

Artificial Intelligence Deepak Khemani Book

An Integrated Approach to Product Development Reliability Engineering presents an integrated approach to the design, engineering, and management of reliability activities throughout the life cycle of a product, including concept, research and development, design, manufacturing, assembly, sales, and service. Containing illustrative guides that include worked problems, numerical examples, homework problems, a solutions manual, and class-tested materials, it demonstrates to product development and manufacturing professionals how to distribute key reliability practices throughout an organization. The authors explain how to integrate reliability methods and techniques in the Six Sigma process and Design for Six Sigma (DFSS). They also discuss relationships between warranty and reliability, as well as legal and liability issues. Other topics covered include: Reliability engineering in the 21st Century Probability life distributions for reliability analysis Process control and process capability Failure modes, mechanisms, and effects analysis Health monitoring and prognostics Reliability tests and reliability estimation Reliability Engineering provides a comprehensive list of references on the topics covered in each chapter. It is an invaluable resource for those interested in gaining fundamental knowledge of the practical aspects of reliability in design, manufacturing, and testing. In addition, it is useful for implementation and management of reliability programs.

Normal adults do not have any difficulty in recognizing their homes. But can artificial systems do in the same way as humans? This book collects interdisciplinary evidences and presents an answer from the perspective of computing, namely, the theory of cognitive prism. To recognize an environment, an intelligent system only needs to classify objects, structures them based on the connection relation (not through measuring!), subjectively orders the objects, and compares with the target environment, whose knowledge is similarly structured. The intelligent system works, therefore, like a prism: when a beam of light (a scene) reaches (is perceived) to an optical prism (by an intelligent system), some light (objects) is reflected (are neglected), those passed through (the recognized objects) are distorted (are ordered differently). So comes the term 'cognitive prism'! Two fundamental propositions used in the theory can be informally stated as follow: an orientation relation is a kind of distance comparison relation -- you being in front of me means you being nearer to my face than to my other sides; a pair of objects being connected means any object, precisely the space occupied by the object, can be moved to a place where it connects with the pair.

About the Book Compounding was taught to us in school, but some-how we have forgotten about it, or we tend to associate it with Interest on our Investments only. We were taught about the formula of calculating Compound Interest in School, but there's so much more to Compounding than only Interest. Compounding touches and affects us, at each and

every stage in our life, while growing up, in our habits, in education, In Sports, in our Relationships, what we eat, drink and what the person we finally become. Everything we do, is either a result of Compounding or is compounding into something that will affect us later on. The book explains it in Simple language even a child can understand. You will be amazed by what it can do to your Wealth, your Health, your relationships and consequently your life.No Wonder Compounding is the 8th wonder! Happy Reading and Happy Compounding!About the AuthorHe is a Financial Planner, having 35 years of experience in Financial Services, with certifications of CFGP and AFGP from AAFM(USA) and AWARD IN FINANCIAL PLANNING from CII(UK).Having started out as a distributor of Financial Products in 1984, he is considered an Influencer in the Personal Finance space. Having qualified to be a LIFE MEMBER of the Prestigious MILLOIN DOLLAR ROUND TABLE-MDRT(USA), a distinction, very few Insurance Advisors across the world are able to achieve, he is also a life member of the CHAIRMAN'S CLUB, LIC OF INDIA.He has the unique distinction of putting out ONE, 2-minute video on a Personal Finance topic, The 2 Minute Money and Personal Finance Show, every day, on his YouTube channel. At the time of writing this book there were 312 videos already out, which he intends to do continuously for 365 days at least!

Constraint satisfaction is a simple but powerful tool. Constraints identify the impossible and reduce the realm of possibilities to effectively focus on the possible, allowing for a natural declarative formulation of what must be satisfied, without expressing how. The field of constraint reasoning has matured over the last three decades with contributions from a diverse community of researchers in artificial intelligence, databases and programming languages, operations research, management science, and applied mathematics. Today, constraint problems are used to model cognitive tasks in vision, language comprehension, default reasoning, diagnosis, scheduling, temporal and spatial reasoning. In Constraint Processing, Rina Dechter, synthesizes these contributions, along with her own significant work, to provide the first comprehensive examination of the theory that underlies constraint processing algorithms. Throughout, she focuses on fundamental tools and principles, emphasizing the representation and analysis of algorithms. ·Examines the basic practical aspects of each topic and then tackles more advanced issues, including current research challenges ·Builds the reader's understanding with definitions, examples, theory, algorithms and complexity analysis ·Synthesizes three decades of researchers work on constraint processing in AI, databases and programming languages, operations research, management science, and applied mathematics

These original contributions provide a current sampling of AI approaches to problems of biological significance; they are the first to treat the computational needs of the biology community hand-in-hand with appropriate advances in artificial intelligence. The enormous amount of data generated by the Human Genome Project and other large-scale biological

research has created a rich and challenging domain for research in artificial intelligence. These original contributions provide a current sampling of AI approaches to problems of biological significance; they are the first to treat the computational needs of the biology community hand-in-hand with appropriate advances in artificial intelligence. Focusing on novel technologies and approaches, rather than on proven applications, they cover genetic sequence analysis, protein structure representation and prediction, automated data analysis aids, and simulation of biological systems. A brief introductory primer on molecular biology and AI gives computer scientists sufficient background to understand much of the biology discussed in the book. Lawrence Hunter is Director of the Machine Learning Project at the National Library of Medicine, National Institutes of Health.

In this richly illustrated book, deep neural network learning algorithms are explained informally first, followed by detailed mathematical analyses. Written in an informal style, with a comprehensive glossary, tutorial appendices, and further readings, this is an ideal introduction to the algorithmic engines of modern artificial intelligence.

This book provides an in-depth analysis of the current evolutionary machine learning techniques. Discussing the most highly regarded methods for classification, clustering, regression, and prediction, it includes techniques such as support vector machines, extreme learning machines, evolutionary feature selection, artificial neural networks including feed-forward neural networks, multi-layer perceptron, probabilistic neural networks, self-optimizing neural networks, radial basis function networks, recurrent neural networks, spiking neural networks, neuro-fuzzy networks, modular neural networks, physical neural networks, and deep neural networks. The book provides essential definitions, literature reviews, and the training algorithms for machine learning using classical and modern nature-inspired techniques. It also investigates the pros and cons of classical training algorithms. It features a range of proven and recent nature-inspired algorithms used to train different types of artificial neural networks, including genetic algorithm, ant colony optimization, particle swarm optimization, grey wolf optimizer, whale optimization algorithm, ant lion optimizer, moth flame algorithm, dragonfly algorithm, salp swarm algorithm, multi-verse optimizer, and sine cosine algorithm. The book also covers applications of the improved artificial neural networks to solve classification, clustering, prediction and regression problems in diverse fields.

Machine Learning Methods for Planning provides information pertinent to learning methods for planning and scheduling. This book covers a wide variety of learning methods and learning architectures, including analogical, case-based, decision-tree, explanation-based, and reinforcement learning. Organized into 15 chapters, this book begins with an overview of planning and scheduling and describes some representative learning systems that have been developed for these tasks. This text then describes a learning apprentice for calendar management. Other chapters consider the problem of temporal credit assignment and describe tractable classes of problems for which optimal plans can be derived. This book discusses as well how reactive, integrated systems give rise to new requirements and opportunities

for machine learning. The final chapter deals with a method for learning problem decompositions, which is based on an idealized model of efficiency for problem-reduction search. This book is a valuable resource for production managers, planners, scientists, and research workers.

No pleasure lasts long unless there is variety in it. Publilius Syrus, Moral Sayings We've been very fortunate to receive fantastic feedback from our readers during the last four years, since the first edition of How to Solve It: Modern Heuristics was published in 1999. It's heartening to know that so many people appreciated the book and, even more importantly, were using the book to help them solve their problems. One professor, who published a review of the book, said that his students had given the best course reviews he'd seen in 15 years when using our text. There can be hardly any better praise, except to add that one of the book reviews published in a SIAM journal received the best review award as well. We greatly appreciate your kind words and personal comments that you sent, including the few cases where you found some typographical or other errors. Thank you all for this wonderful support.

The concept of Artificial Intelligence (AI) & Machine Learning (ML) has been in practice for over years with the advent of technological progress. Over time, it has blended our lives through nearly every narration of learning, teaching, enjoyment, normal routine operations and what not. The aspect delivers a common understanding of the topics with reference to it making an impact on our lives, with a better framework of technology affecting our lives in particular. Let us look up to science for a change to be brought about in us. Let us create awareness of making technology available to people, in a broader sense. As that happens, people who are responsible need to be told about the use and misuse of the same. As we lead our lives, we come across the fact that AI, Robotics and Learning Machines seem to be the household topic of discussion. Earlier, AI was perceived to be reserved for only 'Geniuses' or 'Researchers' or the 'computer' community, but it very aptly integrates and impacts each and every aspect of our lives. Knowingly or unknowingly, it has become intellectually influential in shaping our thoughts, actions and the day-to-day chores.

Artificial intelligence touches nearly every part of your day. While you may initially assume that technology such as smart speakers and digital assistants are the extent of it, AI has in fact rapidly become a general-purpose technology, reverberating across industries including transportation, healthcare, financial services, and many more. In our modern era, an understanding of AI and its possibilities for your organization is essential for growth and success. Artificial Intelligence Basics has arrived to equip you with a fundamental, timely grasp of AI and its impact. Author Tom Taulli provides an engaging, non-technical introduction to important concepts such as machine learning, deep learning, natural language processing (NLP), robotics, and more. In addition to guiding you through real-world case studies and practical implementation steps, Taulli uses his expertise to expand on the bigger questions that surround AI. These include societal trends, ethics, and future impact AI will have on world governments, company structures, and daily life. Google, Amazon, Facebook, and similar tech giants are far from the only organizations on which artificial intelligence has had—and will continue to have—an incredibly significant result. AI is the present and the future of your business as well as your home life. Strengthening your prowess on the subject will prove invaluable to your preparation for the future of tech, and Artificial Intelligence Basics is the indispensable guide that you've been seeking. What You Will Learn Study the core principles for AI approaches such as machine learning, deep learning, and NLP (Natural Language Processing) Discover the best practices to successfully implement AI by examining case studies including Uber, Facebook, Waymo, UiPath, and Stitch Fix Understand how AI capabilities for robots can improve business Deploy chatbots and Robotic Processing Automation (RPA) to save costs and improve customer service Avoid costly gotchas Recognize ethical concerns and other risk factors of using artificial intelligence Examine the secular

trends and how they may impact your business Who This Book Is For Readers without a technical background, such as managers, looking to understand AI to evaluate solutions.

Dynamic Epistemic Logic is the logic of knowledge change. This book provides various logics to support such formal specifications, including proof systems. Concrete examples and epistemic puzzles enliven the exposition. The book also offers exercises with answers. It is suitable for graduate courses in logic. Many examples, exercises, and thorough completeness proofs and expressivity results are included. A companion web page offers slides for lecturers and exams for further practice.

This bestselling book gives business leaders and executives a foundational education on how to leverage artificial intelligence and machine learning solutions to deliver ROI for your business.

The second edition of a comprehensive introduction to machine learning approaches used in predictive data analytics, covering both theory and practice. Machine learning is often used to build predictive models by extracting patterns from large datasets. These models are used in predictive data analytics applications including price prediction, risk assessment, predicting customer behavior, and document classification. This introductory textbook offers a detailed and focused treatment of the most important machine learning approaches used in predictive data analytics, covering both theoretical concepts and practical applications. Technical and mathematical material is augmented with explanatory worked examples, and case studies illustrate the application of these models in the broader business context. This second edition covers recent developments in machine learning, especially in a new chapter on deep learning, and two new chapters that go beyond predictive analytics to cover unsupervised learning and reinforcement learning.

The 7th Pacific Asia Conference on Knowledge Discovery and Data Mining (PAKDD) was held from April 30 to May 2, 2003 in the Convention and Exhibition Center (COEX), Seoul, Korea. The PAKDD conference is a major forum for academic researchers and industry practitioners in the Pacific Asia region to share original research results and development experiences from different KDD-related areas such as data mining, data warehousing, machine learning, databases, statistics, knowledge acquisition and discovery, data visualization, and knowledge-based systems. The conference was organized by the Advanced Information Technology Research Center (AITrc) at KAIST and the Statistical Research Center for Complex Systems (SRCCS) at Seoul National University. It was sponsored by the Korean Data Mining Society (KDMS), the Korea Information Science Society (KISS), the United States Air Force Office of Scientific Research, the Asian Office of Aerospace Research & Development, and KAIST. It was held with cooperation from ACM's Special Group on Knowledge Discovery and Data Mining (SIGKDD).

This work is based on the experience and notes of the authors while teaching mathematics courses to engineering students at the Indian Institute of Technology, New Delhi. It covers syllabi of two core courses in mathematics for engineering students.

This book constitutes the thoroughly refereed post-conference proceedings of the 20th International Conference on Case-Based Reasoning Research and Development (ICCBR 2012) held in Lyon, France, September 3-6, 2012. The 34 revised

full papers presented were carefully selected from 51 submissions. The presentations and posters covered a wide range of CBR topics of interest to both practitioners and researchers, including foundational issues covering case representation, similarity, retrieval, and adaptation; conversational CBR recommender systems; multi-agent collaborative systems; data mining; time series analysis; Web applications; knowledge management; legal reasoning; healthcare systems and planning and scheduling systems.

Designed as an introductory level textbook on Artificial Neural Networks at the postgraduate and senior undergraduate levels in any branch of engineering, this self-contained and well-organized book highlights the need for new models of computing based on the fundamental principles of neural networks. Professor Yegnanarayana compresses, into the covers of a single volume, his several years of rich experience, in teaching and research in the areas of speech processing, image processing, artificial intelligence and neural networks. He gives a masterly analysis of such topics as Basics of artificial neural networks, Functional units of artificial neural networks for pattern recognition tasks, Feedforward and Feedback neural networks, and Archi-tectures for complex pattern recognition tasks. Throughout, the emphasis is on the pattern processing feature of the neural networks. Besides, the presentation of real-world applications provides a practical thrust to the discussion.

A First Course in Artificial Intelligence McGraw-Hill Education Compounding: The 8th Wonder

The 3-volume set LNAI 12712-12714 constitutes the proceedings of the 25th Pacific-Asia Conference on Advances in Knowledge Discovery and Data Mining, PAKDD 2021, which was held during May 11-14, 2021. The 157 papers included in the proceedings were carefully reviewed and selected from a total of 628 submissions. They were organized in topical sections as follows: Part I: Applications of knowledge discovery and data mining of specialized data; Part II: Classical data mining; data mining theory and principles; recommender systems; and text analytics; Part III: Representation learning and embedding, and learning from data.

Develop real-world applications powered by the latest advances in intelligent systems Key Features Gain real-world contextualization using deep learning problems concerning research and application Get to know the best practices to improve and optimize your machine learning systems and algorithms Design and implement machine intelligence using real-world AI-based examples Book Description This Learning Path offers practical knowledge and techniques you need to create and contribute to machine learning, deep learning, and modern data analysis. You will be introduced to various machine learning and deep learning algorithms from scratch, and show you how to apply them to practical industry challenges using realistic and interesting examples. You will learn to build powerful, robust, and accurate predictive models with the power of TensorFlow, combined with other open-source Python libraries. Throughout the Learning Path,

you'll learn how to develop deep learning applications for machine learning systems. Discover how to attain deep learning programming on GPU in a distributed way. By the end of this Learning Path, you know the fundamentals of AI and have worked through a number of case studies that will help you apply your skills to real-world projects. This Learning Path includes content from the following Packt products: Artificial Intelligence By Example by Denis Rothman Python Deep Learning Projects by Matthew Lamons, Rahul Kumar, and Abhishek Nagaraja Hands-On Artificial Intelligence with TensorFlow by Amir Ziai, Ankit Dixit What you will learn Use adaptive thinking to solve real-life AI case studies Rise beyond being a modern-day factory code worker Understand future AI solutions and adapt quickly to them Master deep neural network implementation using TensorFlow Predict continuous target outcomes using regression analysis Dive deep into textual and social media data using sentiment analysis Who this book is for This Learning Path is for anyone who wants to understand the fundamentals of Artificial Intelligence and implement it practically by devising smart solutions. You will learn to extend your machine learning and deep learning knowledge by creating practical AI smart solutions. Prior experience with Python and statistical knowledge is essential to make the most out of this Learning Path.

"The authors' clear visual style provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." –Tim Urban, author of Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional

nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

One of Mark Cuban's top reads for better understanding A.I. (inc.com, 2021) Your comprehensive entry-level guide to machine learning While machine learning expertise doesn't quite mean you can create your own Turing Test-proof android—as in the movie *Ex Machina*—it is a form of artificial intelligence and one of the most exciting technological means of identifying opportunities and solving problems fast and on a large scale. Anyone who masters the principles of machine learning is mastering a big part of our tech future and opening up incredible new directions in careers that include fraud detection, optimizing search results, serving real-time ads, credit-scoring, building accurate and sophisticated pricing models—and way, way more. Unlike most machine learning books, the fully updated 2nd Edition of *Machine Learning For Dummies* doesn't assume you have years of experience using programming languages such as Python (R source is also included in a downloadable form with comments and explanations), but lets you in on the ground floor, covering the entry-level materials that will get you up and running building models you need to perform practical tasks. It takes a look at the underlying—and fascinating—math principles that power machine learning but also shows that you don't need to be a math whiz to build fun new tools and apply them to your work and study. Understand the history of AI and machine learning Work with Python 3.8 and TensorFlow 2.x (and R as a download) Build and test your own models Use the latest datasets, rather than the worn out data found in other books Apply machine learning to real problems Whether you want to learn for college or to enhance your business or career performance, this friendly beginner's guide is your best introduction to machine learning, allowing you to become quickly confident using this amazing and fast-developing technology that's impacting lives for the better all over the world.

Artificial Intelligence and Intelligent Systems provides a comprehensive coverage of the fundamental concepts and techniques in artificial intelligence. The book discusses current trends in AI and its application to various fields. Intelligent systems such as expert systems, fuzzy systems, artificial neural networks, genetic algorithms, and swarm intelligent systems are discussed in detail with examples to facilitate in-depth understanding of AI. The text emphasizes the solution of real-world problems using the latest AI techniques. Since the ultimate goal of AI is the construction of programs to solve problems, an entire chapter has been devoted to the programming languages used in AI problem solving. Written in a clear and lucid style, this student-friendly book has been specially designed for undergraduate engineering students. With its application oriented approach and inclusion of recent topics, the book would also be useful to postgraduate

students and researchers in this field. Features * Includes real-world examples to illustrate concepts * Contains a separate chapter on programming languages in AI * Includes new topics such as swarm intelligent systems * Explains genetic algorithms and swarm intelligence using examples * Provides numerous illustrations, examples, and end-chapter exercises

Crucial reading for those concerned with education and school reform.

"XQuery Kick Start" delivers a concise introduction to the XQuery standard, and useful implementation advice for developers needing to put it into practice. The book starts by explaining the role of XQuery in the XML family of specifications, and its relationship with XPath. The authors then explain the specification in detail, describing the semantics and data model, before moving to examples using XQuery to manipulate XML databases and document storage systems. Later chapters discuss Java implementations of XQuery and development tools that facilitate the development of Web sites with XQuery. This book is up to date with the latest XQuery specifications, and includes coverage of new features for extending the XQuery language.

Problem-solving strategies and the nature of Heuristic information. Heuristics and problem representations. Basic Heuristic-Search procedures. Formal properties of Heuristic methods. Heuristics viewed as information provided by simplified models. Performance analysis of Heuristic methods. Abstract models for quantitative performance analysis. Complexity versus precision of admissible Heuristics. Searching with nonadmissible Heuristics. Game-playing programs. Strategies and models for game-playing programs. Performance analysis for game-searching strategies. Decision quality in game searching. Bibliography. Index.

This book constitutes the refereed proceedings of the 7th European Conference on Case-Based Reasoning, ECCBR 2004, held in Madrid, Spain in August/September 2004. The 56 revised full papers presented together with an invited paper and the abstract of an invited talk were carefully reviewed and selected from 85 submissions. All current issues in case-based reasoning, ranging from theoretical and methodological issues to advanced applications in various fields are addressed.

Demystify the complexity of machine learning techniques and create evolving, clever solutions to solve your problems
Key Features Master supervised, unsupervised, and semi-supervised ML algorithms and their implementation Build deep learning models for object detection, image classification, similarity learning, and more Build, deploy, and scale end-to-end deep neural network models in a production environment
Book Description This Learning Path is your complete guide to quickly getting to grips with popular machine learning algorithms. You'll be introduced to the most widely used algorithms in supervised, unsupervised, and semi-supervised machine learning, and learn how to use them in the best

possible manner. Ranging from Bayesian models to the MCMC algorithm to Hidden Markov models, this Learning Path will teach you how to extract features from your dataset and perform dimensionality reduction by making use of Python-based libraries. You'll bring the use of TensorFlow and Keras to build deep learning models, using concepts such as transfer learning, generative adversarial networks, and deep reinforcement learning. Next, you'll learn the advanced features of TensorFlow 1.x, such as distributed TensorFlow with TF clusters, deploy production models with TensorFlow Serving. You'll implement different techniques related to object classification, object detection, image segmentation, and more. By the end of this Learning Path, you'll have obtained in-depth knowledge of TensorFlow, making you the go-to person for solving artificial intelligence problems This Learning Path includes content from the following Packt products: Mastering Machine Learning Algorithms by Giuseppe Bonaccorso Mastering TensorFlow 1.x by Armando Fandango Deep Learning for Computer Vision by Rajalingappaa Shanmugamani What you will learn Explore how an ML model can be trained, optimized, and evaluated Work with Autoencoders and Generative Adversarial Networks Explore the most important Reinforcement Learning techniques Build end-to-end deep learning (CNN, RNN, and Autoencoders) models Who this book is for This Learning Path is for data scientists, machine learning engineers, artificial intelligence engineers who want to delve into complex machine learning algorithms, calibrate models, and improve the predictions of the trained model. You will encounter the advanced intricacies and complex use cases of deep learning and AI. A basic knowledge of programming in Python and some understanding of machine learning concepts are required to get the best out of this Learning Path.

Search has been vital to artificial intelligence from the very beginning as a core technique in problem solving. The authors present a thorough overview of heuristic search with a balance of discussion between theoretical analysis and efficient implementation and application to real-world problems. Current developments in search such as pattern databases and search with efficient use of external memory and parallel processing units on main boards and graphics cards are detailed. Heuristic search as a problem solving tool is demonstrated in applications for puzzle solving, game playing, constraint satisfaction and machine learning. While no previous familiarity with heuristic search is necessary the reader should have a basic knowledge of algorithms, data structures, and calculus. Real-world case studies and chapter ending exercises help to create a full and realized picture of how search fits into the world of artificial intelligence and the one around us. Provides real-world success stories and case studies for heuristic search algorithms Includes many AI developments not yet covered in textbooks such as pattern databases, symbolic search, and parallel processing units Cutting through the hype, a practical guide to using artificial intelligence for business benefits and competitive advantage. In The AI Advantage, Thomas Davenport offers a guide to using artificial intelligence in business. He describes what

technologies are available and how companies can use them for business benefits and competitive advantage. He cuts through the hype of the AI craze—remember when it seemed plausible that IBM's Watson could cure cancer?—to explain how businesses can put artificial intelligence to work now, in the real world. His key recommendation: don't go for the “moonshot” (curing cancer, or synthesizing all investment knowledge); look for the “low-hanging fruit” to make your company more efficient. Davenport explains that the business value AI offers is solid rather than sexy or splashy. AI will improve products and processes and make decisions better informed—important but largely invisible tasks. AI technologies won't replace human workers but augment their capabilities, with smart machines to work alongside smart people. AI can automate structured and repetitive work; provide extensive analysis of data through machine learning (“analytics on steroids”), and engage with customers and employees via chatbots and intelligent agents. Companies should experiment with these technologies and develop their own expertise. Davenport describes the major AI technologies and explains how they are being used, reports on the AI work done by large commercial enterprises like Amazon and Google, and outlines strategies and steps to becoming a cognitive corporation. This book provides an invaluable guide to the real-world future of business AI. A book in the Management on the Cutting Edge series, published in cooperation with MIT Sloan Management Review.

Today, a scientific explanation is not meant to ascribe agency to natural phenomena: we would not say a rock falls because it seeks the center of the earth. Even for living things, in the natural sciences and often in the social sciences, the same is true. A modern botanist would not say that plants pursue sunlight. This has not always been the case, nor, perhaps, was it inevitable. Since the seventeenth century, many thinkers have made agency, in various forms, central to science. *The Restless Clock* examines the history of this principle, banning agency, in the life sciences. It also tells the story of dissenters embracing the opposite idea: that agency is essential to nature. The story begins with the automata of early modern Europe, as models for the new science of living things, and traces questions of science and agency through Descartes, Leibniz, Lamarck, and Darwin, among many others. Mechanist science, Jessica Riskin shows, had an associated theology: the argument from design, which found evidence for a designer in the mechanisms of nature. Rejecting such appeals to a supernatural God, the dissenters sought to naturalize agency rather than outsourcing it to a “divine engineer.” Their model cast living things not as passive but as active, self-making machines. The conflict between passive- and active-mechanist approaches maintains a subterranean life in current science, shaping debates in fields such as evolutionary biology, cognitive science, and artificial intelligence. This history promises not only to inform such debates, but also our sense of the possibilities for what it means to engage in science—and even what it means to be alive.

The two volume set CCIS 1030 and 1031 constitutes the refereed proceedings of the Second International Conference on Computational Intelligence, Communications, and Business Analytics, CICBA 2018, held in Kalyani, India, in July 2018. The 76 revised full papers presented in the two volumes were carefully reviewed and selected from 240 submissions. The papers are organized in topical sections on computational intelligence; signal processing and communications; microelectronics, sensors, and intelligent networks; data science & advanced data analytics; intelligent data mining & data warehousing; and computational forensics (privacy and security).

Knowledge representation is at the very core of a radical idea for understanding intelligence. This book talks about the central concepts of knowledge representation developed over the years. It is suitable for researchers and practitioners in database management, information retrieval, object-oriented systems and artificial intelligence.

[Copyright: 837cb3bd69e5ec1225522581c6fc89fe](https://www.researchgate.net/publication/331111111)