

Artificial Intelligence Foundations Of Computational Agents

In this book, the author examines the ethical implications of Artificial Intelligence systems as they integrate and replace traditional social structures in new sociocognitive-technological environments. She discusses issues related to the integrity of researchers, technologists, and manufacturers as they design, construct, use, and manage artificially intelligent systems; formalisms for reasoning about moral decisions as part of the behavior of artificial autonomous systems such as agents and robots; and design methodologies for social agents based on societal, moral, and legal values. Throughout the book the author discusses related work, conscious of both classical, philosophical treatments of ethical issues and the implications in modern, algorithmic systems, and she combines regular references and footnotes with suggestions for further reading. This short overview is suitable for undergraduate students, in both technical and non-technical courses, and for interested and concerned researchers, practitioners, and citizens. Providing a unique approach to machine learning, this text contains fresh and intuitive, yet rigorous, descriptions of all fundamental concepts necessary to conduct research, build products, tinker, and play. By prioritizing geometric intuition, algorithmic thinking, and practical real world applications in disciplines including computer vision, natural language processing, economics, neuroscience, recommender systems, physics, and biology, this text provides readers with both a lucid understanding of foundational material as well as the practical tools needed to solve real-world problems. With in-depth Python and MATLAB/OCTAVE-based computational exercises and a complete treatment of cutting edge numerical optimization techniques, this is an essential resource for students and an ideal reference for researchers and practitioners working in machine learning, computer science, electrical engineering, signal processing, and numerical optimization.

A guide to computational thinking education, with a focus on artificial intelligence literacy and the integration of computing and physical objects. Computing has become an essential part of today's primary and secondary school curricula. In recent years, K-12 computer education has shifted from computer science itself to the broader perspective of computational thinking (CT), which is less about technology than a way of thinking and solving problems—"a fundamental skill for everyone, not just computer scientists," in the words of Jeanette Wing, author of a foundational article on CT. This volume introduces a variety of approaches to CT in K-12 education, offering a wide range of international perspectives that focus on artificial intelligence (AI) literacy and the integration of computing and physical objects. The book first offers an overview of CT and its importance in K-12 education, covering such topics as the rationale for teaching CT; programming as a general problem-solving skill; and the "phenomenon-based learning" approach. It then addresses the educational implications of the explosion in AI research, discussing, among other things, the importance of teaching children to be conscientious designers and consumers of AI. Finally, the book examines the increasing influence of physical devices in CT education, considering the learning opportunities offered by robotics. Contributors Harold Abelson, Cynthia Breazeal, Karen Brennan, Michael E. Caspersen, Christian Dindler, Daniella DiPaola, Nardie Fanchamps, Christina Gardner-McCune, Mark Guzdial, Kai Hakkarainen, Fredrik Heintz, Paul Hennissen, H. Ulrich Hoppe, Ole Sejer Iversen, Siu-Cheung Kong, Wai-Ying Kwok, Sven Manske, Jesus Moreno-León, Blakeley H. Payne, Sini Riikonen, Gregorio Robles, Marcos Román-González, Pirta Seitamaa-Hakkarainen, Ju-Ling Shih, Pasi Silander, Lou Slangen, Rachel Charlotte Smith, Marcus Specht, Florence R. Sullivan, David S. Touretzky *This book provides conceptual understanding of machine learning algorithms though supervised, unsupervised, and advanced learning techniques. The book consists of four parts: foundation, supervised learning, unsupervised learning, and advanced learning. The first part provides the fundamental materials, background, and simple machine learning algorithms, as the preparation for studying machine learning algorithms. The second and the third parts provide understanding of the supervised learning algorithms and the unsupervised learning algorithms as the core parts. The last part provides advanced machine learning algorithms: ensemble learning, semi-supervised learning, temporal learning, and reinforced learning. Provides comprehensive coverage of both learning algorithms: supervised and unsupervised learning; Outlines the computation paradigm for solving classification, regression, and clustering; Features essential techniques for building the a new generation of machine learning.*

Machine Learning, a vital and core area of artificial intelligence (AI), is propelling the AI field ever further and making it one of the most compelling areas of computer science research. This textbook offers a comprehensive and unbiased introduction to almost all aspects of machine learning, from the fundamentals to advanced topics. It consists of 16 chapters divided into three parts: Part 1 (Chapters 1-3) introduces the fundamentals of machine learning, including terminology, basic principles, evaluation, and linear models; Part 2 (Chapters 4-10) presents classic and commonly used machine learning methods, such as decision trees, neural networks, support vector machines, Bayesian classifiers, ensemble methods, clustering, dimension reduction and metric learning; Part 3 (Chapters 11-16) introduces some advanced topics, covering feature selection and sparse learning, computational learning theory, semi-supervised learning, probabilistic graphical models, rule learning, and reinforcement learning. Each chapter includes exercises and further reading, so that readers can explore areas of interest. The book can be used as an undergraduate or postgraduate textbook for computer science, computer engineering, electrical engineering, data science, and related majors. It is also a useful reference resource for researchers and practitioners of machine learning.

The Quest for Artificial Intelligence

Foundations of Knowledge Acquisition

Society Of Mind

Foundations, Algorithms, and Applications

Multiagent Systems

Responsible Artificial Intelligence

Artificial IntelligenceCambridge University Press

This is the first textbook dedicated to explaining how artificial intelligence (AI) techniques can be used in and for games. After introductory chapters that explain the background and key techniques in AI and games, the authors explain how to use AI to play games, to generate content for games and to model players. The book will be suitable for undergraduate and graduate courses in games, artificial intelligence, design, human-computer interaction, and computational intelligence, and also for self-study by industrial game developers and practitioners. The authors have developed a website (<http://www.gameaibook.org>) that complements the material covered in the book with up-to-date exercises, lecture slides and reading.

This book presents a variety of techniques designed to enhance and empower multi-disciplinary and multi-institutional machine learning research in healthcare informatics. It is intended to provide a unique compendium of current and emerging machine learning paradigms for healthcare informatics, reflecting the diversity, complexity, and depth and breadth of this multi-disciplinary area.

The field of computational intelligence has grown tremendously over that past five years, thanks to evolving soft computing and artificial intelligent methodologies, tools and techniques for envisaging the essence of intelligence embedded in real life observations. Consequently, scientists have been able to explain and understand real life processes and practices which previously often remain unexplored by virtue of their underlying imprecision, uncertainties and redundancies, and the unavailability of appropriate methods for describing the incompleteness and vagueness of information represented. With the advent of the field of computational intelligence, researchers are now able to explore and unearth the intelligence, otherwise insurmountable, embedded in the systems under consideration. Computational Intelligence is now not limited to only specific computational fields, it has made inroads in signal processing, smart manufacturing, predictive control, robot navigation, smart cities, and sensor design to name a few. Recent Trends in Computational Intelligence Enabled Research: Theoretical Foundations and Applications explores the use of this computational paradigm across a wide range of applied domains which handle meaningful information. Chapters investigate a broad spectrum of the applications of computational intelligence across different platforms and disciplines, expanding our knowledge base of various research initiatives in this direction. This volume aims to bring together researchers, engineers, developers and practitioners from academia and industry working in all major areas and interdisciplinary areas of computational intelligence, communication systems, computer networks, and soft computing. Provides insights into the theory, algorithms, implementation, and application of computational intelligence techniques Covers a wide range of applications of deep learning across various domains which are researching the applications of computational intelligence Investigates novel techniques and reviews the state-of-the-art in the areas of machine learning, computer vision, soft computing techniques

Paradigms of AI Programming is the first text to teach advanced Common Lisp techniques in the context of building major AI systems. By reconstructing authentic, complex AI programs using state-of-the-art Common Lisp, the book teaches students and professionals how to build and debug robust practical programs, while demonstrating superior programming style and important AI concepts. The author strongly emphasizes the practical performance issues involved in writing real working programs of significant size. Chapters on troubleshooting and efficiency are included, along with a discussion of the fundamentals of object-oriented programming and a description of the main CLOS functions. This volume is an excellent text for a course on AI programming, a useful supplement for general AI courses and an indispensable reference for the professional programmer.

A Pragmatic Approach

Statistical Relational Artificial Intelligence

Logic, Probability, and Computation

An Introduction

Foundations of Machine Learning, second edition

Machine Learning Refined

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

This book provides an introduction to the mathematical and algorithmic foundations of data science, including machine learning, high-dimensional geometry, and analysis of large networks. Topics include the counterintuitive nature of data in high dimensions, important linear algebraic techniques such as singular value decomposition, the theory of random walks and Markov chains, the fundamentals of and important algorithms for machine learning, algorithms and analysis for clustering, probabilistic models for large networks, representation learning including topic modelling and non-negative matrix factorization, wavelets and compressed sensing. Important probabilistic techniques are developed including the law of large numbers, tail inequalities, analysis of random projections, generalization guarantees in machine learning, and moment methods for analysis of phase transitions in large random graphs. Additionally, important structural and complexity measures are discussed such as matrix norms and VC-dimension. This book is suitable for both undergraduate and graduate courses in the design and analysis of algorithms for data.

Intelligent Decision Support Systems have the potential to transform human decision making by combining research in artificial intelligence, information technology, and systems engineering. The field of intelligent decision making is expanding rapidly due, in part, to advances in artificial intelligence and network-centric environments that can deliver the technology. Communication and coordination between dispersed systems can deliver just-in-time information, real-time processing, collaborative environments, and globally up-to-date information to a human decision maker. At the same time, artificial intelligence techniques have demonstrated that they have matured sufficiently to provide computational assistance to humans in practical applications. This book includes contributions from leading researchers in the field beginning with the foundations of human decision making and the complexity of the human cognitive system. Researchers contrast human and artificial intelligence, survey computational intelligence, present pragmatic systems, and discuss future trends. This book will be an invaluable resource to anyone interested in the current state of knowledge and key research gaps in the rapidly developing field of intelligent decision support.

A new edition of a graduate-level machine learning textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern topics in machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes several key aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first four chapters lay the theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct (PAC) learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel methods; boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata and languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise probability review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New material in the appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

Computational Intelligence: An Introduction, Second Edition offers an in-depth exploration into the adaptive mechanisms that enable intelligent behaviour in complex and changing environments. The main focus of this text is centred on the computational modelling of biological and natural intelligent systems, encompassing swarm intelligence, fuzzy systems, artificial neural networks, artificial immune systems and evolutionary computation. Engelbrecht provides readers with a wide knowledge of Computational Intelligence (CI) paradigms and algorithms; inviting readers to implement and problem solve real-world, complex problems within the CI development framework. This implementation framework will enable readers to tackle new problems without any difficulty through a single Java class as part of the CI library. Key features of this second edition include: A tutorial, hands-on based presentation of the material. State-of-the-art coverage of the most recent developments in computational intelligence with more elaborate discussions on intelligence and artificial intelligence (AI). New discussion of Darwinian evolution versus Lamarckian evolution, also including swarm robotics, hybrid systems and artificial immune systems. A section on how to perform empirical studies; topics including statistical analysis of stochastic algorithms, and an open source library of CI algorithms. Tables, illustrations, graphs, examples, assignments, Java code implementing the algorithms, and a complete CI implementation and experimental framework. Computational Intelligence: An Introduction, Second Edition is essential reading for third and fourth year undergraduate and postgraduate students studying CI. The first edition has been prescribed by a number of overseas universities and is thus a valuable teaching tool. In addition, it will also be a useful resource for researchers in Computational Intelligence and Artificial Intelligence, as well as engineers, statisticians, operational researchers, and bioinformaticians with an interest in applying AI or CI to solve problems in their domains. Check out <http://www.ci.cs.up.ac.za> for examples, assignments and Java code implementing the algorithms.

Computational Intelligence

Computational Logic and Human Thinking

Algorithmic Intelligence

From Computer to Brain

Foundations of Computational Agents

Concepts to Implementations

Real-world intelligence includes the ability to handle complex, uncertain, dynamic, multi-modal information in real time. In order to pursue the artificial realization of such "human" or "intelligent" information processing, a novel system of representing and interpreting knowledge must first be developed. This book collects the results of ten years of research at six laboratories, focusing on the theoretical and algorithmic foundations of the intelligence we find in the real world.

This outstanding collection is designed to address the fundamental issues and principles underlying the task of Artificial Intelligence.

The latest advances in Artificial Intelligence and (deep) Machine Learning in particular revealed a major drawback of modern intelligent systems, namely the inability to explain their decisions in a way that humans can easily understand. While eXplainable AI rapidly became an active area of research in response to this need for improved understandability and trustworthiness, the field of Knowledge Representation and Reasoning (KRR) has on the other hand a long-standing tradition in managing information in a symbolic, human-understandable form. This book provides the first comprehensive collection of research contributions on the role of knowledge graphs for eXplainable AI (KG4XAI), and the papers included here present academic and industrial research focused on the theory, methods and implementations of AI systems that use structured knowledge to generate reliable explanations. Introductory material on knowledge graphs is included for those readers with only a minimal background in the field, as well as specific chapters devoted to advanced methods, applications and case-studies that use knowledge graphs as a part of knowledge-based, explainable systems (KBX-systems). The final chapters explore current challenges and future research directions in the area of knowledge graphs for eXplainable AI. The book not only provides a scholarly, state-of-the-art overview of research in this subject area, but also fosters the hybrid combination of symbolic and subsymbolic AI methods, and will be of interest to all those working in the field.

This book presents recent research on computational intelligence (CI) algorithms in the field of sport. In the modern age, information technologies have greatly reduced the need for human effort in the carrying out of many daily tasks. These technologies have radically influenced the lives of humans, and the information society in general. Unfortunately, these advances have brought with them certain negative effects, including the encouragement of sedentary lifestyles and the attendant health problems such as obesity that these engender. Other modern maladies, chiefly cardiovascular disease, diabetes, and cancer, have also been on the increase. Today, sports are virtually the only activity that still connects modern humans to their original lifestyle, which was based on physical motion. This book tears familiarizing sports scientists with the foundations of computational intelligence, while at the same time presenting the problems that have arisen in the training domain to computer scientists. Lastly, the book proposes the use of an Artificial Sports Trainer designed to enhance the training of modern athletes who cannot afford the considerable expense of hiring a human personal trainer. This intelligent system can monitor performance and design and direct appropriate future training, thus promoting both healthy lifestyles and competitive success in athletes.

The central task of future-oriented computational linguistics is the development of cognitive machines which humans can freely speak to in their natural language. This will involve the development of a functional theory of language, an objective method of verification, and a wide range of practical applications. Natural communication requires not only verbal processing, but also non-verbal perception and action. Therefore, the content of this book is organized as a theory of language for the construction of talking robots with a focus on the mechanics of natural language communication in both the listener and the speaker.

Machine Learning Foundations

Artificial Intelligence with Python

Computational Intelligence for Machine Learning and Healthcare Informatics

Tree-Structure based Hybrid Computational Intelligence

Recent Trends in Computational Intelligence Enabled Research

Artificial Intelligence to Solve Pervasive Internet of Things Issues

Artificial Intelligence to Solve Pervasive Internet of Things Issues discusses standards and technologies and wide-ranging technology areas and their applications and challenges, including discussions on architectures, frameworks, applications, best practices, methods and techniques required for integrating AI to resolve IoT issues. Chapters also provide step-by-step measures, practices and solutions to tackle vital decision-making and practical issues affecting IoT technology, including autonomous devices and computerized systems. Such issues range from adopting, mitigating, maintaining, modernizing and protecting AI and IoT infrastructure components such as scalability, sustainability, latency, system decentralization and maintainability. The book enables readers to explore, discover and implement new solutions for integrating AI to solve IoT issues. Resolving these issues will help readers address many real-world applications in areas such as scientific research, healthcare, defense, aeronautics, engineering, social media, and many others. Discusses intelligent techniques for the implementation of Artificial Intelligence in Internet of Things Prepared for researchers and specialists who are interested in the use and integration of IoT and Artificial Intelligence technologies

This book uncovers stakes and possibilities offered by Computational Intelligence and Predictive Analytics to Medical Science. The main focus is on data technologies,classification, analysis and mining, information retrieval, and in the algorithms needed to elaborate the informations. A section with use cases and applications follows the two main parts of the book, respectively dedicated to the foundations and techniques of the discipline.

In this book the author argues that the basis of what we consider computer intelligence has algorithmic roots, and he presents this with a holistic view, showing examples and explaining approaches that encompass theoretical computer science and machine learning via engineered algorithmic solutions. The book will be of value to researchers, practitioners, and students. Research in computational intelligence is directed toward building thinking machines and improving our understanding of intelligence. As evident, the ultimate achievement in this field would be to mimic or exceed human cognitive capabilities including reasoning, recognition, creativity, emotions, understanding, learning and so on. In this book, the authors illustrate an hybrid computational intelligence framework and its applications for various problem solving tasks. Based on tree-structure based encoding and the specific function operators, the models can be flexibly constructed and evolved by using simple computational intelligence techniques. The main idea behind this model is the flexible neural tree, which is very adaptive, accurate and efficient. Based on the pre-defined instruction/operator sets, a flexible neural tree model can be created and evolved. This volume comprises of 6 chapters including an introductory chapter giving the fundamental definitions and the last Chapter provides some important research challenges. Academics, scientists as well as engineers engaged in research, development and application of computational intelligence techniques and data mining will find the comprehensive coverage of this book invaluable.

Computational Intelligence: Concepts to Implementations provides the most complete and practical coverage of computational intelligence tools and techniques to date. This book integrates various natural and engineering disciplines to establish Computational Intelligence. This is the first comprehensive textbook on the subject, supported with lots of practical examples. It asserts that computational intelligence rests on a foundation of evolutionary computation. This refreshing view has set the book apart from other books on computational intelligence. This book lays emphasis on practical applications and computational tools, which are very useful and important for further development of the computational intelligence field. Focusing on evolutionary computation, neural networks, and fuzzy logic, the authors have constructed an approach to thinking about and working with computational intelligence that has, in their extensive experience, proved highly effective. The book moves clearly and efficiently from concepts and paradigms to algorithms and implementation techniques by focusing, in the early chapters, on the specific con. It explores a number of key themes, including self-organization, complex adaptive systems, and emergent computation. It details the metrics and analytical tools needed to assess the performance of computational intelligence tools. The book concludes with a series of case studies that illustrate a wide range of successful applications. This book will appeal to professional and academic researchers in computational intelligence applications, tool development, and systems. Moves clearly and efficiently from concepts and paradigms to algorithms and implementation techniques by focusing, in the early chapters, on the specific concepts and paradigms that inform the authors' methodologies Explores a number of key themes, including self-organization, complex adaptive systems, and emergent computation Details the metrics and analytical tools needed to assess the performance of computational intelligence tools Concludes with a series of case studies that illustrate a wide range of successful applications Presents code examples in C and C++ Provides, at the end of each chapter, review questions and exercises suitable for graduate students, as well as researchers and practitioners engaged in self-study

Artificial Intelligence and Games

Foundations of Deep Learning

Understanding Machine Learning

Foundations of Real-world Intelligence

Theoretical Foundations and Applications

An authority on artificial intelligence introduces a theory that explores the workings of the human mind and the mysteries of thought

Provides an integrated introduction to artificial intelligence. Develops AI representation schemes and describes their uses for diverse applications, from autonomous robots to diagnostic assistants to infobots. DLC:

Artificial intelligence.

Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.

This book discusses issues relating to the application of AI and computational modelling in criminal proceedings from a European perspective. Part one provides a definition of the topics. Rather than focusing on policing or prevention of crime - largely tackled by recent literature - it explores ways in which AI can affect the investigation and adjudication of crime. There are two main areas of application: the first is evidence gathering, which is addressed in Part two. This section examines how traditional evidentiary law is affected by both new ways of investigation - based on automated processes (often using machine learning) - and new kinds of evidence, automatically generated by AI instruments. Drawing on the comprehensive case law of the European Court of Human Rights, it also presents reflections on the reliability and, ultimately, the admissibility of such evidence. Part three investigates the second application area: judicial decision-making, providing an unbiased review of the meaning, benefits, and possible long-term effects of 'predictive justice' in the criminal field. It highlights the prediction of both violent behaviour, or recidivism, and future court decisions, based on precedents. Touching on the foundations of common law and civil law traditions, the book offers insights into the usefulness of 'prediction' in criminal proceedings.

One of the most intriguing questions about the new computer technology that has appeared over the past few decades is whether we humans will ever be able to make computers learn. As is painfully obvious to even the most casual computer user, most current computers do not. Yet if we could devise learning techniques that enable computers to routinely improve their performance through experience, the impact would be enormous. The result would be an explosion of new computer applications that would suddenly become economically feasible (e. g. , personalized computer assistants that automatically tune themselves to the needs of individual users), and a dramatic improvement in the quality of current computer applications (e. g. , imagine an airline scheduling program that improves its scheduling method based on analyzing past delays). And while the potential economic impact of successful learning methods is sufficient reason to invest in research into machine learning, there is a second significant reason: studying machine learning helps us understand our own human learning abilities and disabilities, leading to the possibility of improved methods in education. While many open questions remain about the methods by which machines and humans might learn, significant progress has been made.

From Theory to Algorithms

Man-Machine Communication in Natural Language

Machine Learning

Case Studies in Common Lisp

Intelligent Decision Making: An AI-Based Approach

Towards an Algorithmic Foundation for Artificial Intelligence

Multiagent systems combine multiple autonomous entities, each having diverging interests or different information. This overview of the field offers a computer science perspective, but also draws on ideas from game theory, economics, operations research, logic, philosophy and linguistics. It will serve as a reference for researchers in each of these fields, and be used as a text for advanced undergraduate or graduate courses. The authors emphasize foundations to create a broad and rigorous treatment of their subject, with thorough presentations of distributed problem solving, game theory, multiagent communication and learning, social choice, mechanism design, auctions, cooperative game theory, and modal logics of knowledge and belief. For each topic, basic concepts are introduced, examples are given, proofs of key results are offered, and algorithmic considerations are examined. An appendix covers background material in probability theory, classical logic, Markov decision processes and mathematical programming.

Artificial Intelligence presents a practical guide to AI, including agents, machine learning and problem-solving simple and complex domains.

Neural networks and fuzzy systems are different approaches to introducing human-like reasoning into expert systems. This text is the first to combine the study of these two subjects, their basics and their use, along with symbolic AI methods to build comprehensive artificial intelligence systems. In a clear and accessible style, Kasabov describes rule-based and connectionist techniques and then their combinations, with fuzzy logic included, showing the application of the different techniques to a set of simple prototype problems, which makes comparisons possible. A particularly strong feature of the text is that it is filled with applications in engineering, business, and finance. AI problems that cover most of the application-oriented research in the field (pattern recognition, speech and image processing, classification, planning, optimization, prediction, control, decision making, and game simulations) are discussed and illustrated with concrete examples. Intended both as a text for advanced undergraduate and postgraduate students as well as a reference for researchers in the field of knowledge engineering, Foundations of Neural Networks, Fuzzy Systems, and Knowledge Engineering has chapters structured for various levels of teaching and includes original work by the author along with the classic material. Data sets for the examples in the book as well as an integrated software environment that can be used to solve the problems and do the exercises at the end of each chapter are available free through anonymous ftp.

In theory, there is no difference between theory and practice. But, in practice, there is. Jan L. A. van de Snepscheut The pow of academic ideas in the area of computational intelligence has penetrated industry with tremendous speed and persistence. Thousands of applications have proved the practical potential of fuzzy logic, neural networks, evolutionary com- tation, swarm intelligence, and intelligent agents even before their theoretical foundation is completely understood. And the popularity is rising. Some software vendors have pronounced the new machine learning gold rush to "Transfer Data into Gold". New buzzwords like "data mining", "genetic algorithms", and "swarm optimization" have enriched the top executives' vocabulary to make them look more "visionary" for the 21st century. The phrase "fuzzy math" became political jargon after being used by US President George W. Bush in one of the election debates in the campaign in 2000. Even process operators are discussing the perf- mance of neural networks with the same passion as the performance of the Dallas Cowboys. However, for most of the engineers and scientists introducing computational intelligence technologies into practice, looking at the growing number of new approaches, and understanding their theoretical principles and potential for value creation becomes a more and more dif?cult task.

Biology undergraduates, medical students and life-science graduate students often have limited mathematical skills. Similarly, physics, math and engineering students have little patience for the detailed facts that make up much of biological knowledge. Teaching computational neuroscience as an integrated discipline requires that both groups be brought forward onto common ground. This book does this by making ancillary material available in an appendix and providing basic explanations without becoming bogged down in unnecessary details. The book will be suitable for undergraduates and beginning graduate students taking a computational neuroscience course and also to anyone with an interest in the uses of the computer in modeling the nervous system.

Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges

Foundations of Neural Networks, Fuzzy Systems, and Knowledge Engineering

Algorithmic, Game-Theoretic, and Logical Foundations

Artificial Intelligence Literacy and Physical Computing

Applying Computational Intelligence

A Sourcebook

Recent decades have witnessed the emergence of artificial intelligence as a serious science and engineering discipline. This textbook, aimed at junior to senior undergraduate students and first-year graduate students, presents artificial intelligence (AI) using a coherent framework to study the design of intelligent computational agents. By showing how basic approaches fit into a multidimensional design space, readers can learn the fundamentals without losing sight of the bigger picture. The book balances theory and experiment, showing how to link them intimately together, and develops the science of AI together with its engineering applications. Although structured as a textbook, the book's straightforward, self-contained style will also appeal to a wide audience of professionals, researchers, and independent learners. AI is a rapidly developing field: this book encapsulates the latest results without being exhaustive and encyclopedic. The text is supported by an online learning environment, AIspace, <http://aispace.org>, so that students can experiment with the main AI algorithms plus problems, animations, lecture slides, and a knowledge representation system, Allog, for experimentation and problem solving.

Deep learning has significantly reshaped a variety of technologies, such as image processing, natural language processing, and audio processing. The excellent generalizability of deep learning is like a "cloud" to conventional complexity-based learning theory: the over-parameterization of deep learning makes almost all existing tools vacuous. This irreconciliation considerably undermines the confidence of deploying deep learning to security-critical areas, including autonomous vehicles and medical diagnosis, where small algorithmic mistakes can lead to fatal disasters. This book seeks to explaining the excellent generalizability, including generalization analysis via the size-independent complexity measures, the role of optimization in understanding the generalizability, and the relationship between generalizability and ethical/security issues. The efforts to understand the excellent generalizability are following two major paths: (1) developing size-independent complexity measures, which can evaluate the "effective" hypothesis complexity that can be learned, instead of the whole hypothesis space; and (2) modelling the learned hypothesis through stochastic gradient methods, the dominant optimizers in deep learning, via stochastic differential functions and the geometry of the associated loss functions. Related works discover that over-parameterization surprisingly bring many good properties to the loss functions. Rising concerns of deep learning are seen on the ethical and security issues, including privacy preservation and adversarial robustness. Related works also reveal an interplay between them and generalizability: a good generalizability usually means a good privacy-preserving ability; and more robust algorithms might have a worse generalizability. We expect readers can have a big picture of the current knowledge in deep learning theory, understand how the deep learning theory can guide new algorithm designing, and identify future research directions. Readers need knowledge of calculus, linear algebra, probability, statistics, and statistical learning theory.

The practical benefits of computational logic need not be limited to mathematics and computing. As this book shows, ordinary people in their everyday lives can profit from the recent advances that have been developed for artificial intelligence. The book draws upon related developments in various fields from philosophy to psychology and law. It pays special attention to the integration of logic with decision theory, and the use of logic to improve the clarity and coherence of communication in natural languages such as English. This book is essential reading for teachers and researchers who may be out of touch with the latest developments in computational logic. It will also be useful in any undergraduate course that teaches practical thinking, problem solving or communication skills. Its informal presentation makes the book accessible to readers from any background, but optional, more formal, chapters are also included for those who are more technically oriented.

Artificial intelligence (AI) is a field within computer science that is attempting to build enhanced intelligence into computer systems. This book traces the history of the subject, from the early dreams of eighteenth-century (and earlier) pioneers to the more successful work of today's AI engineers. AI is becoming more and more a part of everyone's life. The technology is already embedded in face-recognizing cameras, speech-recognition software, Internet search engines, and health-care robots, among other applications. The book's many diagrams and easy-to-understand descriptions of AI programs will help the casual reader gain an understanding of how these and other AI systems actually work. Its thorough (but unobtrusive) end-of-chapter notes containing citations to important source materials will be of great use to AI scholars and researchers. This book promises to be the definitive history of a field that has captivated the imaginations of scientists, philosophers, and writers for centuries.

An intelligent agent interacting with the real world will encounter individual people, courses, test results, drugs prescriptions, chairs, boxes, etc., and needs to reason about properties of these individuals and relations among them as well as cope with uncertainty.

Uncertainty has been studied in probability theory and graphical models, and relations have been studied in logic, in particular in the predicate calculus and its extensions. This book examines the foundations of combining logic and probability into what are called relational probabilistic models. It introduces representations, inference, and learning techniques for probability, logic, and their combinations. The book focuses on two representations in detail: Markov logic networks, a relational extension of undirected graphical models and weighted first-order predicate calculus formula, and Problog, a probabilistic extension of logic programs that can also be viewed as a Turing-complete relational extension of Bayesian networks.

Paradigms of Artificial Intelligence Programming

Computational Thinking Education in K-12

Foundations of Computational Linguistics

How to Be Artificially Intelligent

The Foundations of Artificial Intelligence

Supervised, Unsupervised, and Advanced 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.

Foundations of Data Science

A Logical Approach

Computational Intelligence and Predictive Analysis for Medical Science

How to Develop and Use AI in a Responsible Way

Artificial Intelligence, Computational Modelling and Criminal Proceedings

How to Create Value