

grokking deep learning pdf

grokking deep learning pdf is a highly sought-after resource for individuals aiming to gain a solid understanding of deep learning concepts and techniques. This valuable document provides a comprehensive introduction to deep learning, covering fundamental principles, neural network architectures, and practical applications. Whether you are a beginner or an experienced practitioner, the grokking deep learning pdf offers clear explanations, illustrative examples, and hands-on exercises designed to enhance learning. This article explores the key features of the grokking deep learning pdf, its benefits, and how it can serve as an essential tool for mastering deep learning. Additionally, the article discusses where to find this resource legally and how to utilize it effectively for both academic and professional purposes.

- Overview of Grokking Deep Learning PDF
- Key Features and Content Structure
- Benefits of Using Grokking Deep Learning PDF
- How to Access the Grokking Deep Learning PDF
- Tips for Effective Learning with Grokking Deep Learning PDF

Overview of Grokking Deep Learning PDF

The grokking deep learning pdf is a digital version of a popular deep learning book designed to demystify complex concepts for learners at various skill levels. It focuses on providing an intuitive grasp of deep learning fundamentals, including neural networks, backpropagation, and optimization techniques. The resource is structured to facilitate progressive learning, starting from basic ideas and advancing to more sophisticated topics. The pdf format ensures easy accessibility and portability, allowing users to study deep learning concepts anytime and anywhere. This makes it an ideal choice for students, developers, and data scientists seeking flexible learning options.

Author and Pedagogical Approach

The grokking deep learning pdf is authored by an expert in machine learning who emphasizes clarity and simplicity in teaching. The pedagogical approach revolves around "grokking" — deeply understanding concepts rather than memorizing formulas. This approach is reflected in the book's use of analogies, visual explanations, and step-by-step walkthroughs of algorithms. The pdf includes practical coding examples, usually implemented in Python, to bridge theory with real-world application. This methodical, hands-on approach helps learners internalize deep learning principles effectively.

Target Audience

This resource is well-suited for a wide audience, including beginners with limited prior knowledge of machine learning, intermediate practitioners looking to solidify their understanding, and professionals aiming to update their skills. Its clear language and structured content make it accessible for self-learners and classroom settings alike. The grokking deep learning pdf also supports readers interested in both academic research and industry applications, providing foundational knowledge essential for advanced studies and practical projects.

Key Features and Content Structure

The grokking deep learning pdf is distinguished by its organized content and comprehensive coverage of essential topics in the deep learning domain. The document is divided into distinct chapters that systematically build the learner's expertise. Each chapter focuses on a specific theme, such as artificial neurons, activation functions, model training, and deep learning architectures.

Core Topics Covered

The pdf includes detailed discussions on the following core topics:

- Introduction to neural networks and perceptrons
- Activation functions and their roles
- Forward propagation and backpropagation algorithms
- Optimization techniques like gradient descent
- Common deep learning architectures such as CNNs and RNNs
- Regularization methods to prevent overfitting
- Practical implementation tips and code examples

Interactive Exercises and Code Samples

To reinforce learning, the grokking deep learning pdf incorporates interactive exercises and coding challenges. These practical components allow readers to apply concepts through hands-on programming, typically using popular libraries such as NumPy and TensorFlow. Code samples are explained in detail to ensure understanding of each step. This active learning strategy enhances retention and prepares users for real-world deep learning tasks.

Benefits of Using Grokking Deep Learning PDF

Utilizing the grokking deep learning pdf offers numerous advantages for individuals pursuing expertise in artificial intelligence and machine learning. The resource's structured and clear presentation makes complex topics more approachable. Its portability as a pdf ensures immediate access on multiple devices, facilitating continuous learning.

Enhanced Conceptual Understanding

The grokking deep learning pdf focuses on building a deep, conceptual understanding rather than rote memorization. This leads to better problem-solving skills and the ability to adapt knowledge to novel challenges. The intuitive explanations and visual aids support learners in internalizing key principles effectively.

Flexibility and Convenience

As a downloadable pdf, this resource can be used offline, making it convenient for learners who prefer studying without constant internet access. It allows users to highlight, annotate, and bookmark important sections for personalized study sessions. The format also supports printing for those who favor physical copies.

Cost-Effectiveness

Many versions of the grokking deep learning pdf are available at affordable prices or as part of educational packages, providing excellent value for the depth of information offered. This makes it an attractive option compared to costly courses or extensive textbook collections.

How to Access the Grokking Deep Learning PDF

Accessing the grokking deep learning pdf can be done through various legitimate channels. It is important to obtain the document through authorized platforms to ensure compliance with copyright laws and to receive the most updated content.

Official Publishers and Retailers

The most reliable source for the grokking deep learning pdf is the official publisher's website or recognized online retailers. Purchasing from these sources guarantees a high-quality, error-free version and often includes additional learning resources such as companion websites or code repositories.

Educational Institutions and Libraries

Some academic institutions provide access to the grokking deep learning pdf through their

digital libraries. Students and faculty members may benefit from institutional subscriptions that offer free or discounted access to the material. Public libraries with digital collections might also have this resource available for borrowing.

Ensuring Legal and Safe Downloads

To avoid counterfeit or pirated copies, it is crucial to verify the authenticity of the download source. Legal copies support the authors and publishers, encouraging the continued production of quality educational content. Users should refrain from downloading the grokking deep learning pdf from unauthorized file-sharing sites.

Tips for Effective Learning with Grokking Deep Learning PDF

Maximizing the benefits of the grokking deep learning pdf requires strategic study practices and engagement with the material. The following tips can enhance the learning experience and facilitate mastery of deep learning concepts.

Create a Structured Study Schedule

Consistent study sessions with defined goals help in gradually building knowledge without feeling overwhelmed. Allocating time to read chapters, attempt exercises, and review concepts improves retention and confidence.

Engage with Coding Exercises

Actively working through the included code samples and exercises is essential. Writing and modifying code deepens understanding and prepares learners for practical applications in machine learning projects.

Supplement with External Resources

While the grokking deep learning pdf is comprehensive, complementing it with additional materials such as video tutorials, forums, and research papers can provide broader perspectives and updated insights.

Join Study Groups or Online Communities

Collaborating with peers or participating in online forums dedicated to deep learning fosters knowledge sharing and problem-solving. Discussing challenging topics and exchanging ideas can accelerate learning progress.

Frequently Asked Questions

Where can I download the Grokking Deep Learning PDF?

The Grokking Deep Learning PDF can be purchased or accessed through authorized platforms such as Manning Publications' official website or other legitimate eBook retailers. Free unauthorized downloads are illegal and not recommended.

Is there an official Grokking Deep Learning PDF available for free?

No, there is no official free PDF version of Grokking Deep Learning. The book is a paid resource, and obtaining it through official channels supports the authors and publishers.

What topics are covered in the Grokking Deep Learning PDF?

Grokking Deep Learning covers foundational concepts of deep learning including neural networks, backpropagation, training models, convolutional networks, recurrent networks, and practical implementations with clear explanations suitable for beginners.

Can I use the Grokking Deep Learning PDF for offline study?

Yes, if you purchase the Grokking Deep Learning PDF from an authorized source, you can download it and use it for offline study, making it convenient to learn without internet access.

Are there updated editions of Grokking Deep Learning PDF available?

As of now, the main edition of Grokking Deep Learning remains the primary resource. Check the publisher's website or official sources for any announcements on updated editions or supplementary materials.

How does the Grokking Deep Learning PDF help beginners?

The Grokking Deep Learning PDF uses simple language, visual explanations, and practical examples to help beginners understand complex deep learning concepts without requiring advanced math or prior experience.

Can I get sample chapters of Grokking Deep Learning PDF before buying?

Yes, Manning Publications and other retailers often provide sample chapters or previews of Grokking Deep Learning PDF on their websites, allowing potential readers to evaluate the content before making a purchase.

Additional Resources

1. *Grokking Deep Learning*

This book by Andrew W. Trask offers an accessible introduction to deep learning concepts. It breaks down complex ideas into simple, intuitive explanations, making it ideal for beginners. The book also includes practical exercises and Python code to reinforce learning.

2. *Deep Learning with Python*

Authored by François Chollet, the creator of Keras, this book provides a hands-on approach to deep learning using the Python programming language. It covers fundamental concepts and advanced techniques, blending theory with practical examples. Readers will learn how to build, train, and deploy deep learning models efficiently.

3. *Neural Networks and Deep Learning*

Michael Nielsen's book is a free online resource that dives deep into the mathematics and intuition behind neural networks. It is well-suited for those who want a thorough understanding of how deep learning algorithms work under the hood. The text includes interactive elements and code snippets for better comprehension.

4. *Deep Learning*

Written by Ian Goodfellow, Yoshua Bengio, and Aaron Courville, this is considered the definitive textbook on deep learning. It covers a broad range of topics from the basics to cutting-edge research, providing both theoretical foundations and practical insights. Suitable for students, researchers, and practitioners alike.

5. *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*

This book by Aurélien Géron focuses on practical machine learning and deep learning techniques using popular Python libraries. It guides readers through real-world projects and workflows, making it easier to apply deep learning concepts. The clear explanations and code examples help build a robust understanding.

6. *Deep Learning for Computer Vision*

This book emphasizes the application of deep learning techniques to computer vision tasks. It covers convolutional neural networks, image classification, object detection, and segmentation. Readers will gain practical skills to implement and experiment with vision-based deep learning models.

7. *Pattern Recognition and Machine Learning*

Christopher M. Bishop's classic text offers a comprehensive introduction to machine learning, including foundational concepts relevant to deep learning. Although more general, it provides essential mathematical tools and frameworks to understand learning algorithms. The book is well-suited for those seeking a rigorous academic approach.

8. *Practical Deep Learning for Cloud, Mobile, and Edge*

This book explores how to deploy deep learning models across various platforms, such as cloud services, mobile devices, and edge computing environments. It focuses on optimization, scalability, and real-world applications. Readers interested in production-level deep learning will find this resource valuable.

9. *Deep Reinforcement Learning Hands-On*

By Maxim Lapan, this book introduces deep reinforcement learning, a cutting-edge area combining deep learning with reinforcement learning principles. It includes practical implementations using Python and PyTorch, covering algorithms like DQN, A3C, and PPO. Suitable for readers looking to expand beyond supervised learning into autonomous decision-making systems.

[Grokking Deep Learning Pdf](#)

Grokking Deep Learning PDF

Ebook Title: Deep Learning Demystified: A Grokking Approach

Ebook Outline:

Introduction: What is Deep Learning? Why is it important? Overview of the book's structure and approach.

Chapter 1: Foundational Concepts: Linear Algebra essentials, Calculus basics, Probability and Statistics refresher.

Chapter 2: Neural Networks: The Building Blocks: Perceptrons, Multilayer Perceptrons (MLPs), Activation Functions, Backpropagation.

Chapter 3: Convolutional Neural Networks (CNNs): Architectures, Applications (image classification, object detection), Practical examples.

Chapter 4: Recurrent Neural Networks (RNNs): Architectures (LSTMs, GRUs), Applications (NLP, time series analysis), Challenges and solutions.

Chapter 5: Advanced Deep Learning Techniques: Autoencoders, Generative Adversarial Networks (GANs), Transfer Learning, Model Optimization.

Chapter 6: Deep Learning Frameworks: TensorFlow, Keras, PyTorch - a comparative overview and practical introduction.

Conclusion: The future of Deep Learning, further learning resources, and concluding remarks.

Deep Learning Demystified: A Grokking Approach

The field of artificial intelligence (AI) has witnessed a phenomenal surge in recent years, largely propelled by advancements in deep learning. Deep learning, a subfield of machine learning, employs artificial neural networks with multiple layers to extract higher-level features from raw data. This

ability to automatically learn complex patterns from data has revolutionized various sectors, from image recognition and natural language processing to self-driving cars and medical diagnosis. Our comprehensive guide, "Deep Learning Demystified: A Grokking Approach," aims to provide a clear and accessible understanding of this transformative technology. This guide, available as a convenient PDF, is designed to equip you with the fundamental knowledge and practical skills to navigate the world of deep learning.

1. Introduction: Unveiling the Power of Deep Learning

Deep learning isn't just another buzzword; it's a paradigm shift in how we approach problem-solving using computers. Unlike traditional machine learning algorithms that rely heavily on feature engineering (manual selection and extraction of relevant features from data), deep learning algorithms automatically learn these features from raw data through multiple layers of interconnected nodes. This hierarchical learning process allows deep learning models to handle incredibly complex data and achieve unprecedented accuracy.

This introduction sets the stage by defining deep learning, explaining its core principles, and highlighting its transformative impact across numerous industries. We will explore the reasons behind deep learning's recent dominance, emphasizing its ability to tackle previously intractable problems and its potential for future innovation. The structure of the ebook, including the flow of chapters and the learning objectives for each section, will also be clearly outlined to guide the reader's learning journey. Finally, we'll dispel common misconceptions about deep learning, making it accessible even to those with limited prior knowledge of AI or machine learning.

2. Chapter 1: Foundational Concepts - Building the Blocks

Before diving into the intricacies of neural networks, it's crucial to establish a solid foundation in the mathematical and statistical concepts that underpin deep learning. This chapter serves as a refresher course, covering essential elements of linear algebra (vectors, matrices, operations), calculus (derivatives, gradients), and probability and statistics (distributions, hypothesis testing). While we won't delve into rigorous mathematical proofs, we'll provide intuitive explanations and practical examples to ensure a clear understanding of these fundamental concepts. The focus will be on the application of these concepts within the context of deep learning, emphasizing their relevance to neural network training and optimization. Practical examples illustrating these concepts in Python will be included to enhance comprehension and provide a hands-on learning experience.

3. Chapter 2: Neural Networks: The Building Blocks

This chapter introduces the core components of deep learning: artificial neural networks. We begin with the simplest unit, the perceptron – a single-layer neural network capable of performing basic binary classification. Building upon this foundation, we introduce multilayer perceptrons (MLPs), which consist of multiple layers of interconnected perceptrons, allowing for the learning of more

complex patterns. We will delve into the crucial role of activation functions, explaining their purpose in introducing non-linearity and enabling the network to learn intricate relationships within the data. The backpropagation algorithm, the cornerstone of neural network training, will be explained in detail, elucidating how the network learns from its errors and adjusts its weights to improve its performance. We'll use clear visualizations and intuitive explanations to make even complex concepts easily understandable.

4. Chapter 3: Convolutional Neural Networks (CNNs) - Image Understanding

Convolutional Neural Networks (CNNs) are a specialized type of neural network particularly well-suited for processing grid-like data, such as images and videos. This chapter explores the architecture of CNNs, focusing on convolutional layers, pooling layers, and fully connected layers. We'll explain how these layers work together to extract meaningful features from images, enabling tasks such as image classification, object detection, and image segmentation. We'll explore various successful CNN architectures, including LeNet, AlexNet, VGGNet, and ResNet, highlighting their strengths and limitations. Practical applications and real-world examples will be used to illustrate the power of CNNs and their wide-ranging impact.

5. Chapter 4: Recurrent Neural Networks (RNNs) - Understanding Sequences

Recurrent Neural Networks (RNNs) are designed to process sequential data, such as text, speech, and time series. Unlike CNNs, RNNs have loops in their architecture, allowing them to maintain a "memory" of past inputs. This chapter will delve into the intricacies of RNNs, exploring their architecture and the challenges associated with training them, such as the vanishing gradient problem. We'll discuss Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs), which are designed to mitigate these challenges. The chapter will conclude with several applications demonstrating the power of RNNs in various fields, including Natural Language Processing (NLP) and time series forecasting. We'll illustrate how to use RNNs for tasks such as machine translation and sentiment analysis.

6. Chapter 5: Advanced Deep Learning Techniques - Expanding the Horizons

This chapter explores more advanced deep learning techniques that have further pushed the boundaries of what's possible with AI. We'll cover autoencoders, which are used for dimensionality reduction and feature extraction, and generative adversarial networks (GANs), which can generate new data samples that resemble the training data. We'll also discuss transfer learning, a powerful technique that allows you to leverage pre-trained models to improve the performance of your own models, even with limited data. Finally, we'll discuss various model optimization techniques,

including regularization and hyperparameter tuning, to help you build robust and high-performing deep learning models.

7. Chapter 6: Deep Learning Frameworks - Putting it all Together

This chapter provides a practical introduction to popular deep learning frameworks, including TensorFlow, Keras, and PyTorch. We'll compare and contrast these frameworks, highlighting their strengths and weaknesses. We'll provide practical examples and tutorials demonstrating how to build and train deep learning models using these frameworks. This hands-on approach will empower readers to apply the theoretical concepts learned in previous chapters to real-world problems. The chapter will focus on providing practical guidance on choosing the right framework for specific tasks and building efficient deep learning pipelines.

Conclusion: Looking Ahead

Deep learning continues to evolve at a rapid pace, with new architectures, techniques, and applications emerging constantly. This conclusion summarizes the key takeaways from the ebook and looks towards the future of deep learning, highlighting exciting new research directions and potential breakthroughs. We'll also provide resources for further learning, including online courses, research papers, and communities where you can continue to expand your knowledge and engage with the deep learning community. Finally, we'll offer encouragement and guidance to readers as they embark on their journey to mastering this transformative technology.

FAQs

1. What is the prerequisite knowledge required to understand this ebook? Basic understanding of mathematics and programming is helpful but not strictly required. The ebook provides a refresher on essential concepts.
2. What programming languages are used in the ebook? Primarily Python, due to its popularity in the deep learning community.
3. Are there any coding exercises included? Yes, practical examples and coding snippets are provided throughout the ebook to reinforce learning.
4. Is this ebook suitable for beginners? Yes, the ebook is designed to be accessible to beginners with a clear and concise writing style.

5. What kind of deep learning models are covered? The ebook covers a range of models, including MLPs, CNNs, RNNs, autoencoders, and GANs.
6. Which deep learning frameworks are discussed? TensorFlow, Keras, and PyTorch are introduced and compared.
7. What is the ebook's format? It is available as a downloadable PDF.
8. Is there any support available after purchasing the ebook? While direct support might not be included, the ebook provides extensive resources and links for further learning.
9. What are the practical applications of deep learning discussed in the ebook? The ebook explores applications in image recognition, natural language processing, time series analysis, and more.

Related Articles

1. Introduction to Neural Networks: A beginner-friendly overview of the fundamental concepts of neural networks.
2. Backpropagation Explained Simply: A clear explanation of the backpropagation algorithm, crucial for training neural networks.
3. Understanding Activation Functions: A detailed exploration of various activation functions and their roles in deep learning.
4. Convolutional Neural Networks for Image Classification: A focused guide on using CNNs for image recognition tasks.
5. Recurrent Neural Networks for Natural Language Processing: An in-depth look at applying RNNs to text and language data.
6. Generative Adversarial Networks (GANs) and Their Applications: A comprehensive overview of GANs and their use in generating synthetic data.
7. Transfer Learning in Deep Learning: A guide to leveraging pre-trained models to improve model performance.
8. Choosing the Right Deep Learning Framework: A comparative analysis of popular deep learning frameworks, assisting in framework selection.
9. Optimizing Deep Learning Models for Performance: Techniques for improving the accuracy and efficiency of deep learning models.

grokking deep learning pdf: [Grokking Deep Learning](#) Andrew W. Trask, 2019-01-23 Summary
Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology
Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book
Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself

every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise: introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide

grokking deep learning pdf: Grokking Deep Reinforcement Learning Miguel Morales, 2020-11-10 Grokking Deep Reinforcement Learning uses engaging exercises to teach you how to build deep learning systems. This book combines annotated Python code with intuitive explanations to explore DRL techniques. You'll see how algorithms function and learn to develop your own DRL agents using evaluative feedback. Summary We all learn through trial and error. We avoid the things that cause us to experience pain and failure. We embrace and build on the things that give us reward and success. This common pattern is the foundation of deep reinforcement learning: building machine learning systems that explore and learn based on the responses of the environment. Grokking Deep Reinforcement Learning introduces this powerful machine learning approach, using examples, illustrations, exercises, and crystal-clear teaching. You'll love the perfectly paced teaching and the clever, engaging writing style as you dig into this awesome exploration of reinforcement learning fundamentals, effective deep learning techniques, and practical applications in this emerging field. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology We learn by interacting with our environment, and the rewards or punishments we experience guide our future behavior. Deep reinforcement learning brings that same natural process to artificial intelligence, analyzing results to uncover the most efficient ways forward. DRL agents can improve marketing campaigns, predict stock performance, and beat grand masters in Go and chess. About the book Grokking Deep Reinforcement Learning uses engaging exercises to teach you how to build deep learning systems. This book combines annotated Python code with intuitive explanations to explore DRL techniques. You'll see how algorithms function and learn to develop your own DRL agents using evaluative feedback. What's inside An introduction to reinforcement learning DRL agents with human-like behaviors Applying DRL to complex situations About the reader For developers with basic deep learning experience. About the author Miguel Morales works on reinforcement learning at Lockheed Martin and is an instructor for the Georgia Institute of Technology's Reinforcement Learning and Decision Making course. Table of Contents 1 Introduction to deep reinforcement learning 2 Mathematical foundations of reinforcement learning 3 Balancing immediate and long-term goals 4 Balancing the gathering and use of information 5 Evaluating agents' behaviors 6 Improving agents' behaviors 7 Achieving goals more effectively and efficiently 8 Introduction to value-based deep reinforcement learning 9 More stable value-based methods 10 Sample-efficient value-based methods 11 Policy-gradient and

actor-critic methods 12 Advanced actor-critic methods 13 Toward artificial general intelligence

grokking deep learning pdf: *Grokking Machine Learning* Luis Serrano, 2021-12-14 Grokking Machine Learning presents machine learning algorithms and techniques in a way that anyone can understand. This book skips the confused academic jargon and offers clear explanations that require only basic algebra. As you go, you'll build interesting projects with Python, including models for spam detection and image recognition. You'll also pick up practical skills for cleaning and preparing data.

grokking deep learning pdf: *Grokking Deep Learning* Andrew Trask, 2019-01-25 Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise: introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide

grokking deep learning pdf: *Grokking Algorithms* Aditya Bhargava, 2016-05-12 This book does the impossible: it makes math fun and easy! - Sander Rossel, COAS Software Systems Grokking Algorithms is a fully illustrated, friendly guide that teaches you how to apply common algorithms to the practical problems you face every day as a programmer. You'll start with sorting and searching and, as you build up your skills in thinking algorithmically, you'll tackle more complex concerns such as data compression and artificial intelligence. Each carefully presented example includes helpful diagrams and fully annotated code samples in Python. Learning about algorithms doesn't have to be boring! Get a sneak peek at the fun, illustrated, and friendly examples you'll find in Grokking Algorithms on Manning Publications' YouTube channel. Continue your journey into the world of algorithms with Algorithms in Motion, a practical, hands-on video course available exclusively at Manning.com (www.manning.com/livevideo/algorithms-?in-motion). Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the

Technology An algorithm is nothing more than a step-by-step procedure for solving a problem. The algorithms you'll use most often as a programmer have already been discovered, tested, and proven. If you want to understand them but refuse to slog through dense multipage proofs, this is the book for you. This fully illustrated and engaging guide makes it easy to learn how to use the most important algorithms effectively in your own programs. About the Book *Grokking Algorithms* is a friendly take on this core computer science topic. In it, you'll learn how to apply common algorithms to the practical programming problems you face every day. You'll start with tasks like sorting and searching. As you build up your skills, you'll tackle more complex problems like data compression and artificial intelligence. Each carefully presented example includes helpful diagrams and fully annotated code samples in Python. By the end of this book, you will have mastered widely applicable algorithms as well as how and when to use them. What's Inside Covers search, sort, and graph algorithms Over 400 pictures with detailed walkthroughs Performance trade-offs between algorithms Python-based code samples About the Reader This easy-to-read, picture-heavy introduction is suitable for self-taught programmers, engineers, or anyone who wants to brush up on algorithms. About the Author Aditya Bhargava is a Software Engineer with a dual background in Computer Science and Fine Arts. He blogs on programming at adit.io. Table of Contents Introduction to algorithms Selection sort Recursion Quicksort Hash tables Breadth-first search Dijkstra's algorithm Greedy algorithms Dynamic programming K-nearest neighbors

grokking deep learning pdf: *Deep Reinforcement Learning in Action* Alexander Zai, Brandon Brown, 2020-04-28 Summary Humans learn best from feedback—we are encouraged to take actions that lead to positive results while deterred by decisions with negative consequences. This reinforcement process can be applied to computer programs allowing them to solve more complex problems that classical programming cannot. Deep Reinforcement Learning in Action teaches you the fundamental concepts and terminology of deep reinforcement learning, along with the practical skills and techniques you'll need to implement it into your own projects. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Deep reinforcement learning AI systems rapidly adapt to new environments, a vast improvement over standard neural networks. A DRL agent learns like people do, taking in raw data such as sensor input and refining its responses and predictions through trial and error. About the book Deep Reinforcement Learning in Action teaches you how to program AI agents that adapt and improve based on direct feedback from their environment. In this example-rich tutorial, you'll master foundational and advanced DRL techniques by taking on interesting challenges like navigating a maze and playing video games. Along the way, you'll work with core algorithms, including deep Q-networks and policy gradients, along with industry-standard tools like PyTorch and OpenAI Gym. What's inside Building and training DRL networks The most popular DRL algorithms for learning and problem solving Evolutionary algorithms for curiosity and multi-agent learning All examples available as Jupyter Notebooks About the reader For readers with intermediate skills in Python and deep learning. About the author Alexander Zai is a machine learning engineer at Amazon AI. Brandon Brown is a machine learning and data analysis blogger. Table of Contents PART 1 - FOUNDATIONS 1. What is reinforcement learning? 2. Modeling reinforcement learning problems: Markov decision processes 3. Predicting the best states and actions: Deep Q-networks 4. Learning to pick the best policy: Policy gradient methods 5. Tackling more complex problems with actor-critic methods PART 2 - ABOVE AND BEYOND 6. Alternative optimization methods: Evolutionary algorithms 7. Distributional DQN: Getting the full story 8. Curiosity-driven exploration 9. Multi-agent reinforcement learning 10. Interpretable reinforcement learning: Attention and relational models 11. In conclusion: A review and roadmap

grokking deep learning pdf: *Grokking Artificial Intelligence Algorithms* Rishal Hurbans, 2020-07-20 From start to finish, the best book to help you learn AI algorithms and recall why and how you use them. - Linda Ristevski, York Region District School Board "This book takes an impossibly broad area of computer science and communicates what working developers need to understand in a clear and thorough way." - David Jacobs, Product Advance Local Key Features

Master the core algorithms of deep learning and AI Build an intuitive understanding of AI problems and solutions Written in simple language, with lots of illustrations and hands-on examples Creative coding exercises, including building a maze puzzle game and exploring drone optimization About The Book “Artificial intelligence” requires teaching a computer how to approach different types of problems in a systematic way. The core of AI is the algorithms that the system uses to do things like identifying objects in an image, interpreting the meaning of text, or looking for patterns in data to spot fraud and other anomalies. Mastering the core algorithms for search, image recognition, and other common tasks is essential to building good AI applications Grokking Artificial Intelligence Algorithms uses illustrations, exercises, and jargon-free explanations to teach fundamental AI concepts. You’ll explore coding challenges like detecting bank fraud, creating artistic masterpieces, and setting a self-driving car in motion. All you need is the algebra you remember from high school math class and beginning programming skills. What You Will Learn Use cases for different AI algorithms Intelligent search for decision making Biologically inspired algorithms Machine learning and neural networks Reinforcement learning to build a better robot This Book Is Written For For software developers with high school-level math skills. About the Author Rishal Hurbans is a technologist, startup and AI group founder, and international speaker. Table of Contents 1 Intuition of artificial intelligence 2 Search fundamentals 3 Intelligent search 4 Evolutionary algorithms 5 Advanced evolutionary approaches 6 Swarm intelligence: Ants 7 Swarm intelligence: Particles 8 Machine learning 9 Artificial neural networks 10 Reinforcement learning with Q-learning

grokking deep learning pdf: Deep Learning Interviews Shlomo Kashani, 2020-12-09 The book's contents is a large inventory of numerous topics relevant to DL job interviews and graduate level exams. That places this work at the forefront of the growing trend in science to teach a core set of practical mathematical and computational skills. It is widely accepted that the training of every computer scientist must include the fundamental theorems of ML, and AI appears in the curriculum of nearly every university. This volume is designed as an excellent reference for graduates of such programs.

grokking deep learning pdf: Deep Learning with Python Francois Chollet, 2017-11-30 Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing

Systems (NIPS), the International Conference on Learning Representations (ICLR), and others.

Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

grokking deep learning pdf: *AI and Machine Learning for Coders* Laurence Moroney, 2020-10-01 If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow Serving

grokking deep learning pdf: *Deep Learning for Search* Tommaso Teofili, 2019-06-02 Summary Deep Learning for Search teaches you how to improve the effectiveness of your search by implementing neural network-based techniques. By the time you're finished with the book, you'll be ready to build amazing search engines that deliver the results your users need and that get better as time goes on! Foreword by Chris Mattmann. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning handles the toughest search challenges, including imprecise search terms, badly indexed data, and retrieving images with minimal metadata. And with modern tools like DL4J and TensorFlow, you can apply powerful DL techniques without a deep background in data science or natural language processing (NLP). This book will show you how. About the Book Deep Learning for Search teaches you to improve your search results with neural networks. You'll review how DL relates to search basics like indexing and ranking. Then, you'll walk through in-depth examples to upgrade your search with DL techniques using Apache Lucene and Deeplearning4j. As the book progresses, you'll explore advanced topics like searching through images, translating user queries, and designing search engines that improve as they learn! What's inside Accurate and relevant rankings Searching across languages Content-based image search Search with recommendations About the Reader For developers comfortable with Java or a similar language and search basics. No experience with deep learning or NLP needed. About the Author Tommaso Teofili is a software engineer with a passion for open source and machine learning. As a member of the Apache Software Foundation, he contributes to a number of open source projects, ranging from topics like information retrieval (such as Lucene and Solr) to natural language processing and machine translation (including OpenNLP, Joshua, and UIMA). He currently works at Adobe, developing search and indexing infrastructure components, and researching the areas of natural language processing, information retrieval, and deep learning. He has presented search and machine learning talks at conferences including BerlinBuzzwords, International Conference on Computational Science, ApacheCon, EclipseCon, and others. You can find him on Twitter at @tteofili. Table of Contents PART 1 - SEARCH MEETS DEEP LEARNING Neural search Generating synonyms PART 2 - THROWING NEURAL NETS AT A SEARCH ENGINE From plain retrieval to text generation More-sensitive query suggestions Ranking search results with word embeddings Document embeddings for rankings and recommendations PART 3 - ONE STEP BEYOND Searching across languages Content-based image search A peek at performance

grokking deep learning pdf: *Deep Learning from Scratch* Seth Weidman, 2019-09-09 With the

resurgence of neural networks in the 2010s, deep learning has become essential for machine learning practitioners and even many software engineers. This book provides a comprehensive introduction for data scientists and software engineers with machine learning experience. You'll start with deep learning basics and move quickly to the details of important advanced architectures, implementing everything from scratch along the way. Author Seth Weidman shows you how neural networks work using a first principles approach. You'll learn how to apply multilayer neural networks, convolutional neural networks, and recurrent neural networks from the ground up. With a thorough understanding of how neural networks work mathematically, computationally, and conceptually, you'll be set up for success on all future deep learning projects. This book provides: Extremely clear and thorough mental models—accompanied by working code examples and mathematical explanations—for understanding neural networks Methods for implementing multilayer neural networks from scratch, using an easy-to-understand object-oriented framework Working implementations and clear-cut explanations of convolutional and recurrent neural networks Implementation of these neural network concepts using the popular PyTorch framework

grokking deep learning pdf: Reinforcement Learning, second edition Richard S. Sutton, Andrew G. Barto, 2018-11-13 The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

grokking deep learning pdf: Grokking Bitcoin Kalle Rosenbaum, 2019-04-17 Summary If you think Bitcoin is just an alternative currency for geeks, it's time to think again. Grokking Bitcoin opens up this powerful distributed ledger system, exploring the technology that enables applications both for Bitcoin-based financial transactions and using the blockchain for registering physical property ownership. With this fully illustrated, easy-to-read guide, you'll finally understand how Bitcoin works, how you can use it, and why you can trust the blockchain. Foreword by David A. Harding, Contributor to Bitcoin documentation. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Inflation, depressed economies, debased currencies ... these are just a few of the problems centralized banking has caused throughout history. Bitcoin, a digital currency created with the ambition to shift control away from change-prone governments, has the potential to bring an end to those problems once and for all. It's time to find out how it can help you. About the Book Grokking Bitcoin explains why Bitcoin's supporters trust it so deeply, and why you can too. This approachable book will introduce you to Bitcoin's groundbreaking technology, which is the key to this world-changing system. This illustrated, easy-to-read guide prepares you for a new way of thinking with easy-to-follow diagrams and exercises. You'll discover how Bitcoin mining works, how to accept Bitcoin, how to participate in the Bitcoin network, and how to set up a digital wallet. What's inside Bitcoin transactions The blockchain Bitcoin mining Bitcoin wallets About the Reader Intended for

anyone interested in learning about Bitcoin technology. While a basic understanding of technical concepts is beneficial, no programming skills are necessary. About the Author Kalle Rosenbaum is a computer scientist, an avid Bitcoin supporter, and the founder of Propeller, a Bitcoin consultancy. Table of Contents Introduction to Bitcoin Cryptographic hash functions and digital signatures Addresses Wallets Transactions The blockchain Proof of work Peer-to-peer network Transactions revisited Segregated witness Bitcoin upgrades

grokking deep learning pdf: Deep Learning with PyTorch Luca Pietro Giovanni Antiga, Eli Stevens, Thomas Viehmann, 2020-07-01 “We finally have the definitive treatise on PyTorch! It covers the basics and abstractions in great detail. I hope this book becomes your extended reference document.” —Soumith Chintala, co-creator of PyTorch Key Features Written by PyTorch’s creator and key contributors Develop deep learning models in a familiar Pythonic way Use PyTorch to build an image classifier for cancer detection Diagnose problems with your neural network and improve training with data augmentation Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About The Book Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands. Instantly familiar to anyone who knows Python data tools like NumPy and Scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It’s great for building quick models, and it scales smoothly from laptop to enterprise. Deep Learning with PyTorch teaches you to create deep learning and neural network systems with PyTorch. This practical book gets you to work right away building a tumor image classifier from scratch. After covering the basics, you’ll learn best practices for the entire deep learning pipeline, tackling advanced projects as your PyTorch skills become more sophisticated. All code samples are easy to explore in downloadable Jupyter notebooks. What You Will Learn Understanding deep learning data structures such as tensors and neural networks Best practices for the PyTorch Tensor API, loading data in Python, and visualizing results Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Methods for training networks with limited inputs Sifting through unreliable results to diagnose and fix problems in your neural network Improve your results with augmented data, better model architecture, and fine tuning This Book Is Written For For Python programmers with an interest in machine learning. No experience with PyTorch or other deep learning frameworks is required. About The Authors Eli Stevens has worked in Silicon Valley for the past 15 years as a software engineer, and the past 7 years as Chief Technical Officer of a startup making medical device software. Luca Antiga is co-founder and CEO of an AI engineering company located in Bergamo, Italy, and a regular contributor to PyTorch. Thomas Viehmann is a Machine Learning and PyTorch speciality trainer and consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

grokking deep learning pdf: Grokking Streaming Systems Josh Fischer, Ning Wang, 2022-04-19 A friendly, framework-agnostic tutorial that will help you grok how streaming systems work—and how to build your own! In Grokking Streaming Systems you will learn how to: Implement and troubleshoot streaming systems Design streaming systems for complex functionalities Assess parallelization requirements Spot networking bottlenecks and resolve back pressure Group data for high-performance systems Handle delayed events in real-time systems Grokking Streaming Systems is a simple guide to the complex concepts behind streaming systems. This friendly and framework-agnostic tutorial teaches you how to handle real-time events, and even design and build

your own streaming job that's a perfect fit for your needs. Each new idea is carefully explained with diagrams, clear examples, and fun dialogue between perplexed personalities! About the technology Streaming systems minimize the time between receiving and processing event data, so they can deliver responses in real time. For applications in finance, security, and IoT where milliseconds matter, streaming systems are a requirement. And streaming is hot! Skills on platforms like Spark, Heron, and Kafka are in high demand. About the book *Grokking Streaming Systems* introduces real-time event streaming applications in clear, reader-friendly language. This engaging book illuminates core concepts like data parallelization, event windows, and backpressure without getting bogged down in framework-specific details. As you go, you'll build your own simple streaming tool from the ground up to make sure all the ideas and techniques stick. The helpful and entertaining illustrations make streaming systems come alive as you tackle relevant examples like real-time credit card fraud detection and monitoring IoT services. What's inside Implement and troubleshoot streaming systems Design streaming systems for complex functionalities Spot networking bottlenecks and resolve backpressure Group data for high-performance systems About the reader No prior experience with streaming systems is assumed. Examples in Java. About the author Josh Fischer and Ning Wang are Apache Committers, and part of the committee for the Apache Heron distributed stream processing engine. Table of Contents PART 1 GETTING STARTED WITH STREAMING 1 Welcome to Grokking Streaming Systems 2 Hello, streaming systems! 3 Parallelization and data grouping 4 Stream graph 5 Delivery semantics 6 Streaming systems review and a glimpse ahead PART 2 STEPPING UP 7 Windowed computations 8 Join operations 9 Backpressure 10 Stateful computation 11 Wrap-up: Advanced concepts in streaming systems

grokking deep learning pdf: *Deep Learning for Coders with fastai and PyTorch* Jeremy Howard, Sylvain Gugger, 2020-06-29 Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

grokking deep learning pdf: Dive Into Deep Learning Joanne Quinn, Joanne McEachen, Michael Fullan, Mag Gardner, Max Drummy, 2019-07-15 The leading experts in system change and learning, with their school-based partners around the world, have created this essential companion to their runaway best-seller, *Deep Learning: Engage the World Change the World*. This hands-on guide provides a roadmap for building capacity in teachers, schools, districts, and systems to design deep learning, measure progress, and assess conditions needed to activate and sustain innovation. *Dive Into Deep Learning: Tools for Engagement* is rich with resources educators need to construct and drive meaningful deep learning experiences in order to develop the kind of mindset and know-how that is crucial to becoming a problem-solving change agent in our global society. Designed in full color, this easy-to-use guide is loaded with tools, tips, protocols, and real-world examples. It includes: • A framework for deep learning that provides a pathway to develop the six global competencies needed to flourish in a complex world — character, citizenship, collaboration, communication, creativity, and critical thinking. • Learning progressions to help educators analyze student work and measure progress. • Learning design rubrics, templates and examples for incorporating the four elements of learning design: learning partnerships, pedagogical practices, learning environments, and leveraging digital. • Conditions rubrics, teacher self-assessment tools,

and planning guides to help educators build, mobilize, and sustain deep learning in schools and districts. Learn about, improve, and expand your world of learning. Put the joy back into learning for students and adults alike. Dive into deep learning to create learning experiences that give purpose, unleash student potential, and transform not only learning, but life itself.

grokking deep learning pdf: *Grokking Simplicity* Eric Normand, 2021-05-18 Distributed across servers, difficult to test, and resistant to modification—modern software is complex. *Grokking Simplicity* is a friendly, practical guide that will change the way you approach software design and development. It introduces a unique approach to functional programming that explains why certain features of software are prone to complexity, and teaches you the functional techniques you can use to simplify these systems so that they're easier to test and debug. Available in PDF (ePub, kindle, and liveBook formats coming soon). about the technology Even experienced developers struggle with software systems that sprawl across distributed servers and APIs, are filled with redundant code, and are difficult to reliably test and modify. Adopting ways of thinking derived from functional programming can help you design and refactor your codebase in ways that reduce complexity, rather than encouraging it. *Grokking Simplicity* lays out how to use functional programming in a professional environment to write a codebase that's easier to test and reuse, has fewer bugs, and is better at handling the asynchronous nature of distributed systems. about the book In *Grokking Simplicity*, you'll learn techniques and, more importantly, a mindset that will help you tackle common problems that arise when software gets complex. Veteran functional programmer Eric Normand guides you to a crystal-clear understanding of why certain features of modern software are so prone to complexity and introduces you to the functional techniques you can use to simplify these systems so that they're easier to read, test, and debug. Through hands-on examples, exercises, and numerous self-assessments, you'll learn to organize your code for maximum reusability and internalize methods to keep unwanted complexity out of your codebase. Regardless of the language you're using, the ways of thinking in this book will help recognize problematic code and tame even the most complex software. what's inside Apply functional programming principles to reduce codebase complexity Work with data transformation pipelines for code that's easier to test and reuse Tools for modeling time to simplify asynchrony 60 exercises and 100 questions to test your knowledge about the reader For experienced programmers. Examples are in JavaScript. about the author Eric Normand has been a functional programmer since 2001 and has been teaching functional programming online and in person since 2007. Visit LispCast.com to see more of his credentials.

grokking deep learning pdf: *Practical Deep Learning* Ronald T. Kneusel, 2021-02-23 *Practical Deep Learning* teaches total beginners how to build the datasets and models needed to train neural networks for your own DL projects. If you've been curious about machine learning but didn't know where to start, this is the book you've been waiting for. Focusing on the subfield of machine learning known as deep learning, it explains core concepts and gives you the foundation you need to start building your own models. Rather than simply outlining recipes for using existing toolkits, *Practical Deep Learning* teaches you the why of deep learning and will inspire you to explore further. All you need is basic familiarity with computer programming and high school math—the book will cover the rest. After an introduction to Python, you'll move through key topics like how to build a good training dataset, work with the scikit-learn and Keras libraries, and evaluate your models' performance. You'll also learn: How to use classic machine learning models like k-Nearest Neighbors, Random Forests, and Support Vector Machines How neural networks work and how they're trained How to use convolutional neural networks How to develop a successful deep learning model from scratch You'll conduct experiments along the way, building to a final case study that incorporates everything you've learned. The perfect introduction to this dynamic, ever-expanding field, *Practical Deep Learning* will give you the skills and confidence to dive into your own machine learning projects.

grokking deep learning pdf: *Six Septembers: Mathematics for the Humanist* Patrick Juola, Stephen Ramsay, 2017 Scholars of all stripes are turning their attention to materials that represent

enormous opportunities for the future of humanistic inquiry. The purpose of this book is to impart the concepts that underlie the mathematics they are likely to encounter and to unfold the notation in a way that removes that particular barrier completely. This book is a primer for developing the skills to enable humanist scholars to address complicated technical material with confidence. This book, to put it plainly, is concerned with the things that the author of a technical article knows, but isn't saying. Like any field, mathematics operates under a regime of shared assumptions, and it is our purpose to elucidate some of those assumptions for the newcomer. The individual subjects we tackle are (in order): logic and proof, discrete mathematics, abstract algebra, probability and statistics, calculus, and differential equations.

grokking deep learning pdf: Machine Learning Jason Bell, 2020-02-17 Dig deep into the data with a hands-on guide to machine learning with updated examples and more! Machine Learning: Hands-On for Developers and Technical Professionals provides hands-on instruction and fully-coded working examples for the most common machine learning techniques used by developers and technical professionals. The book contains a breakdown of each ML variant, explaining how it works and how it is used within certain industries, allowing readers to incorporate the presented techniques into their own work as they follow along. A core tenant of machine learning is a strong focus on data preparation, and a full exploration of the various types of learning algorithms illustrates how the proper tools can help any developer extract information and insights from existing data. The book includes a full complement of Instructor's Materials to facilitate use in the classroom, making this resource useful for students and as a professional reference. At its core, machine learning is a mathematical, algorithm-based technology that forms the basis of historical data mining and modern big data science. Scientific analysis of big data requires a working knowledge of machine learning, which forms predictions based on known properties learned from training data. Machine Learning is an accessible, comprehensive guide for the non-mathematician, providing clear guidance that allows readers to: Learn the languages of machine learning including Hadoop, Mahout, and Weka Understand decision trees, Bayesian networks, and artificial neural networks Implement Association Rule, Real Time, and Batch learning Develop a strategic plan for safe, effective, and efficient machine learning By learning to construct a system that can learn from data, readers can increase their utility across industries. Machine learning sits at the core of deep dive data analysis and visualization, which is increasingly in demand as companies discover the goldmine hiding in their existing data. For the tech professional involved in data science, Machine Learning: Hands-On for Developers and Technical Professionals provides the skills and techniques required to dig deeper.

grokking deep learning pdf: Deep Learning Josh Patterson, Adam Gibson, 2017-07-28 Although interest in machine learning has reached a high point, lofty expectations often scuttle projects before they get very far. How can machine learning—especially deep neural networks—make a real difference in your organization? This hands-on guide not only provides the most practical information available on the subject, but also helps you get started building efficient deep learning networks. Authors Adam Gibson and Josh Patterson provide theory on deep learning before introducing their open-source Deeplearning4j (DL4J) library for developing production-class workflows. Through real-world examples, you'll learn methods and strategies for training deep network architectures and running deep learning workflows on Spark and Hadoop with DL4J. Dive into machine learning concepts in general, as well as deep learning in particular Understand how deep networks evolved from neural network fundamentals Explore the major deep network architectures, including Convolutional and Recurrent Learn how to map specific deep networks to the right problem Walk through the fundamentals of tuning general neural networks and specific deep network architectures Use vectorization techniques for different data types with DataVec, DL4J's workflow tool Learn how to use DL4J natively on Spark and Hadoop

grokking deep learning pdf: Cracking The Machine Learning Interview Nitin Suri, 2018-12-18 A breakthrough in machine learning would be worth ten Microsofts. -Bill Gates Despite being one of the hottest disciplines in the Tech industry right now, Artificial Intelligence and Machine Learning

remain a little elusive to most. The erratic availability of resources online makes it extremely challenging for us to delve deeper into these fields. Especially when gearing up for job interviews, most of us are at a loss due to the unavailability of a complete and uncondensed source of learning. *Cracking the Machine Learning Interview* Equips you with 225 of the best Machine Learning problems along with their solutions. Requires only a basic knowledge of fundamental mathematical and statistical concepts. Assists in learning the intricacies underlying Machine Learning concepts and algorithms suited to specific problems. Uniquely provides a manifold understanding of both statistical foundations and applied programming models for solving problems. Discusses key points and concrete tips for approaching real life system design problems and imparts the ability to apply them to your day to day work. This book covers all the major topics within Machine Learning which are frequently asked in the Interviews. These include: Supervised and Unsupervised Learning Classification and Regression Decision Trees Ensembles K-Nearest Neighbors Logistic Regression Support Vector Machines Neural Networks Regularization Clustering Dimensionality Reduction Feature Extraction Feature Engineering Model Evaluation Natural Language Processing Real life system design problems Mathematics and Statistics behind the Machine Learning Algorithms Various distributions and statistical tests This book can be used by students and professionals alike. It has been drafted in a way to benefit both, novices as well as individuals with substantial experience in Machine Learning. Following *Cracking The Machine Learning Interview* diligently would equip you to face any Machine Learning Interview.

grokking deep learning pdf: *Generative Deep Learning* David Foster, 2019-06-28 Generative modeling is one of the hottest topics in AI. It's now possible to teach a machine to excel at human endeavors such as painting, writing, and composing music. With this practical book, machine-learning engineers and data scientists will discover how to re-create some of the most impressive examples of generative deep learning models, such as variational autoencoders, generative adversarial networks (GANs), encoder-decoder models and world models. Author David Foster demonstrates the inner workings of each technique, starting with the basics of deep learning before advancing to some of the most cutting-edge algorithms in the field. Through tips and tricks, you'll understand how to make your models learn more efficiently and become more creative. Discover how variational autoencoders can change facial expressions in photos Build practical GAN examples from scratch, including CycleGAN for style transfer and MuseGAN for music generation Create recurrent generative models for text generation and learn how to improve the models using attention Understand how generative models can help agents to accomplish tasks within a reinforcement learning setting Explore the architecture of the Transformer (BERT, GPT-2) and image generation models such as ProGAN and StyleGAN

grokking deep learning pdf: Grokking Continuous Delivery Christie Wilson, 2022-12-06 Build and use systems that safely automate software delivery from testing through release with this jargon-busting guide to continuous delivery pipelines. In *Grokking Continuous Delivery* you will learn how to: Design effective CD pipelines for new and legacy projects Keep your software projects release-ready Maintain effective tests Scale CD across multiple applications Ensure pipelines give the right signals at the right time Use version control as the source of truth Safely automate deployments with metrics Describe CD in a way that makes sense to your colleagues *Grokking Continuous Delivery* teaches you the design and purpose of continuous delivery systems that you can use with any language or stack. You'll learn directly from your mentor Christie Wilson, Google engineer and co-creator of the Tekton CI/CD framework. Using crystal-clear, well-illustrated examples, Christie lays out the practical nuts and bolts of continuous delivery for developers and pipeline designers. In each chapter, you'll uncover the proper approaches to solve the real-world challenges of setting up a CD pipeline. With this book as your roadmap, you'll have a clear plan for bringing CD to your team without the need for costly trial-and-error experimentation. About the technology Keep your codebase release-ready. A continuous delivery pipeline automates version control, testing, and deployment with minimal developer intervention. Master the tools and practices of continuous delivery, and you'll be able to add features and push updates quickly and consistently.

About the book *Grokking Continuous Delivery* is a friendly guide to setting up and working with a continuous delivery pipeline. Each chapter takes on a different scenario you'll face when setting up a CD system, with real-world examples like automated scaling and testing legacy applications. Taking a tool-agnostic approach, author Christie Wilson guides you each step of the way with illustrations, crystal-clear explanations, and practical exercises to lock in what you're learning. What's inside

Design effective CD pipelines for new and legacy projects
 Ensure your pipelines give the right signals at the right times
 Version control as the source of truth
 Safely automate deployments

About the reader
 For software engineers who want to add CD to their development process.

About the author
 Christie Wilson is a software engineer at Google, where she co-created Tekton, a cloud-native CI/CD platform built on Kubernetes.

Table of Contents

PART 1 Introducing continuous delivery

1 Welcome to *Grokking Continuous Delivery*

2 A basic pipeline

PART 2 Keeping software in a deliverable state at all times

3 Version control is the only way to roll

4 Use linting effectively

5 Dealing with noisy tests

6 Speeding up slow test suites

7 Give the right signals at the right times

PART 3 Making delivery easy

8 Easy delivery starts with version control

9 Building securely and reliably

10 Deploying confidently

PART 4 CD design

11 Starter packs: From zero to CD

12 Scripts are code, too

13 Pipeline design

grokking deep learning pdf: *Deep Reinforcement Learning Hands-On* Maxim Lapan, 2020-01-31 Revised and expanded to include multi-agent methods, discrete optimization, RL in robotics, advanced exploration techniques, and more

Key Features

Second edition of the bestselling introduction to deep reinforcement learning, expanded with six new chapters

Learn advanced exploration techniques including noisy networks, pseudo-count, and network distillation methods

Apply RL methods to cheap hardware robotics platforms

Book Description

Deep Reinforcement Learning Hands-On, Second Edition is an updated and expanded version of the bestselling guide to the very latest reinforcement learning (RL) tools and techniques. It provides you with an introduction to the fundamentals of RL, along with the hands-on ability to code intelligent learning agents to perform a range of practical tasks. With six new chapters devoted to a variety of up-to-the-minute developments in RL, including discrete optimization (solving the Rubik's Cube), multi-agent methods, Microsoft's TextWorld environment, advanced exploration techniques, and more, you will come away from this book with a deep understanding of the latest innovations in this emerging field. In addition, you will gain actionable insights into such topic areas as deep Q-networks, policy gradient methods, continuous control problems, and highly scalable, non-gradient methods. You will also discover how to build a real hardware robot trained with RL for less than \$100 and solve the Pong environment in just 30 minutes of training using step-by-step code optimization. In short, *Deep Reinforcement Learning Hands-On, Second Edition*, is your companion to navigating the exciting complexities of RL as it helps you attain experience and knowledge through real-world examples.

What you will learn

Understand the deep learning context of RL and implement complex deep learning models

Evaluate RL methods including cross-entropy, DQN, actor-critic, TRPO, PPO, DDPG, D4PG, and others

Build a practical hardware robot trained with RL methods for less than \$100

Discover Microsoft's TextWorld environment, which is an interactive fiction games platform

Use discrete optimization in RL to solve a Rubik's Cube

Teach your agent to play Connect 4 using AlphaGo Zero

Explore the very latest deep RL research on topics including AI chatbots

Discover advanced exploration techniques, including noisy networks and network distillation techniques

Who this book is for

Some fluency in Python is assumed. Sound understanding of the fundamentals of deep learning will be helpful. This book is an introduction to deep RL and requires no background in RL

grokking deep learning pdf: *The Principles of Deep Learning Theory* Daniel A. Roberts, Sho Yaida, Boris Hanin, 2022-05-26 This volume develops an effective theory approach to understanding deep neural networks of practical relevance.

grokking deep learning pdf: *Artificial Intelligence* Melanie Mitchell, 2019-10-15 Melanie Mitchell separates science fact from science fiction in this sweeping examination of the current state of AI and how it is remaking our world No recent scientific enterprise has proved as alluring,

terrifying, and filled with extravagant promise and frustrating setbacks as artificial intelligence. The award-winning author Melanie Mitchell, a leading computer scientist, now reveals AI's turbulent history and the recent spate of apparent successes, grand hopes, and emerging fears surrounding it. In *Artificial Intelligence*, Mitchell turns to the most urgent questions concerning AI today: How intelligent—really—are the best AI programs? How do they work? What can they actually do, and when do they fail? How humanlike do we expect them to become, and how soon do we need to worry about them surpassing us? Along the way, she introduces the dominant models of modern AI and machine learning, describing cutting-edge AI programs, their human inventors, and the historical lines of thought underpinning recent achievements. She meets with fellow experts such as Douglas Hofstadter, the cognitive scientist and Pulitzer Prize-winning author of the modern classic *Gödel, Escher, Bach*, who explains why he is “terrified” about the future of AI. She explores the profound disconnect between the hype and the actual achievements in AI, providing a clear sense of what the field has accomplished and how much further it has to go. Interweaving stories about the science of AI and the people behind it, *Artificial Intelligence* brims with clear-sighted, captivating, and accessible accounts of the most interesting and provocative modern work in the field, flavored with Mitchell's humor and personal observations. This frank, lively book is an indispensable guide to understanding today's AI, its quest for “human-level” intelligence, and its impact on the future for us all.

grokking deep learning pdf: Grokking Functional Programming Michal Plachta, 2023-02-07 There's no need to fear going functional! This friendly, lively, and engaging guide is perfect for any perplexed programmer. It lays out the principles of functional programming in a simple and concise way that will help you grok what FP is really all about. In *Grokking Functional Programming* you will learn: Designing with functions and types instead of objects Programming with pure functions and immutable values Writing concurrent programs using the functional style Testing functional programs Multiple learning approaches to help you grok each new concept If you've ever found yourself rolling your eyes at functional programming, this is the book for you. Open up *Grokking Functional Programming* and you'll find functional ideas mapped onto what you already know as an object-oriented programmer. The book focuses on practical aspects from page one. Hands-on examples apply functional principles to everyday programming tasks like concurrency, error handling, and improving readability. Plus, puzzles and exercises let you think and practice what you're learning. You'll soon reach an amazing “aha” moment and start seeing code in a completely new way. About the technology Finally, there's an easy way to learn functional programming! This unique book starts with the familiar ideas of OOP and introduces FP step-by-step using relevant examples, engaging exercises, and lots of illustrations. You'll be amazed at how quickly you'll start seeing software tasks from this valuable new perspective. About the book *Grokking Functional Programming* introduces functional programming to imperative developers. You'll start with small, comfortable coding tasks that expose basic concepts like writing pure functions and working with immutable data. Along the way, you'll learn how to write code that eliminates common bugs caused by complex distributed state. You'll also explore the FP approach to IO, concurrency, and data streaming. By the time you finish, you'll be writing clean functional code that's easy to understand, test, and maintain. What's inside Designing with functions and types instead of objects Programming with pure functions and immutable values Writing concurrent programs using the functional style Testing functional programs About the reader For developers who know an object-oriented language. Examples in Java and Scala. About the author Michal Plachta is an experienced software developer who regularly speaks and writes about creating maintainable applications. Table of Contents Part 1 The functional toolkit 1 Learning functional programming 2 Pure functions 3 Immutable values 4 Functions as values Part 2 Functional programs 5 Sequential programs 6 Error handling 7 Requirements as types 8 IO as values 9 Streams as values 10 Concurrent programs Part 3 Applied functional programming 11 Designing functional programs 12 Testing functional programs

grokking deep learning pdf: Advanced Algorithms and Data Structures Marcello La

Rocca, 2021-08-10 Advanced Algorithms and Data Structures introduces a collection of algorithms for complex programming challenges in data analysis, machine learning, and graph computing.

Summary As a software engineer, you'll encounter countless programming challenges that initially seem confusing, difficult, or even impossible. Don't despair! Many of these "new" problems already have well-established solutions. Advanced Algorithms and Data Structures teaches you powerful approaches to a wide range of tricky coding challenges that you can adapt and apply to your own applications. Providing a balanced blend of classic, advanced, and new algorithms, this practical guide upgrades your programming toolbox with new perspectives and hands-on techniques.

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Can you improve the speed and efficiency of your applications without investing in new hardware? Well, yes, you can: Innovations in algorithms and data

structures have led to huge advances in application performance. Pick up this book to discover a collection of advanced algorithms that will make you a more effective developer. About the book Advanced Algorithms and Data Structures introduces a collection of algorithms for complex programming challenges in data analysis, machine learning, and graph computing. You'll discover cutting-edge approaches to a variety of tricky scenarios. You'll even learn to design your own data structures for projects that require a custom solution. What's inside Build on basic data structures you already know Profile your algorithms to speed up application Store and query strings efficiently Distribute clustering algorithms with MapReduce Solve logistics problems using graphs and optimization algorithms About the reader For intermediate programmers. About the author Marcello La Rocca is a research scientist and a full-stack engineer. His focus is on optimization algorithms, genetic algorithms, machine learning, and quantum computing. Table of Contents 1 Introducing data structures PART 1 IMPROVING OVER BASIC DATA STRUCTURES 2 Improving priority queues: d-way heaps 3 Treaps: Using randomization to balance binary search trees 4 Bloom filters: Reducing the memory for tracking content 5 Disjoint sets: Sub-linear time processing 6 Trie, radix trie: Efficient string search 7 Use case: LRU cache PART 2 MULTIDIMENSIONAL QUERIES 8 Nearest neighbors search 9 K-d trees: Multidimensional data indexing 10 Similarity Search Trees: Approximate nearest neighbors search for image retrieval 11 Applications of nearest neighbor search 12 Clustering 13 Parallel clustering: MapReduce and canopy clustering PART 3 PLANAR GRAPHS AND MINIMUM CROSSING NUMBER 14 An introduction to graphs: Finding paths of minimum distance 15 Graph embeddings and planarity: Drawing graphs with minimal edge intersections 16 Gradient descent: Optimization problems (not just) on graphs 17 Simulated annealing: Optimization beyond local minima 18 Genetic algorithms: Biologically inspired, fast-converging optimization

grokking deep learning pdf: Deep Learning and the Game of Go Kevin Ferguson, Max Pumperla, 2019-01-06 Summary Deep Learning and the Game of Go teaches you how to apply the power of deep learning to complex reasoning tasks by building a Go-playing AI. After exposing you to the foundations of machine and deep learning, you'll use Python to build a bot and then teach it the rules of the game. Foreword by Thore Graepel, DeepMind Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology The ancient strategy game of Go is an incredible case study for AI. In 2016, a deep learning-based system shocked the Go world by defeating a world champion. Shortly after that, the upgraded AlphaGo Zero crushed the original bot by using deep reinforcement learning to master the game. Now, you can learn those same deep learning techniques by building your own Go bot! About the Book Deep Learning and the Game of Go introduces deep learning by teaching you to build a Go-winning bot. As you progress, you'll apply increasingly complex training techniques and strategies using the Python deep learning library Keras. You'll enjoy watching your bot master the game of Go, and along the way, you'll discover how to apply your new deep learning skills to a wide range of other scenarios! What's inside Build and teach a self-improving game AI Enhance classical game AI systems with deep learning Implement neural networks for deep learning About the Reader All you need are basic Python skills and high school-level math. No deep learning experience

required. About the Author Max Pumperla and Kevin Ferguson are experienced deep learning specialists skilled in distributed systems and data science. Together, Max and Kevin built the open source bot BetaGo. Table of Contents PART 1 - FOUNDATIONS Toward deep learning: a machine-learning introduction Go as a machine-learning problem Implementing your first Go bot PART 2 - MACHINE LEARNING AND GAME AI Playing games with tree search Getting started with neural networks Designing a neural network for Go data Learning from data: a deep-learning bot Deploying bots in the wild Learning by practice: reinforcement learning Reinforcement learning with policy gradients Reinforcement learning with value methods Reinforcement learning with actor-critic methods PART 3 - GREATER THAN THE SUM OF ITS PARTS AlphaGo: Bringing it all together AlphaGo Zero: Integrating tree search with reinforcement learning

grokking deep learning pdf: Machine Learning with TensorFlow, Second Edition

Mattmann A. Chris, 2021-02-02 Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape

grokking deep learning pdf: Machine Learning Design Patterns Valliappa Lakshmanan, Sara Robinson, Michael Munn, 2020-10-15 The design patterns in this book capture best practices and solutions to recurring problems in machine learning. The authors, three Google engineers, catalog proven methods to help data scientists tackle common problems throughout the ML process. These design patterns codify the experience of hundreds of experts into straightforward, approachable advice. In this book, you will find detailed explanations of 30 patterns for data and problem

representation, operationalization, repeatability, reproducibility, flexibility, explainability, and fairness. Each pattern includes a description of the problem, a variety of potential solutions, and recommendations for choosing the best technique for your situation. You'll learn how to: Identify and mitigate common challenges when training, evaluating, and deploying ML models Represent data for different ML model types, including embeddings, feature crosses, and more Choose the right model type for specific problems Build a robust training loop that uses checkpoints, distribution strategy, and hyperparameter tuning Deploy scalable ML systems that you can retrain and update to reflect new data Interpret model predictions for stakeholders and ensure models are treating users fairly

grokking deep learning pdf: Reinforcement Learning Phil Winder Ph.D., 2020-11-06

Reinforcement learning (RL) will deliver one of the biggest breakthroughs in AI over the next decade, enabling algorithms to learn from their environment to achieve arbitrary goals. This exciting development avoids constraints found in traditional machine learning (ML) algorithms. This practical book shows data science and AI professionals how to learn by reinforcement and enable a machine to learn by itself. Author Phil Winder of Winder Research covers everything from basic building blocks to state-of-the-art practices. You'll explore the current state of RL, focus on industrial applications, learn numerous algorithms, and benefit from dedicated chapters on deploying RL solutions to production. This is no cookbook; doesn't shy away from math and expects familiarity with ML. Learn what RL is and how the algorithms help solve problems Become grounded in RL fundamentals including Markov decision processes, dynamic programming, and temporal difference learning Dive deep into a range of value and policy gradient methods Apply advanced RL solutions such as meta learning, hierarchical learning, multi-agent, and imitation learning Understand cutting-edge deep RL algorithms including Rainbow, PPO, TD3, SAC, and more Get practical examples through the accompanying website

grokking deep learning pdf: Deep Learning with PyTorch Vishnu Subramanian, 2018-02-23

Build neural network models in text, vision and advanced analytics using PyTorch Key Features Learn PyTorch for implementing cutting-edge deep learning algorithms. Train your neural networks for higher speed and flexibility and learn how to implement them in various scenarios; Cover various advanced neural network architecture such as ResNet, Inception, DenseNet and more with practical examples; Book Description Deep learning powers the most intelligent systems in the world, such as Google Voice, Siri, and Alexa. Advancements in powerful hardware, such as GPUs, software frameworks such as PyTorch, Keras, Tensorflow, and CNTK along with the availability of big data have made it easier to implement solutions to problems in the areas of text, vision, and advanced analytics. This book will get you up and running with one of the most cutting-edge deep learning libraries—PyTorch. PyTorch is grabbing the attention of deep learning researchers and data science professionals due to its accessibility, efficiency and being more native to Python way of development. You'll start off by installing PyTorch, then quickly move on to learn various fundamental blocks that power modern deep learning. You will also learn how to use CNN, RNN, LSTM and other networks to solve real-world problems. This book explains the concepts of various state-of-the-art deep learning architectures, such as ResNet, DenseNet, Inception, and Seq2Seq, without diving deep into the math behind them. You will also learn about GPU computing during the course of the book. You will see how to train a model with PyTorch and dive into complex neural networks such as generative networks for producing text and images. By the end of the book, you'll be able to implement deep learning applications in PyTorch with ease. What you will learn Use PyTorch for GPU-accelerated tensor computations Build custom datasets and data loaders for images and test the models using torchvision and torchtext Build an image classifier by implementing CNN architectures using PyTorch Build systems that do text classification and language modeling using RNN, LSTM, and GRU Learn advanced CNN architectures such as ResNet, Inception, Densenet, and learn how to use them for transfer learning Learn how to mix multiple models for a powerful ensemble model Generate new images using GAN's and generate artistic images using style transfer Who this book is for This book is for machine learning engineers, data analysts, data scientists interested in deep learning and are looking to explore implementing advanced algorithms in

PyTorch. Some knowledge of machine learning is helpful but not a mandatory need. Working knowledge of Python programming is expected.

grokking deep learning pdf: Building Machine Learning Powered Applications

Emmanuel Ameisen, 2020-01-21 Learn the skills necessary to design, build, and deploy applications powered by machine learning (ML). Through the course of this hands-on book, you'll build an example ML-driven application from initial idea to deployed product. Data scientists, software engineers, and product managers—including experienced practitioners and novices alike—will learn the tools, best practices, and challenges involved in building a real-world ML application step by step. Author Emmanuel Ameisen, an experienced data scientist who led an AI education program, demonstrates practical ML concepts using code snippets, illustrations, screenshots, and interviews with industry leaders. Part I teaches you how to plan an ML application and measure success. Part II explains how to build a working ML model. Part III demonstrates ways to improve the model until it fulfills your original vision. Part IV covers deployment and monitoring strategies. This book will help you: Define your product goal and set up a machine learning problem Build your first end-to-end pipeline quickly and acquire an initial dataset Train and evaluate your ML models and address performance bottlenecks Deploy and monitor your models in a production environment

grokking deep learning pdf: Advanced Deep Learning with Python Ivan Vasilev, 2019-12-12

Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem Key FeaturesGet to grips with building faster and more robust deep learning architecturesInvestigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorchApply deep neural networks (DNNs) to computer vision problems, NLP, and GANsBook Description In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks, such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world. What you will learnCover advanced and state-of-the-art neural network architecturesUnderstand the theory and math behind neural networksTrain DNNs and apply them to modern deep learning problemsUse CNNs for object detection and image segmentationImplement generative adversarial networks (GANs) and variational autoencoders to generate new imagesSolve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence modelsUnderstand DL techniques, such as meta-learning and graph neural networksWho this book is for This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

grokking deep learning pdf: Machine Learning in Action Peter Harrington, 2012-04-03

Summary Machine Learning in Action is unique book that blends the foundational theories of machine learning with the practical realities of building tools for everyday data analysis. You'll use the flexible Python programming language to build programs that implement algorithms for data

classification, forecasting, recommendations, and higher-level features like summarization and simplification. About the Book A machine is said to learn when its performance improves with experience. Learning requires algorithms and programs that capture data and ferret out the interesting or useful patterns. Once the specialized domain of analysts and mathematicians, machine learning is becoming a skill needed by many. Machine Learning in Action is a clearly written tutorial for developers. It avoids academic language and takes you straight to the techniques you'll use in your day-to-day work. Many (Python) examples present the core algorithms of statistical data processing, data analysis, and data visualization in code you can reuse. You'll understand the concepts and how they fit in with tactical tasks like classification, forecasting, recommendations, and higher-level features like summarization and simplification. Readers need no prior experience with machine learning or statistical processing. Familiarity with Python is helpful. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside A no-nonsense introduction Examples showing common ML tasks Everyday data analysis Implementing classic algorithms like Apriori and Adaboos Table of Contents PART 1 CLASSIFICATION Machine learning basics Classifying with k-Nearest Neighbors Splitting datasets one feature at a time: decision trees Classifying with probability theory: naïve Bayes Logistic regression Support vector machines Improving classification with the AdaBoost meta algorithm PART 2 FORECASTING NUMERIC VALUES WITH REGRESSION Predicting numeric values: regression Tree-based regression PART 3 UNSUPERVISED LEARNING Grouping unlabeled items using k-means clustering Association analysis with the Apriori algorithm Efficiently finding frequent itemsets with FP-growth PART 4 ADDITIONAL TOOLS Using principal component analysis to simplify data Simplifying data with the singular value decomposition Big data and MapReduce

grokking deep learning pdf: Probabilistic Deep Learning Oliver Duerr, Beate Sick, Elvis Murina, 2020-11-10 Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. Summary Probabilistic Deep Learning: With Python, Keras and TensorFlow Probability teaches the increasingly popular probabilistic approach to deep learning that allows you to refine your results more quickly and accurately without much trial-and-error testing. Emphasizing practical techniques that use the Python-based Tensorflow Probability Framework, you'll learn to build highly-performant deep learning applications that can reliably handle the noise and uncertainty of real-world data. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology The world is a noisy and uncertain place. Probabilistic deep learning models capture that noise and uncertainty, pulling it into real-world scenarios. Crucial for self-driving cars and scientific testing, these techniques help deep learning engineers assess the accuracy of their results, spot errors, and improve their understanding of how algorithms work. About the book Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. What's inside Explore maximum likelihood and the statistical basis of deep learning Discover probabilistic models that can indicate possible outcomes Learn to use normalizing flows for modeling and generating complex distributions Use Bayesian neural networks to access the uncertainty in the model About the reader For experienced machine learning developers. About the author Oliver Dürr is a professor at the University of Applied Sciences in Konstanz, Germany. Beate Sick holds a chair for applied statistics at ZHAW and works as a researcher and lecturer at the University of Zurich. Elvis Murina is a data scientist. Table of Contents PART 1 - BASICS OF DEEP LEARNING 1 Introduction to probabilistic deep learning 2 Neural network architectures 3 Principles of curve fitting PART 2 - MAXIMUM LIKELIHOOD APPROACHES FOR PROBABILISTIC DL MODELS 4 Building loss functions with the likelihood approach 5 Probabilistic deep learning models with

TensorFlow Probability 6 Probabilistic deep learning models in the wild PART 3 - BAYESIAN APPROACHES FOR PROBABILISTIC DL MODELS 7 Bayesian learning 8 Bayesian neural networks

Related Articles

- [hampton bay 1-855-hd-hampton](#)
- [gizmo mouse genetics two traits](#)
- [hobbing machine gear calculation pdf](#)

[Back to Home](#)