
Make Your Own Neural Network

Make Your Own Neural Network

Downloaded from api.delirest.hu by Gina & Mya

DOWNLOAD AND INSTALL MAKE YOUR OWN NEURAL NETWORK AND DISCOVER A MYRIAD OF LITERATURE WITHIN YOUR REACHES

Invite to our website, where you can quickly **download Make Your Own Neural Network publication** selections that accommodate your **analysis taste** - done in one convenient place. With simply a couple of clicks, you can immediately access a diverse series of **Make Your Own Neural Network literature** and take pleasure in hours of reading pleasure.

Gone are the days of searching multiple web sites or heading to the bookstore to locate your next read. Our website supplies a problem-free experience that places a myriad of publications at your **fingertips**. Say goodbye to the lengthy process of searching for your favored books like Make Your Own Neural Network and hello there to the comfort of downloading them easily.

Explore our website's substantial collection of fiction, non-fiction, love, secret, and various other categories that match your **reading preference** by seeing us today. Discover new authors or discover the current launches all in one area at our **api.delirest.hu**. Start your publication journey currently and allow us be your go-to for all your literary needs.

EXPLORE A MYRIAD OF LITERARY WORKS

Are you tired of checking out heaps of books, looking for your following read? Look no more than our website for a vast choice of literature that caters to your reading taste. We offer a varied variety of categories, from timeless literary works to contemporary fiction, non-fiction, romance, mystery, and a lot more.

Our downloadable Make Your Own Neural Network period a multitude of subjects, making certain that there's something for everybody. From biographies to sci-fi, from background to self-help, our collection has all of it. With simply a couple of clicks, you can check out the various classifications and discover the ideal publication like Make Your Own Neural Network to download.

And the best part? You can access every one of this literary works from the convenience of your own home. No more driving to the bookstore or waiting in line at the library. With our website, you can download and install Make Your Own Neural Network directly to your gadget and begin reading instantly. So why wait? Discover your next preferred read today!

EASY DOWNLOAD PROCESS OF MAKE YOUR OWN NEURAL NETWORK

Are you all set to begin downloading and install Make Your Own Neural Network? Our site provides a simple and hassle-free download procedure that you can begin today. Initially, develop an account

with us by signing up on our site. Once you're logged in, you can search our huge collection of books and locate the excellent literature that suits your reading taste.

Once you've located guide Make Your Own Neural Network you wish to download, merely click the download button. Our site makes sure that the downloading and install process is quick and reliable, so you can start reading your favorite books in no time at all.

Our straightforward platform is designed to supply you with a smooth experience, making it easy for you to download Make Your Own Neural Network and begin reading right now. You do not require to be tech-savvy to use our website - we give step-by-step directions to assist you navigate with the process.

So what are you waiting on? Begin your book trip today by downloading **Make Your Own Neural Network** from our website. With our simple download procedure, you'll have the ability to access your analysis product quickly. Happy reading!

WIDE OPTION OF PUBLICATION STYLES

At our site, we recognize the relevance of accommodating your analysis choices. That's why we offer a wide selection of Make Your Own Neural Network publication styles for you to pick from. Whether you choose the classic PDF, the versatile EPUB, or the practical MOBI, we have actually obtained you covered. Not just that, we likewise support various other popular styles to make certain compatibility across different tools.

With our substantial variety of styles, you can appreciate your downloaded Make Your Own Neural Network book effortlessly on your e-reader, tablet computer, or mobile phone without any trouble. So, go on and choose the format that matches your reading preference and begin downloading your favored literary works today!

STAY GOTTEN IN TOUCH WITH NEW RELEASES

Do not lose out on the most recent literary treasures! By staying gotten in touch with us, you can uncover brand-new releases and stay up to date with your preferred writers.

To make certain you never ever miss a beat, register for our newsletter or follow us on social media sites - you'll be the very first to learn about exciting book launches, writer meetings, and special offers.

Our selection of downloadable Make Your Own Neural Network is constantly expanding, so see to it to remain linked to locate your next terrific read that fits your unique analysis taste.

Join our area today and begin your trip into the globe of literature with simple downloads of all your

avored publications like **Make Your Own Neural Network!**

Make Your Own Neural Network *Createspace Independent Publishing Platform*

A step-by-step gentle journey through the mathematics of neural networks, and making your own using the Python computer language. Neural networks are a key element of deep learning and artificial intelligence, which today is capable of some truly impressive feats. Yet too few really understand how neural networks actually work. This guide will take you on a fun and unhurried journey, starting from very simple ideas, and gradually building up an understanding of how neural networks work. You won't need any mathematics beyond secondary school, and an accessible introduction to calculus is also included. The ambition of this guide is to make neural networks as accessible as possible to as many readers as possible - there are enough texts for advanced readers already! You'll learn to code in Python and make your own neural network, teaching it to recognise human handwritten numbers, and performing as well as professionally developed networks. Part 1 is about ideas. We introduce the mathematical ideas underlying the neural networks, gently with lots of illustrations and examples. Part 2 is practical. We introduce the popular and easy to learn Python programming language, and gradually builds up a neural network which can learn to recognise human handwritten numbers, easily getting it to perform as well as networks made by professionals. Part 3 extends these ideas further. We push the performance of our neural network to an industry leading 98% using only simple ideas and code, test the network on your own handwriting, take a privileged peek inside the mysterious mind of a neural network, and even get it all working on a Raspberry Pi. All the code in this has been tested to work on a Raspberry Pi Zero.

Make Your Own Neural Network: An In-Depth Visual Introduction for Beginners *Independently Published*

A step-by-step visual journey through the mathematics of neural networks, and making your own using Python and Tensorflow. What you will gain from this book: * A deep understanding of how a Neural Network works. * How to build a Neural Network from scratch using Python. Who this book is for: * Beginners who want to fully understand how networks work, and learn to build two step-by-step examples in Python. * Programmers who need an easy to read, but solid refresher, on the math of neural networks. What's Inside - 'Make Your Own Neural Network: An Indepth Visual Introduction For Beginners' What Is a Neural Network? Neural networks have made a gigantic comeback in the last few decades and you likely make use of them everyday without realizing it, but what exactly is a neural network? What is it used for and how does it fit within the broader arena of machine learning? we gently explore these topics so that we can be prepared to dive deep further on. To start, we'll begin with a high-level overview of machine learning and then drill down into the specifics of a neural network. The Math of Neural Networks On a high level, a network learns just like we do, through trial and error. This is true regardless if the network is supervised, unsupervised, or semi-supervised. Once we dig a bit deeper though, we discover that a handful of mathematical functions play a major role in the trial and error process. It also becomes clear that a grasp of the underlying mathematics helps clarify how a network learns. * Forward Propagation * Calculating The Total Error

* Calculating The Gradients * Updating The Weights Make Your Own Artificial Neural Network: Hands on Example You will learn to build a simple neural network using all the concepts and functions we learned in the previous few chapters. Our example will be basic but hopefully very intuitive. Many examples available online are either hopelessly abstract or make use of the same data sets, which can be repetitive. Our goal is to be crystal clear and engaging, but with a touch of fun and uniqueness. This section contains the following eight chapters. Building Neural Networks in Python There are many ways to build a neural network and lots of tools to get the job done. This is fantastic, but it can also be overwhelming when you start, because there are so many tools to choose from. We are going to take a look at what tools are needed and help you nail down the essentials. To build a neural network Tensorflow and Neural Networks There is no single way to build a feedforward neural network with Python, and that is especially true if you throw Tensorflow into the mix. However, there is a general framework that exists that can be divided into five steps and grouped into two parts. We are going to briefly explore these five steps so that we are prepared to use them to build a network later on. Ready? Let's begin. Neural Network: Distinguish Handwriting We are going to dig deep with Tensorflow and build a neural network that can distinguish between handwritten numbers. We'll use the same 5 steps we covered in the high-level overview, and we are going to take time exploring each line of code. Neural Network: Classify Images 10 minutes. That's all it takes to build an image classifier thanks to Google! We will provide a high-level overview of how to classify images using a convolutional neural network (CNN) and Google's Inception V3 model. Once finished, you will be able to tweak this code to classify any type of image sets! Cats, bats, super heroes - the sky's the limit.

Neural Network Projects with Python

The ultimate guide to using Python to explore the true power of neural networks through six projects *Packt Publishing Ltd*

Build your Machine Learning portfolio by creating 6 cutting-edge Artificial Intelligence projects using neural networks in Python Key Features Discover neural network architectures (like CNN and LSTM) that are driving recent advancements in AI Build expert neural networks in Python using popular libraries such as Keras Includes projects such as object detection, face identification, sentiment analysis, and more Book Description Neural networks are at the core of recent AI advances, providing some of the best resolutions to many real-world problems, including image recognition, medical diagnosis, text analysis, and more. This book goes through some basic neural network and deep learning concepts, as well as some popular libraries in Python for implementing them. It contains practical demonstrations of neural networks in domains such as fare prediction, image classification, sentiment analysis, and more. In each case, the book provides a problem statement, the specific neural network architecture required to tackle that problem, the reasoning behind the algorithm used, and the associated Python code to implement the solution from scratch. In the process, you will gain hands-on experience with using popular Python libraries such as Keras to build and train your own neural networks from scratch. By the end of this book, you will have mastered the different neural network architectures and created cutting-edge AI projects in Python that will

immediately strengthen your machine learning portfolio. What you will learn Learn various neural network architectures and its advancements in AI Master deep learning in Python by building and training neural network Master neural networks for regression and classification Discover convolutional neural networks for image recognition Learn sentiment analysis on textual data using Long Short-Term Memory Build and train a highly accurate facial recognition security system Who this book is for This book is a perfect match for data scientists, machine learning engineers, and deep learning enthusiasts who wish to create practical neural network projects in Python. Readers should already have some basic knowledge of machine learning and neural networks.

Tensorflow in 1 Day: Make Your Own Neural Network *Independently Published*

Tensorflow is the most popular Deep Learning Library out there. It has fantastic graph computations feature which helps data scientist to visualize his designed neural network using TensorBoard. This Machine learning library supports both Convolution as well as Recurrent Neural network. It supports parallel processing on CPU as well as GPU. Prominent machine learning algorithms supported by TensorFlow are Deep Learning Classification, wip & deep, Boston Tree amongst others. The book is very hands-on and gives you industry ready deep learnings practices. Here is what is covered in the book

Artificial Intelligence with Python *Packt Publishing Ltd*

Build real-world Artificial Intelligence applications with Python to intelligently interact with the world around you About This Book Step into the amazing world of intelligent apps using this comprehensive guide Enter the world of Artificial Intelligence, explore it, and create your own applications Work through simple yet insightful examples that will get you up and running with Artificial Intelligence in no time Who This Book Is For This book is for Python developers who want to build real-world Artificial Intelligence applications. This book is friendly to Python beginners, but being familiar with Python would be useful to play around with the code. It will also be useful for experienced Python programmers who are looking to use Artificial Intelligence techniques in their existing technology stacks. What You Will Learn Realize different classification and regression techniques Understand the concept of clustering and how to use it to automatically segment data See how to build an intelligent recommender system Understand logic programming and how to use it Build automatic speech recognition systems Understand the basics of heuristic search and genetic programming Develop games using Artificial Intelligence Learn how reinforcement learning works Discover how to build intelligent applications centered on images, text, and time series data See how to use deep learning algorithms and build applications based on it In Detail Artificial Intelligence is becoming increasingly relevant in the modern world where everything is driven by technology and data. It is used extensively across many fields such as search engines, image recognition, robotics, finance, and so on. We will explore various real-world scenarios in this book and you'll learn about various algorithms that can be used to build Artificial Intelligence applications. During the course of this book, you will find out how to make informed decisions about what algorithms to use in a given context. Starting from the basics of Artificial Intelligence, you will learn how to develop various building blocks using different data mining techniques. You will see how to implement different

algorithms to get the best possible results, and will understand how to apply them to real-world scenarios. If you want to add an intelligence layer to any application that's based on images, text, stock market, or some other form of data, this exciting book on Artificial Intelligence will definitely be your guide! Style and approach This highly practical book will show you how to implement Artificial Intelligence. The book provides multiple examples enabling you to create smart applications to meet the needs of your organization. In every chapter, we explain an algorithm, implement it, and then build a smart application.

Make Your Own Neural Network

Tariq Rashid

An Introduction to Neural Networks *CRC Press*

Though mathematical ideas underpin the study of neural networks, the author presents the fundamentals without the full mathematical apparatus. All aspects of the field are tackled, including artificial neurons as models of their real counterparts; the geometry of network action in pattern space; gradient descent methods, including back-propagation; associative memory and Hopfield nets; and self-organization and feature maps. The traditionally difficult topic of adaptive resonance theory is clarified within a hierarchical description of its operation. The book also includes several real-world examples to provide a concrete focus. This should enhance its appeal to those involved in the design, construction and management of networks in commercial environments and who wish to improve their understanding of network simulator packages. As a comprehensive and highly accessible introduction to one of the most important topics in cognitive and computer science, this volume should interest a wide range of readers, both students and professionals, in cognitive science, psychology, computer science and electrical engineering.

Neural Networks for Beginners

An Easy Textbook for Machine Learning Fundamentals to Guide You Implementing Neural Networks with Python and Deep Learning

Do you want to understand Neural Networks and learn everything about them but it looks like it is an exclusive club? Are you fascinated by Artificial Intelligence but you think that it would be too difficult for you to learn? If you think that Neural Networks and Artificial Intelligence are the present and, even more, the future of technology, and you want to be part of it... well you are in the right place, and you are looking at the right book. If you are reading these lines you have probably already noticed this: Artificial Intelligence is all around you. Your smartphone that suggests you the next word you want to type, your Netflix account that recommends you the series you may like or Spotify's personalised playlists. This is how machines are learning from you in everyday life. And these examples are only the surface of this technological revolution. Either if you want to start your own AI enterprise, to empower your business or to work in the greatest and most innovative companies, Artificial Intelligence is the future, and Neural Networks programming is the skill you

want to have. The good news is that there is no exclusive club, you can easily (if you commit, of course) learn how to program and use neural networks, and to do that *Neural Networks for Beginners* is the perfect way. In this book you will learn: The types and components of neural networks The smartest way to approach neural network programming Why Algorithms are your friends The "three Vs" of Big Data (plus two new Vs) How machine learning will help you making predictions The three most common problems with Neural Networks and how to overcome them Even if you don't know anything about programming, *Neural Networks* is the perfect place to start now. Still, if you already know about programming but not about how to do it in Artificial Intelligence, neural networks are the next thing you want to learn. And *Neural Networks for Beginners* is the best way to do it. Download *Neural Network for Beginners* now to get the best start for your journey to Artificial Intelligence. Scroll to the top of the page and click the BUY NOW button.

The Math of Neural Networks *Independently Published*

There are many reasons why neural networks fascinate us and have captivated headlines in recent years. They make web searches better, organize photos, and are even used in speech translation. Heck, they can even generate encryption. At the same time, they are also mysterious and mind-bending: how exactly do they accomplish these things? What goes on inside a neural network? On a high level, a network learns just like we do, through trial and error. This is true regardless if the network is supervised, unsupervised, or semi-supervised. Once we dig a bit deeper though, we discover that a handful of mathematical functions play a major role in the trial and error process. It also becomes clear that a grasp of the underlying mathematics helps clarify how a network learns. In the following chapters we will unpack the mathematics that drive a neural network. To do this, we will use a feedforward network as our model and follow input as it moves through the network.

Neural Networks and Deep Learning

A Textbook *Springer*

This book covers both classical and modern models in deep learning. The primary focus is on the theory and algorithms of deep learning. The theory and algorithms of neural networks are particularly important for understanding important concepts, so that one can understand the important design concepts of neural architectures in different applications. Why do neural networks work? When do they work better than off-the-shelf machine-learning models? When is depth useful? Why is training neural networks so hard? What are the pitfalls? The book is also rich in discussing different applications in order to give the practitioner a flavor of how neural architectures are designed for different types of problems. Applications associated with many different areas like recommender systems, machine translation, image captioning, image classification, reinforcement-learning based gaming, and text analytics are covered. The chapters of this book span three categories: The basics of neural networks: Many traditional machine learning models can be understood as special cases of neural networks. An emphasis is placed in the first two chapters on understanding the relationship between traditional machine learning and neural networks. Support vector machines, linear/logistic regression, singular value decomposition, matrix factorization, and

recommender systems are shown to be special cases of neural networks. These methods are studied together with recent feature engineering methods like word2vec. Fundamentals of neural networks: A detailed discussion of training and regularization is provided in Chapters 3 and 4. Chapters 5 and 6 present radial-basis function (RBF) networks and restricted Boltzmann machines. Advanced topics in neural networks: Chapters 7 and 8 discuss recurrent neural networks and convolutional neural networks. Several advanced topics like deep reinforcement learning, neural Turing machines, Kohonen self-organizing maps, and generative adversarial networks are introduced in Chapters 9 and 10. The book is written for graduate students, researchers, and practitioners. Numerous exercises are available along with a solution manual to aid in classroom teaching. Where possible, an application-centric view is highlighted in order to provide an understanding of the practical uses of each class of techniques.

Deep Learning Illustrated

A Visual, Interactive Guide to Artificial Intelligence *Addison-Wesley Professional*

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

Neural Networks

A Systematic Introduction *Springer Science & Business Media*

Neural networks are a computing paradigm that is finding increasing attention among computer scientists. In this book, theoretical laws and models previously scattered in the literature are brought together into a general theory of artificial neural nets. Always with a view to biology and starting with the simplest nets, it is shown how the properties of models change when more general computing elements and net topologies are introduced. Each chapter contains examples, numerous illustrations, and a bibliography. The book is aimed at readers who seek an overview of the field or who wish to deepen their knowledge. It is suitable as a basis for university courses in neurocomputing.

Deep Learning with Python *Simon and Schuster*

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

Introduction to Deep Learning**From Logical Calculus to Artificial Intelligence** *Springer*

This textbook presents a concise, accessible and engaging first introduction to deep learning, offering a wide range of connectionist models which represent the current state-of-the-art. The text explores the most popular algorithms and architectures in a simple and intuitive style, explaining the mathematical derivations in a step-by-step manner. The content coverage includes convolutional networks, LSTMs, Word2vec, RBMs, DBNs, neural Turing machines, memory networks and autoencoders. Numerous examples in working Python code are provided throughout the book, and the code is also supplied separately at an accompanying website. Topics and features: introduces the fundamentals of machine learning, and the mathematical and computational prerequisites for deep learning; discusses feed-forward neural networks, and explores the modifications to these which can be applied to any neural network; examines convolutional neural networks, and the recurrent connections to a feed-forward neural network; describes the notion of distributed representations, the concept of the autoencoder, and the ideas behind language processing with deep learning; presents a brief history of artificial intelligence and neural networks, and reviews interesting open research problems in deep learning and connectionism. This clearly written and lively primer on deep learning is essential reading for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as fields such as linguistics, logic, philosophy, and psychology.

Grokking Deep Learning *Simon and Schuster*

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

Deep Learning *MIT Press*

An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. “Written by three experts in the field, Deep Learning is the only comprehensive book on the subject.” —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX

Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

Neural Network Design

Deep Learning and the Game of Go *Simon and Schuster*

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

Deep Learning for Natural Language Processing

Creating Neural Networks with Python *Apress*

Discover the concepts of deep learning used for natural language processing (NLP), with full-fledged examples of neural network models such as recurrent neural networks, long short-term memory networks, and sequence-2-sequence models. You'll start by covering the mathematical prerequisites and the fundamentals of deep learning and NLP with practical examples. The first three chapters of the book cover the basics of NLP, starting with word-vector representation before moving onto advanced algorithms. The final chapters focus entirely on implementation, and deal with sophisticated architectures such as RNN, LSTM, and Seq2seq, using Python tools: TensorFlow, and Keras. Deep Learning for Natural Language Processing follows a progressive approach and combines all the knowledge you have gained to build a question-answer chatbot system. This book is a good starting point for people who want to get started in deep learning for NLP. All the code presented in

the book will be available in the form of IPython notebooks and scripts, which allow you to try out the examples and extend them in interesting ways. What You Will Learn Gain the fundamentals of deep learning and its mathematical prerequisites Discover deep learning frameworks in Python Develop a chatbot Implement a research paper on sentiment classification Who This Book Is For Software developers who are curious to try out deep learning with NLP.

Hands-On Neural Networks

Learn How to Build and Train Your First Neural Network Model Using Python

Design and create neural networks with deep learning and artificial intelligence principles using OpenAI Gym, TensorFlow, and Keras Key Features Explore neural network architecture and understand how it functions Learn algorithms to solve common problems using back propagation and perceptrons Understand how to apply neural networks to applications with the help of useful illustrations Book Description Neural networks play a very important role in deep learning and artificial intelligence (AI), with applications in a wide variety of domains, right from medical diagnosis, to financial forecasting, and even machine diagnostics. Hands-On Neural Networks is designed to guide you through learning about neural networks in a practical way. The book will get you started by giving you a brief introduction to perceptron networks. You will then gain insights into machine learning and also understand what the future of AI could look like. Next, you will study how embeddings can be used to process textual data and the role of long short-term memory networks (LSTMs) in helping you solve common natural language processing (NLP) problems. The later chapters will demonstrate how you can implement advanced concepts including transfer learning, generative adversarial networks (GANs), autoencoders, and reinforcement learning. Finally, you can look forward to further content on the latest advancements in the field of neural networks. By the end of this book, you will have the skills you need to build, train, and optimize your own neural network model that can be used to provide predictable solutions. What you will learn Learn how to train a network by using backpropagation Discover how to load and transform images for use in neural networks Study how neural networks can be applied to a varied set of applications Solve common challenges faced in neural network development Understand the transfer learning concept to solve tasks using Keras and Visual Geometry Group (VGG) network Get up to speed with advanced and complex deep learning concepts like LSTMs and NLP Explore innovative algorithms like GANs and deep reinforcement learning Who this book is for If you are interested in artificial intelligence and deep learning and want to further your skills, then this intermediate-level book is for you. Some knowledge of statistics will help you get the most out of this book.

Hands-On Neural Network Programming with C#

Add powerful neural network capabilities to your C# enterprise applications Packt Publishing Ltd

Create and unleash the power of neural networks by implementing C# and .Net code Key Features Get a strong foundation of neural networks with access to various machine learning and deep

learning libraries Real-world case studies illustrating various neural network techniques and architectures used by practitioners Cutting-edge coverage of Deep Networks, optimization algorithms, convolutional networks, autoencoders and many more Book Description Neural networks have made a surprise comeback in the last few years and have brought tremendous innovation in the world of artificial intelligence. The goal of this book is to provide C# programmers with practical guidance in solving complex computational challenges using neural networks and C# libraries such as CNTK, and TensorFlowSharp. This book will take you on a step-by-step practical journey, covering everything from the mathematical and theoretical aspects of neural networks, to building your own deep neural networks into your applications with the C# and .NET frameworks. This book begins by giving you a quick refresher of neural networks. You will learn how to build a neural network from scratch using packages such as Encog, Aforge, and Accord. You will learn about various concepts and techniques, such as deep networks, perceptrons, optimization algorithms, convolutional networks, and autoencoders. You will learn ways to add intelligent features to your .NET apps, such as facial and motion detection, object detection and labeling, language understanding, knowledge, and intelligent search. Throughout this book, you will be working on interesting demonstrations that will make it easier to implement complex neural networks in your enterprise applications. What you will learn Understand perceptrons and how to implement them in C# Learn how to train and visualize a neural network using cognitive services Perform image recognition for detecting and labeling objects using C# and TensorFlowSharp Detect specific image characteristics such as a face using Accord.Net Demonstrate particle swarm optimization using a simple XOR problem and Encog Train convolutional neural networks using ConvNetSharp Find optimal parameters for your neural network functions using numeric and heuristic optimization techniques. Who this book is for This book is for Machine Learning Engineers, Data Scientists, Deep Learning Aspirants and Data Analysts who are now looking to move into advanced machine learning and deep learning with C#. Prior knowledge of machine learning and working experience with C# programming is required to take most out of this book

Generative Deep Learning

Teaching Machines to Paint, Write, Compose, and Play "O'Reilly Media, Inc."

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

Ultimate Step by Step Guide to Deep Learning Using Python

Artificial Intelligence and Neural Network Concepts Explained in Simple Terms

*Start your Data Science career using Python today!*Are you ready to start your new exciting career? Ready to master artificial intelligence and deep learning concepts?Are you overwhelmed with complexity of the books on this subject?Then let this breezy and fun little book on Python, Machine Learning and Deep Learning models make you a Data Scientist in 7 days!This book continues from where the first book in the series, Ultimate Step by Step Guide to Machine Learning Using Python, left of. In the first book you were introduced to Python concepts such as: -Data Structures like Pandas -Foundational libraries like Numpy, Seaborn and Scikit-Learn-Regression analysis-Classification-Clustering-Association Learning-Dimension ReductionThis book builds on those concepts to expand on Machine Learning algorithms like: -Linear and Logistical regression-Decision tree-Support vector machines (SVM)After that, this book takes you on a journey into Deep Learning and Neural Networks with important concepts and libraries like: -Convolutional and Recurrent Neural Networks-TensorFlow-Keras-PyTorch-Keras-Apache MXNet-Microsoft Cognitive Toolkit (CNTK)The final part of the book covers all foundational concepts that are required for Amazon Web Services (AWS) Certified Machine Learning Specialization by explaining how to deploy your models at scale on Cloud technologies. While AWS is used in the book for illustrative purposes, Microsoft Azure and Google Cloud are also introduced as alternative cloud technologies. After reading this book you will be able to: -Code in Python with confidence-Build new machine learning and deep learning models from scratch-Know how to clean and prepare your data for analytics-Speak confidently about statistical analysis techniquesData Science was ranked the fast-growing field by LinkedIn and Data Scientist is one of the most highly sought after and lucrative careers in the world!If you are on the fence about making the leap to a new and lucrative career, this is the book for you!What sets this book apart from other books on the topic of Python and Machine learning: -Step by step code examples and explanation-Complex concepts explained visually-Real world applicability of the machine learning and deep learning models introducedWhat do I need to get started?You will have a step by step action plan in place once you finish this book and finally feel that you, can master data science and artificial intelligence and start a lucrative and rewarding career! Ready to dive in to the exciting world of Python and Deep Learning?Then scroll up to the top and hit that BUY BUTTON!

Data Structures and Algorithm Analysis in C++, Third Edition Courier Corporation

Comprehensive treatment focuses on creation of efficient data structures and algorithms and selection or design of data structure best suited to specific problems. This edition uses C++ as the programming language.

The Hundred-page Machine Learning Book

Provides a practical guide to get started and execute on machine learning within a few days without necessarily knowing much about machine learning.The first five chapters are enough to get you started and the next few chapters provide you a good feel of more advanced topics to pursue.

A Practical Guide to Neural Nets Prentice Hall

Based on a course given to internal managers at Texas Instruments, this book is an introduction to neural nets for computer science, artificial intelligence and R & D professionals, as well as MIS or DP managers.

Neural Networks and Deep Learning

Neural Networks and Deep Learning, Deep Learning Explained to Your Granny Createspace Independent Publishing Platform

Ready to crank up a neural network to get your self-driving car pick up the kids from school? Want to add 'Deep Learning' to your LinkedIn profile? Well, hold on there... Before you embark on your epic journey into the world of deep learning, there is basic theory to march through first! Take a step-by-step journey through the basics of Neural Networks and Deep Learning, made so simple that...even your granny could understand it! What you will gain from this book: * A deep understanding of how a Neural Network and Deep Learning work * A basics comprehension on how to build a Deep Neural Network from scratch Who this book is for: * Beginners who want to approach the topic, but are too afraid of complex math to start! What's Inside? * A brief introduction to Machine Learning * Two main Types of Machine Learning Algorithms * A practical example of Unsupervised Learning * What are Neural Networks? * McCulloch-Pitts's Neuron * Types of activation function * Types of network architectures * Learning processes * Advantages and disadvantages * Let us give a memory to our Neural Network * The example of book writing Software * Deep learning: the ability of learning to learn * How does Deep Learning work? * Main architectures and algorithms * Main types of DNN * Available Frameworks and libraries * Convolutional Neural Networks * Tunnel Vision * Convolution * The right Architecture for a Neural Network * Test your Neural Network * A general overview of Deep Learning * What are the limits of Deep Learning? * Deep Learning: the basics * Layers, Learning paradigms, Training, Validation * Main architectures and algorithms * Models for Deep Learning * Probabilistic graphic models * Restricted Boltzmann Machines * Deep Belief Networks * Available Frameworks and libraries * TensorFlow Hit download. Now!

Python Coding

An Introduction to Neural Networks and a Brief Overview of the Processes You Need to Know when Programming Computers and Coding with Python Independently Published

python codingWandering how to learn everything on Python Programming right from the beginning? The next few lines can tell you something! Learning Python is one of the 21st century specialties you

can have right now. You know how to code with Python, you become one of the most relevant citizens of the computer age. You can access neural networks, interpret, understand, code and decode certain special languages of a computer. So in order to be relevant, you need a program like python. And what is Python? Python is a specialized computer program that is used to code data into a computer. It supplies information into the computer in a specialized language. The computer then works on that information to produce desired response. This is exactly what is called Coding. So, Python is a very famous program used to code a computer. It can be used to perform various activities, ranging from basic mathematical calculations to website coding and data processing. Also computing an A. I. software. This field used to be restricted to Computer scientists, Engineers, Technicians and related fields originally. But today, everyone programs a computer and you can't afford not belonging to that class for long. Learning Python programming is your pathway to understanding neural networks and coding information into a computer. But learning the basic coding processes requires a lot of technicalities. That's something most trained specialists in the field find difficult to teach others, probably because it is hard to go to the basics if you are already an expert in the more complex fields. You should remember it is a very creative field too. It's a field where you can become so creative you'd code a complex program that would look mysterious to even the head of coding specialists. But remember, that is only if you understand the basic programming lessons itself. Since even specialists find it difficult to teach newbies, how can you learn? It is not difficult. Do you understand every lesson mentioned on programming so far? Then you will find it very easy, interactive and fascinating to learn python programming if you download: **DOWNLOAD: python coding.** Python coding is an introduction to neural networks and a brief overview of the processes that you need to know when programming computers and coding with python. It is a detailed book that introduced you to the techniques of python programming in a simple way. At your own pace, understand the basics of python programming. Click here to download your copy. This eBook is one sure bet to learn the complex lessons of coding without getting bored with technicalities. You probably won't finish it before the urge to yank off everything else and jump straight at your computer, coding with compassion. What specialties do you stand to learn? -Introduction to python machine.-The process of neural networks and a brief overview-Learn coding with python in computer programming-Organize data using effective pre-processing techniques-Get grips to a deeper textual and social media data-To optimize your machine learning systems and algorithms.You even get to learn how to organize data using various techniques. C'mon you can't afford not to have a copy. Be sure you scroll up and use the Download button!

Fundamentals of Neural Networks

Architectures, Algorithms, and Applications *Prentice Hall*

Providing detailed examples of simple applications, this new book introduces the use of neural networks. It covers simple neural nets for pattern classification; pattern association; neural networks based on competition; adaptive-resonance theory; and more. For professionals working with neural networks.

Machine Learning For Dummies *John Wiley & Sons*

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

Machine Learning with Neural Networks

An Introduction for Scientists and Engineers *Cambridge University Press*

This modern and self-contained book offers a clear and accessible introduction to the important topic of machine learning with neural networks. In addition to describing the mathematical principles of the topic, and its historical evolution, strong connections are drawn with underlying methods from statistical physics and current applications within science and engineering. Closely based around a well-established undergraduate course, this pedagogical text provides a solid understanding of the key aspects of modern machine learning with artificial neural networks, for students in physics, mathematics, and engineering. Numerous exercises expand and reinforce key concepts within the book and allow students to hone their programming skills. Frequent references to current research develop a detailed perspective on the state-of-the-art in machine learning research.

Better Deep Learning

Train Faster, Reduce Overfitting, and Make Better Predictions *Machine Learning Mastery*

Deep learning neural networks have become easy to define and fit, but are still hard to configure. Discover exactly how to improve the performance of deep learning neural network models on your

predictive modeling projects. With clear explanations, standard Python libraries, and step-by-step tutorial lessons, you'll discover how to better train your models, reduce overfitting, and make more accurate predictions.

Supervised Machine Learning for Text Analysis in R *CRC Press*

Text data is important for many domains, from healthcare to marketing to the digital humanities, but specialized approaches are necessary to create features for machine learning from language. Supervised Machine Learning for Text Analysis in R explains how to preprocess text data for modeling, train models, and evaluate model performance using tools from the tidyverse and tidymodels ecosystem. Models like these can be used to make predictions for new observations, to understand what natural language features or characteristics contribute to differences in the output, and more. If you are already familiar with the basics of predictive modeling, use the comprehensive, detailed examples in this book to extend your skills to the domain of natural language processing. This book provides practical guidance and directly applicable knowledge for data scientists and analysts who want to integrate unstructured text data into their modeling pipelines. Learn how to use text data for both regression and classification tasks, and how to apply more straightforward algorithms like regularized regression or support vector machines as well as deep learning approaches. Natural language must be dramatically transformed to be ready for computation, so we explore typical text preprocessing and feature engineering steps like tokenization and word embeddings from the ground up. These steps influence model results in ways we can measure, both in terms of model metrics and other tangible consequences such as how fair or appropriate model results are.

Decision Trees and Random Forests

A Visual Introduction for Beginners *Independently Published*

If you want to learn how decision trees and random forests work, plus create your own, this visual book is for you. The fact is, decision tree and random forest algorithms are powerful and likely touch your life everyday. From online search to product development and credit scoring, both types of algorithms are at work behind the scenes in many modern applications and services. They are also used in countless industries such as medicine, manufacturing and finance to help companies make better decisions and reduce risk. Whether coded or scratched out by hand, both algorithms are powerful tools that can make a significant impact. This book is a visual introduction for beginners that unpacks the fundamentals of decision trees and random forests. If you want to dig into the basics with a visual twist plus create your own algorithms in Python, this book is for you.

Efficient Processing of Deep Neural Networks *Morgan & Claypool Publishers*

This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high computational

complexity. Therefore, techniques that enable efficient processing of deep neural networks to improve metrics—such as energy-efficiency, throughput, and latency—without sacrificing accuracy or increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems. The book includes background on DNN processing; a description and taxonomy of hardware architectural approaches for designing DNN accelerators; key metrics for evaluating and comparing different designs; features of the DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as a formalization and organization of key concepts from contemporary works that provides insights that may spark new ideas.

Learn TensorFlow in 24 Hours

Make your own Neural Network *Guru99*

Tensorflow is the most popular Deep Learning Library out there. It has fantastic graph computations feature which helps data scientist to visualize his designed neural network using TensorBoard. This Machine learning library supports both Convolution as well as Recurrent Neural network. It supports parallel processing on CPU as well as GPU. Prominent machine learning algorithms supported by TensorFlow are Deep Learning Classification, wip & deep, Boston Tree amongst others. The book is very hands-on and gives you industry ready deep learnings practices. Here is what is covered in the book – Table Of Content Chapter 1: What is Deep learning? Chapter 2: Machine Learning vs Deep Learning Chapter 3: What is TensorFlow? Chapter 4: Comparison of Deep Learning Libraries Chapter 5: How to Download and Install TensorFlow Windows and Mac Chapter 6: Jupyter Notebook Tutorial Chapter 7: Tensorflow on AWS Chapter 8: TensorFlow Basics: Tensor, Shape, Type, Graph, Sessions & Operators Chapter 9: Tensorboard: Graph Visualization with Example Chapter 10: NumPy Chapter 11: Pandas Chapter 12: Scikit-Learn Chapter 13: Linear Regression Chapter 14: Linear Regression Case Study Chapter 15: Linear Classifier in TensorFlow Chapter 16: Kernel Methods Chapter 17: TensorFlow ANN (Artificial Neural Network) Chapter 18: ConvNet(Convolutional Neural Network): TensorFlow Image Classification Chapter 19: Autoencoder with TensorFlow Chapter 20: RNN(Recurrent Neural Network) TensorFlow

Make Your First GAN With PyTorch *Independently Published*

A gentle introduction to Generative Adversarial Networks, and a practical step-by-step tutorial on making your own with PyTorch. This beginner-friendly guide will give you hands-on experience: * understanding PyTorch basics * developing your first PyTorch neural network * exploring neural network refinements to improve performance * introduce CUDA GPU acceleration It will introduce GANs, one of the most exciting areas of machine learning: * introducing the concept step-by-step, in plain English * coding the simplest GAN to develop a good workflow * growing our confidence with an MNIST GAN * progressing to develop a GAN to generate full-colour human faces * experiencing how GANs fail, exploring remedies and improving GAN performance and stability Beyond the very basics, readers can explore more sophisticated GANs: * convolutional GANs for generated higher

quality images * conditional GANs for generated images of a desired class The appendices will be useful for students of machine learning as they explain themes often skipped over in many courses: * calculating ideal loss values for balanced GANs * probability distributions and sampling them to create images * carefully chosen examples illustrating how convolutions work * a brief explanation of why gradient descent isn't suited to adversarial machine learning

Deep Learning for Coders with fastai and PyTorch *"O'Reilly Media, Inc."*

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

Dive Into Deep Learning

Tools for Engagement *Corwin Press*

Create learning experiences that transform not only learning, but life itself. Learn about, improve, and expand your world of learning. This hands-on companion to the runaway best-seller, *Deep Learning: Engage the World Change the World*, provides an essential 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. Loaded with tips, tools, protocols, and real-world examples, the easy-to-use guide has everything educators need to construct and drive meaningful deep learning experiences that give purpose, unleash student potential, and prepare students to become problem-solving change agents in a global society.

In the Agora

The Public Face of Canadian Philosophy *University of Toronto Press*

A spirited and engaging read, *In the Agora* effectively illustrates how Canadian philosophers have contributed to public discourse and enriched our world. It is a collection that is sure to prompt both interest and debate.