intended application. It makes no attempt to retain the learned knowledge and use it in subsequent learning. Unlike this isolated system, humans learn effectively with only a few examples precisely because our learning is very knowledge-driven: the knowledge learned in the past helps us learn new things with little data or effort. Lifelong learning aims to emulate this capability, because without it, an AI system cannot be considered truly intelligent. Research in lifelong learning has developed significantly in the relatively short time since the first edition of this book was published. The purpose of this second edition is to expand the definition of lifelong learning, update the content of several chapters, and add a new chapter about continual learning in deep neural networks—which has been actively researched over the past two or three years. A few chapters have also been reorganized to make each of them more coherent for the reader. Moreover, the authors want to propose a unified framework for the research area. Currently, there are several research topics in machine learning that are closely related to lifelong learning—most notably, multi-task learning, transfer learning, and meta-learning—because they also employ the idea of knowledge sharing and transfer. This book brings all these topics under one roof and discusses their similarities and differences. Its goal is to introduce this emerging machine learning paradigm and present a comprehensive survey and review of the important research results and latest ideas in the area. This book is thus suitable for students, researchers, and practitioners who are interested in machine learning, data mining, natural language processing, or pattern recognition. Lecturers can readily use the book for courses in any of these related fields.

From driving, flying, and swimming, to digging for unknown objects in space exploration, autonomous robots take on varied shapes and sizes. In part, autonomous robots are designed to perform tasks that are too dirty, dull, or dangerous for humans. With nontrivial autonomy and volition, they may soon claim their own place in human society. These robots will be our allies as we strive for understanding our natural and man-made environments and build positive synergies around us. Although we may never perfect replication of biological capabilities in robots, we must harness the inevitable emergence of robots that synchronizes with our own capacities to live, learn, and grow. This book is a snapshot of motivations and methodologies for our collective attempts to transform our lives and enable us to cohabit with

