Artificial Intelligence 3rd Edition Text By Winston

The two-volume set LNAI 9119 and LNAI 9120 constitutes the refereed proceedings of the 14th International Conference on Artificial Intelligence and Soft Computing, ICAISC 2015, held in Zakopane, Poland in June 2015. The 142 revised full papers presented in the volumes, were carefully reviewed and selected from 322 submissions. These proceedings present both traditional artificial intelligence methods and soft computing techniques. The goal is to bring together scientists representing both areas of research. The first volume covers topics as follows neural networks and their applications, fuzzy systems and their applications, evolutionary algorithms and their applications, classification and estimation, computer vision, image and speech analysis and the workshop: large-scale visual recognition and machine learning. The second volume has the focus on the following subjects: data mining, bioinformatics, biometrics and medical applications, concurrent and parallel processing, agent systems, robotics and control, artificial intelligence in modeling and simulation and various problems of artificial intelligence. For one or two-semester, undergraduate or graduate-level courses in Artificial Intelligence. The long-anticipated revision of this best-selling text offers the most comprehensive, up-to-date introduction to the theory and practice of artificial intelligence. This book constitutes the refereed proceedings of the

28th European Conference on Information Retrieval Research, ECIR 2006, held in London, April 2006. The 37 revised full papers and 28 revised poster papers presented are organized in topical sections on formal models, document and query representation and text understanding, topic identification and news retrieval, clustering and classification, refinement and feedback, performance and peer-to-peer networks, Web search, cross-language retrieval, genomic IR, and much more. This open access book examines recent advances in how artificial intelligence (AI) and robotics have elicited widespread debate over their benefits and drawbacks for humanity. The emergent technologies have for instance implications within medicine and health care, employment, transport, manufacturing, agriculture, and armed conflict. While there has been considerable attention devoted to robotics/AI applications in each of these domains, a fuller picture of their connections and the possible consequences for our shared humanity seems needed. This volume covers multidisciplinary research, examines current research frontiers in AI/robotics and likely impacts on societal well-being, human – robot relationships, as well as the opportunities and risks for sustainable development and peace. The attendant ethical and religious dimensions of these technologies are addressed and implications for regulatory policies on the use and future development of AI/robotics technologies are elaborated.
This book constitutes the refereed proceedings of the 15th International Conference on Advanced Information Systems Engineering, CaiSE 2003, held in Klagenfurt,

Austria in June 2003. The 45 revised full papers presented together with 3 invited contributions were carefully reviewed and selected from 219 submissions. The papers are organized in topical sections on XML, methods and models for information systems, UML, Internet business and social modeling, peer-to-peer systems, ontology-based methods, advanced design of information systems, knowledge, knowledge management, Web services, data warehouses, electronic agreements and workflow, requirements engineering, metrics and method engineering, and agent technologies and advanced environments.
This is an eagerly awaited revision of the single bestselling introduction to Artificial Intelligence ever published. It retains the best features of the earlier works including superior readability, currency, and excellence in the selection of the examples.
About the book:? Do you know what AI is doing to improve our health and wellbeing?? Does this new technology concern you, or impress you?? Do you want to know more about the future of AI in healthcare?Technology continues to advance at a pace that can seem bewildering. Nowhere else is it moving faster than in the health sector, where ?AI is now being used to improve millions of lives?.In this book, ? Artificial Intelligence in Healthcare: AI, Machine Learning, and Deep and Intelligent Medicine Simplified for Everyone ?, you can discover the great improvements that AI is making, with chapters covering:? The current applications and future of AI in healthcare and all major medical specialties? The benefits and risks weighed up?

The ethics involved? Machine learning and data science simplified? AI's role in medical research and education, health insurance, drug discovery, electronic health records, and the fight against COVID-19? The roles that major corporations and start-up companies are playing? The implementation of AI in clinical practice? And lots more?Quite simply the most authoritative text on the subject, Artificial Intelligence in Healthcare - 3rd Edition, is an absorbing and compelling read for anyone who wants to know more. It is packed with more updated information than any other book currently available, written in easy-to-understand language, and accessible to all.

New Scientist magazine was launched in 1956 "for all those men and women who are interested in scientific discovery, and in its industrial, commercial and social consequences". The brand's mission is no different today - for its consumers, New Scientist reports, explores and interprets the results of human endeavour set in the context of society and culture.

The present book is the product of conferences held in Bielefeld at the Center for interdisciplinary Sturlies (ZiF) in connection with a year-long ZiF Research Group with the theme "Prerational intelligence". The premise ex plored by the research group is that traditional notions of intelligent behav ior, which form the basis for much work in artificial intelligence and cog nitive science, presuppose many basic capabilities which are not trivial, as more recent work in robotics and neuroscience has shown, and that these capabilities may be best understood as ernerging from interaction and coop

eration in systems of simple agents, elements that accept inputs from and act upon their surroundings. The main focus is on the way animals and artificial systems process in formation about their surroundings in order to move and act adaptively. The analysis of the collective properties of systems of interacting agents, how ever, is a problem that occurs repeatedly in many disciplines. Therefore, contributions from a wide variety of areas have been included in order to obtain a broad overview of phenomena that demoostrate complexity arising from simple interactions or can be described as adaptive behavior arising from the collective action of groups of agents. To this end we have invited contributions on topics ranging from the development of complex structures and functions in systems ranging from cellular automata, genetic codes, and neural connectivity to social behavior and evolution. Additional contribu tions discuss traditional concepts of intelligence and adaptive behavior. 1.

Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural

networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets

computing techniques.

Provides an ontological characterization of texts, explores the issues raised by the identity of various texts, and presents a view of the function of authors and audiences, and of their relations to texts.

This book presents recent work on healthcare management and engineering using artificial intelligence and data mining techniques. Specific topics covered in the contributed chapters include predictive mining, decision support, capacity management, patient flow optimization, image compression, data clustering, and feature selection. The content will be valuable for researchers and postgraduate students in computer science, information technology, industrial engineering, and applied mathematics.

Artificial IntelligenceA Modern Approach, Global Edition

Enabling information interoperability, fostering legal knowledge usability and reuse, enhancing legal information search, in short, formalizing the complexity of

legal knowledge to enhance legal knowledge management are challenging tasks, for which different solutions and lines of research have been proposed. During the last decade, research and applications based on the use of legal ontologies as a technique to represent legal knowledge has raised a very interesting debate about their capacity and limitations to represent conceptual structures in the legal domain. Making conceptual legal knowledge explicit would support the development of a web of legal knowledge, improve communication, create trust and enable and support open data, e-government and e-democracy activities. Moreover, this explicit knowledge is also relevant to the formalization of software agents and the shaping of virtual institutions and multi-agent systems or environments. This book explores the use of ontologism in legal knowledge representation for semantically-enhanced legal knowledge systems or web-based applications. In it, current methodologies, tools and languages used for ontology development are revised, and the book includes an exhaustive revision of existing ontologies in the legal domain. The development of the Ontology of Professional Judicial Knowledge (OPJK) is presented as a case study.

CICLing 2004 was the 5th Annual Conference on Intelligent Text Processing and Computational Linguistics; see www.CICLing.org. CICLing conferences are intended to provide a balanced view of the cutting-edge developments in both theoretical foundations of computational linguistics and the practice of natural language text processing with its numerous applications.

A feature of CICLing conferences is their wide scope that covers nearly all areas of computational linguistics and all aspects of natural language processing applications. These conferences are a forum for dialogue between the specialists working in the two areas. This year we were honored by the presence of our invited speakers Martin KayofStanfordUniversity,PhilipResnikoftheUniversityofMaryland,Ricardo Baeza-Yates of the University of Chile, and Nick Campbell of the ATR Spoken Language Translation Research Laboratories. They delivered excellent extended lectures and organized vivid discussions. Of129submissionsreceived(74fullpapersand44shortpapers),aftercareful international reviewing 74 papers were selected for presentation (40 full papers and 35shortpapers),writtenby176authorsfrom21countries:Korea(37),Spain (34), Japan (22), Mexico (15), China (11), Germany (10), Ireland (10), UK (10), Singapore (6), Canada (3), Czech Rep. (3), France (3), Brazil (2), Sweden (2), Taiwan (2), Turkey (2), USA (2), Chile (1), Romania (1), Thailand (1), and The Netherlands (1); the ?gures in parentheses stand for the number of authors from the corresponding country.

This comprehensive reference text discusses the fundamental concepts of artificial intelligence and its applications in a single volume. Artificial Intelligence: Fundamentals and Applications presents a detailed discussion of basic aspects and ethics in the field of artificial intelligence and its applications in areas, including electronic devices and systems, consumer electronics, automobile engineering, manufacturing, robotics and automation, agriculture, banking, and predictive analysis. Aimed at senior undergraduate and graduate students in the field of electrical

engineering, electronics engineering, manufacturing engineering, pharmacy, and healthcare, this text: Discusses advances in artificial intelligence and its applications. Presents the predictive analysis and data analysis using artificial intelligence. Covers the algorithms and pseudo-codes for different domains. Discusses the latest development of artificial intelligence in the field of practical speech recognition, machine translation, autonomous vehicles, and household robotics. Covers the applications of artificial intelligence in fields, including pharmacy and healthcare, electronic devices and systems, manufacturing, consumer electronics, and robotics.

The focus of prerational intelligence is on the way animals and artificial systems utilize information about their surroundings in order to behave intelligently; the premise is that logic and symbolic reasoning are neither necessary nor, possibly, sufficient. Experts in the fields of biology, psychology, robotics, AI, mathematics, engineering, computer science, and philosophy review the evidence that intelligent behaviour can arise in systems of simple agents interacting according to simple rules; that self-organization and interaction with the environment are critical; and that quick approximations may replace logical analyses. It is argued that a better understanding of the intelligence inherent in procedure like those illustrated will eventually shed light on how rational intelligence is realised in humans. Readership: Scientifically literate general readers and scientists in all fields interested in understanding and duplicating biological intelligence.

Artificial intelligence: A Modern Approach, 3e,is ideal for one or two-semester, undergraduate or graduate-level courses in Artificial Intelligence. It is also a valuable resource for computer professionals, linguists, and cognitive scientists interested in artificial intelligence. The revision of this best-

selling text offers the most comprehensive, up-to-date introduction to the theory and practice of artificial intelligence. This book constitutes the refereed proceedings of the 17th Australian Conference on Artificial Intelligence, AI 2004, held in Cairns, Australia, in December 2004. The 78 revised full papers and 62 revised short papers presented were carefully reviewed and selected from 340 submissions. The papers are organized in topical sections on agents; biomedical applications; computer vision, image processing, and pattern recognition; ontologies, knowledge discovery and data mining; natural language and speech processing; problem solving and reasoning; robotics; and soft computing. Keeping the maths to a minimum, Negnevitsky explains the principles of AI, demonstrates how systems are built, what they are useful for and how to choose the right tool for the job.

Given the exponential growth of Artificial Intelligence (AI) over the past few decades, AI and its related applications have become part of daily life in ways that we could never have dreamt of only a century ago. Our routines have been changed beyond measure by robotics and AI, which are now used in a vast array of services. Though AI is still in its infancy, we have already benefited immensely. This book introduces readers to basic Artificial Intelligence concepts, and helps them understand the relationship between AI and daily life. In the interest of clarity, the content is divided into four major parts. Part I (AI Concepts) presents fundamental concepts of and information on AI; while Part II (AI Technology) introduces readers to the five core AI Technologies that provide the building blocks for various AI applications, namely: Machine Learning (ML), Data Mining (DM), Computer Vision (CV), Natural Languages Processing (NLP), and Ontology-based Search Engine (OSE). In turn, Part III (AI Applications) reviews major contemporary

applications that are impacting our ways of life, working styles and environment, ranging from intelligent agents and robotics to smart campus and smart city projects. Lastly, Part IV (Beyond AI) addresses related topics that are vital to the future development of AI. It also discusses a number of critical issues, such as AI ethics and privacy, the development of a conscious mind, and autonomous robotics in our daily lives.

This is primarily a business book that discusses the research and associated practical application of artificial intelligence (AI) and machine learning (ML) in order to achieve business optimization (BO). AI comprises a wide range of technologies, databases, algorithms, and devices. This book aims for a holistic approach to AI by focusing on developing business strategies that will not only automate but also optimize business functions, processes, and people's behaviors. Artificial Intelligence for Business Optimization: Research and Applications explores AI and ML from a business viewpoint with the key purpose of enhancing customer value. It applies research methods and fundamentals from a practitioner's viewpoint and incorporates discussions around risks and changes associated with the utilization of AI in business. Furthermore, governance risks, privacy, and security are also addressed in this book to ensure compliance with AI/ML applications. Readers should find direct and practical applications of the discussions in this book quite useful in their work environment. Researchers will find many ideas to further explore the applications of AI to business.

Deep learning methods are achieving state-of-the-art results on challenging machine learning problems such as describing photos and translating text from one language to another. In this new laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about natural language processing. Using clear explanations, standard

Python libraries and step-by-step tutorial lessons you will discover what natural language processing is, the promise of deep learning in the field, how to clean and prepare text data for modeling, and how to develop deep learning models for your own natural language processing projects.

This book constitutes the refereed proceedings of the Second Hellenic Conference on Artificial Intelligence, SETN 2002, held in Thessaloniki, Greece, in April 2002. The 42 revised full papers presented together with two invited contributions were carefully reviewed and selected for inclusion in the book. The papers are organized in topical sections on knowledge representation and reasoning, logic programming and constraint satisfaction, planning and scheduling, natural language processing, human-computer interaction, machine learning, intelligent Internet and multiagent systems, and intelligent applications.

This book discusses in detail the great historical and social significance of the development of Artificial Intelligence (AI). It consists of seven chapters, each focusing on a specific issue related to AI, such as ethical principles, legal regulations, education, employment and security. Adopting a multidisciplinary approach, it appeals to wide readership, ranging from experts and government officials to the general public.

Data Mining: Practical Machine Learning Tools and Techniques, Third Edition, offers a thorough

grounding in machine learning concepts as well as practical advice on applying machine learning tools and techniques in real-world data mining situations. This highly anticipated third edition of the most acclaimed work on data mining and machine learning will teach you everything you need to know about preparing inputs, interpreting outputs, evaluating results, and the algorithmic methods at the heart of successful data mining. Thorough updates reflect the technical changes and modernizations that have taken place in the field since the last edition, including new material on Data Transformations, Ensemble Learning, Massive Data Sets, Multi-instance Learning, plus a new version of the popular Weka machine learning software developed by the authors. Witten, Frank, and Hall include both tried-and-true techniques of today as well as methods at the leading edge of contemporary research. The book is targeted at information systems practitioners, programmers, consultants, developers, information technology managers, specification writers, data analysts, data modelers, database R&D professionals, data warehouse engineers, data mining professionals. The book will also be useful for professors and students of upper-level undergraduate and graduate-level data mining and machine learning courses who want to incorporate data mining as part of their data management knowledge base and expertise.

Provides a thorough grounding in machine learning concepts as well as practical advice on applying the tools and techniques to your data mining projects Offers concrete tips and techniques for performance improvement that work by transforming the input or output in machine learning methods Includes downloadable Weka software toolkit, a collection of machine learning algorithms for data mining tasks—in an updated, interactive interface. Algorithms in toolkit cover: data pre-processing, classification, regression, clustering, association rules, visualization

Graphics in this book are printed in black and white. Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—scikit-learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine

learning landscape, particularly neural nets Use scikit-learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets Apply practical code examples without acquiring excessive machine learning theory or algorithm details

Use Unity 2017 to create fun and unbelievable AI entities in your games with A*, Fuzzy logic and NavMesh Key Features Explore the brand-new Unity 2017 features that makes implementing Artificial Intelligence in your game easier than ever Use fuzzy logic concepts in your AI decision-making to make your characters more engaging Build exciting and richer games by mastering advanced Artificial Intelligence concepts such as Neural Networks Book Description Unity 2017 provides game and app developers with a variety of tools to implement Artificial Intelligence. Leveraging these tools via Unity's API or built-in features allows limitless possibilities when it comes to creating your game's worlds and characters. This third edition with Unity will help you break down Artificial Intelligence into simple concepts to give you a fundamental

understanding of the topic to build upon. Using a variety of examples, the book then takes those concepts and walks you through actual implementations designed to highlight key concepts, and features related to game AI in Unity 5. Further on you will learn to distinguish the state machine pattern and implement one of your own. This is followed by learning how to implement a basic sensory system for your AI agent and coupling it with a Finite State Machine (FSM). Next you'll learn how to use Unity's built-in NavMesh feature and implement your own A* pathfinding system. You will then learn how to implement simple flocks and crowd's dynamics, key AI concepts. Moving on, you will learn how to implement a behavior tree through a game-focused example. Lastly, you'll combine fuzzy logic concepts with state machines and apply all the concepts in the book to build a simple tank game. What you will learn Understand the basic terminology and concepts in game AI Explore advanced AI Concepts such as Neural Networks Implement a basic finite state machine using state machine behaviors in Unity 2017 Create sensory systems for your AI and couple it with a Finite State Machine Wok with Unity 2017's built-in NavMesh features in your game Build believable and highly-efficient artificial flocks and crowds Create a basic behavior tree to drive a character's actions Who this book is for This book is intended for Unity

developers with a basic understanding of C# and the Unity editor. Whether you're looking to build your first game or are looking to expand your knowledge as a game programmer, you will find plenty of exciting information and examples of game AI in terms of concepts and implementation.

This book covers anaphora resolution for the English language from a linguistic and computational point of view. First, a definition of anaphors that applies to linguistics as well as information technology is given. On this foundation, all types of anaphors and their characteristics for English are outlined. To examine how frequent each type of anaphor is, a corpus of different hypertexts has been established and analysed with regard to anaphors. The most frequent type are non-finite clause anaphors - a type which has never been investigated so far. Therefore, the potential of non-finite clause anaphors are further explored with respect to anaphora resolution. After presenting the fundamentals of computational anaphora resolution and its application in text retrieval, rules for resolving non-finite clause anaphors are established. Therefore, this book shows that a truly interdisciplinary approach can achieve results which would not have been possible otherwise. Open Access: In July 2019, this volume was retroactively turned into an Open Access publication thanks to the support of the Fachinformationsdienst

Linguistik.https://www.linguistik.de/
The theory and practice of AI and ML in marketing saving time, money
This is the first text on pattern recognition to present the Bayesian viewpoint, one that has become increasing popular in the last five years. It presents approximate inference algorithms that permit fast approximate answers in situations where exact answers are not feasible. It provides the first text to use graphical models to describe probability distributions when there are no other books that apply graphical models to machine learning. It is also the first four-color book on pattern recognition. The book is suitable for courses on machine learning, statistics, computer science, signal processing, computer vision, data mining, and bioinformatics. Extensive support is provided for course instructors, including more than 400 exercises, graded according to difficulty. Example solutions for a subset of the exercises are available from the book web site, while solutions for the remainder can be obtained by instructors from the publisher.
This volume presents state-of-the-art tools and techniques for automatically detecting, diagnosing, and predicting the effects of adverse events in an engineered system. It emphasizes the importance of these techniques in managing the intricate interactions within and between engineering systems to maintain a high degree of reliability. Reflecting the

interdisciplinary nature of the field, the book explains how the fundamental algorithms and methods of both physics-based and data-driven approaches effectively address systems health management in application areas such as data centers, aircraft, and software systems.

For one or two-semester, undergraduate or graduate-level courses in Artificial Intelligence. The long-anticipated revision of this best-selling text offers the most comprehensive, up-to-date introduction to the theory and practice of artificial intelligence. View chapters 3 and 4 from the Third Edition. Artificial Intelligence: A Modern Approach, 3e is available to purchase as an eText for your Kindle™, NOOK™, and the iPhone®/iPad®. You can also purchase the eText for 180 days through CourseSmart http://www.mypearsonstore.com/bookstore/product.asp?isbn=0 136067336

This book constitutes the refereed proceedings of the 13th International Conference on Artificial General Intelligence, AGI 2020, held in St. Petersburg, Russia, in September 2020. The 30 full papers and 8 short papers presented in this book were carefully reviewed and selected from 60 submissions. The papers cover topics such as AGI architectures, artificial creativity and AI safety, transfer learning, AI unification and benchmarks for AGI.

This book provides a perspective on the application

of machine learning-based methods in knowledge discovery from natural languages texts. By analysing various data sets, conclusions which are not normally evident, emerge and can be used for various purposes and applications. The book provides explanations of principles of time-proven machine learning algorithms applied in text mining together with step-by-step demonstrations of how to reveal the semantic contents in real-world datasets using the popular R-language with its implemented machine learning algorithms. The book is not only aimed at IT specialists, but is meant for a wider audience that needs to process big sets of text documents and has basic knowledge of the subject, e.g. e-mail service providers, online shoppers, librarians, etc. The book starts with an introduction to text-based natural language data processing and its goals and problems. It focuses on machine learning, presenting various algorithms with their use and possibilities, and reviews the positives and negatives. Beginning with the initial data pre-processing, a reader can follow the steps provided in the R-language including the subsuming of various available plug-ins into the resulting software tool. A big advantage is that R also contains many libraries implementing machine learning algorithms, so a reader can concentrate on the principal target without the need to implement the details of the algorithms her- or himself. To make sense of the

results, the book also provides explanations of the algorithms, which supports the final evaluation and interpretation of the results. The examples are demonstrated using realworld data from commonly accessible Internet sources.

This book constitutes the refereed proceedings of the 5th European Conference on Research and Advanced Technology for Digital Libraries, ECDL 2001, held in Darmstadt, Germany, in September 2001. The 38 revised full papers presented were carefully reviewed and selected from 79 submissions. The papers are organized in topical sections on user modeling; digitization, interpretation, and annotation of documents; knowledge management; data and metadata models; integration in user communities; information retrieval and filtering; multimedia system; and multilinguality.

"To design future networks that are worthy of society's trust, we must put the 'discipline' of computer networking on a much stronger foundation. This book rises above the considerable minutiae of today's networking technologies to emphasize the long-standing mathematical underpinnings of the field." –Professor Jennifer Rexford, Department of Computer Science, Princeton University "This book is exactly the one I have been waiting for the last couple of years. Recently, I decided most students were already very familiar with the way the net works

but were not being taught the fundamentals–the math. This book contains the knowledge for people who will create and understand future communications systems." –Professor Jon Crowcroft, The Computer Laboratory, University of Cambridge The Essential Mathematical Principles Required to Design, Implement, or Evaluate Advanced Computer Networks Students, researchers, and professionals in computer networking require a firm conceptual understanding of its foundations. Mathematical Foundations of Computer Networking provides an intuitive yet rigorous introduction to these essential mathematical principles and techniques. Assuming a basic grasp of calculus, this book offers sufficient detail to serve as the only reference many readers will need. Each concept is described in four ways: intuitively; using appropriate mathematical notation; with a numerical example carefully chosen for its relevance to networking; and with a numerical exercise for the reader. The first part of the text presents basic concepts, and the second part introduces four theories in a progression that has been designed to gradually deepen readers' understanding. Within each part, chapters are as self-contained as possible. The first part covers probability; statistics; linear algebra; optimization; and signals, systems, and transforms. Topics range from Bayesian networks to hypothesis testing, and eigenvalue

computation to Fourier transforms. These preliminary chapters establish a basis for the four theories covered in the second part of the book: queueing theory, game theory, control theory, and information theory. The second part also demonstrates how mathematical concepts can be applied to issues such as contention for limited resources, and the optimization of network responsiveness, stability, and throughput.

Copyright: 76824522916fe1446f85fda486ddb810