

Fp Dnn An Automated Framework For Mapping Deep Neural

Explores the principles of automatic partial evaluation, provides simple and complete algorithms, and demonstrates via examples that specialization can increase efficiency. Covers partial evaluation of programming languages from C and Prolog to Scheme and the lambda calculus. For researchers, programmers, and students in advanced programming languages.

This book captures the state of the art research in the area of malicious code detection, prevention and mitigation. It contains cutting-edge behavior-based techniques to analyze and detect obfuscated malware. The book analyzes current trends in malware activity online, including botnets and malicious code for profit, and it proposes effective models for detection and prevention of attacks using. Furthermore, the book introduces novel techniques for creating services that protect their own integrity and safety, plus the data they manage.

This book gathers selected high-quality papers presented at the International Conference on Machine Learning and Computational Intelligence (ICMLCI-2019), jointly organized by Kuninging University of Science and Technology and the Interference Research Network, Bhubaneswar, India, from April 6 to 7, 2019. Addressing virtually all aspects of intelligent systems, soft computing and machine learning, the topics covered include: prediction; data mining; information retrieval; game playing; robotics; learning methods; pattern visualization; automated knowledge acquisition; fuzzy, stochastic and probabilistic computing; neural computing; big data; social networks and applications of soft computing in various areas. An Introduction to Variational Autoencoders provides a quick summary for of a topic that has become an important tool in modern-day deep learning techniques.

The book Green, Energy-Efficient and Sustainable Networks provides insights and solutions for a range of problems in the field of obtaining greener, energy-efficient, and sustainable networks. The book contains the outcomes of the Special Issue on " Green, Energy-Efficient and Sustainable Networks " of the Sensors journal. Seventeen high-quality papers published in the Special Issue have been collected and reproduced in this book, demonstrating significant achievements in the field.

Among the published papers, one paper is an editorial and one is a review, while the remaining 15 works are research articles. The published papers are self-contained peer-reviewed scientific works that are authored by more than 75 different contributors with both academic and industry backgrounds. The editorial paper gives an introduction to the problem of information and communication technology (ICT) energy consumption and greenhouse gas emissions, presenting the state of the art and future trends in terms of improving the energy-efficiency of wireless networks and data centers, as the major energy consumers in the ICT sector. In addition, the published articles aim to improve energy efficiency in the fields of software-defined networking, Internet of things, machine learning, authentication, energy harvesting, wireless relay systems, routing metrics, wireless sensor networks, device-to-device communications, heterogeneous wireless networks, and image sensing. The last paper is a review that gives a detailed overview of energy-efficiency improvements and methods for the implementation of fifth-generation networks and beyond. This book can serve as a source of information in industrial, teaching, and/or research and development activities. The book is a valuable source of information, since it presents recent advances in different fields related to greening and improving the energy-efficiency and sustainability of those ICTs particularly addressed in this book.

This book constitutes the proceedings of the 21st International Conference on Speech and Computer, SPECOM 2019, held in Istanbul, Turkey, in August 2019. The 57 papers presented were carefully reviewed and selected from 86 submissions. The papers present current research in the area of computer speech processing including audio signal processing, automatic speech recognition, speaker recognition, computational paralinguistics, speech synthesis, sign language and multimodal processing, and speech and language resources.

The book offers a timely snapshot of neural network technologies as a significant component of big data analytics platforms. It promotes new advances and research directions in efficient and innovative algorithmic approaches to analyzing big data (e.g. deep networks, nature-inspired and brain-inspired algorithms), implementations on different computing platforms (e.g. neuromorphic, graphics processing units (GPUs), clouds, clusters); and big data analytics applications to solve real-world problems (e.g. weather prediction, transportation, energy management). The book, which reports on the second edition of the INNS Conference on Big Data, held on October 23 – 25, 2016, in Thessaloniki, Greece, depicts an interesting collaborative adventure of neural networks with big data and other learning technologies.

During the 1980s and early 1990s there was significant work in the design and implementation of hardware neurocomputers. Nevertheless, most of these efforts may be judged to have been unsuccessful, at no time have hardware neurocomputers been in wide use. This lack of success may be largely attributed to the fact that earlier work was almost entirely aimed at developing custom neurocomputers, based on ASIC technology, but for such niche - as this technology was never sufficiently developed or competitive enough to justify large-scale adoption. On the other hand, gate-arrays of the period mentioned were never large enough nor fast enough for serious artificial-neur- network (ANN) applications. But technology has now improved: the capacity and performance of current FPGAs are such that they present a much more realistic alternative. Consequently neurocomputers based on FPGAs are now a much more practical proposition than they have been in the past. This book summarizes some work towards this goal and consists of 12 papers that were selected, after review, from a number of submissions. The book is nominally divided into three parts: Chapters 1 through 4 deal with foundational issues; Chapters 5 through 11 deal with a variety of implementations; and Chapter 12 looks at the lessons learned from a large-scale project and also reconsiders design issues in light of current and future technology.

[Proceedings of the 2nd INNS Conference on Big Data, October 23-25, 2016, Thessaloniki, Greece](#)

[Machine Learning with R](#)

[8th International Conference, TPNC 2019, Kingston, ON, Canada, December 9 – 11, 2019, Proceedings](#)

[Proceedings of ICWSNLUCA 2021](#)

[Advances in Computing and Data Sciences](#)

[Applied Reconfigurable Computing, Architectures, Tools, and Applications](#)

[Medical Image Registration](#)

[Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow](#)

[Speech and Computer](#)

[An Introduction to Variational Autoencoders](#)

[Efficient Processing of Deep Neural Networks](#)

[17th International Symposium, ARC 2021, Virtual Event, June 29 – 30, 2021, Proceedings](#)

[A comprehensive introduction to the tools, techniques and applications of convex optimization.](#)

This book reviews the state of the art in deep learning approaches to high-performance robust disease detection, robust and accurate organ segmentation in medical image computing (radiological and pathological imaging modalities), and the construction and mining of large-scale radiology databases. It particularly focuses on the application of convolutional neural networks, and on recurrent neural networks like LSTM, using numerous practical examples to complement the theory. The book's chief features are as follows: It highlights how deep neural networks can be used to address new questions and protocols, and to tackle current challenges in medical image computing; presents a comprehensive review of the latest research and literature; and describes a range of different methods that employ deep learning for object or landmark detection tasks in 2D and 3D medical imaging. In addition, the book examines a broad selection of techniques for semantic segmentation using deep learning principles in medical imaging; introduces a novel approach to text and image deep embedding for a large-scale chest x-ray image database; and discusses how deep learning relational graphs can be used to organize a stable collection of radiology findings from real clinical practice, allowing semantic similarity-based retrieval. The intended reader of this edited book is a professional engineer, scientist or a graduate student who is able to comprehend general concepts of image processing, computer vision and medical image analysis. They can apply computer science and mathematical principles into problem solving practices. It may be necessary to have a certain level of familiarity with a number of more advanced subjects: image formation and enhancement, image understanding, visual recognition in medical applications, statistical learning, deep neural networks, structured prediction and image segmentation.

In the past decade, social media has become increasingly popular for news consumption due to its easy access, fast dissemination, and low cost. However, social media also enables the wide propagation of "fake news," i.e., news with intentionally false information. Fake news on social media can have significant negative societal effects. Therefore, fake news detection on social media has recently become an emerging research area that is attracting tremendous attention. This book, from a data mining perspective, introduces the basic concepts and characteristics of fake news across disciplines, reviews representative fake news detection methods in a principled way, and illustrates challenging issues of fake news detection on social media. In particular, we discussed the value of news content and social context, and important extensions to handle early detection, weakly-supervised detection, and explainable detection. The concepts, algorithms, and methods described in this lecture can help harness the power of social media to build effective and intelligent fake news detection systems. This book is an accessible introduction to the study of detecting fake news on social media. It is an essential reading for students, researchers, and practitioners to understand, manage, and excel in this area. This book is supported by additional materials, including lecture slides, the complete set of figures, key references, datasets, tools used in this book, and the source code of representative algorithms.

Image registration is the process of systematically placing separate images in a common frame of reference so that the information they contain can be optimally integrated or compared. This is becoming the central tool for image analysis, understanding, and visualization in both medical and scientific applications. Medical Image Registration provid

This book constitutes the refereed proceedings of the 4th International Conference on Smart Computing and Communications, SmartCom 2019, held in Birmingham, UK, in October 2019. The 40 papers presented in this volume were carefully reviewed and selected from 286 submissions. They focus on both smart computing and communications fields and aimed to collect recent academic work to improve the research and practical application in the field.

This book constitutes the proceedings of the 4th International Conference on Computational Intelligence, Cyber Security, and Computational Models, ICC3 2019, which was held in Coimbatore, India, in December 2019. The 9 papers presented in this volume were carefully reviewed and selected from 38 submissions. They were organized in topical sections named: computational intelligence; cyber security; and computational models.

Applied Reconfigurable Computing, Architectures, Tools, and Applications17th International Symposium, ARC 2021, Virtual Event, June 29-30, 2021, ProceedingsSpringer NatureApplied Reconfigurable Computing, Architectures, Tools, and Applications16th International Symposium, ARC 2020, Toledo, Spain, April 1-3, 2020, ProceedingsSpringer Nature

[Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers](#)

[Topology Optimization](#)

[Towards Ubiquitous Low-power Image Processing Platforms](#)

[21st International Conference, SPECOM 2019, Istanbul, Turkey, August 20-25, 2019, Proceedings](#)

[IoT and Analytics for Sensor Networks](#)

[Smart Computing and Communication](#)

[Advances in Big Data](#)

[Elements of Artificial Neural Networks](#)

[FPGA Implementations of Neural Networks](#)

[Intelligent Data Engineering and Automated Learning - IDEAL 2018](#)

[Theory, Methods, and Applications](#)

[Intel Sgx Security Analysis and Mit Sanctum Architecture](#)

[Deep Learning and Convolutional Neural Networks for Medical Imaging and Clinical Informatics](#)

This book constitutes the refereed joint proceedings of the Third International Workshop on Deep Learning in Medical Image Analysis, DLMIA 2017, and the 6th International Workshop on Multimodal Learning for Clinical Decision Support, ML-CDS 2017, held in conjunction with the 20th International Conference on Medical Imaging and Computer-Assisted Intervention, MICCAI 2017, in Québec City, QC, Canada, in September 2017. The 38 full papers presented at DLMIA 2017 and the 5 full papers presented at ML-CDS 2017 were carefully reviewed and selected. The DLMIA papers focus on the design and use of deep learning methods in medical imaging. The ML-CDS papers discuss new techniques of multimodal mining/retrieval and

[Mobile devices, such as smart phones, have achieved computing and networking capabilities comparable to traditional personal computers. Their successful consumerization has also become a source of pain for adopting users and organizations. In particular, the widespread presence of information-stealing applications and other types of mobile malware raises](#)

[substantial security and privacy concerns. Android Malware presents a systematic view on state-of-the-art mobile malware that targets the popular Android mobile platform. Covering key topics like the Android malware history, malware behavior and classification, as well as, possible defense techniques.](#)

[Deep learning algorithms have brought a revolution to the computer vision community by introducing non-traditional and efficient solutions to several image-related problems that had long remained unsolved or partially addressed. This book presents a collection of eleven chapters where each individual chapter explains the deep learning principles of a specific topic,](#)

[introduces reviews of up-to-date techniques, and presents research findings to the computer vision community. The book covers a broad scope of topics in deep learning concepts and applications such as accelerating the convolutional neural network inference on field-programmable gate arrays, fire detection in surveillance applications, face recognition, action and activity recognition, semantic segmentation for autonomous driving, aerial imagery registration, robot vision, tumor detection, and skin lesion segmentation as well as skin melanoma classification. The content of this book has been organized such that each chapter can be read independently from the others. The book is a valuable companion for researchers, for postgraduate and possibly senior undergraduate students who are taking an advanced course in related topics, and for those who are interested in deep learning with applications in computer vision, image processing, and pattern recognition.](#)

[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 is about the Zynq-7000 All Programmable System on Chip, the family of devices from Xilinx that combines an application-grade ARM Cortex-A9 processor with traditional FPGA logic fabric. Catering for both new and experienced readers, it covers fundamental issues in an accessible way, starting with a clear overview of the device architecture, and an introduction to the design tools and processes for developing a Zynq SoC. Later chapters progress to more advanced topics such as embedded systems development, IP block design and operating systems. Maintaining a 'real-world' perspective, the book also compares Zynq with other device alternatives, and considers end-user applications. The Zynq Book is](#)

[accompanied by a set of practical tutorials hosted on a companion website. These tutorials will guide the reader through first steps with Zynq, following on to a complete, audio-based embedded systems design.](#)

[This book constitutes the refereed proceedings of the 8th International Conference on Theory and Practice of Natural Computing, TPNC 2019, held in Kingston, ON, Canada, in December 2019. The 15 full papers presented in this book, together with two invited talk, were carefully reviewed and selected from 38 submissions. The papers are organized in topical sections named: Applications of Natural Computing; Evolutionary Computation; Genetic Algorithms, Swarm Intelligence, and Heuristics; Quantum Computing and Information.](#)

[This book constitutes the proceedings of the 16th International Symposium on Applied Reconfigurable Computing, ARC 2020, held in Toledo, Spain, in April 2020. The 18 full papers and 11 poster presentations presented in this volume were carefully reviewed and selected from 40 submissions. The papers are organized in the following topical sections: design methods & tools; design space exploration & estimation techniques; high-level synthesis; architectures; applications.](#)

[This book constitutes the post-conference proceedings of the 4th International Conference on Advances in Computing and Data Sciences, ICACDS 2020, held in Valletta, Malta, in April 2020. * The 46 full papers were carefully reviewed and selected from 354 submissions. The papers are centered around topics like advanced computing, data sciences, distributed systems organizing principles, development frameworks and environments, software verification and validation, computational complexity and cryptography, machine learning theory, database theory, probabilistic representations. * The conference was held virtually due to the COVID-19 pandemic.](#)

[The Zynq Book](#)

[19th International Conference, Madrid, Spain, November 21–23, 2018, Proceedings](#)

[Theory and Practice of Natural Computing](#)

[Partial Evaluation and Automatic Program Generation](#)

[Convex Optimization](#)

[Deep Learning in Computer Vision](#)

[Innovation in Information Systems and Technologies to Support Learning Research](#)

[A Problem-Solver's Guide to Building Real-World Intelligent Systems](#)

[Concepts, Tools, and Techniques to Build Intelligent Systems](#)

[Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support](#)

[Technical Advances of Machine Learning in Healthcare](#)

[Practical Machine Learning with Python](#)

[Neural Network Design](#)

The topology optimization method solves the basic enginee- ring problem of distributing a limited amount of material in a design space. The first edition of this book has become the standard text on optimal design which is concerned with the optimization of structural topology, shape and material. This edition, has been substantially revised and updated to reflect progress made in modelling and computational procedures. It also encompasses a comprehensive and unified description of the state-of-the-art of the so-called material distribution method, based on the use of mathematical programming and finite elements. Applications treated include not only structures but also materials and MEMS.

Deep learning networks are getting smaller. Much smaller. The Google Assistant team can detect words with a model just 14 kilobytes in size—small enough to run on a microcontroller. With this practical book you'll enter the field of TinyML, where deep learning and embedded systems combine to make astounding things possible with tiny devices. Pete Warden and Daniel Situnayake explain how you can train models small enough to fit into any environment. Ideal for software and hardware developers who want to build embedded systems using machine learning, this guide walks you through creating a series of TinyML projects, step-by-step. No machine learning or microcontroller experience is necessary. Build a speech recognizer, a camera that detects people, and a magic wand that responds to gestures Work with Arduino and ultra-low-power microcontrollers Learn the essentials of ML and how to train your own models Train models to understand audio, image, and accelerometer data Explore TensorFlow Lite for Microcontrollers, Google's toolkit for TinyML Debug applications and provide safeguards for privacy and security Optimize latency, energy usage, and model and binary size

This book summarizes the key scientific outcomes of the Horizon 2020 research project TULIPP: Towards Ubiquitous Low-power Image Processing Platforms. The main focus lies on the development of high-performance, energy-efficient embedded systems for the growing range of increasingly complex image processing applications. The holistic TULIPP approach is described in the book, which addresses hardware platforms, programming tools and embedded operating systems. Several of the results are available as open-source hardware/software for the community. The results are evaluated with several use cases taken from real-world applications in key domains such as Unmanned Aerial Vehicles (UAVs), robotics, space and medicine. Discusses the development of high-performance, energy-efficient embedded systems for the growing range of increasingly complex image processing applications; Covers the hardware architecture of embedded image processing systems, novel methods, tools and libraries for programming those systems as well as embedded operating systems to manage those systems; Demonstrates results with several challenging applications, such as medical systems, robotics, drones and automotive.

This book presents high-quality, original contributions (both theoretical and experimental) on software engineering, cloud computing, computer networks & internet technologies, artificial intelligence, information security, and database and distributed computing. It gathers papers presented at ICRIC 2019, the 2nd International Conference on Recent Innovations in Computing, which was held in Jammu, India, in March 2019. This conference series represents a targeted response to the growing need for research that reports on and assesses the practical implications of IoT and network technologies, AI and machine learning, cloud-based e-Learning and big data, security and privacy, image processing and computer vision, and next-generation computing technologies.

This two-volume set LNCS 11314 and 11315 constitutes the thoroughly refereed conference proceedings of the 19th International Conference on Intelligent Data Engineering and Automated Learning, IDEAL 2018, held in Madrid, Spain, in November 2018. The 125 full papers presented were carefully reviewed and selected from 204 submissions. These papers provided a timely sample of the latest advances in data engineering and automated learning, from methodologies, frameworks and techniques to applications. In addition to various topics such as evolutionary algorithms, deep learning neural networks, probabilistic modelling, particle swarm intelligence, big data analytics, and applications in image recognition, regression, classification, clustering, medical and biological modelling and prediction, text processing and social media analysis.

This book includes high-quality research papers presented at the 1st International Conference on Wireless Sensor Networks, Ubiquitous Computing and Applications (ICWSNUCA, 2021), which is held at Gokaraju Rangaraju Institute of Engineering and Technology, Hyderabad, India, during 26-27 February, 2021. This volume includes applications, use-cases, architectures, deployments, and recent advances of wireless sensor networks as well as ubiquitous computing. Different research topics are illustrated in this book, like wireless sensor networks for the Internet of Things; IoT applications for eHealth; smart cities; architectures for WSNs and IoT, WSNs hardware and new devices; low-power wireless technologies; wireless ad hoc sensor networks; routing and data transfer in WSNs; multicast communication in WSNs; security management in WSNs and in IoT systems; and power consumption optimization in WSNs.

This Synthesis Lecture focuses on techniques for efficient data orchestration within DNN accelerators. The End of Moore's Law, coupled with the increasing growth in deep learning and other AI applications has led to the emergence of custom Deep Neural Network (DNN) accelerators for energy-efficient inference on edge devices. Modern DNNs have millions of hyper parameters and involve billions of computations; this necessitates extensive data movement from memory to on-chip processing engines. It is well known that the cost of data movement today surpasses the cost of the actual computation; therefore, DNN accelerators require careful orchestration of data across on-chip compute, network, and memory elements to minimize the number of accesses to external DRAM. The book covers DNN dataflows, data reuse, buffer hierarchies, networks-on-chip, and automated design-space exploration. It concludes with data orchestration challenges with compressed and sparse DNNs and future trends. The target audience is students, engineers, and researchers interested in designing high-performance and low-energy accelerators for DNN inference.

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 Learn techniques for training and scaling deep neural nets

[Android Malware](#)

[Deep Learning for Natural Language Processing](#)

[Computational Intelligence, Cyber Security and Computational Models, Models and Techniques for Intelligent Systems and Automation](#)

[Proceedings of ICRIC 2019](#)

[Principles and Applications](#)

[4th International Conference, ICC3 2019, Coimbatore, India, December 19-21, 2019, Revised Selected Papers](#)

[Secure Processors Part II](#)

[2018 IEEE ACM International Conference on Computer Aided Design \(ICCAD\)](#)

[Data Orchestration in Deep Learning Accelerators](#)

[Proceedings of EMENA-ISTL 2019](#)

[Proceedings of ICMLCI 2019](#)

[Green, Energy-Efficient and Sustainable Networks](#)

[4th International Conference, ICACDS 2020, Valletta, Malta, April 24-25, 2020, Revised Selected Papers](#)

This book provides glimpses into contemporary research in information systems & technology, learning, artificial intelligence (AI), machine learning, and security and how it applies to the real world, but the ideas presented also span the domains of telehealth, computer vision, the role and use of mobile devices, brain–computer interfaces, virtual reality, language and image processing and big data analytics and applications. Great research arises from asking pertinent research questions. This book reveals some of the authors' "beautiful questions" and how they develop the subsequent "what if" and "how" questions, offering readers food for thought and whetting their appetite for further research by the same authors.

This manuscript is the second in a two part survey and analysis of the state of the art in secure processor systems, with a specific focus on remote software attestation and software isolation. The first part established the taxonomy and prerequisite concepts relevant to an examination of the state of the art in trusted remote computation: attested software isolation containers (enclaves). This second part extends Part I's description of Intel's Software Guard Extensions (SGX), an available and documented enclave-capable system, with a rigorous security analysis of SGX as a system for trusted remote computation. This part documents the authors' concerns over the shortcomings of SGX as a secure system and introduces the MIT Sanctum processor developed by the authors: a system designed to offer stronger security guarantees, lend itself better to analysis and formal verification, and offer a more straightforward and complete threat model than the Intel system, all with an equivalent programming model. This two part work advocates a principled, transparent, and wellscrutinized approach to system design, and argues that practical guarantees of privacy and integrity for remote computation are achievable at a reasonable design cost and performance overhead.

Elements of Artificial Neural Networks provides a clearly organized general Introduction, focusing on a broad range of algorithms, for students and others who want to use neural networks rather than simply study them. The authors, who have been developing and team teaching the material in a one-semester course over the past six years, describe most of the basic neural network models (with several detailed worked examples) and discuss the rationale and advantages of the models, as well as their limitations. The approach is practical and open-minded and requires very little mathematical or technical background. Written from a computer science and statistics point of view, the text stresses links to contiguous fields and can easily serve as a first course for students in economics and management. The opening chapter sets the stage, presenting the basic concepts in a clear and objective way and tackling important – yet rarely addressed – questions related to the use of neural networks in practical situations.

Subsequent chapters on supervised learning (single layer and multilayer networks), unsupervised learning, and associative models are structured around classes of problems to which networks can be applied. Applications are discussed along with the algorithms. A separate chapter takes up optimization methods. The most frequently used algorithms, such as backpropagation, are introduced early on, right after perceptions, so that these can form the basis for initiating course projects. Algorithms published as late as 1995 are also included. All of the algorithms are presented using block-structured pseudo-code, and exercises are provided throughout. Software implementing many commonly used neural network algorithms is available at the book's website. Transparency masters, including abbreviated text and figures for the entire book, are available for instructors using the text.

Written as a tutorial to explore and understand the power of R for machine learning. This practical guide that covers all of the need to know topics in a very systematic way. For each machine learning approach, each step in the process is detailed, from preparing the data for analysis to evaluating the results. These steps will build the knowledge you need to apply them to your own data science tasks.Intended for those who want to learn how to use R's machine learning capabilities and gain insight from your data. Perhaps you already know a bit about machine learning, but have never used R; or perhaps you know a little R but are new to machine learning. In either case, this book will get you up and running quickly. It would be helpful to have a bit of familiarity with basic programming concepts, but no prior experience is required.

This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high computational complexity. Therefore, techniques that enable efficient processing of deep neural networks to improve key metrics—such as energy-efficiency, throughput, and latency—without sacrificing accuracy or increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems.

The book includes background on DNN processing; a description and taxonomy of hardware architectural approaches for designing DNN accelerators; key metrics for evaluating and comparing different designs; features of DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as formalization and organization of key concepts from contemporary work that provide insights that may spark new ideas.

Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

[Malware Detection](#)

[TinyML](#)

[Develop Deep Learning Models for your Natural Language Problems](#)

[Detecting Fake News on Social Media](#)

[Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings](#)

[Advances in Machine Learning and Computational Intelligence](#)

[Recent Innovations in Computing](#)

[16th International Symposium, ARC 2020, Toledo, Spain, April 1–3, 2020, Proceedings](#)

[4th International Conference, SmartCom 2019, Birmingham, UK, October 11–13, 2019, Proceedings](#)

[Embedded Processing with the Arm Cortex-A9 on the Xilinx Zynq-7000 All Programmable SoC](#)