

study, Steinhoff develops a Marxist analysis to argue that the popular theory of immaterial labour, which holds that information technologies increase the autonomy of workers from capital, tending towards a post-capitalist economy, does not adequately describe the situation of high-tech digital labour today. In the AI industry, digital labour remains firmly under the control of capital. Steinhoff argues that theories discerning therein an emergent autonomy of labour are in fact witnessing labour's increasing automation.

The book presents a remarkable collection of chapters covering a wide range of topics in the areas of intelligent systems and artificial intelligence, and their real-world applications. It gathers the proceedings of the Intelligent Systems Conference 2019, which attracted a total of 546 submissions from pioneering researchers, scientists, industrial engineers, and students from all around the world. These submissions underwent a double-blind peer-review process, after which 190 were selected for inclusion in these proceedings. As intelligent systems continue to replace and sometimes outperform human intelligence in decision-making processes, they have made it possible to tackle a host of problems more effectively. This branching out of computational intelligence in several directions and use of intelligent systems in everyday applications have created the need for an international conference as a venue for reporting on the latest innovations and trends. This book collects both theory and application based chapters on virtually all aspects of artificial intelligence; presenting state-of-the-art intelligent methods and techniques for solving real-world problems, along with a vision for future research, it represents a unique and valuable asset.

If machine learning transforms the nature of knowledge, does it also transform the practice of critical thought? Machine learning—programming computers to learn from data—has spread across scientific disciplines, media, entertainment, and government. Medical research, autonomous vehicles, credit transaction processing, computer gaming, recommendation systems, finance, surveillance, and robotics use machine learning. Machine learning devices (sometimes understood as scientific models, sometimes as operational algorithms) anchor the field of data science. They have also become mundane mechanisms deeply embedded in a variety of systems and gadgets. In contexts from the everyday to the esoteric, machine learning is said to transform the nature of knowledge. In this book, Adrian Mackenzie investigates whether machine learning also transforms the practice of critical thinking. Mackenzie focuses on machine learners—either humans and machines or human-machine relations—situated among settings, data, and devices. The settings range from fMRI to Facebook; the data anything from cat images to DNA sequences; the devices include neural networks, support vector machines, and decision trees. He examines specific learning algorithms—writing code and writing about code—and develops an archaeology of operations that, following Foucault, views machine learning as a form of knowledge production and a strategy of power. Exploring layers of abstraction, data infrastructures, coding practices, diagrams, mathematical formalisms, and the social organization of machine learning, Mackenzie traces the mostly invisible architecture of one of the central zones of contemporary technological cultures. Mackenzie's account of machine learning locates places in which a sense of agency can take root. His archaeology of the operational formation of machine learning does not unearth the footprint of a strategic monolith but reveals the local tributaries of force that feed into the generalization and plurality of the field.

A new edition of an introductory text in machine learning that gives a unified treatment of machine learning problems and solutions. The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Many successful applications of machine learning exist already, including systems that analyze past sales data to predict customer behavior, optimize robot behavior so that a task can be completed using minimum resources, and extract knowledge from bioinformatics data. The second edition of Introduction to Machine Learning is a comprehensive textbook on the subject, covering a broad array of topics not usually included in introductory machine learning texts. In order to present a unified treatment of machine learning problems and solutions, it discusses many methods from different fields, including statistics, pattern recognition, neural networks, artificial intelligence, signal processing, control, and data mining. All learning algorithms are explained so that the student can easily move from the equations in the book to a computer program. The text covers such topics as supervised learning, Bayesian decision theory, parametric methods, multivariate methods, multilayer perceptrons, local models, hidden Markov models, assessing and comparing classification algorithms, and reinforcement learning. New to the second edition are chapters on kernel machines, graphical models, and Bayesian estimation; expanded coverage of statistical tests in a chapter on design and analysis of machine learning experiments; case studies available on the Web (with downloadable results for instructors); and many additional exercises. All chapters have been revised and updated. Introduction to Machine Learning can be used by advanced undergraduates and graduate students who have completed courses in computer programming, probability, calculus, and linear algebra. It will also be of interest to engineers in the field who are concerned with the application of machine learning methods.

Now that people are aware that data can make the difference in an election or a business model, data science as an occupation is gaining ground. But how can you get started working in a wide-ranging, interdisciplinary field that's so clouded in hype? This insightful book, based on Columbia University's Introduction to Data Science class, tells you what you need to know. In many of these chapter-long lectures, data scientists from companies such as Google, Microsoft, and eBay share new algorithms, methods, and models by presenting case studies and the code they use. If you're familiar with linear algebra, probability, and statistics, and have programming experience, this book is an ideal introduction to data science. Topics include: Statistical inference, exploratory data analysis, and the data science process Algorithms Spam filters, Naive Bayes, and data wrangling Logistic regression Financial modeling Recommendation engines and causality Data visualization Social networks and data journalism Data engineering, MapReduce, Pregel, and Hadoop Doing Data Science is collaboration between course instructor Rachel Schutt, Senior VP of Data Science at News Corp, and data science consultant Cathy O'Neil, a senior data scientist at Johnson Research Labs, who attended

and blogged about the course.

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.

Introduction to Machine Learning, fourth edition MIT Press

A laboratory study that investigates how algorithms come into existence. Algorithms--often associated with the terms big data, machine learning, or artificial intelligence--underlie the technologies we use every day, and disputes over the consequences, actual or potential, of new algorithms arise regularly. In this book, Florian Jatton offers a new way to study computerized methods, providing an account of where algorithms come from and how they are constituted, investigating the practical activities by which algorithms are progressively assembled rather than what they may suggest or require once they are assembled.

This book provides a comprehensive overview of the concepts and approaches used for sequence, structure, and phylogenetic analysis. Starting with an introduction to the subject and intellectual property protection for bioinformatics, it guides readers through the latest sequencing technologies, sequence analysis, genomic variations, metagenomics, epigenomics, molecular evolution and phylogenetics, structural bioinformatics, protein folding, structure analysis and validation, drug discovery, reverse vaccinology, machine learning, application of R programming in biological data analysis, and the use of Linux in handling large data files.

A concise overview of machine learning--computer programs that learn from data--the basis of such applications as voice recognition and driverless cars. Today, machine learning underlies a range of applications we use every day, from product recommendations to voice recognition--as well as some we don't yet use everyday, including driverless cars. It is the basis for a new approach to artificial intelligence that aims to program computers to use example data or past experience to solve a given problem. In this volume in the MIT Press Essential Knowledge series, Ethem Alpaydin offers a concise and accessible overview of "the new AI." This expanded edition offers new material on such challenges facing machine learning as privacy, security, accountability, and bias. Alpaydin, author of a popular textbook on machine learning, explains that as "Big Data" has gotten bigger, the theory of machine learning--the foundation of efforts to process that data into knowledge--has also advanced. He describes the evolution of the field, explains important learning algorithms, and presents example applications. He discusses the use of machine learning algorithms for pattern recognition; artificial neural networks inspired by the human brain; algorithms that learn associations between instances; and reinforcement learning, when an autonomous agent learns to take actions to maximize reward. In a new chapter, he considers transparency, explainability, and fairness, and the ethical and legal implications of making decisions based on data.

A comprehensive review of an area of machine learning that deals with the use of unlabeled data in classification problems: state-of-the-art algorithms, a taxonomy of the field, applications, benchmark experiments, and directions for future research. In the field of machine learning, semi-supervised learning (SSL) occupies the middle ground, between supervised learning (in which all training examples are labeled) and unsupervised learning (in which no label data are given). Interest in SSL has increased in recent years, particularly because of application domains in which unlabeled data are plentiful, such as images, text, and bioinformatics. This first comprehensive overview of SSL presents state-of-the-art algorithms, a taxonomy of the field, selected applications, benchmark experiments, and perspectives on ongoing and future research. Semi-Supervised Learning first presents the key assumptions and ideas underlying the field: smoothness, cluster or low-density separation, manifold structure, and transduction. The core of the book is the presentation of SSL methods, organized according to algorithmic strategies. After an examination of generative models, the book describes algorithms that implement the low-density separation assumption, graph-based methods, and algorithms that perform two-step learning. The book then discusses SSL applications and offers guidelines for SSL practitioners by analyzing the results of extensive benchmark experiments. Finally, the book looks at interesting directions for SSL research. The book closes with a discussion of the relationship between semi-supervised learning and transduction.

A comprehensive introduction to machine learning that uses probabilistic models and inference as a unifying approach. Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides

these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package—PMTK (probabilistic modeling toolkit)—that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

This book attempts to provide a unified overview of the broad field of Machine Learning and its Practical implementation. This book is a survey of the state of art. It breaks this massive subject into comprehensible parts piece by piece. The objective is to focus on basic principles of machine learning with some leading edge topics. This book addresses a full spectrum of machine learning programming. The emphasis is to solve lot many programming examples using step-by step practical implementation of machine learning algorithms. To facilitate easy understanding of machine learning, this book has been written in such a simple style that a student thinks as if a teacher is sitting behind him and guiding him. This book is written as per the new syllabus of different Universities of India. It also Cover the syllabus of B.Tech.(CSE/IT), MCA, BCA of Delhi University, Delhi. GGSIPU, MDU, RGTU, Nagpur University, UTU, APJ Abdul Kalam University so on. The book is intended for both academic and professional audience.

An accessible introduction and essential reference for an approach to machine learning that creates highly accurate prediction rules by combining many weak and inaccurate ones. Boosting is an approach to machine learning based on the idea of creating a highly accurate predictor by combining many weak and inaccurate “rules of thumb.” A remarkably rich theory has evolved around boosting, with connections to a range of topics, including statistics, game theory, convex optimization, and information geometry. Boosting algorithms have also enjoyed practical success in such fields as biology, vision, and speech processing. At various times in its history, boosting has been perceived as mysterious, controversial, even paradoxical. This book, written by the inventors of the method, brings together, organizes, simplifies, and substantially extends two decades of research on boosting, presenting both theory and applications in a way that is accessible to readers from diverse backgrounds while also providing an authoritative reference for advanced researchers. With its introductory treatment of all material and its inclusion of exercises in every chapter, the book is appropriate for course use as well. The book begins with a general introduction to machine learning algorithms and their analysis; then explores the core theory of boosting, especially its ability to generalize; examines some of the myriad other theoretical viewpoints that help to explain and understand boosting; provides practical extensions of boosting for more complex learning problems; and finally presents a number of advanced theoretical topics. Numerous applications and practical illustrations are offered throughout.

A good book is like a teacher who sits behind the reader and guides him/her accordingly. Deep Learning has been an area of current research. After toiling through the various concepts of Deep Learning, the book slithers around all principles of deep learning. This book highlights in deep the concepts of deep learning so that new projects and researchers can be done. The book serves, both as textbook and as a reference book. Some of the highlights of the book are: Simple Language, Recent Concepts of Machine and Deep Learning explained, MCQ's, Conceptual Short Questions & Answers, Case Studies, Case Tools (like TensorFlow, H2O etc).

A general framework for constructing and using probabilistic models of complex systems that would enable a computer to use available information for making decisions. Most tasks require a person or an automated system to reason—to reach conclusions based on available information. The framework of probabilistic graphical models, presented in this book, provides a general approach for this task. The approach is model-based, allowing interpretable models to be constructed and then manipulated by reasoning algorithms. These models can also be learned automatically from data, allowing the approach to be used in cases where manually constructing a model is difficult or even impossible. Because uncertainty is an inescapable aspect of most real-world applications, the book focuses on probabilistic models, which make the uncertainty explicit and provide models that are more faithful to reality. Probabilistic Graphical Models discusses a variety of models, spanning Bayesian networks, undirected Markov networks, discrete and continuous models, and extensions to deal with dynamical systems and relational data. For each class of models, the text describes the three fundamental cornerstones: representation, inference, and learning, presenting both basic concepts and advanced techniques. Finally, the book considers the use of the proposed framework for causal reasoning and decision making under uncertainty. The main text in each chapter provides the detailed technical development of the key ideas. Most chapters also include boxes with additional material: skill boxes, which describe techniques; case study boxes, which discuss empirical cases related to the approach described in the text, including applications in computer vision, robotics, natural language understanding, and computational biology; and concept boxes, which present significant concepts drawn from the material in the chapter. Instructors (and readers) can group chapters in various combinations, from core topics to more technically advanced material, to suit their particular needs.

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.

Advanced statistical modeling and knowledge representation techniques for a newly emerging area of machine learning and probabilistic reasoning; includes introductory material, tutorials for different proposed approaches, and applications. Handling inherent uncertainty and

