
Neural Networks And Deep Learning

Making AI Less Susceptible to Adversarial Trickery

Applied Neural Networks with TensorFlow 2

Neural networks in TensorFlow.js

Artificial Intelligence for Humans, Volume 3

Precision Medicine, High Performance and Large-Scale Datasets

Neural Networks and Deep Learning, Deep Learning Explained to Your Granny

Deep Learning with JavaScript

Advanced Applied Deep Learning

An Introduction to Neural Networks

Learning Deep Architectures for AI

Deep Learning

The Perceptron

A Textbook

MATLAB Deep Learning

Introduction to Deep Learning

Deep Learning and Convolutional Neural Networks for Medical Image Computing
Deep Learning on Graphs
Neural Networks with R
Design and Case Studies
Neural Networks and Statistical Learning
Deep Learning
Neural Networks and Deep Learning
Applied Artificial Intelligence
Neural Networks and Deep Learning
Deep Learning and Neural Networks
Neural Networks and Deep Learning
Implement advanced deep learning models using Python
Grokking Deep Learning
Concepts, Methodologies, Tools, and Applications
Neural Networks and Deep Learning
Evolutionary Approach to Machine Learning and Deep Neural Networks
Deep Learning and Machine Learning Outlined
Strengthening Deep Neural Networks
Deep Learning and Neural Networks: Concepts, Methodologies, Tools, and Applications

A Fast-Track Approach to Modern Deep Learning with Python
Machine Learning with Neural Networks
The Principles of Deep Learning Theory
With Machine Learning, Neural Networks and Artificial Intelligence
Deep Learning with PyTorch

Neural Networks And Deep Learning Downloaded from ecobankpayservices.ecobank.com by guest

CHASE JACK

Making AI Less Susceptible to Adversarial Trickery Springer
Develop and optimize deep learning models with advanced architectures. This book teaches you the intricate details and subtleties of the algorithms that are at the

core of convolutional neural networks. In Advanced Applied Deep Learning, you will study advanced topics on CNN and object detection using Keras and TensorFlow. Along the way, you will look at the fundamental operations in CNN, such as convolution and pooling, and then look at more advanced architectures such as

inception networks, resnets, and many more. While the book discusses theoretical topics, you will discover how to work efficiently with Keras with many tricks and tips, including how to customize logging in Keras with custom callback classes, what is eager execution, and how to use it in your models. Finally, you will study how

object detection works, and build a complete implementation of the YOLO (you only look once) algorithm in Keras and TensorFlow. By the end of the book you will have implemented various models in Keras and learned many advanced tricks that will bring your skills to the next level. What You Will Learn See how convolutional neural networks and object detection work Save weights and models on disk Pause training and restart it at a later stage Use hardware

acceleration (GPUs) in your code Work with the Dataset TensorFlow abstraction and use pre-trained models and transfer learning Remove and add layers to pre-trained networks to adapt them to your specific project Apply pre-trained models such as Alexnet and VGG16 to new datasets Who This Book Is For Scientists and researchers with intermediate-to-advanced Python and machine learning know-how. Additionally, intermediate knowledge of Keras and

TensorFlow is expected. [Applied Neural Networks with TensorFlow 2](#) Simon and Schuster Uncover the power of artificial neural networks by implementing them through R code. About This Book Develop a strong background in neural networks with R, to implement them in your applications Build smart systems using the power of deep learning Real-world case studies to illustrate the power of neural network models Who This Book Is For This book is intended for

anyone who has a statistical background with knowledge in R and wants to work with neural networks to get better results from complex data. If you are interested in artificial intelligence and deep learning and you want to level up, then this book is what you need! What You Will Learn Set up R packages for neural networks and deep learning Understand the core concepts of artificial neural networks Understand neurons, perceptrons, bias, weights, and activation

functions Implement supervised and unsupervised machine learning in R for neural networks Predict and classify data automatically using neural networks Evaluate and fine-tune the models you build. In Detail Neural networks are one of the most fascinating machine learning models for solving complex computational problems efficiently. Neural networks are used to solve wide range of problems in different areas of AI and machine

learning. This book explains the niche aspects of neural networking and provides you with foundation to get started with advanced topics. The book begins with neural network design using the neural net package, then you'll build a solid foundation knowledge of how a neural network learns from data, and the principles behind it. This book covers various types of neural network including recurrent neural networks and convoluted neural networks. You will not only learn how to train

neural networks, but will also explore generalization of these networks. Later we will delve into combining different neural network models and work with the real-world use cases. By the end of this book, you will learn to implement neural network models in your applications with the help of practical examples in the book. Style and approach A step-by-step guide filled with real-world practical examples. *Neural networks in TensorFlow.js* "O'Reilly Media, Inc."

A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets,

convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author

writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to Tensorflow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses;

practitioners will find it an essential reference. [Artificial Intelligence for Humans, Volume 3](#) World Scientific Publishing Company This book provides theoretical and practical knowledge about a methodology for evolutionary algorithm-based search strategy with the integration of several machine learning and deep learning techniques. These include convolutional neural networks, Gröbner bases, relevance vector machines, transfer

learning, bagging and boosting methods, clustering techniques (affinity propagation), and belief networks, among others. The development of such tools contributes to better optimizing methodologies. Beginning with the essentials of evolutionary algorithms and covering interdisciplinary research topics, the contents of this book are valuable for different classes of readers: novice, intermediate, and also expert readers from related fields. Following

the chapters on introduction and basic methods, Chapter 3 details a new research direction, i.e., neuro-evolution, an evolutionary method for the generation of deep neural networks, and also describes how evolutionary methods are extended in combination with machine learning techniques. Chapter 4 includes novel methods such as particle swarm optimization based on affinity propagation (PSOAP), and transfer learning for differential evolution (TRADE),

another machine learning approach for extending differential evolution. The last chapter is dedicated to the state of the art in gene regulatory network (GRN) research as one of the most interesting and active research fields. The author describes an evolving reaction network, which expands the neuro-evolution methodology to produce a type of genetic network suitable for biochemical systems and has succeeded in designing genetic circuits in synthetic biology. The author also presents real-

world GRN application to several artificial intelligent tasks, proposing a framework of motion generation by GRNs (MONGERN), which evolves GRNs to operate a real humanoid robot.

**Precision Medicine,
High Performance and
Large-Scale Datasets**

Addison-Wesley

Professional

Introduction to Deep

Learning and Neural

Networks with Python™:

A Practical Guide is an

intensive step-by-step

guide for neuroscientists

to fully understand,

practice, and build neural networks. Providing math and Python™ code examples to clarify neural network calculations, by book's end readers will fully understand how neural networks work starting from the simplest model $Y=X$ and building from scratch. Details and explanations are provided on how a generic gradient descent algorithm works based on mathematical and Python™ examples, teaching you how to use the gradient descent algorithm to manually perform all calculations in

both the forward and backward passes of training a neural network. Examines the practical side of deep learning and neural networks Provides a problem-based approach to building artificial neural networks using real data Describes Python™ functions and features for neuroscientists Uses a careful tutorial approach to describe implementation of neural networks in Python™ Features math and code examples (via companion website) with helpful

instructions for easy implementation
[Neural Networks and Deep Learning, Deep Learning Explained to Your Granny](#) Academic Press

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, providing a comfortable Python experience that gets you

started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Summary 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, providing a comfortable Python

experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Foreword by Soumith Chintala, Cocreator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Although many deep learning tools use Python, the PyTorch library is truly

Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead

without sacrificing the access to advanced features, meaning you can focus on what matters the most - building and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class players like FAIR, OpenAI, FastAI and Purdue. About the book

Deep Learning with PyTorch teaches you to create neural networks and deep learning systems with PyTorch. This practical book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The

centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with limited inputs and start processing data to get some results. You'll sift through the unreliable initial results and focus on how to diagnose and fix the problems in your neural network. Finally, you'll look at ways to improve your results by training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside

Training deep neural networks Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Exploring code samples in Jupyter Notebooks About the reader For Python programmers with an interest in machine learning. About the author Eli Stevens had roles from software engineer to CTO, and is currently working on machine learning in the self-driving-car industry. Luca Antiga is cofounder of an AI engineering company and

an AI tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. 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
Deep Learning with JavaScript Cambridge University Press
As deep neural networks (DNNs) become increasingly common in real-world applications, the potential to deliberately "fool" them with data that wouldn't trick a human presents a new attack vector. This practical book examines real-world scenarios where DNNs—the

algorithms intrinsic to much of AI—are used daily to process image, audio, and video data. Author Katy Warr considers attack motivations, the risks posed by this adversarial input, and methods for increasing AI robustness to these attacks. If you're a data scientist developing DNN algorithms, a security architect interested in how to make AI systems more resilient to attack, or someone fascinated by the differences between artificial and biological

perception, this book is for you. Delve into DNNs and discover how they could be tricked by adversarial input Investigate methods used to generate adversarial input capable of fooling DNNs Explore real-world scenarios and model the adversarial threat Evaluate neural network robustness; learn methods to increase resilience of AI systems to adversarial data Examine some ways in which AI might become better at mimicking human perception in years to

come
Advanced Applied Deep Learning Springer Nature
 This textbook establishes a theoretical framework for understanding deep learning models of practical relevance. With an approach that borrows from theoretical physics, Roberts and Yaida provide clear and pedagogical explanations of how realistic deep neural networks actually work. To make results from the theoretical forefront accessible, the authors eschew the subject's traditional emphasis on

intimidating formality without sacrificing accuracy. Straightforward and approachable, this volume balances detailed first-principle derivations of novel results with insight and intuition for theorists and practitioners alike. This self-contained textbook is ideal for students and researchers interested in artificial intelligence with minimal prerequisites of linear algebra, calculus, and informal probability theory, and it can easily fill a semester-long course on deep learning theory.

For the first time, the exciting practical advances in modern artificial intelligence capabilities can be matched with a set of effective principles, providing a timeless blueprint for theoretical research in deep learning.

An Introduction to Neural Networks

Manning Publications
 Ready to crank up a deep neural network to get your self-driving car pick up the kids from school? Want to add 'Neural Networks' and '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! Check out this exceptional bundle of 3 books... This bundle contains 3 books: Book 1: Neural Networks & Deep Learning: Deep Learning explained to your granny - A visual introduction for beginners who want to make their own Deep Learning Neural Network... What you will gain from this book: * A deep understanding of

how Deep Learning works
* 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!
* 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
 Book2: Deep Learning:
 Deep Learning explained
 to your granny - A guide
 for Beginners... What's
 Inside? * 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 * Convolutional
 Neural Networks * Models
 for Deep Learning *
 Probabilistic graphic
 models * Restricted
 Boltzmann Machines *
 Deep Belief Networks *
 Available Frameworks and
 libraries * TensorFlow
 Book 3: Blockchain
 Blueprint: The ultimate
 guide to understanding
 blockchain,
 cryptocurrencies, smart
 contracts and the future

of money The current
 emerging innovation of
 this decade may be the
 connected world of
 computing relying on
 blockchain encryption.
 The attention given to this
 technology by global giant
 players suggests that it
 will become the
 operational philosophy of
 the economic system of
 the future, ranging across
 many industries.
 Blockchain can become
 the solution we needed
 for speeding up the
 economy and transactions
 in order to keep up with
 our multi-device

connected world. In this book, high tech expert Pat Nakamoto answers your questions concerning the future of Blockchain technology along with addressing different major developments linked to it, like Smart Contracts, Fintech and Ethereum. Hit download. Now!

Learning Deep Architectures for AI
Createspace Independent Publishing Platform
Deep Learning Neural Networks is the fastest growing field in machine learning. It serves as a powerful computational

tool for solving prediction, decision, diagnosis, detection and decision problems based on a well-defined computational architecture. It has been successfully applied to a broad field of applications ranging from computer security, speech recognition, image and video recognition to industrial fault detection, medical diagnostics and finance. This comprehensive textbook is the first in the new emerging field. Numerous case studies are succinctly demonstrated

in the text. It is intended for use as a one-semester graduate-level university text and as a textbook for research and development establishments in industry, medicine and financial research.

Deep Learning Wolfgang Beer
Implement deep learning applications using TensorFlow while learning the “why” through in-depth conceptual explanations. You’ll start by learning what deep learning offers over other machine learning models.

Then familiarize yourself with several technologies used to create deep learning models. While some of these technologies are complementary, such as Pandas, Scikit-Learn, and Numpy—others are competitors, such as PyTorch, Caffe, and Theano. This book clarifies the positions of deep learning and Tensorflow among their peers. You'll then work on supervised deep learning models to gain applied experience with the technology. A single-layer of multiple

perceptrons will be used to build a shallow neural network before turning it into a deep neural network. After showing the structure of the ANNs, a real-life application will be created with Tensorflow 2.0 Keras API. Next, you'll work on data augmentation and batch normalization methods. Then, the Fashion MNIST dataset will be used to train a CNN. CIFAR10 and Imagenet pre-trained models will be loaded to create already advanced CNNs. Finally, move into theoretical applications

and unsupervised learning with auto-encoders and reinforcement learning with tf-agent models. With this book, you'll delve into applied deep learning practical functions and build a wealth of knowledge about how to use TensorFlow effectively. What You'll Learn Compare competing technologies and see why TensorFlow is more popular Generate text, image, or sound with GANs Predict the rating or preference a user will give to an item Sequence data with recurrent neural

networks Who This Book
Is For Data scientists and
programmers new to the
fields of deep learning
and machine learning
APIs.

The Perceptron

Createspace Independent
Publishing Platform
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 *A Textbook* Cambridge University Press 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! **MATLAB Deep Learning** Springer
Would you achieve more if you could envision your success? A neural network is a computing system made up of a number of simple, highly interconnected processing elements, which process

information by their dynamic state response to external inputs. All of this sounds fancy, but what does it mean for computer intelligence, or for the future? In this book, you will find answers to many practical and theoretical questions related to neural networks, from insights about nodes and hidden layers to error spaces, network analyses, and computing influences. Topics will be discussed, such as: What the definition of neural networks encompasses and what all the elements

pertaining to them mean. The main advantages of neural networks and how to leverage and apply them. Limitations to neural networks. How neural networks differ from conventional computing systems. Neural Network applications for medical diagnostics, smart computers, artificial intelligence, and forex or stock trading. Troubleshooting tips for when neural networks stop functioning. If you are even in the least interested in computer

technology, artificial intelligence, or what the technological future will bring, you need to read this book and get a better understanding of neural networks and their many applications. This book will bring you to the core of how they function and what you can do with them. Add this book to your cart. [Introduction to Deep Learning](#) Createspace Independent Publishing Platform This book presents a detailed review of the state of the art in deep

learning approaches for semantic object detection and segmentation in medical image computing, and large-scale radiology database mining. A particular focus is placed on the application of convolutional neural networks, with the theory supported by practical examples. Features: highlights how the use of deep neural networks can address new questions and protocols, as well as improve upon existing challenges in medical image computing; discusses the insightful

research experience of Dr. Ronald M. Summers; presents a comprehensive review of the latest research and literature; describes a range of different methods that make use of deep learning for object or landmark detection tasks in 2D and 3D medical imaging; examines a varied selection of techniques for semantic segmentation using deep learning principles in medical imaging; introduces a novel approach to interleaved text and image deep

mining on a large-scale radiology image database.

Deep Learning and Convolutional Neural Networks for Medical Image Computing

Addison-Wesley Professional

A comprehensive text on foundations and techniques of graph neural networks with applications in NLP, data mining, vision and healthcare.

Deep Learning on Graphs
Apress

Learn, understand, and implement deep neural

networks in a math- and programming-friendly approach using Keras and Python. The book focuses on an end-to-end approach to developing supervised learning algorithms in regression and classification with practical business-centric use-cases implemented in Keras. The overall book comprises three sections with two chapters in each section. The first section prepares you with all the necessary basics to get started in deep learning. Chapter 1 introduces you to the world of deep

learning and its difference from machine learning, the choices of frameworks for deep learning, and the Keras ecosystem. You will cover a real-life business problem that can be solved by supervised learning algorithms with deep neural networks. You'll tackle one use case for regression and another for classification leveraging popular Kaggle datasets. Later, you will see an interesting and challenging part of deep learning: hyperparameter tuning; helping you further improve your

models when building robust deep learning applications. Finally, you'll further hone your skills in deep learning and cover areas of active development and research in deep learning. At the end of *Learn Keras for Deep Neural Networks*, you will have a thorough understanding of deep learning principles and have practical hands-on experience in developing enterprise-grade deep learning solutions in Keras. What You'll Learn Master fast-paced practical deep learning

concepts with math- and programming-friendly abstractions. Design, develop, train, validate, and deploy deep neural networks using the Keras framework Use best practices for debugging and validating deep learning models Deploy and integrate deep learning as a service into a larger software service or product Extend deep learning principles into other popular frameworks Who This Book Is For Software engineers and data engineers with basic programming skills in any

language and who are keen on exploring deep learning for a career move or an enterprise project. *Neural Networks with R* "O'Reilly Media, Inc." About This Book Step into the amazing world of Artificial Intelligence and Machine Learning using this compact and easy to understand book. Dive into Neural Networks and Deep Learning and create your own production ready AI models by using TensorFlow and Keras. Work through simple yet insightful examples that will get you up and

running with Artificial Intelligence, TensorFlow and Keras in no time. Who This Book Is For This book is for Python developers who want to understand Neural Networks from ground up and 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. What You Will Learn The basic structure and functionality of a Neuron The basic math behind the Neural Network learning process

See how to build a simple character recognition model from ground up What classification, regression and clustering is How to use TensorFlow to build production ready models Build a first model with the Keras framework How to predict the survival chance for Titanic passengers How to build a simple book recommender How to detect toxic language with an AI model In Detail Artificial Intelligence became one of the hottest topics in the modern economy, where

everything is driven by software, network and data. There exists nearly no startup nor traditional business where Artificial Intelligence is not used extensively across many fields such as search engines, image recognition, robotics or finance. This book gives a ground up, step by step introduction about how a Neural Network is used to learn a given function and to make intelligent data-driven decisions. The book explains how to identify typical use-cases such as classification,

regression and clustering in terms of practical and well known use-cases. This book comes with an introduction into the state-of-the-art Google TensorFlow framework that allows developers to roll out their models in production. On top of TensorFlow, the Keras library is used to simplify the design and training of complex deep-learning models. This book comes with multiple examples that show how to apply Artificial Intelligence and Machine Learning models for use-cases such as

handwriting recognition, decision making, text analysis and toxic comment identification as well as the use of AI to recommend products to customers.

Design and Case Studies

Apress 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

Neural Networks and Statistical Learning MIT

Press
Theoretical results suggest that in order to learn the kind of complicated functions that can represent high-level abstractions (e.g. in vision, language, and other AI-level tasks), one may need deep architectures. Deep architectures are composed of multiple levels of non-linear operations, such as in

neural nets with many hidden layers or in complicated propositional formulae re-using many sub-formulae. Searching the parameter space of deep architectures is a difficult task, but learning algorithms such as those for Deep Belief Networks have recently been proposed to tackle this problem with notable success, beating the

state-of-the-art in certain areas. This paper discusses the motivations and principles regarding learning algorithms for deep architectures, in particular those exploiting as building blocks unsupervised learning of single-layer models such as Restricted Boltzmann Machines, used to construct deeper models such as Deep Belief Networks.

Related with Neural Networks And Deep Learning:

[© Neural Networks And Deep Learning Hogwarts Legacy Gear Guide](#)

[© Neural Networks And Deep Learning Hoa Loan Vs Special Assessment](#)

[© Neural Networks And Deep Learning Hogwarts Legacy Field Guide](#)