
1 Introduction Mit Press

Knowledge and Mind

Assetization

Readings on Color: The philosophy of color

Annotation

Introduction to Deep Learning

Numbered Lives

Introduction to Embedded Systems, Second Edition

Deep Learning

Introduction To Algorithms

Neural Network Learning and Expert Systems

A Semantic Web Primer, third edition

An Introduction to the Event-Related Potential Technique, second edition

Rethinking Public Key Infrastructures and Digital Certificates

Introduction to Industrial Organization, second edition

AI and Humanity

Marine Hydrodynamics, 40th anniversary edition

Introduction to Static Analysis

Introduction to Computation and Programming Using Python, second edition

Deep Learning

Data Science

AI Ethics

Introduction to AI Robotics, second edition

A Gentle Introduction to Effective Computing in Quantitative Research

Reinforcement Learning, second edition

Introduction to Natural Language Processing

Cloud Computing

An Introduction to Natural Computation
Introduction to Machine Learning
Content
Living Books
Machine Learning
An Introduction to Computational Learning Theory
Artificial Unintelligence
Open Access
An Introduction to Genetic Algorithms
An Introduction to Statistical Genetic Data Analysis
Intellectual Property Strategy
The Soar Cognitive Architecture
Computational Thinking
Data Feminism

1 Introduction Mit Press

Downloaded from <ftp.wtvq.com> by guest

ISSAC LAYLAH

Knowledge and Mind MIT Press

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.

Assetization MIT Press

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

Readings on Color: The philosophy of color MIT Press

Why cloud computing represents a paradigm shift for business,

and how business users can best take advantage of cloud services. Most of the information available on cloud computing is either highly technical, with details that are irrelevant to non-technologists, or pure marketing hype, in which the cloud is simply a selling point. This book, however, explains the cloud from the user's viewpoint—the business user's in particular. Nayan Ruparelia explains what the cloud is, when to use it (and when not to), how to select a cloud service, how to integrate it with other technologies, and what the best practices are for using cloud computing. Cutting through the hype, Ruparelia cites the simple and basic definition of cloud computing from the National Institute of Science and Technology: a model enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources. Thus with cloud computing, businesses can harness information technology resources usually available only to large enterprises. And this, Ruparelia demonstrates, represents a paradigm shift for business. It will ease funding for startups, alter business plans, and allow big businesses greater agility. Ruparelia discusses the key issues for any organization considering cloud computing: service level agreements, business service delivery and consumption, finance, legal jurisdiction, security, and social responsibility. He introduces novel concepts made possible by cloud computing: cloud cells, or specialist clouds for specific uses; the personal cloud; the cloud of things; and cloud service exchanges. He examines use case patterns in terms of infrastructure and platform, software information, and business process; and he explains how to transition to a cloud service. Current and future users will find this book an indispensable guide to the cloud.

Annotation MIT Press

An accessible introduction to the artificial intelligence technology that enables computer vision, speech recognition, machine translation, and driverless cars. Deep learning is an artificial intelligence technology that enables computer vision, speech recognition in mobile phones, machine translation, AI games, driverless cars, and other applications. When we use consumer products from Google, Microsoft, Facebook, Apple, or Baidu, we are often interacting with a deep learning system. In this volume in the MIT Press Essential Knowledge series, computer scientist John Kelleher offers an accessible and concise but comprehensive introduction to the fundamental technology at the heart of the artificial intelligence revolution. Kelleher explains that deep learning enables data-driven decisions by identifying and extracting patterns from large datasets; its ability to learn from complex data makes deep learning ideally suited to take advantage of the rapid growth in big data and computational power. Kelleher also explains some of the basic concepts in deep learning, presents a history of advances in the field, and discusses the current state of the art. He describes the most important deep learning architectures, including autoencoders, recurrent neural networks, and long short-term networks, as well as such recent developments as Generative Adversarial Networks and capsule networks. He also provides a comprehensive (and comprehensible) introduction to the two fundamental algorithms in deep learning: gradient descent and backpropagation. Finally, Kelleher considers the future of deep learning—major trends, possible developments, and significant challenges.

Introduction to Deep Learning MIT Press

A new way of thinking about data science and data ethics that is informed by the ideas of intersectional feminism. Today, data science is a form of power. It has been used to expose injustice, improve health outcomes, and topple governments. But it has also been used to discriminate, police, and surveil. This potential for good, on the one hand, and harm, on the other, makes it essential to ask: Data science by whom? Data science for whom? Data science with whose interests in mind? The narratives around big data and data science are overwhelmingly white, male, and techno-heroic. In *Data Feminism*, Catherine D'Ignazio and Lauren Klein present a new way of thinking about data science and data ethics—one that is informed by intersectional feminist thought. Illustrating data feminism in action, D'Ignazio and Klein show how challenges to the male/female binary can help challenge other hierarchical (and empirically wrong) classification systems. They explain how, for example, an understanding of emotion can expand our ideas about effective data visualization, and how the concept of invisible labor can expose the significant human efforts required by our automated systems. And they show why the data never, ever “speak for themselves.” *Data Feminism* offers strategies for data scientists seeking to learn how feminism can help them work toward justice, and for feminists who want to focus their efforts on the growing field of data science. But *Data Feminism* is about much more than gender. It is about power, about who has it and who doesn't, and about how those differentials of power can be challenged and changed.

Numbered Lives MIT Press

Emphasizing issues of computational efficiency, Michael Kearns and Umesh Vazirani introduce a number of central topics in

computational learning theory for researchers and students in artificial intelligence, neural networks, theoretical computer science, and statistics. Emphasizing issues of computational efficiency, Michael Kearns and Umesh Vazirani introduce a number of central topics in computational learning theory for researchers and students in artificial intelligence, neural networks, theoretical computer science, and statistics. Computational learning theory is a new and rapidly expanding area of research that examines formal models of induction with the goals of discovering the common methods underlying efficient learning algorithms and identifying the computational impediments to learning. Each topic in the book has been chosen to elucidate a general principle, which is explored in a precise formal setting. Intuition has been emphasized in the presentation to make the material accessible to the nontheoretician while still providing precise arguments for the specialist. This balance is the result of new proofs of established theorems, and new presentations of the standard proofs. The topics covered include the motivation, definitions, and fundamental results, both positive and negative, for the widely studied L. G. Valiant model of Probably Approximately Correct Learning; Occam's Razor, which formalizes a relationship between learning and data compression; the Vapnik-Chervonenkis dimension; the equivalence of weak and strong learning; efficient learning in the presence of noise by the method of statistical queries; relationships between learning and cryptography, and the resulting computational limitations on efficient learning; reducibility between learning problems; and algorithms for learning finite automata from active experimentation.

Introduction to Embedded Systems, Second Edition MIT Press
A self-contained introduction to abstract interpretation-based static analysis, an essential resource for students, developers, and users. Static program analysis, or static analysis, aims to discover semantic properties of programs without running them. It plays an important role in all phases of development, including verification of specifications and programs, the synthesis of optimized code, and the refactoring and maintenance of software applications. This book offers a self-contained introduction to static analysis, covering the basics of both theoretical foundations and practical considerations in the use of static analysis tools. By offering a quick and comprehensive introduction for nonspecialists, the book fills a notable gap in the literature, which until now has consisted largely of scientific articles on advanced topics. The text covers the mathematical foundations of static analysis, including semantics, semantic abstraction, and computation of program invariants; more advanced notions and techniques, including techniques for enhancing the cost-accuracy balance of analysis and abstractions for advanced programming features and answering a wide range of semantic questions; and techniques for implementing and using static analysis tools. It begins with background information and an intuitive and informal introduction to the main static analysis principles and techniques. It then formalizes the scientific foundations of program analysis techniques, considers practical aspects of implementation, and presents more advanced applications. The book can be used as a textbook in advanced undergraduate and graduate courses in static analysis and program verification, and as a reference for users,

developers, and experts.

Deep Learning MIT Press

An introduction to annotation as a genre--a synthesis of reading, thinking, writing, and communication--and its significance in scholarship and everyday life. Annotation--the addition of a note to a text--is an everyday and social activity that provides information, shares commentary, sparks conversation, expresses power, and aids learning. It helps mediate the relationship between reading and writing. This volume in the MIT Press Essential Knowledge series offers an introduction to annotation and its literary, scholarly, civic, and everyday significance across historical and contemporary contexts. It approaches annotation as a genre--a synthesis of reading, thinking, writing, and communication--and offer examples of annotation that range from medieval rubrication and early book culture to data labeling and online reviews.

Introduction To Algorithms MIT Press

The definitive presentation of Soar, one AI's most enduring architectures, offering comprehensive descriptions of fundamental aspects and new components. In development for thirty years, Soar is a general cognitive architecture that integrates knowledge-intensive reasoning, reactive execution, hierarchical reasoning, planning, and learning from experience, with the goal of creating a general computational system that has the same cognitive abilities as humans. In contrast, most AI systems are designed to solve only one type of problem, such as playing chess, searching the Internet, or scheduling aircraft departures. Soar is both a software system for agent development and a theory of what computational structures are

necessary to support human-level agents. Over the years, both software system and theory have evolved. This book offers the definitive presentation of Soar from theoretical and practical perspectives, providing comprehensive descriptions of fundamental aspects and new components. The current version of Soar features major extensions, adding reinforcement learning, semantic memory, episodic memory, mental imagery, and an appraisal-based model of emotion. This book describes details of Soar's component memories and processes and offers demonstrations of individual components, components working in combination, and real-world applications. Beyond these functional considerations, the book also proposes requirements for general cognitive architectures and explicitly evaluates how well Soar meets those requirements.

Neural Network Learning and Expert Systems MIT Press

A concise introduction to the emerging field of data science, explaining its evolution, relation to machine learning, current uses, data infrastructure issues, and ethical challenges. The goal of data science is to improve decision making through the analysis of data. Today data science determines the ads we see online, the books and movies that are recommended to us online, which emails are filtered into our spam folders, and even how much we pay for health insurance. This volume in the MIT Press Essential Knowledge series offers a concise introduction to the emerging field of data science, explaining its evolution, current uses, data infrastructure issues, and ethical challenges. It has never been easier for organizations to gather, store, and process data. Use of data science is driven by the rise of big data and social media, the development of high-performance computing,

and the emergence of such powerful methods for data analysis and modeling as deep learning. Data science encompasses a set of principles, problem definitions, algorithms, and processes for extracting non-obvious and useful patterns from large datasets. It is closely related to the fields of data mining and machine learning, but broader in scope. This book offers a brief history of the field, introduces fundamental data concepts, and describes the stages in a data science project. It considers data infrastructure and the challenges posed by integrating data from multiple sources, introduces the basics of machine learning, and discusses how to link machine learning expertise with real-world problems. The book also reviews ethical and legal issues, developments in data regulation, and computational approaches to preserving privacy. Finally, it considers the future impact of data science and offers principles for success in data science projects.

[A Semantic Web Primer, third edition](#) MIT Press

A concise introduction to content and the content industry, from the early internet to the Instagram egg. From the time we roll out of bed to check overnight updates to our last posts, likes, and views of the previous day, we're consuming and producing content. But what does the term "content" even mean? When did it become ubiquitous? And at what cost? In this volume in the MIT Press Essential Knowledge series, Kate Eichhorn offers a concise introduction to content and the content industry, examining the far-reaching effects content has on culture, politics, and labor in a digital age. Eichhorn traces the evolution of our current understanding of content from the early internet to the current social mediaverse. The quintessential example of content, she

says, is the Instagram egg—an image that imparted no information or knowledge and circulated simply for the sake of circulation. Eichhorn explores what differentiates user-generated content from content produced by compensated (although often undercompensated) workers; examines how fields from art and literature to journalism and politics have weathered the rise of the content industry; and investigates the increasing importance of artists' "content capital"—the ability of artists, writers, and performers to produce content not about their work but about their status as artists.

An Introduction to the Event-Related Potential Technique, second edition MIT Press

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information

theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

[Rethinking Public Key Infrastructures and Digital Certificates](#) MIT Press

A concise introduction to the basics of open access, describing what it is (and isn't) and showing that it is easy, fast, inexpensive, legal, and beneficial. The Internet lets us share perfect copies of our work with a worldwide audience at virtually no cost. We take advantage of this revolutionary opportunity when we make our work "open access": digital, online, free of charge, and free of most copyright and licensing restrictions. Open access is made possible by the Internet and copyright-holder consent, and many authors, musicians, filmmakers, and other creators who depend on royalties are understandably unwilling to give their consent.

But for 350 years, scholars have written peer-reviewed journal articles for impact, not for money, and are free to consent to open access without losing revenue. In this concise introduction, Peter Suber tells us what open access is and isn't, how it benefits authors and readers of research, how we pay for it, how it avoids copyright problems, how it has moved from the periphery to the mainstream, and what its future may hold. Distilling a decade of Suber's influential writing and thinking about open access, this is the indispensable book on the subject for researchers, librarians, administrators, funders, publishers, and policy makers.

Introduction to Industrial Organization, second edition MIT Press
 A feminist media history of quantification, uncovering the stories behind the tools and technologies we use to count, measure, and weigh our lives and realities. Anglo-American culture has used media to measure and quantify lives for centuries. Historical journal entries map the details of everyday life, while death registers put numbers to life's endings. Today we count our daily steps with fitness trackers and quantify births and deaths with digitized data. How are these present-day methods for measuring ourselves similar to those used in the past? In this book, Jacqueline Wernimont presents a new media history of western quantification, uncovering the stories behind the tools and technologies we use to count, measure, and weigh our lives and realities. *Numbered Lives* is the first book of its kind, a feminist media history that maps connections not only between past and present-day "quantum media" but between media tracking and long-standing systemic inequalities. Wernimont explores the history of the pedometer, mortality statistics, and the census in England and the United States to illuminate the entanglement of

Anglo-American quantification with religious, imperial, and patriarchal paradigms. In Anglo-American culture, Wernimont argues, counting life and counting death are sides of the same coin—one that has always been used to render statistics of life and death more valuable to corporate and state organizations. *Numbered Lives* enumerates our shared media history, helping us understand our digital culture and inheritance.

AI and Humanity MIT Press

A practical guide to using modern software effectively in quantitative research in the social and natural sciences. This book offers a practical guide to the computational methods at the heart of most modern quantitative research. It will be essential reading for research assistants needing hands-on experience; students entering PhD programs in business, economics, and other social or natural sciences; and those seeking quantitative jobs in industry. No background in computer science is assumed; a learner need only have a computer with access to the Internet. Using the example as its principal pedagogical device, the book offers tried-and-true prototypes that illustrate many important computational tasks required in quantitative research. The best way to use the book is to read it at the computer keyboard and learn by doing. The book begins by introducing basic skills: how to use the operating system, how to organize data, and how to complete simple programming tasks. For its demonstrations, the book uses a UNIX-based operating system and a set of free software tools: the scripting language Python for programming tasks; the database management system SQLite; and the freely available R for statistical computing and graphics. The book goes on to describe particular tasks: analyzing data, implementing

commonly used numerical and simulation methods, and creating extensions to Python to reduce cycle time. Finally, the book describes the use of LaTeX, a document markup language and preparation system.

Marine Hydrodynamics, 40th anniversary edition MIT Press

An essential guide to designing, conducting, and analyzing event-related potential (ERP) experiments, completely updated for this edition. The event-related potential (ERP) technique, in which neural responses to specific events are extracted from the EEG, provides a powerful noninvasive tool for exploring the human brain. This volume describes practical methods for ERP research along with the underlying theoretical rationale. It offers researchers and students an essential guide to designing, conducting, and analyzing ERP experiments. This second edition has been completely updated, with additional material, new chapters, and more accessible explanations. Freely available supplementary material, including several online-only chapters, offer expanded or advanced treatment of selected topics. The first half of the book presents essential background information, describing the origins of ERPs, the nature of ERP components, and the design of ERP experiments. The second half of the book offers a detailed treatment of the main steps involved in conducting ERP experiments, covering such topics as recording the EEG, filtering the EEG and ERP waveforms, and quantifying amplitudes and latencies. Throughout, the emphasis is on rigorous experimental design and relatively simple analyses. New material in the second edition includes entire chapters devoted to components, artifacts, measuring amplitudes and latencies, and statistical analysis; updated coverage of recording technologies;

concrete examples of experimental design; and many more figures. Online chapters cover such topics as overlap, localization, writing and reviewing ERP papers, and setting up and running an ERP lab.

Introduction to Static Analysis MIT Press

A survey of computational methods for understanding, generating, and manipulating human language, which offers a synthesis of classical representations and algorithms with contemporary machine learning techniques. This textbook provides a technical perspective on natural language processing—methods for building computer software that understands, generates, and manipulates human language. It emphasizes contemporary data-driven approaches, focusing on techniques from supervised and unsupervised machine learning. The first section establishes a foundation in machine learning by building a set of tools that will be used throughout the book and applying them to word-based textual analysis. The second section introduces structured representations of language, including sequences, trees, and graphs. The third section explores different approaches to the representation and analysis of linguistic meaning, ranging from formal logic to neural word embeddings. The final section offers chapter-length treatments of three transformative applications of natural language processing: information extraction, machine translation, and text generation. End-of-chapter exercises include both paper-and-pencil analysis and software implementation. The text synthesizes and distills a broad and diverse research literature, linking contemporary machine learning techniques with the field's linguistic and computational foundations. It is suitable for use in advanced

undergraduate and graduate-level courses and as a reference for software engineers and data scientists. Readers should have a background in computer programming and college-level mathematics. After mastering the material presented, students will have the technical skill to build and analyze novel natural language processing systems and to understand the latest research in the field.

Introduction to Computation and Programming Using Python, second edition MIT Press

A comprehensive introduction to modern applied statistical genetic data analysis, accessible to those without a background in molecular biology or genetics. Human genetic research is now relevant beyond biology, epidemiology, and the medical sciences, with applications in such fields as psychology, psychiatry, statistics, demography, sociology, and economics. With advances in computing power, the availability of data, and new techniques, it is now possible to integrate large-scale molecular genetic information into research across a broad range of topics. This book offers the first comprehensive introduction to modern applied statistical genetic data analysis that covers theory, data preparation, and analysis of molecular genetic data, with hands-on computer exercises. It is accessible to students and researchers in any empirically oriented medical, biological, or social science discipline; a background in molecular biology or genetics is not required. The book first provides foundations for statistical genetic data analysis, including a survey of fundamental concepts, primers on statistics and human evolution, and an introduction to polygenic scores. It then covers the practicalities of working with genetic data, discussing such

topics as analytical challenges and data management. Finally, the book presents applications and advanced topics, including polygenic score and gene-environment interaction applications, Mendelian Randomization and instrumental variables, and ethical issues. The software and data used in the book are freely available and can be found on the book's website.

Deep Learning MIT Press

An introduction to the engineering principles of embedded systems, with a focus on modeling, design, and analysis of cyber-physical systems. The most visible use of computers and software is processing information for human consumption. The vast majority of computers in use, however, are much less visible. They run the engine, brakes, seatbelts, airbag, and audio system in your car. They digitally encode your voice and construct a radio signal to send it from your cell phone to a base station. They command robots on a factory floor, power generation in a power plant, processes in a chemical plant, and traffic lights in a city. These less visible computers are called

embedded systems, and the software they run is called embedded software. The principal challenges in designing and analyzing embedded systems stem from their interaction with physical processes. This book takes a cyber-physical approach to embedded systems, introducing the engineering concepts underlying embedded systems as a technology and as a subject of study. The focus is on modeling, design, and analysis of cyber-physical systems, which integrate computation, networking, and physical processes. The second edition offers two new chapters, several new exercises, and other improvements. The book can be used as a textbook at the advanced undergraduate or introductory graduate level and as a professional reference for practicing engineers and computer scientists. Readers should have some familiarity with machine structures, computer programming, basic discrete mathematics and algorithms, and signals and systems.

Data Science MIT Press

An extensively revised edition of a mathematically rigorous yet accessible introduction to algorithms.