

Elements of Causal Inference Foundations And Learning Algorithms Adaptive Computation And Machine Learning Series

Emphasizing causation as a functional relationship between variables that describe objects, Linear Causal Modeling with Structural Equations integrates a general philosophical theory of causation with structural equation modeling (SEM) that concerns the special case of linear causal relations. In addition to describing how the functional relation concept may be generalized to treat probabilistic causation, the book reviews historical treatments of causation and explores recent developments in experimental psychology on studies of the perception of causation. It looks at how to perceive causal relations directly by perceiving quantities in magnitudes and motions of causes that are conserved in the effects of causal exchanges. The author surveys the basic concepts of graph theory useful in the formulation of structural models. Focusing on SEM, he shows how to write a set of structural equations corresponding to the path diagram, describes two ways of computing variances and covariances of variables in a structural equation model, and introduces matrix equations for the general structural equation model. The text then discusses the problem of identifying a model, parameter estimation, issues involved in designing structural equation models, the application of confirmatory factor analysis, equivalent models, the use of instrumental variables to resolve issues of causal direction and mediated causation, longitudinal modeling, and nonrecursive models with loops. It also evaluates models on several dimensions and examines the polychoric and polyserial correlation coefficients and their derivation. Covering the fundamentals of algebra and the history of causality, this book provides a solid understanding of causation, linear causal modeling, and SEM. It takes readers through the process of identifying, estimating, analyzing, and evaluating a range of models.

Outside of randomized experiments, association does not imply causation, and yet there is nothing defective about our knowledge that smoking causes lung cancer, a conclusion reached in the absence of randomized experimentation with humans. How is that possible? If observed associations do not identify causal effects in observational studies, how can a sequence of such associations become decisive? Two or more associations may each be susceptible to unmeasured biases, yet not susceptible to the same biases. An observational study has two evidence factors if it provides two comparisons susceptible to different biases that may be combined as if from independent studies of different data by different investigators, despite using the same data twice. If the two factors concur, then they may exhibit greater insensitivity to unmeasured biases than either factor exhibits on its own. Replication and Evidence Factors in Observational Studies includes four parts: A concise introduction to causal inference, making the book self-contained Practical examples of evidence factors from the health and social sciences with analyses in R The theory of evidence factors Study design with evidence factors A companion R package evident is available from CRAN.

"This book presents a new approach to causal inference and explanation, addressing both the timing and complexity of relationships. The method's feasibility and success is demonstrated through theoretical and experimental case studies"--

A pioneer of artificial intelligence shows how the study of causality revolutionized science and the world "Correlation does not imply causation." This mantra was invoked by scientists for decades in order to avoid taking positions as to whether one thing caused another, such as smoking and cancer and carbon dioxide and global warming. But today, that taboo is dead. The causal revolution, sparked by world-renowned computer scientist Judea Pearl and his colleagues, has cut through a century of confusion and placed cause and effect on a firm scientific basis. Now, Pearl and science journalist Dana Mackenzie explain causal thinking to general readers for the first time, showing how it allows us to explore the world that is and the worlds that could have been. It is the essence of human and artificial intelligence. And just as Pearl's discoveries have enabled machines to think better, The Book of Why explains how we can think better.

Foundations
A Primer
Causation, Prediction, and Search
Replication and Evidence Factors in Observational Studies
22nd International Conference, FOSSACS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings
Empirical Inference
A hands-on approach to tasks and techniques in data stream mining and real-time analytics, with examples in MOA, a popular freely available open-source software framework. Today many information sources—including sensor networks, financial markets, social networks, and healthcare monitoring—are so-called data streams, arriving sequentially and at high speed. Analysis must take place in real time, with partial data and without the capacity to store the entire data set. This book presents algorithms and techniques used in data stream mining and real-time analytics. Taking a hands-on approach, the book demonstrates the techniques using MOA (Massive Online Analysis), a popular, freely available open-source software framework, allowing readers to try out the techniques after reading the explanations. The book first offers a brief introduction to the topic, covering big data mining, basic methodologies for mining data streams, and a simple example of MOA. More detailed discussions follow, with chapters on sketching techniques, change, classification, ensemble methods, regression, clustering, and frequent pattern mining. Most of these chapters include exercises, an MOA-based lab session, or both. Finally, the book discusses the MOA software, covering the MOA graphical user interface, the command line, use of its API, and the development of new methods within MOA. The book will be an essential reference for readers who want to use data stream mining as a tool, researchers in innovation or data stream mining, and programmers who want to create new algorithms for MOA.

Why do ideas of how mechanisms relate to causality and probability differ so much across the sciences? Can progress in understanding the tools of causal inference in some sciences lead to progress in others? This book tackles these questions and others concerning the use of causality in the sciences. The application of causal inference methods is growing exponentially in fields that deal with observational data. Written by pioneers in the field, this practical book presents an authoritative yet accessible overview of the methods and applications of causal inference. With a wide range of detailed, worked examples using real epidemiologic data as well as software for replicating the analyses, the text provides a thorough introduction to the basics of the theory for non-time-varying treatments and the generalization to complex longitudinal data.

Elements of Causal InferenceFoundations and Learning AlgorithmsMIT Press
Causal Inference in Statistics
Bayesian Reasoning and Machine Learning
Machine Learning for Data Streams
Biological Sequence Analysis in the Era of High-Throughput Sequencing
Causal Inference for Observational and Experimental Data
Causality in the Sciences
A new approach for defining causality and such related notions as degree of responsibility, degrees of blame, and causal explanation. Causality plays a central role in the way people structure the world; we constantly seek causal explanations for our observations. But what does it even mean that an event C “actually caused” event E? The problem of defining actual causation goes beyond mere philosophical speculation. For example, in many legal arguments, it is precisely what needs to be established in order to determine responsibility. The philosophy literature has been struggling with the problem of defining causality since Hume. In this book, Joseph Halpern explores actual causality, and such related notions as degree of responsibility, degree of blame, and causal explanation. The goal is to arrive at a definition of causality that matches our natural language usage and is helpful, for example, to a jury deciding a legal case, a programmer looking for the line of code that cause some software to fail, or an economist trying to determine whether austerity caused a subsequent depression. Halpern applies and expands an approach to causality that he and Judea Pearl developed, based on structural equations. He carefully formulates a definition of causality, and building on this, defines degree of responsibility, degree of blame, and causal explanation. He concludes by discussing how these ideas can be applied to such practical problems as accountability and program verification. Technical details are generally confined to the final section of each chapter and can be skipped by non-mathematical readers.
In the face of conflicting claims about some treatments, behaviors, and policies, the question arises: What is the most scientifically rigorous way to draw conclusions about cause and effect in the study of humans? In this introduction to causal inference, Paul Rosenbaum explains key concepts and methods through real-world examples.
In this Element and its accompanying second Element, A Practical Introduction to Regression Discontinuity Designs: Extensions, Matias Cattaneo, Nicolás Idrrobo, and Rocio Titiunik provide an accessible and practical guide for the analysis and interpretation of regression discontinuity (RD) designs that encourages the use of a common set of practices and facilitates the accumulation of the evidence. In this Element, the authors discuss the foundations of the canonical Sharp RD design, which has the following features: (i) the score is continuously distributed and has only one dimension, (ii) there is only one cutoff, and (iii) compliance with the treatment assignment is perfect. In the second Element, the authors discuss practical and conceptual extensions to this basic RD setup. Causality offers the first comprehensive coverage of causal analysis in many sciences, including recent advances using graphical methods. Pearl presents a unified account of the probabilistic, manipulative, counterfactual and structural approaches to causation, and devises simple mathematical tools for analyzing the relationships between causal connections, statistical associations, actions and observations. The book will open the way for including causal analysis in the standard curriculum of statistics, artificial intelligence ...
with Practical Examples in MOA
Designing Social Inquiry
Measuring Racial Discrimination
Fundamentals of Causal Inference
Bayesian Nets and Causality: Philosophical and Computational Foundations

This book explores why, regarding practical reasoning, humans are sometimes still faster than artificial intelligence systems. It is the first to offer a self-contained presentation of neural network models for many computer science logics. The statistician profession is at a unique point in history. The need for valid statistical tools is greater than ever; data sets are massive, often measuring hundreds of thousands of measurements for a single subject. The field is ready to move towards clear objective benchmarks under which tools can be evaluated. Targeted learning allows (1) the full generalization and utilization of cross-validation as an estimator selection tool so that the subjective choices made by humans are now made by the machine, and (2) targeting the fitting of the probability distribution of the data toward the target parameter representing the scientific question of interest. This book is aimed at both statisticians and applied researchers interested in causal inference and general effect estimation for observational and experimental data. Part I is an accessible introduction to super learning and the targeted maximum likelihood estimator, including related concepts necessary to understand and apply these methods. Parts II-IX handle complex data structures and topics applied researchers will immediately recognize from their own research, including time-to-event outcomes, direct and indirect effects, positivity violations, case-control studies, censored data, longitudinal data, and genomic studies.

This book presents the most recent and advanced techniques for creating autonomous AI systems capable of planning and acting effectively.

This textbook for graduate students in statistics, data science, and public health deals with the practical challenges that come with big, complex, and dynamic data. It presents a scientific roadmap to translate real-world data science applications into formal statistical estimation problems by using the general template of targeted maximum likelihood estimators. These targeted machine learning algorithms estimate quantities of interest while still providing valid inference. Targeted learning methods within data science area critical component for solving scientific problems in the modern age. The techniques can answer complex questions including optimal rules for assigning treatment based on longitudinal data with time-dependent confounding, as well as other estimands in dependent data structures, such as networks. Included in Targeted Learning in Data Science are demonstrations with soft ware packages and real data sets that present a case that targeted learning is crucial for the next generation of statisticians and data scientists. Th is book is a sequel to the first textbook on machine learning for causal inference, Targeted Learning, published in 2011. Mark van der Laan, PhD, is Jiann-Ping Hsu/Karl E. Peace Professor of Biostatistics and Statistics at UC Berkeley. His research interests include statistical methods in genomics, survival analysis, censored data, machine learning, semiparametric models, causal inference, and targeted learning. Dr. van der Laan received the 2004 Mortimer Spiegelman Award, the 2005 Van Dantzig Award, the 2005 COPSS Snedecor Award, the 2005 COPSS Presidential Award, and has graduated over 40 PhD students in biostatistics and statistics. Sherri Rose, PhD, is Associate Professor of Health Care Policy (Biostatistics) at Harvard Medical School. Her work is centered on developing and integrating innovative statistical approaches to advance human health. Dr. Rose's methodological research focuses on nonparametric machine learning for causal inference and prediction. She co-leads the Health Policy Data Science Lab and currently serves as an associate editor for the Journal of the American Statistical Association and Biostatistics.

An Introduction to Causal Inference
Mathematical Analysis for Machine Learning and Data Mining
Machine-learning Techniques in Economics
Genome-Scale Algorithm Design
Foundations of Software Science and Computation Structures
Festschrift in Honor of Vladimir N. Vapnik

Configural Frequency Analysis (CFA) provides an up-to-the-minute comprehensive introduction to its techniques, models, and applications. Written in a formal yet accessible style, actual empirical data examples are used to illustrate key concepts. Step-by-step program sequences are used to show readers how to employ CFA methods using commercial software packages, such as SAS, SPSS, SYSTAT, S-Plus, or those written specifically to perform CFA. CFA is an important method for analyzing results involved with categorical and longitudinal data. It allows one to answer the question of whether individual cells or groups of cells of cross-classifications differ significantly from expectations. The expectations are calculated using methods employed in log-linear modeling or a priori information. It is the only statistical method that allows one to make statements about empty areas in the data space. Applied and or person-oriented researchers, statisticians, and advance students interested in CFA and categorical and longitudinal data will find this book to be a valuable resource. Developed since 1969, this method is now used by a large number of researchers around the world in a variety of disciplines, including psychology, education, medicine, and sociology. Configural Frequency Analysis will serve as an excellent text for courses on configural frequency analysis, categorical variable analysis, or analysis of contingency tables. Prerequisites include an understanding of descriptive statistics, hypothesis testing, statistical model fitting, and some understanding of categorical data analysis and matrix algebra.

This book summarizes recent advances in causal inference and underscores the paradigmatic shifts that must be undertaken in moving from traditional statistical analysis to causal analysis of multivariate data. Special emphasis is placed on the assumptions that underlie all causal inferences, the languages used in formulating those assumptions, the conditional nature of all causal and counterfactual claims, and the methods that have been developed for the assessment of such claims.

These advances are illustrated using a general theory of causation based on the Structural Causal Model (SCM), which subsumes and unifies other approaches to causation, and provides a coherent mathematical foundation for the analysis of causes and counterfactuals. In particular, the paper surveys the development of mathematical tools for inferring (from a combination of data and assumptions) answers to three types of causal queries: those about (1) the effects of potential treatments, (2) the structural and potential-outcome frameworks and pretests tools for a symbolic analysis that uses the strong features of both. The tools are demonstrated in the analyses of mediation, causes of effects, and probabilities of causation.

This book develops a machine-learning framework for predicting economic growth. It can also be considered as a primer for using machine learning (also known as data mining or data analytics) to answer economic questions. While machine learning itself is not a new idea, advances in computing technology combined with a dawning realization of its applicability to economic questions makes it a new tool for economists.

This compendium provides a self-contained introduction to mathematical analysis in the field of machine learning and data mining. The mathematical analysis component of the typical mathematical curriculum for computer science students omits these very important ideas and techniques which are indispensable for approaching specialized area of machine learning centered around optimization such as support vector machines, neural networks, various types of regression, feature selection, and clustering. The book is of special interest to researchers and graduate students who will benefit from these application areas discussed in the book.

Methods, Models, and Applications

The Causal Structure of Natural Selection

Methods for Applied Empirical Research

How People Think About the World and its Alternatives

Configural Frequency Analysis

New Tools for Predicting Economic Growth

This book is intended for anyone, regardless of discipline, who is interested in the use of statistical methods to help obtain scientific explanations or to predict the outcomes of actions, experiments or policies. Much of G. Udny Yule's work illustrates a vision of statistics whose goal is to investigate when and how causal influences may be reliably inferred, and their comparative strengths estimated, from statistical samples. Yule's enterprise has been largely replaced by Ronald Fisher's conception, in which there is a fundamental cleavage between experimental and non experimental inquiry, and statistics is largely unable to aid in causal inference without randomized experimental trials. Every now and then members of the statistical community express misgivings about this turn of events, and, in our view, rightly so. Our work represents a return to something like Yule's conception of the enterprise of theoretical statistics and its potential practical benefits. If intellectual history in the 20th century had gone otherwise, there might have been a discipline to which our work belongs. As it happens, there is not. We develop material that belongs to statistics, to computer science, and to philosophy; the combination may not be entirely satisfactory for specialists in any of these subjects. We hope it is nonetheless satisfactory for its purpose.

Probabilistic Reasoning in Intelligent Systems is a complete and accessible account of the theoretical foundations and computational methods that underlie plausible reasoning under uncertainty. The author provides a coherent explication of probability as a language for reasoning with partial belief and offers a unifying perspective on other AI approaches to uncertainty, such as the Dempster-Shafer formalism, truth maintenance systems, and nonmonotonic logic. The author distinguishes syntactic and semantic approaches to uncertainty--and offers techniques, based on belief networks, that provide a mechanism for making semantics-based systems operational. Specifically, network-propagation techniques serve as a mechanism for combining the theoretical coherence of probability theory with modern demands of reasoning-systems technology: modular declarative inputs, conceptually meaningful inferences, and parallel distributed representations. Applications include diagnosis, forecasting, image interpretation, multi-sensor fusion, decision support systems, plan recognition, planning, speech recognition--in short, almost every task requiring that conclusions be drawn from uncertain clues and incomplete information. Probabilistic Reasoning in Intelligent Systems will be of special interest to scholars and researchers in AI, decision theory, statistics, logic, philosophy, cognitive psychology, and the management sciences. Professionals in the areas of knowledge-based systems, operations research, engineering, and statistics will find theoretical and computational tools of immediate practical use. The book can also be used as an excellent text for graduate-level courses in AI, operations research, or applied probability.

Presents the Terminology and Methods of Mendelian Randomization for Epidemiological StudiesMendelian randomization uses genetic instrumental variables to make inferences about causal effects based on observational data. It, therefore, can be a reliable way of assessing the causal nature of risk factors, such as biomarkers, for a wide range of disea

'The editors of the new SAGE Handbook of Regression Analysis and Causal Inference have assembled a wide-ranging, high-quality, and timely collection of articles on topics of central importance to quantitative social research, many written by leaders in the field. Everyone engaged in statistical analysis of social-science data will find something of interest in this book.' - John Fox, Professor, Department of Sociology, McMaster University
'The authors do a great job in explaining the various statistical methods in a clear and simple way - focussing on fundamental understanding, interpretation of results, and practical application - yet being precise in their exposition.' - Ben Jann, Executive Director, Institute of Sociology, University of Bern
'Best and Wolf have put together a powerful collection, especially valuable in its separate discussions of uses for both cross-sectional and panel data analysis.' -Tom Smith, Senior Fellow, NORC, University of Chicago
Edited and written by a team of leading international social scientists, this Handbook provides a comprehensive introduction to multivariate methods. The Handbook focuses on regression analysis of cross-sectional and longitudinal data with an emphasis on causal analysis, thereby covering a large number of different techniques including selection models, complex samples, and regression discontinuity designs. Starting with a non-mathematical introduction to the method covered in that section, giving readers a basic knowledge of the method's logic, scope and unique features. Next, the mathematical and statistical basis of each method is presented along with advanced aspects. Using real-world data from the European Social Survey (ESS) and the Socio-Economic Panel (GSOEP), the book provides a comprehensive discussion of each method's application, making this an ideal text for PhD students and researchers embarking on their own data analysis.

Actual Causality
Mendelian Randomization
Networks of Plausible Inference
Elements of Causal Inference
Targeted Learning
The Book of Why

This book honours the outstanding contributions of Vladimir Vapnik, a rare example of a scientist for whom the following statements hold true simultaneously: his work led to the inception of a new field of research, the theory of statistical learning and empirical inference; he has lived to see the field blossom; and he is still as active as ever. He started analyzing learning algorithms in the 1960s and he invented the first version of the generalized portrait algorithm. He later developed one of the most successful methods in machine learning, the support vector machine (SVM) - more than just an algorithm, this was a new approach to learning problems, pioneering the use of functional analysis and convex optimization in machine learning. Part I of this book contains three chapters describing and witnessing some of Vladimir Vapnik's contributions to science. In the first chapter, Léon Bottou discusses the seminal paper published in 1968 by Vapnik and Chervonenkis that lay the foundations of statistical learning theory, and the second chapter is an English-language translation of that original paper. In the third chapter, Alexey Chervonenkis presents a first-hand account of the early history of SVMs and valuable insights into the first steps in the development of the SVM in the framework of the generalised portrait method. The remaining chapters, by leading scientists in domains such as statistics, theoretical computer science, and mathematics, address substantial topics in the theory and practice of statistical learning theory, including SVMs and other kernel-based methods, boosting, PAC-Bayesian theory, online and transductive learning, loss functions, learnable function classes, notions of complexity for function classes, multitask learning, and hypothesis selection. These contributions include historical and context notes, short surveys, and comments on future research directions. This book will be of interest to researchers, engineers, and graduate students engaged with all aspects of statistical learning.

A rigorous and comprehensive textbook covering the major approaches to knowledge graphs, an active and interdisciplinary area within artificial intelligence. The field of knowledge graphs, which allows us to model, process, and derive insights from complex real-world data, has emerged as an active and interdisciplinary area of artificial intelligence over the last decade, drawing on such fields as natural language processing, data mining, and the semantic web. Current projects involve predicting cyberattacks, recommending products, and even gleanig insights from thousands of papers on COVID-19. This textbook offers rigorous and comprehensive coverage of the field. It focuses systematically on the major approaches, both those that have stood the test of time and the latest deep learning methods.

Many of the concepts and terminology surrounding modern causal inference can be quite intimidating to the novice. Judea Pearl presents a book ideal for beginners in statistics, providing a comprehensive introduction to the field of causality. Examples from classical statistics are presented throughout to demonstrate the need for causality in resolving decision-making dilemmas posed by data. Causal methods are also compared to traditional statistical methods, whilst questions are provided at the end of each section to aid student learning.

Many racial and ethnic groups in the United States, including blacks, Hispanics, Asians, American Indians, and others, have historically faced severe discriminationa€pervasive and open denial of civil, social, political, educational, and economic opportunities. Today, large differences among racial and ethnic groups continue to exist in employment, income and wealth, housing, education, criminal justice, health, and other areas. While many factors may contribute to such differences, their size and extent suggest that various forms of discriminatory treatment persist in U.S. society and serve to undercut the achievement of equal opportunity. Measuring Racial Discrimination considers the definition of race and racial discrimination, reviews the existing techniques used to measure racial discrimination, and identifies new tools and areas for future research. The book conducts a thorough evaluation of current methodologies for a wide range of circumstances in which racial discrimination may occur, and makes recommendations on how to better assess the presence and effects of discrimination.

Causal Models

Foundations and Learning Algorithms

Experimental and Quasi-experimental Designs for Generalized Causal Inference

Causality, Probability, and Time

Causal Inference

Automated Planning and Acting

A practical introduction perfect for final-year undergraduate and graduate students without a solid background in linear algebra and calculus.

Process-tracing in social science is a method for studying causal mechanisms linking causes with outcomes. This enables the researcher to make strong inferences about how a cause (or set of causes) contributes to producing an outcome. Derek Beach and Rasmus Brn Pedersen introduce a refined definition of process-tracing, differentiating it into three distinct variants and explaining the applications and limitations of each. The authors develop the underlying logic of process-tracing, including how one should understand causal mechanisms and how Bayesian logic enables strong within-case inferences. They provide instructions for identifying the variant of process-tracing most appropriate for the research question at hand and a set of guidelines for each stage of the research process.

Sections include: experiments and generalised causal inference; statistical conclusion validity and internal validity; construct validity and external validity; quasi-experimental designs that either lack a control group or lack pretest observations on the outcome; quasi-experimental designs that use both control groups and pretests; quasi-experiments: interrupted time-series designs; regression discontinuity designs; randomised experiments: rationale, designs, and conditions conducive to doing them; practical problems 1 ethics, participation recruitment and random assignment; practical problems 2: treatment implementation and attrition; generalised causal inference: a grounded theory; generalised causal inference: methods for single studies; generalised causal inference: methods for multiple studies; a critical assessment of our assumptions.

This open access book constitutes the proceedings of the 22nd International Conference on Foundations of Software Science and Computational Structures, FOSSACS 2019, which took place in Prague, Czech Republic, in April 2019, held as part of the European Joint Conference on Theory and Practice of Software, ETAPS 2019.The 29 papers presented in this volume were carefully reviewed and selected from 85 submissions. They deal with foundational research with a clear significance for software science.

Causal Inference for Complex Longitudinal Studies
Linear Causal Modeling with Structural Equations
Neural-Symbolic Cognitive Reasoning
Statistics and Causality
The New Science of Cause and Effect
Knowledge Graphs

The classic work on qualitative methods in political science Designing Social Inquiry presents a unified approach to qualitative and quantitative research in political science, showing how the same logic of inference underlies both. This stimulating book discusses issues related to framing research questions, measuring the accuracy of data and the uncertainty of empirical inferences, discovering causal effects, and getting the most out of qualitative research. It addresses topics such as interpretation and inference, comparative case studies, constructing causal theories, dependent and explanatory variables, the limits of random selection, selection bias, and errors in measurement. The book only uses mathematical notation to clarify concepts, and assumes no prior knowledge of mathematics or statistics. Featuring a new preface by Robert O. Keohane and Gary King, this edition makes an influential work available to new generations of qualitative researchers in the social sciences.

Human beings are active agents who can think. To understand how thought serves action requires understanding how people conceive of the relation between cause and effect, between action and outcome. In cognitive terms, how do people construct and reason with the causal models we use to represent our world? A revolution is occurring in how statisticians, philosophers, and computer scientists answer this question. Those fields have ushered in new insights about causal models by thinking about how to represent causal structure mathematically, in a framework that uses graphs and probability theory to develop what are called causal Bayesian networks. The framework starts with the idea that the purpose of causal structure is to understand and predict the effects of intervention. How does intervening on one thing affect other things? This is not a question merely about probability (or logic), but about action. The framework offers a new understanding of mind: Thought is about the effects of intervention and cognition is thus intimately tied to actions that take place either in the actual physical world or in imagination, in counterfactual worlds. The book offers a conceptual introduction to the key mathematical ideas, presenting them in a non-technical way, focusing on the intuitions rather than the theorems. It tries to show why the ideas are important to understanding how people explain things and why thinking not only about the world as it is but the world as it could be is so central to human action. The book reviews the role of causality, causal models, and intervention in the basic human cognitive functions: decision making, reasoning, judgment, categorization, inductive inference, language, and learning. In short, the book offers a discussion about how people think, talk, learn, and explain things in causal terms, in terms of action and manipulation.

A concise and self-contained introduction to causal inference, increasingly important in data science and machine learning. The mathematization of causality is a relatively recent development, and has become increasingly important in data science and machine learning. This book offers a self-contained and concise introduction to causal models and how to learn them from data. After explaining the need for causal models and discussing some of the principles underlying causal inference, the book teaches readers how to use causal models: how to compute intervention distributions, how to infer causal models from observational and interventional data, and how causal ideas could be exploited for classical machine learning problems. All of these topics are discussed first in terms of two variables and then in a more general multivariate case. The bivariate case turns out to be a particularly hard problem for causal learning because there are no conditional independences as used by classical methods for solving multivariate cases. The authors consider analyzing statistical asymmetries between cause and effect to be highly instructive, and they report on their decade of intensive research into this problem. The book is accessible to readers with a background in machine learning or statistics, and can be used in graduate courses or as a reference for researchers. The text includes code snippets that can be copied and pasted, exercises, and an appendix with a summary of the most important technical concepts.

High-throughput sequencing has revolutionised the field of biological sequence analysis. Its application has enabled researchers to address important biological questions, often for the first time. This book provides an integrated presentation of the fundamental algorithms and data structures that power modern sequence analysis workflows. The topics covered range from the foundations of biological sequence analysis (alignments and hidden Markov models), to classical index structures (k-mer indexes, suffix arrays and suffix

trees), Burrows–Wheeler indexes, graph algorithms and a number of advanced omics applications. The chapters feature numerous examples, algorithm visualisations, exercises and problems, each chosen to reflect the steps of large-scale sequencing projects, including read alignment, variant calling, haplotyping, fragment assembly, alignment-free genome comparison, transcript prediction and analysis of metagenomic samples. Each biological problem is accompanied by precise formulations, providing graduate students and researchers in bioinformatics and computer science with a powerful toolkit for the emerging applications of high-throughput sequencing.

Causality

Fundamentals, Techniques, and Applications

Targeted Learning in Data Science

Observation and Experiment

A Practical Introduction to Regression Discontinuity Designs

With R

Recent arguments concerning the nature of causation in evolutionary theory, now often known as the debate between the 'causalists' and 'statisticalist' positions, have involved answers to a variety of independent questions – definitions of key evolutionary concepts like natural selection, fitness, and genetic drift; causation in multi-level systems; or the nature of evolutionary explanations, among others. This Element offers a way to disentangle one set of these questions surrounding the causal structure of natural selection. Doing so allows us to clearly reconstruct the approach that some of these major competing interpretations of evolutionary theory have to this causal structure, highlighting particular features of philosophical interest within each. Further, those features concern problems not exclusive to the philosophy of biology. Connections between them and, in two case studies, contemporary metaphysics and philosophy of physics demonstrate the potential value of broader collaboration in the understanding of evolution.

One of the primary motivations for clinical trials and observational studies of humans is to infer cause and effect. Disentangling causation from confounding is of utmost importance. Fundamentals of Causal Inference explains and relates different methods of confounding adjustment in terms of potential outcomes and graphical models, including standardization, difference-in-differences estimation, the front-door method, instrumental variables estimation, and propensity score methods. It also covers effect-measure modification, precision variables, mediation analyses, and time-dependent confounding. Several real data examples, simulation studies, and analyses using R motivate the methods throughout. The book assumes familiarity with basic statistics and probability, regression, and R and is suitable for seniors or graduate students in statistics, biostatistics, and data science as well as PhD students in a wide variety of other disciplines, including epidemiology, pharmacy, the health sciences, education, and the social, economic, and behavioral sciences. Beginning with a brief history and a review of essential elements of probability and statistics, a unique feature of the book is its focus on real and simulated datasets with all binary variables to reduce complex methods down to their fundamentals. Calculus is not required, but a willingness to tackle mathematical notation, difficult concepts, and intricate logical arguments is essential. While many real data examples are included, the book also features the Double What-If Study, based on simulated data with known causal mechanisms, in the belief that the methods are best understood in circumstances where they are known to either succeed or fail. Datasets, R code, and solutions to odd-numbered exercises are available at www.routledge.com.

Bayesian nets are used in artificial intelligence as a calculus for causal reasoning, enabling machines to make predictions perform diagnoses, take decisions and even to discover causal relationships. This book brings together how to automate reasoning in artificial intelligence, and the nature of causality and probability in philosophy.

A one-of-a-kind guide to identifying and dealing with modern statistical developments in causality Written by a group of well-known experts, Statistics and Causality: Methods for Applied Empirical Research focuses on the most up-to-date developments in statistical methods in respect to causality. Illustrating the properties of statistical methods to theories of causality, the book features a summary of the latest developments in methods for statistical analysis of causality hypotheses. The book is divided into five accessible and independent parts. The first part introduces the foundations of causal structures and discusses issues associated with standard mechanistic and difference-making theories of causality. The second part features novel generalizations of methods designed to make statements concerning the direction of effects. The third part illustrates advances in Granger-causality testing and related issues. The fourth part focuses on counterfactual approaches and propensity score analysis. Finally, the fifth part presents designs for causal inference with an overview of the research designs commonly used in epidemiology. Statistics and Causality: Methods for Applied Empirical Research also includes: New statistical methodologies and approaches to causal analysis in the context of the continuing development of philosophical theories End-of-chapter bibliographies that provide references for further discussions and additional research topics Discussions on the use and applicability of software when appropriate Statistics and Causality: Methods for Applied Empirical Research is an ideal reference for practicing statisticians, applied mathematicians, psychologists, sociologists, logicians, medical professionals, epidemiologists, and educators who want to learn more about new methodologies in causal analysis. The book is also an excellent textbook for graduate-level courses in causality and qualitative logic.

Probabilistic Reasoning in Intelligent Systems

The SAGE Handbook of Regression Analysis and Causal Inference

Process-Tracing Methods

Foundations and Guidelines

Methods for Using Genetic Variants in Causal Estimation