Information And Randomness An Algorithmic Perspective Texts In Theoretical Computer Science An Eatcs Series

"Algorithmic information theory (AIT) is the result of putting Shannon's information theory and Turing's computability theory into a cocktail shaker and shaking vigorously", says G.J. Chaitin, one of the fathers of this theory of complexity and randomness, which is also known as Kolmogorov complexity. It is relevant for logic (new light is shed on Gödel's incompleteness results), physics (chaotic motion), biology (how likely is life to appear and evolve?), and metaphysics (how ordered is the universe?). This book, benefiting from the author's research and teaching experience in Algorithmic Information Theory (AIT), should help to make the detailed mathematical techniques of AIT accessible to a much wider audience.

Briefly, we review the basic elements of computability theory and prob ability theory that are required. Finally, in order to place the subject in the appropriate historical and conceptual context we trace the main roots of Kolmogorov complexity. This way the stage is set for Chapters 2 and 3, where we introduce the notion of optimal effective descriptions of objects. The length of such a description (or the number of bits of information in it) is its Kolmogorov complexity. We treat all aspects of the elementary mathematical theory of Kolmogorov complexity. This body of knowledge may be called algo rithmic complexity theory. The theory of Martin-Lof tests for random ness of finite objects and infinite sequences is inextricably intertwined with the theory of Kolmogorov complexity and is completely treated. We also investigate the statistical properties of finite strings with high Kolmogorov complexity. Both of these topics are eminently useful in the applications part of the book. We also investigate the recursion theoretic properties of Kolmogorov complexity (relations with Godel's incompleteness result), and the Kolmogorov complexity version of infor mation theory, which we may call "algorithmic information theory" or "absolute information theory. " The treatment of algorithmic probability theory in Chapter 4 presup poses Sections 1. 6, 1. 11. 2, and Chapter 3 (at least Sections 3. 1 through 3. 4).

Constructibility and complexity play central roles in recent research in computer science, mathematics and physics. For example, scientists are investigating the complexity of computer programs, constructive proofs in mathematics and the randomness of physical processes. But there are different approaches to the explication of these concepts. This volume presents important research on the state of this discussion, especially as it refers to quantum mechanics. This `foundational debate' in computer science, mathematics and physics was already fully developed in 1930 in the Vienna Circle. A special section is devoted to its real founder Hans Hahn, referring to his contribution to the history and philosophy of science. The documentation section presents articles on the early Philipp Frank and on the Vienna Circle in exile. Reviews cover important recent literature on logical empiricism and related topics.

This book presents the scientific outcome of a joint effort of the computer science departments of the universities of Berne, Fribourg and Neuchâtel. Within an initiative devoted to "Information and Knowledge", these research groups collaborated over several years on issues of logic, probability, inference, and deduction. The goal of this volume is to examine whether there is any common ground between the different approaches to the concept of information. The structure of this book could be represented by a circular model, with an innermost syntactical circle, comprising statistical and algorithmic approaches; a second, larger circle, the semantical one, in which "meaning" enters the stage; and finally an outermost circle, the pragmatic one, casting light on real-life logical reasoning. These articles are complemented by two philosophical contributions exploring the wide conceptual field as well as taking stock of the articles on the various formal theories of information.

Algorithmic Learning in a Random World describes recent theoretical and experimental developments in building computable approximations to Kolmogorov's algorithmic notion of randomness. Based on these approximations, a new set of machine learning algorithms have been developed that can be used to make predictions and to estimate their confidence and credibility in high-dimensional spaces under the usual assumption that the data are independent and identically distributed (assumption of randomness). Another aim of this unique monograph is to outline some limits of predictions: The approach based on algorithmic theory of randomness allows for the proof of impossibility of prediction in certain situations. The book describes how several important machine learning problems, such as density estimation in high-dimensional spaces, cannot be solved if the only assumption is randomness.

The research presented in Aspects of Kolmogorov Complexity addresses the fundamental standard of defining randomness as measured by a Martin-Lof level of randomness as found in random sequential binary strings. A classical study of statistics that addresses both a fundamental standard of statistics as well as an applied measure for statistical communication theory. The research points to compression levels in a random state that are greater than is found in current literature. A historical overview of the field of Kolmogorov Complexity and Algorithmic Information Theory, a subfield of Information Theory, is given as well as examples using a radix 3, radix 4, and radix 5 base numbers for both random and non-random sequential strings. The text also examines monochromatic and chromatic symbols and both theoretical and applied aspects of data compression as they relate to the transmission and storage of information. The appendix contains papers on the subject given at conferences and the references are current.

"This textbook is designed to accompany a one- or two-semester course for advanced undergraduates or beginning graduate students in computer science and applied mathematics. - It gives an excellent introduction to the probabilistic techniques and paradigms used in the development of probabilistic algorithms and analyses. - It assumes only an elementary background in discrete mathematics and gives a rigorous yet accessible treatment of the material, with numerous examples and applications."--Jacket.

Algorithmic information theory (AIT), or Kolmogorov complexity as it is known to mathematicians, can provide a useful tool for scientists to look at natural systems, however, some critical conceptual issues need to be understood and the advances already made collated and put in a form accessible to scientists. This book has been written in the hope that readers will be able to absorb the key ideas behind AIT so that they are in a better position to access the mathematical developments and to apply the ideas to their own areas of interest. The theoretical underpinning of AIT is outlined in the earlier chapters, while later chapters focus on the applications, drawing attention to the thermodynamic commonality between ordered physical systems such as the alignment of magnetic spins, the maintenance of a laser distant from equilibrium, and ordered living systems such as bacterial systems, an ecology, and an economy. Key Features Presents a mathematically complex subject in language accessible to scientists Provides rich insights into modelling far-from-equilibrium systems Emphasises applications across range of fields, including physics, biology and econophysics Empowers scientists to apply these mathematical tools to their own research

New and classical results in computational complexity, including interactive proofs, PCP, derandomization, and quantum computation. Ideal for graduate students.

Cryptography is one of the most active areas in current mathematics research and applications. This book focuses on cryptography along with two related areas: the study of probabilistic proof systems, and the theory of computational pseudorandomness. Following a common theme that explores the interplay between randomness and computation, the important notions in each field are covered, as well as novel ideas and insights.

Information and RandomnessAn Algorithmic PerspectiveSpringer Science & Business Media

This book is the first one that provides a solid bridge between algorithmic information theory and statistical mechanics. Algorithmic information theory (AIT) is a theory of program size and recently is also known as algorithmic randomness. AIT provides a framework for characterizing the notion of randomness for an individual object and for studying it closely and comprehensively. In this book, a statistical mechanical interpretation of AIT is introduced while explaining the basic notions and results of AIT to the reader who has an acquaintance with an elementary theory of computation. A simplification of the setting of AIT is the noiseless source coding in information theory. First, in the book, a statistical mechanical interpretation of the noiseless source coding scheme is introduced. It can be seen that the notions in statistical mechanics such as entropy, temperature, and thermal equilibrium are translated into the context of noiseless source coding in a natural manner. Then, the framework of AIT is introduced. On this basis, the introduction of a statistical mechanical interpretation of AIT is begun. Namely, the notion of thermodynamic quantities, such as free energy, energy, and entropy, is introduced into AIT. In the interpretation, the temperature is shown to be equal to the partial randomness of the values of all these thermodynamic quantities, where the notion of partial randomness is a stronger representation of the compression rate measured by means of program-size complexity. Additionally, it is demonstrated that this situation holds for the temperature itself as a thermodynamic quantity. That is, for each of all the thermodynamic quantities above, the computability of its value at temperature T gives a sufficient condition for T to be a fixed point on partial randomness. In this groundbreaking book, the current status of the interpretation from both mathematical and physical points of view is reported. For example, a total statistical mechanical interpretation of AIT that actualizes a perfect correspondence to normal statistical mechanics can be developed by identifying a microcanonical ensemble in the framework of AIT. As a result, the statistical mechanical meaning of the thermodynamic quantities of AIT is clarified. In the book, the close relationship of the interpretation to Landauer's principle is pointed out.

Recent findings in the computer sciences, discrete mathematics, formal logics and metamathematics have opened up a royal road for the investigation of undecidability and randomness in physics. A translation of these formal concepts yields a fresh look into diverse features of physical modelling such as quantum complementarity and the measurement problem, but also stipulates questions related to the necessity of the assumption of continua.Conversely, any computer may be perceived as a physical system: not only in the immediate sense of the physical properties of its hardware. Computers are a medium to virtual realities. The foreseeable importance of such virtual realities stimulates the investigation of an ?inner description?, a ?virtual physics? of these universes of computation. Indeed, one may consider our own universe as just one particular realisation of an enormous number of virtual realities, most of them awaiting discovery.One motive of this book is the recognition that what is often referred to as ?randomness? in physics might actually be a signature of undecidability for systems whose evolution is computable on a step-by-step basis. To give a flavour of the type of questions envisaged: Consider an arbitrary algorithmic system which is computable on a step-by-step basis. Then it is in general impossible to specify a second algorithmic procedure, including itself, which, by experimental input-output analysis, is capable of finding the deterministic law of the first system. But even if such a law is specified beforehand, it is in general impossible to predict the system behaviour in the ?distant future?. In other words: no ?speedup? or ?computational shortcut? is available. In this approach, classical paradoxes can be formally translated into no-go theorems concerning intrinsic physical perception.It is suggested that complementarity can be modelled by experiments on finite automata, where measurements of one observable of the automaton destroys the possibility to measure another observable of the same automaton and it vice versa.Besides undecidability, a great part of the book is dedicated to a formal definition of randomness and entropy measures based on algorithmic information theory.

From the bestselling author of the acclaimed Chaos and Genius comes a thoughtful and provocative exploration of the big ideas of the modern era: Information, communication, and information theory. Acclaimed science writer James Gleick presents an eye-opening vision of how our relationship to information has transformed the very nature of human consciousness. A fascinating intellectual journey through the history of communication and information, from the language of Africa’s talking drums to the invention of written alphabets; from the electronic transmission of code to the origins of information theory, into the new information age and the current deluge of news, tweets, images, and blogs. Along the way, Gleick profiles key innovators, including Charles Babbage, Ada Lovelace, Samuel Morse, and Claude Shannon, and reveals how our understanding of information is transforming not only how we look at the world, but how we live. A New York Times Notable Book A Los Angeles Times and Cleveland Plain Dealer Best Book of the Year Winner of the PEN/E. O. Wilson Literary Science Writing Award

Chaitin, the inventor of algorithmic information theory, presents in this book the strongest possible version of Gödel's incompleteness theorem, using an information theoretic approach based on the size of computer programs. One half of the book is concerned with studying the halting probability of a universal computer if its program is chosen by tossing a coin. The other half is concerned with encoding the halting probability as an algebraic equation in integers, a so-called exponential diophantine equation.

To truly understand how the Internet and Web are organized and function requires knowledge of mathematics and computation theory. Mathematical and Algorithmic Foundations of the Internet introduces the concepts and methods upon which computer networks rely and explores their applications to the Internet and Web. The book offers a unique approach to mathematical and algorithmic concepts, demonstrating their universality by presenting ideas and examples from various fields, including literature, history, and art. Progressing from fundamental concepts to more specific topics and applications, the text covers computational complexity and randomness, networks and graphs, parallel and distributed computing, and search engines. While the mathematical treatment is rigorous, it is presented at a level that can be grasped by readers with an elementary mathematical background. The authors also present a lighter side to this complex subject by illustrating how many of the mathematical concepts have counterparts in everyday life. The book provides in-depth coverage of the mathematical prerequisites and assembles a complete presentation of how computer networks function. It is a useful resource for anyone interested in the inner functioning, design, and organization of the Internet.

This book contains in easily accessible form all the main ideas of the creator and principal architect of algorithmic information theory. This expanded second edition has added thirteen abstracts, a 1988 Scientific American Article, a transcript of a EUROPALIA 89 lecture, an essay on biology, and an extensive bibliography. Its new larger format makes it easier to read. Chaitin's ideas are a fundamental extension of those of G”del and Turning and have exploded some basic assumptions of mathematics and thrown new light on the scientific method, epistemology, probability theory, and of course computer science and information theory.

Algorithmic learning theory is mathematics about computer programs which learn from experience. This involves considerable interaction between various mathematical disciplines including theory of computation, statistics, and c- binatorics. There is also considerable interaction with the practical, empirical ?elds of machine and statistical learning in which a principal aim is to predict, from past data about phenomena, useful features of future data from the same phenomena. The papers in this volume cover a broad range of topics of current research in the ?eld of algorithmic learning theory. We have divided the 29 technical, contributed papers in this volume into eight categories (corresponding to eight sessions) re?ecting this broad range. The categories featured are Inductive Inf- ence, Approximate Optimization Algorithms, Online Sequence Prediction, S- tistical Analysis of Unlabeled Data, PAC Learning & Boosting, Statistical - pervisedLearning,LogicBasedLearning,andQuery&ReinforcementLearning. Below we give a brief overview of the ?eld, placing each of these topics in the general context of the ?eld. Formal models of automated learning re?ect various facets of the wide range of activities that can be viewed as learning. A ?rst dichotomy is between viewing learning as an inde?nite process and viewing it as a ?nite activity with a de?ned termination. Inductive Inference models focus on inde?nite learning processes, requiring only eventual success of the learner to converge to a satisfactory conclusion.

This book brings together historical notes, reviews of research developments, fresh ideas on how to make VC (Vapnik–Chervonenkis) guarantees tighter, and new technical contributions in the areas of machine learning, statistical inference, classification, algorithmic statistics, and pattern recognition. The contributors are leading scientists in domains such as statistics, mathematics, and theoretical computer science, and the book will be of interest to researchers and graduate students in these domains.

This Festschrift volume has been published in honor of Cristian Calude on the occasion of his 60th birthday and contains contributions from invited speakers and regular papers presented at the International Workshop on Theoretical Computer Science, WTCS 2012, held in Auckland, New Zealand, in February 2012. Cristian Calude has made a significant contribution to research in computer science theory. Along with early work by Chaitin, Ku?era, Kurtz, Solovay, and Terwijn his papers published in the mid-1990s jointly with Khoussainov, Hertling, and Wang laid the foundation for the development of modern theory of algorithmic randomness. His work was essential for establishing the leading role of New Zealand in this area. The research interests of Cristian Calude are reflected in the topics covered by the 32 papers included in this book, namely: algorithmic information theory, algorithms, automata and formal languages, computing and natural sciences, computability and applications, logic and applications, philosophy of computation, physics and computation, and unconventional models of computation. They have been organized into four parts. The first part consists of papers discussing his life achievements. This is followed by papers in the three general areas of complexity, computability, and randomness; physics, philosophy (and logic), and computation; and algorithms, automata, and formal models (including unconventional computing).

Computability and complexity theory are two central areas of research in theoretical computer science. This book provides a systematic, technical development of "algorithmic randomness" and complexity for scientists from diverse fields.

The last two decades have seen a wave of exciting new developments in the theory of algorithmic randomness and its applications to other areas of mathematics. This volume surveys much of the recent work that has not been included in published volumes until now. It contains a range of articles on algorithmic randomness and its interactions with closely related topics such as computability theory and computational complexity, as well as wider applications in areas of mathematics including analysis, probability, and ergodic theory. In addition to being an indispensable reference for researchers in algorithmic randomness, the unified view of the theory presented here makes this an excellent entry point for graduate students and other newcomers to the field.

The interplay between computability and randomness has been an active area of research in recent years, reflected by ample funding in the USA, numerous workshops, and publications on the subject. The complexity and the randomness aspect of a set of natural numbers are closely related. Traditionally, computability theory is concerned with the complexity aspect. However, computability theoretic tools can also be used to introduce mathematical counterparts for the intuitive notion of randomness of a set. Recent research shows that, conversely, concepts and methods originating from randomness enrich computability theory. The book covers topics such as lowness and highness properties, Kolmogorov complexity, betting strategies and higher computability. Both the basics and recent research results are desribed, providing a very readable introduction to the exciting interface of computability and randomness for graduates and researchers in computability theory, theoretical computer science, and measure theory.

Exploring the theory of computability and complexity over the real numbers, this book provides a systematic, technical development of "algorithmic randomness" and complexity for scientists working in diverse fields concerned with theoretical computer science.

Uses of Randomness in Algorithms and Protocolsmakes fundamental contributions to two different fields of complexity theory: computational number theory and cryptography. The most famous result is Goldwasser and Kilian's invention of a new approach to distinguish prime numbers from composites, using methods from the theory of elliptic curves over finite fields. The Goldwasser-Kilian algorithm is the first to yield a polynomial size proof of its assertions, ensuring correctness while still provably running fast on most inputs. This new primality test implies for the first time and without any assumptions that large certified primes can be generated in expected polynomial time under a distribution that is close to uniform. It provides a provocative new link between algebraic geometry and primality testing, one of the most ancient algorithmic problems in number theory. Heuristic implementations of the algorithm are currently considered to be the fastest existing methods to certify primes. Kilian also provides two elegant and original contributions to theoretical cryptography. He shows how to base general two-party protocols on a simple protocol, known as "oblivious transfer," proving the first completeness result of this kind. He also introduces a generalization of interactive proof systems, known as "multi-prover interactive proof systems," and shows that anything provable in this model is provable in zero knowledge. Joe Kilian is a National Science Foundation Postdoctoral Fellow at MIT and Harvard. Contents: Introduction. New Techniques in Primality Testing. Committing Bits Using Oblivious Transfer. Circuit Evaluation Using Oblivious Transfer: The NC1 Circuit Base. Oblivious Evaluation of Arbitrary Circuits. Interactive Proof Systems with Multiple Provers.

The papers gathered in this book were published over a period of more than twenty years in widely scattered journals. They led to the discovery of randomness in arithmetic which was presented in the recently published monograph on “Algorithmic Information Theory” by the author. There the strongest possible version of Gödel's incompleteness theorem, using an information-theoretic approach based on the size of computer programs, was discussed. The present book is intended as a companion volume to the monograph and it will serve as a stimulus for work on complexity, randomness and unpredictability, in physics and biology as well as in metamathematics.

Contents:Introductory/Tutorial/Survey PapersApplications to MetamathematicsApplications to BiologyTechnical Papers on Self-Delimiting ProgramsTechnical Papers on Blank-Endmarker ProgramsTechnical Papers on Turing Machines Readership: Computer scientists, mathematicians, physicists and philosophers. Keywords:Omega;Randomness;Godel Incompleteness;Algorithmic Information Theory;Program-Size Complexity;Kolmogorov ComplexityReview: “Many of Chaitin's results are discussed in a delightful collection of his published articles.” JosephFord American Scientist, 1989 “Chaitin advances the cause of truths whose time have come; he is preparing a roadmap to ease our voyage into a truly uncertain future. Those who embark on this great adventure will most assuredly find sustenance in the books reviewed here.” JosephFord Foundations of Physics, 1989

This volume presents some exciting new developments occurring on the interface between set theory and computability as well as their applications in algebra, analysis and topology. These include effective versions of Borel equivalence, Borel reducibility and Borel determinacy. It also covers algorithmic randomness and dimension, Ramsey sets and Ramsey spaces. Many of these topics are being discussed in the NSF-supported annual Southeastern Logic Symposium.

This review volume consists of an indispensable set of chapters written by leading scholars, scientists and researchers in the field of Randomness, including related subfields specially but not limited to the strong developed connections to the Computability and Recursion Theory. Highly respected, indeed renowned in their areas of specialization, many of these contributors are the founders of their fields. The scope of Randomness Through Computation is novel. Each contributor shares his personal views and anecdotes on the various reasons and motivations which led him to the study of the subject. They share their visions from their vantage and distinctive viewpoints. In summary, this is an opportunity to learn about the topic and its various angles from the leading thinkers.

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

In this fascinating book, mathematician Ed Beltrami takes a close enough look at randomness to make it mysteriously disappear. The results of coin tosses, it turns out, are determined from the start, and only our incomplete knowledge makes them look random. "Random" sequences of numbers are more elusive, but Godels undecidability theorem informs us that we will never know. Those familiar with quantum indeterminacy assert that order is an illusion, and that the world is fundamentally random. Yet randomness is also an illusion. Perhaps order and randomness, like waves and particles, are only two sides of the same (tossed) coin.

A fascinating exploration of how insights from computer algorithms can be applied to our everyday lives, helping to solve common decision-making problems and illuminate the workings of the human mind All our lives are constrained by limited space and time, limits that give rise to a particular set of problems. What should we do, or leave undone, in a day or a lifetime? How much messiness should we accept? What balance of new activities and familiar favorites is the most fulfilling? These may seem like uniquely human quandaries, but they are not: computers, too, face the same constraints, so computer scientists have been grappling with their version of such issues for decades. And the solutions they've found have much to teach us. In a dazzlingly interdisciplinary work, acclaimed author Brian Christian and cognitive scientist Tom Griffiths show how the algorithms used by computers can also untangle very human questions. They explain how to have better hunches and when to leave things to chance, how to deal with overwhelming choices and how best to connect with others. From finding a spouse to finding a parking spot, from organizing one's inbox to understanding the workings of memory, Algorithms to Live By transforms the wisdom of computer science into strategies for human living.

This essential companion to Chaitin's successful books The Unknowable and The Limits of Mathematics, presents the technical core of his theory of program-size complexity. The two previous volumes are more concerned with applications to meta-mathematics. LISP is used to present the key algorithms and to enable computer users to interact with the authors proofs and discover for themselves how they work. The LISP code for this book is available at the author's Web site together with a Java applet LISP interpreter. "No one has looked deeper and farther into the abyss of randomness and its role in mathematics than Greg Chaitin. This book tells you everything hes seen. Don miss it." John Casti, Santa Fe Institute, Author of Goedel: A Life of Logic.'

Looking at a sequence of zeros and ones, we often feel that it is not random, that is, it is not plausible as an outcome of fair coin tossing. Why? The answer is provided by algorithmic information theory: because the sequence is compressible, that is, it has small complexity or, equivalently, can be produced by a short program. This idea, going back to Solomonoff, Kolmogorov, Chaitin, Levin, and others, is now the starting point of algorithmic information theory. The first part of this book is a textbook-style exposition of the basic notions of complexity and randomness; the second part covers some recent work done by participants of the “Kolmogorov seminar” in Moscow (started by Kolmogorov himself in the 1980s) and their colleagues. This book contains numerous exercises (embedded in the text) that will help readers to grasp the material.

In the sixteenth and seventeenth centuries, gamblers and mathematicians transformed the idea of chance from a mystery into the discipline of probability, setting the stage for a series of breakthroughs that enabled or transformed innumerable fields, from gambling, mathematics, statistics, economics, and finance to physics and computer science. This book tells the story of ten great ideas about chance and the thinkers who developed them, tracing the philosophical implications of these ideas as well as their mathematical impact.

The present book is based on a course developed as partofthe large NSF-funded GatewayCoalitionInitiativeinEngineeringEducationwhichincludedCaseWest ern Reserve University, Columbia University, Cooper Union, Drexel University, Florida International University, New Jersey Institute ofTechnology, Ohio State University, University ofPennsylvania, Polytechnic University, and Universityof South Carolina. The Coalition aimed to restructure the engineering curriculum by incorporating the latest technological innovations and tried to attract more and betterstudents to engineering and science. Draftsofthis textbookhave been used since 1992instatisticscoursestaughtatCWRU, IndianaUniversity, Bloomington, and at the universities in Gottingen, Germany, and Grenoble, France. Another purpose of this project was to develop a courseware that would take advantage ofthe Electronic Learning Environment created by CWRUnet-the all fiber-optic Case Western Reserve University computer network, and its ability to let students run Mathematica experiments and projects in their dormitory rooms, and interactpaperlessly with the instructor. Theoretically, onecould try togothroughthisbook withoutdoing Mathematica experimentsonthecomputer, butitwouldbelikeplayingChopin's Piano Concerto in E-minor, or Pink Floyd's The Wall, on an accordion. One would get an idea ofwhatthe tune was without everexperiencing the full richness andpowerofthe entire composition, and the whole ambience would be miscued.

Introduces cutting-edge research on machine learning theory and practice, providing an accessible, modern algorithmic toolkit.

Over the course of a generation, algorithms have gone from mathematical abstractions to powerful mediators of daily life. Algorithms have made our lives more efficient, more entertaining, and, sometimes, better informed. At the same time, complex algorithms are increasingly violating the basic rights of individual citizens. Allegedly anonymized datasets routinely leak our most sensitive personal information; statistical models for everything from mortgages to college admissions reflect racial and gender bias. Meanwhile, users manipulate algorithms to "game" search engines, spam filters, online reviewing services, and navigation apps. Understanding and improving the science behind the algorithms that run our lives is rapidly becoming one of the most pressing issues of this century. Traditional fixes, such as laws, regulations and watchdog groups, have proven woefully inadequate. Reporting from the cutting edge of scientific research, The Ethical Algorithm offers a new approach: a set of principled solutions based on the emerging and exciting science of socially aware algorithm design. Michael Kearns and Aaron Roth explain how we can better embed human principles into machine code - without halting the advance of data-driven scientific exploration. Weaving together innovative research with stories of citizens, scientists, and activists on the front lines, The Ethical Algorithm offers a compelling vision for a future, one in which we can better protect humans from the unintended impacts of algorithms while continuing to inspire wondrous advances in technology.