Javascript Artificial Intelligence Made Easy W Essential Programming Create Your Problem Solving Algorithms Today W Machine Learning Data Structures Artificial Intelligence Series

"What does AI mean for your business? Read this book to find out." -- Hal Varian, Chief Economist, Google Artificial intelligence does the seemingly impossible, magically bringing machines to life--driving cars, trading stocks, and teaching children. But facing the sea change that AI will bring can be paralyzing. How should companies set strategies, governments design policies, and people plan their lives for a world so different from what we know? In the face of such uncertainty, many analysts either cower in fear or predict an impossibly sunny future. But in Prediction Machines, three eminent economists recast the rise of AI as a drop in the cost of prediction. With this single, masterful stroke, they lift the curtain on the AI-is-magic hype and show how basic tools from economics provide clarity about the AI revolution and a basis for action by CEOs, managers, policy makers, investors, and entrepreneurs. When AI is framed as cheap prediction, its extraordinary potential becomes clear: Prediction is at the heart of making decisions under uncertainty. Our businesses and personal lives are riddled with such decisions. Prediction tools increase productivity--operating machines, handling documents, communicating with customers. Uncertainty constrains strategy. Better prediction creates opportunities for new business structures and strategies to compete. Penetrating, fun, and always insightful and practical, Prediction Machines follows its inescapable logic to explain how to navigate the changes on the horizon. The impact of AI will be profound, but the economic framework for understanding it is surprisingly simple.

Given the demand for AI and the ubiquity of JavaScript, TensorFlow.js was inevitable. With this Google framework, seasoned AI veterans and web developers alike can help propel the future of AI-driven websites. In this guide, author Gant Laborde--Google Developer Expert in machine learningand the web--provides a hands-on end-to-end approach to TensorFlow.js fundamentals for a broad technical audience that includes data scientists, engineers, web developers, students, and researchers. You'll begin by working through some basic examples in TensorFlow.js before diving deeper into neural network architectures, DataFrames, TensorFlow Hub, model conversion, transfer learning, and more. Once you finish this book, you'll know how to build and deploy production-readydeep learning systems with TensorFlow.js. Explore tensors, the most fundamental structure of machine learning Convert data into tensors and back with a real-world example Combine AI with the web using TensorFlow.js Use resources to convert, train, and manage machine learning data Build and train your own training models from scratch

Get hands-on with the browser-based JavaScript library for training and deploying machine learning models effectively Key Features Build, train and run machine learning models in the browser using TensorFlow.js Create smart web applications from scratch with the help of useful examples Use flexible and intuitive APIs from TensorFlow.js to understand how machine learning algorithms function Book Description TensorFlow.js is a framework that enables you to create performant machine learning (ML) applications that run smoothly in a web browser. With this book, you will learn how to use TensorFlow.js to implement various ML models through an example-based approach. Starting with the basics, you'll understand how ML models can be built on the web. Moving on, you will get to grips with the TensorFlow.js ecosystem to develop applications more efficiently. The book will then guide you through implementing ML techniques and algorithms such as regression, clustering, fast Fourier transform (FFT), and dimensionality reduction. You will later cover the Bellman equation to solve Markov decision process (MDP) problems and understand how it is related to reinforcement learning. Finally, you will explore techniques for deploying ML-based web applications and training models with TensorFlow Core. Throughout this ML book, you'll discover useful tips and tricks that will build on your knowledge. By the end of this book, you will be equipped with the skills you need to create your own web-based ML applications and fine-tune models to achieve high performance. What you will learn Use the t-SNE algorithm in TensorFlow.js to reduce dimensions in an input dataset Deploy tfjs-converter to convert Keras models and load them into TensorFlow.js Apply the Bellman equation to solve MDP problems Use the k-means algorithm in TensorFlow.js to visualize prediction results Create tf.js packages with Parcel, Webpack, and Rollup to deploy web apps Implement tf.js backend frameworks to tune and accelerate app performance Who this book is for This book is for web developers who want to learn how to integrate machine learning techniques with web-based applications from scratch. This book will also appeal to data scientists, machine learning practitioners, and deep learning enthusiasts who are looking to perform accelerated, browser-based machine learning on Web using TensorFlow.js. Working knowledge of JavaScript programming language is all you need to get started.

This book shows readers how they can successfully analyze data using only two core machine learning algorithms---and how to do so using the popular Python programming language. These algorithms deal with common scenarios faced by all data analysts and data scientists. This book focuses on two algorithm families (linear methods and ensemble methods) that effectively predict outcomes. This type of problem covers a multitude of use cases (what ad to place on a web page, predicting prices in securities markets, detecting credit card fraud, etc.). The focus on two families gives enough room for full descriptions of the mechanisms at work in the algorithms. Then the code examples serve to illustrate the workings of the machinery with specific hackable code. The author will explain in simple terms, using no complex math, how these algorithms work, and will then show how to apply them in Python. He will also provide advice on how to select from among these algorithms, and will show how to prepare the data, and how to use the trained models in practice. The author begins with an overview of the two core algorithms, explaining the types of problems solved by each one. He then introduces a core set of Python programming techniques that can be used to apply these algorithms. The author shows various techniques for building predictive models that solve a range of problems, from simple to complex; he also shows how to measure the performance of each model to ensure you use the right one. The following chapters provide a deep dive into each of the two algorithms: penalized linear regression and ensemble methods. Chapters will show how to apply each algorithm in Python. Readers can directly use the sample code to build their own solutions.

Apply Artificial Intelligence techniques in the browser or on resource constrained computing devices. Machine learning (ML) can be an intimidating subject until you know the essentials and for what applications it works. This book takes advantage of the intricacies of the ML processes by using a simple, flexible and portable programming language such as JavaScript to work with more approachable, fundamental coding ideas. Using JavaScript programming features along with standard libraries, you'll first learn to design and develop interactive graphics applications. Then move further into neural systems and human pose estimation strategies. For training and deploying your ML models in the browser, TensorFlow.js libraries will be emphasized. After

conquering the fundamentals, you'll dig into the wilderness of ML. Employ the ML and Processing (P5) libraries for Human Gait analysis. Building up Gait recognition with themes, you'll come to understand a variety of ML implementation issues. For example, you'll learn about the classification of normal and abnormal Gait patterns. With Beginning Machine Learning in the Browser, you'll be on your way to becoming an experienced Machine Learning developer. What You'll Learn Work with ML models, calculations, and information gathering Implement TensorFlow.js libraries for ML models Perform Human Gait Analysis using ML techniques in the browser Who This Book Is For Computer science students and research scholars, and novice programmers/web developers in the domain of Internet Technologies

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

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

Ready to learn how to code a game? Get an introduction to programming with this fun and accessible guide. Learn HTML and JavaScript. Design and build five interactive computer games. Create cool graphics. Code simple artificial intelligence. This appealing guide, covering essential coding concepts, offers an ideal introduction to all these activities and more. By following simple step-by-step instructions and completing five exciting missions, aspiring programmers are invited to code well-known games such as tic-tac-toe and table tennis, then customize their projects to test their skills.

A definitive guide to creating an intelligent web application with the best of machine learning and JavaScript Key Features Solve complex computational problems in browser with JavaScript Teach your browser how to learn from rules using the power of machine learning Understand discoveries on web interface and API in machine learning Book Description In over 20 years of existence, JavaScript has been pushing beyond the boundaries of web evolution with proven existence on servers, embedded devices, Smart TVs, IoT, Smart Cars, and more. Today, with the added advantage of machine learning research and support for JS libraries, JavaScript makes your browsers smarter than ever with the ability to learn patterns and reproduce them to become a part of innovative products and applications. Hands-on Machine Learning with JavaScript presents various avenues of machine learning in a practical and objective way, and helps implement them using the JavaScript language. Predicting behaviors, analyzing feelings, grouping data, and building neural models are some of the skills you will build from this book. You will learn how to train your machine learning models and work with different kinds of data. During this journey, you will come across use cases such as face detection, spam filtering, recommendation systems, character recognition, and more. Moreover, you will learn how to work with deep neural networks and guide your applications to gain insights from data. By the end of this book, you'll have gained hands-on knowledge on evaluating and implementing the right model, along with choosing from different JS libraries, such as NaturalNode, brain, harthur, classifier, and many more to design smarter applications. What you will learn Get an overview of state-of-the-art machine learning Understand the pre-processing of data handling, cleaning, and preparation Learn Mining and Pattern Extraction with JavaScript Build your own model for classification, clustering, and prediction Identify the most appropriate model for each type of problem Apply machine learning techniques to real-world applications Learn how JavaScript can be a powerful language for machine learning Who this book is for This book is for you if you are a JavaScript developer who wants to implement machine learning to make applications smarter, gain insightful information from the data, and enter the field of machine learning without switching to another language. Working knowledge of JavaScript language is expected to get the most out of the book.

One of the most persistent concerns about the future is whether it will be dominated by the predictive algorithms of AI – and, if so, what this will mean for our behaviour, for our institutions and for what it means to be human. AI changes our experience of time and the future and challenges our identities, yet we are blinded by its efficiency and fail to understand how it affects us. At the heart of our trust in AI lies a paradox: we leverage AI to increase our control over the future and uncertainty, while at the same time the performativity of AI, the power it has to make us act in the ways it predicts, reduces our agency over the future. This happens when we forget that that we humans have created the digital technologies to which we attribute agency. These developments also challenge the narrative of progress, which played such a central role in modernity and is based on the hubris of total control. We are now moving into an era where this control is limited as AI monitors our actions, posing the threat of surveillance, but also offering the opportunity to reappropriate control and transform it into care. As we try to adjust to a world in which algorithms, robots and avatars play an ever-increasing role, we need to understand better the limitations of AI and how their predictions affect our agency, while at the same time having the courage to embrace the uncertainty of the future.

Two leaders in the field offer a compelling analysis of the current state of the art and reveal the steps we must take to achieve a truly robust artificial intelligence. Despite the hype surrounding AI, creating an intelligence that rivals or exceeds human levels is far more complicated than we have been led to believe. Professors Gary Marcus and Ernest Davis have spent their careers at the forefront of AI research and have witnessed some of the greatest milestones in the field, but they argue that a computer beating a human in Jeopardy! does not signal that we are on the doorstep of fully autonomous cars or superintelligent machines. The achievements in the field thus far have occurred in closed systems with fixed sets of rules, and these approaches are too narrow to achieve genuine intelligence. The real world, in contrast, is wildly complex and open-ended. How can we bridge this gap? What will the consequences be when we do? Taking inspiration from the human mind, Marcus and Davis explain what we need to advance AI to the next level, and suggest that if we are wise along the way, we won't need to worry about a future of machine overlords. If we focus on endowing machines with common sense and deep understanding, rather than simply focusing on statistical analysis and gatherine ever larger collections of data, we will be able to create an AI we can trust--in our homes, our cars, and our doctors' offices. Rebooting AI provides a lucid, clear-eyed assessment of the current science and offers an inspiring vision of how a new generation of AI can make our lives better.

Completely revised and updated, this best-selling introduction to programming in JavaScript focuses on writing real applications. JavaScript lies at the heart of almost every modern web application, from social apps like Twitter to browser-based game frameworks like Phaser and Babylon. Though simple for beginners to pick up and play with, JavaScript is a flexible, complex language that you can use to build full-scale applications. This much anticipated and thoroughly revised third edition of Eloquent JavaScript dives deep into the JavaScript language to show you how to write beautiful, effective code. It has been updated to reflect the current state of Java¬Script and web browsers and includes brand-new material on features like class notation, arrow functions, iterators, async functions, template strings, and block scope. A host of new exercises have also been added to test your skills and keep you on track. As with previous editions, Haverbeke continues to teach through extensive examples and immerses you in code from the start, while exercises and full-chapter projects give you hands-on experience with writing your own programs. You start by learning the basic structure of the JavaScript language as well as control structures, functions, and data structures to help you write basic programs. Then you'll learn about error handling and bug fixing, modularity, and asynchronous programming before moving on to web browsers and how JavaScript is used to program them. As you build projects such as an artificial life simulation, a simple programming language, and a paint program, you'll learn how to: - Understand the essential elements of programming, including syntax, control, and data - Organize and clarify your code with object-oriented and functional programming techniques - Script the browser and make basic web applications - Use the DOM effectively to interact with browsers - Harness Node.js to build servers and utilities Isn't it time you became fluent in the language of the Web? * All source code is available online in an inter¬active sandbox, where you can edit the code, run it, and see its output instantly.

The Road to Firebase is your personal journey to master advanced React for business web applications in JavaScript whereas Firebase is used to replace everything that you would want from a backend application. Firebase enables you to connect your React application to a database, to authenticated users with your application with a login, logout and register mechanisms, and to authorize only certain users to access your application. It also comes with hosting capabilities and with social logins via Google, Facebook and more. Everything will be explained in the book while building a business web application yourself. I wrote the The Road to React with Firebase over the last two years. During this time, I came to understand the practical genius of Firebase, and how it dramatically reduces the tech stack to focus on getting sh*t done. Once you have set up your starter kit project -- that's what we are going to do together in this book -- you are ready to iterate fast on your personal ideas. There is no need to complicate things by adding a backend application with a database to your frontend application, because Firebase takes care of it with a well-designed API. I applied the same principles as my other books: Stay pragmatic Keep it simple Answer the why, not just the how Experience a problem, solve a problem This book is not intended to be an end-all reference for the Firebase API nor an in-depth guide about the internals of Firebase. Instead, its purpose is to journey through learning Firebase with React the pragmatic way, building an entire application on this tech stack yourself. The end result is the foundation to make your business application a reality. Requirements To get the most out of this book, you should be familiar with the basics of web development, which includes knowledge of HTML, CSS and JavaScript. You will also need to be familiar with the term API, because APIs are used frequently for the applications in this book. Editor/Terminal or IDE For the development environment, use a running editor/terminal (command line tool) or IDE with integrated terminal. I will provide a setup guide if you're unsure about which tools to use. The guide is set up for MacOS users, but you can find a Windows setup guide there as well. Node and NPM You will need to have node and npm installed, which are used to run the applications we'll build and manage the libraries we'll use along the way. In this book, you will install external node packages via npm (node package manager). These node packages can be libraries or whole frameworks. You can verify which node and npm versions you have in the command line: node --version v10.11.0 npm --version v6.5.0 These are the versions used for this publication. If you don't see output in your terminal, you will need to install node and npm. React My other book, called The Road to learn React, teaches the fundamentals about React by building a real world application. It is available for free, and after having read it, you should possess all the understanding necessary to work with the application(s) from this book. Also there will be many sidenotes to React articles that may be helpful.

Get up and running with the latest numerical computing library by Google and dive deeper into your data!About This Book- Get the first book on the market that shows you the key aspects TensorFlow, how it works, and how to use it for the second generation of machine learning- Want to perform faster and more accurate computations in the field of data science? This book will acquaint you with an all-new refreshing library-TensorFlow!- Dive into the next generation of numerical computing and get the most out of your data with this quick guideWho This Book Is ForThis book is dedicated to all the machine learning and deep learning enthusiasts, data scientists, researchers, and even students who want to perform more accurate, fast machine learning operations with TensorFlow. Those with basic knowledge of programming (Python and C/C++) and math concepts who want to be introduced to the topics of machine learning will find this book useful.What You Will Learn- Install and adopt TensorFlow in your Python environment to solve mathematical problems- Get to know the basic machine and deep learning concepts- Train and test neural networks to fit your data model- Make predictions using regression algorithms- Analyze your data with a clustering procedure- Develop algorithms for clustering and data classification- Use GPU computing to analyze big dataIn DetailGoogle's TensorFlow engine, after much fanfare, has evolved in to a robust, user-friendly, and customizable, application-grade software library of machine learning (ML) code for numerical computation and neural networks.This book takes you through the practical software implementation of various machine learning techniques with TensorFlow. In the first few chapters, you'll gain familiarity with the framework and perform the mathematical operations required for data analysis. As you progress further, you'll learn to implement various machine learning techniques such as classification, clustering, neural networks, and deep learning through practical examples.By the end of this book, you'll have gained hands-on experience of using TensorFlow and building classification, image recognition systems, language processing, and information retrieving systems for your application.Style and approachGet quickly up and running with TensorFlow using this fast-paced guide. You will get to know everything that can be done with TensorFlow and we'll show you how to implement it in your environment. The examples in the book are from the core of the computation industry-something you can connect to and will find familiar.

Human-in-the-Loop Machine Learning lays out methods for humans and machines to work together effectively. Summary Most machine learning systems that are deployed in the world today learn from human feedback. However, most machine learning courses focus almost exclusively on the algorithms, not the human-computer interaction part of the systems. This can leave a big knowledge gap for data scientists working in real-world machine learning, where data scientists spend more time on data management than on building algorithms. Human-in-the-Loop Machine Learning is a practical guide to optimizing the entire machine learning process, including techniques for annotation, active learning, transfer learning, and using machine learning to optimize every step of the process. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Machine learning applications perform better with human feedback. Keeping the right people in the loop improves the accuracy of models, reduces errors in data, lowers costs, and helps you ship models faster. About the book Human-in-the-Loop Machine Learning lays out methods for humans and machines to work together effectively. You'll find best practices on selecting sample data for human feedback, quality control for human annotations, and designing annotation interfaces. You'll learn to create training data for labeling, object detection, and semantic segmentation, sequence labeling, and more. The book starts with the basics and progresses to advanced techniques like transfer learning and self-supervision within annotation workflows. What's inside Identifying the right training and evaluation data Finding and managing people to annotate data Selecting annotation quality control strategies Designing interfaces to improve accuracy and efficiency About the author Robert (Munro) Monarch is a data scientist and engineer who has built machine learning data for companies such as Apple, Amazon, Google, and IBM. He holds a PhD from Stanford. Robert holds a PhD from Stanford focused on Human-in-the-Loop machine learning for healthcare and disaster response, and is a disaster response professional in addition to being a machine learning professional. A worked example throughout this text is classifying disaster-related messages from real disasters that Robert has helped respond to in the past. Table of Contents PART 1 - FIRST STEPS 1 Introduction to human-in-the-loop machine learning 2 Getting started with human-in-the-loop machine learning PART 2 - ACTIVE LEARNING 3 Uncertainty sampling 4 Diversity sampling 5 Advanced active learning 6 Applying active learning to different machine learning tasks PART 3 - ANNOTATION 7 Working with the people annotating your data 8 Quality control for data annotation 9 Advanced data annotation and augmentation 10 Annotation quality for different machine learning tasks PART 4 - HUMAN–COMPUTER INTERACTION FOR MACHINE LEARNING 11 Interfaces for data annotation 12 Human-in-the-loop machine learning products

A human-inspired, linguistically sophisticated model of language understanding for intelligent agent systems. One of the original goals of artificial intelligence research was to endow intelligent agents with human-level natural language capabilities. Recent AI research, however, has focused on applying statistical and machine learning approaches to big data rather than attempting to model what people do and how they do it. In this book, Marjorie McShane and Sergei Nirenburg return to the original goal of recreating human-level intelligence in a machine. They present a human-inspired, linguistically sophisticated model of language understanding for intelligent agent systems that emphasizes meaning--the deep, context-sensitive meaning that a person derives from spoken or written language.

Unlock the groundbreaking advances of deep learning with this extensively revised new edition of the bestselling original. Learn directly from the creator of Keras and master practical Python deep learning techniques that are easy to apply in the real world. In Deep Learning with Python, Second Edition you will learn: Deep learning from first principles Image classification and image segmentation Timeseries forecasting Text classification and machine translation Text generation, neural style transfer, and image generation Deep Learning with Python has taught thousands of readers how to put the full capabilities of deep learning into action. This extensively revised second edition introduces deep learning using Python and Keras, and is loaded with insights for both novice and experienced ML practitioners. You'll learn practical techniques that are easy to apply in the real world, and important theory for perfecting neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Recent innovations in deep learning unlock exciting new software capabilities like automated language translation, image recognition, and more. Deep learning is quickly becoming essential knowledge for every software developer, and modern tools like Keras and TensorFlow put it within your reach—even if you have no background in mathematics or data science. This book shows you how to get started. About the book Deep Learning with Python, Second Edition introduces the field of deep learning using Python and the powerful Keras library. In this revised and expanded new edition, Keras creator François Chollet offers insights for both novice and experienced machine learning practitioners. As you move through this book, you'll build your understanding through intuitive explanations, crisp illustrations, and

clear examples. You'll quickly pick up the skills you need to start developing deep-learning applications. What's inside Deep learning from first principles Image classification and image segmentation Time series forecasting Text classification and machine translation Text generation, neural style transfer, and image generation About the reader For readers with intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the author François Chollet is a software engineer at Google and creator of the Keras deep-learning library. Table of Contents 1 What is deep learning? 2 The mathematical building blocks of neural networks 3 Introduction to Keras and TensorFlow 4 Getting started with neural networks: Classification and regression 5 Fundamentals of machine learning 6 The universal workflow of machine learning 7 Working with Keras: A deep dive 8 Introduction to deep learning for computer vision 9 Advanced deep learning for computer vision 10 Deep learning for timeseries 11 Deep learning for text 12 Generative deep learning 13 Best practices for the real world 14 Conclusions

Get hands-on with building data-driven applications using Danfo.js in combination with other data analysis tools and techniques Key Features Build microservices to perform data transformation and ML model serving in JavaScript Explore what Danfo.js is and how it helps with data analysis and data visualization Combine Danfo.js and TensorFlow.js for machine learning Book Description Most data analysts use Python and pandas for data processing for the convenience and performance these libraries provide. However, JavaScript developers have always wanted to use machine learning in the browser as well. This book focuses on how Danfo.js brings data processing, analysis, and ML tools to JavaScript developers and how to make the most of this library to build data-driven applications. Starting with an overview of modern JavaScript, you'll cover data analysis and transformation with Danfo.js and Dnotebook. The book then shows you how to load different datasets, combine and analyze them by performing operations such as handling missing values and string manipulations. You'll also get to grips with data plotting, visualization, aggregation, and group operations by combining Danfo.js with Plotly. As you advance, you'll create a no-code data analysis and handling system and create-react-app, react-table, react-chart, Draggable.js, and tailwindcss, and understand how to use TensorFlow.js and Danfo.js to build a recommendation system. Finally, you'll build a Twitter analytics dashboard powered by Danfo.js, Next.js, node-nlp, and Twit.js. By the end of this app development book, you'll be able to build and embed data analytics, visualization, and ML capabilities into any JavaScript app in server-side Node.js or the browser. What you will learn Perform data experimentation and analysis with Danfo.js and Dnotebook Build machine learning applications using Danfo.js integrated with TensorFlow.js Connect Danfo.js with popular database applications to aid data analysis Create a no-code data analysis and handling system using internal libraries Develop a recommendation system with Danfo.js and TensorFlow.js Build a Twitter analytics dashboard for sentiment analysis and other types of data insights Who this book is for This book is for data analysts, data scientists, and JavaScript developers who want to create data-driven applications in the JavaScript/Node.js environment. Intermediate-level knowledge of JavaScript programming and data science using pandas is expected.

If you know HTML and/or CSS and want to take your skills to the next level, or even if you are a complete web novice, you really need to learn JavaScript. Not only is it the language behind the smooth and dynamic operation of Web 2.0 websites like Facebook, Twitter and Gmail, but in conjunction with HTML5 it's also the standard means Microsoft supports for creating Windows 8 apps - JavaScript is definitely the future for Windows! So, whether you want to simply add a little functionality to your website, such as smooth menus that pop up and down, image transition effects, user-friendly form handling and verification, or anything else that's more than a simple, flat HTML/CSS design, JavaScript is the way to go. What's more, JavaScript is easy. If you've ever tried to learn it and been put off by a plethora of jargon and technical mumbo-jumbo then you're in for a real treat, because Robin Nixon's Crash Courses have helped tens of thousands of people learn the new skills they need. From the top-selling author of "Learning PHP, MySQL & JavaScript," and starting from the ground up with no assumption of prior knowledge, every aspect of JavaScript is explained in this book, in logical order with plenty of simple examples, clear explanations, informative figures, and advice on how best to use the new things you learn. If you want to learn JavaScript up to a solid intermediate level, this book will teach you all you need to know, without recourse to other books and materials. Plus all the examples are free to download from the companion website, so you won't have to type them in to follow along and try them out for yourself. This course features the following lectures: Introduction to JavaScript Incorporating JavaScript Code Into a Web Page JavaScript Language Syntax JavaScript Operators JavaScript Arrays Multidimensional Arrays The JavaScript Array Functions Controlling Program Flow Looping Sections of Code JavaScript Functions JavaScript Objects Errors and Expressions The Document Object Model Advanced JavaScript Appendix: 150+ Functions Detailed Reasons why you will learn all you need from this course: No assumption is made of previous knowledge. Every new concept is explained in logical order. Fully-tested examples are provided throughout. Each lecture features several notes offering extra, handy advice. The examples can all be downloaded free from the companion website.

This book covers the crossroads of web development and deep learning. Both technologies are beginning to meet, and this honeymoon will produce new fantastic applications that you cannot even imagine yet. In this book you will see how to concretely use the main JavaScript deep learning frameworks and web programming in the browser with the capture of inputs and the WebGL implementation. Deep learning in the browser is currently at an embryonic stage, but this is the best time to bet on it before it becomes a giant, and this book will get you in on the action. Are you ready to embark on the adventure?

Design the MIND of a Robotic Thinker! " This book will help you get started with this exciting language and gives you an idea of what is possible. " - Melchizedek B, from Amazon.com " The examples it uses are easy to follow and the illustrations bring out the more complex aspects while making them simple. " - C. Brant, from Amazon.com " Such a cool book that covers basic Javascript programming then incorporates tools and components to explore Artificial Intelligence. " - M. Gavel, from Amazon.com * * INCLUDED BONUS: a Quick-start guide to Learning Javascript in less than a Day! * * How would you like to Create the Next SIRI? Artificial Intelligence. One of the most brilliant creations of mankind. No longer a sci-fi fantasy, but a realistic approach to making work more efficient and lives easier.And the best news? It's not that complicated after all Does it require THAT much advanced math? NO!And are you paying THOUSANDS of dollars just to learn this information? NO!Hundreds? Not even close. Within this book's pages, you'll find GREAT coding skills to learn - and more. Just some of the questions and topics include: - Complicated scheduling problem? Here's how to solve it. - How good are your AI algorithms? Analysis for Efficiency- How to interpret a system into logical code for the AI- How would an AI system would diagnose a system? We show you...- Getting an AI agent to solve problems for youand Much, much more!World-Class TrainingThis book breaks your training

down into easy-to-understand modules. It starts from the very essentials of algorithms and program procedures, so you can write great code - even as a beginner!

Build machine learning web applications without having to learn a new language. This book will help you develop basic knowledge of machine learning concepts and applications. You'll learn not only theory, but also dive into code samples and example projects with TensorFlow.js. Using these skills and your knowledge as a web developer, you'll add a whole new field of development to your tool set. This will give you a more concrete understanding of the possibilities offered by machine learning. Discover how ML will impact the future of not just programming in general, but web development specifically. Machine learning is currently one of the most exciting technology fields with the potential to impact industries from health to home automation to retail, and even art. Google has now introduced TensorFlow.js—an iteration of TensorFlow aimed directly at web developers. Practical Machine Learning in JavaScript will help you stay relevant in the tech industry with new tools, trends, and best practices. What You'll Learn Use the JavaScript framework for ML Build machine learning applications for the web Develop dynamic and intelligent web content Who This Book Is For Web developers and who want a hands-on introduction to machine learning in JavaScript. A working knowledge of the JavaScript language is recommended.

Develop and deploy deep learning web apps using the TensorFlow.js library. TensorFlow.?js? is part of a bigger framework named TensorFlow, which has many tools that supplement it, such as TensorBoard?, ?ml5js?, ?tfjs-vis. This book will cover all these technologies and show they integrate with TensorFlow.?js? to create intelligent web apps. The most common and accessible platform users interact with everyday is their web browser, making it an ideal environment to deploy AI systems. TensorFlow.js is a well-known and battle-tested library for creating browser solutions. Working in JavaScript, the so-called language of the web, directly on a browser, you can develop and serve deep learning applications.You'll work with deep learning algorithms such as feedforward neural networks, convolutional neural networks (CNN), recurrent neural networks (RNN), and generative adversarial network (GAN). Through hands-on examples, apply these networks in use cases related to image classification, natural language processing, object detection, dimensionality reduction, image translation, transfer learning, and time series analysis. Also, these topics are very varied in terms of the kind of data they use, their output, and the training phase. Not everything in machine learning is deep networks, there is also what some call shallow or traditional machine learning. While TensorFlow.js is not the most common place to implement these, you'll be introduce them and review the basics of machine learning through TensorFlow.js. What You'll Learn Build deep learning products suitable for web browsers Work with deep learning algorithms such as feedforward neural networks, convolutional neural networks (CNN), recurrent neural networks (RNN), and generative adversarial network (GAN) Develop apps using image classification, natural language processing, object detection, dimensionality reduction, image translation, transfer learning, and time series analysis Who This Book Is For Programmers developing deep learning solutions for the web and those who want to learn TensorFlow.js with at least minimal programming and software development knowledge. No prior JavaScript knowledge is required, but familiarity with it is helpful.

A quick, easy guide to building cutting-edge intelligence into your website Without "smart," interactive content, your website just won't cut it with today's savvy Web surfers. Fortunately, with JavaScript, it's now incredibly easy to build intelligent sites, and this book shows you how. Featuring a rapid skill-building format, this book is organized around goals you want to accomplish with JavaScript rather than the language's technical aspects. Even if you've never programmed a line of code in your life, you'll quickly learn how to enliven your site with the smarts it needs to automatically tailor its content to an individual user's needs, solve problems, make decisions, create new pages on the fly, and even close sales! Nigel Ford eases you into the basics of JavaScript and then, with the help of dozens of vivid examples, teaches you: * Basic and advanced techniques for adding intelligence to your Web pages * How to store different types of knowledge in your Web pages * Techniques for handling requests, accessing files, and using databases in server-side JavaScript * Programming for language processing, searching, problem solving, and game playing * How to program for both the client side and server side with JavaScript.

Combining the demand for AI with the ubiquity of JavaScript was inevitable. With Google's TensorFlow.js framework, seasoned AI veterans and web developers alike can help propel the future of AI-driven websites. In this guide, author Gant Laborde--Google Developer Expert in machine learning and the web--provides a hands-on, end-to-end approach to TensorFlow.js fundamentals for a broad technical audience that includes data scientists, engineers, web developers, students, and researchers. You'll begin by working through some basic examples in TensorFlow.js before diving deeper into neural network architectures, DataFrames, TensorFlow Hub, model conversion, transfer learning, and more. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems with TensorFlow.js. Explore tensors, the most fundamental structure of machine learning Convert data into tensors and back with a real-world example Combine AI with the web using TensorFlow.js and other tools Use resources to convert, train, and manage machine learning data Start building and training your own training models from scratch Learn how to create your own image classification models Examine transfer learning: retraining an advanced model to perform a new task

Are you looking for a super-fast computer programming course? Would you like to learn the Python Programming Language in 7 days? Do you want to increase your business thanks to the web applications? If so, keep reading: this bundle book is for you! Finally on launch the most complete Python guide with 3 Manuscripts in 1 book: 1-Python for beginners 2-Python for Data Science 4-Python Crash Course Python will introduce you many selected practices for coding . You will discover as a beginner the world of data science, machine learning and artificial intelligence. The following list is just a tiny fraction of what you will learn in this collection bundle. 1) Python for beginners ? The basics of Python programming ? Differences among programming languages ? Vba, SQL, R, Python ? Game creation with Pyhton ? Easy-to-follow steps for reading and writing codes. ? Control flow statements and Error handling ? 4 best strategies with NumPy, Pandas, Matplotlib 2) Python for Data science ? 4 reason why Python is fundamental for Data Science ? Python design patterns ? How to use Python Data Analysis in your business ? Data visualization optimal tools and techniques ? Analysis of popular Python projects templates ? How to set up the Python environment for Data Science ? Most important Machine Learning Algorithms ? How to leverage Data Science in the Cloud 3) Python Crash Course * A Proven Method to Write your First Program in 7 Days * 5 Common Mistakes to Avoid when You Start Coding * A Simple Strategy to Write Clean, Understandable and Flexible Codes * The One Thing You Need to Debug your Codes in Python * 5 Practical exercises to start programming Even if you have never written a programming code before, you will quickly grasp the basics thanks to visual charts and guidelines for coding. Examples and step-by-step guides will guide you during the code-writing learning process. The description of each topic is crystal-clear and you can easily practice with related exercises. You will also learn all the best tricks of

writing codes with point by point descriptions of the code elements. If you really wish to to learn Python and master its language, please click the BUY NOW button.

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.

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

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.

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire

The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow Serving

AI as a Service is a practical handbook to building and implementing serverless AI applications, without bogging you down with a lot of theory. Instead, you'll find easy-to-digest instruction and two complete hands-on serverless AI builds in this must-have guide! Summary Companies everywhere are moving everyday business processes over to the cloud, and AI is increasingly being given the reins in these tasks. As this massive digital transformation continues, the combination of serverless computing and AI promises to become the de facto standard for business-to-consumer platform development—and developers who can design, develop, implement, and maintain these systems will be in high demand! AI as a Service is a practical handbook to building and implementing serverless AI applications, without bogging you down with a lot of theory. Instead, you'll find easy-to-digest instruction and two complete hands-on serverless AI builds in this must-have guide! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Cloud-based AI services can automate a variety of labor intensive business tasks in areas such as customer service, data analysis, and financial reporting. The secret is taking advantage of pre-built tools like Amazon Rekognition for image analysis or AWS Comprehend for natural language processing. That way, there's no need to build expensive custom software. Artificial Intelligence (AI), a machine's ability to learn and make predictions based on patterns it identifies, is already being leveraged by businesses around the world in areas like targeted product recommendations, financial forecasting and resource planning, customer service chatbots, healthcare diagnostics, data security, and more. With the exciting combination of serverless computing and AI, software developers now have enormous power to improve their businesses' existing systems and rapidly deploy new AI-enabled platforms. And to get on this fast-moving train, you don't have to invest loads of time and effort in becoming a data scientist or AI expert, thanks to cloud platforms and the readily available off-the-shelf cloud-based AI services! About the book AI as a Service is a fast-paced guide to harnessing the power of cloud-based solutions. You'll learn to build real-world apps—such as chatbots and text-to-speech services—by stitching together cloud components. Work your way from small projects to large data-intensive applications. What's inside - Apply cloud AI services to existing platforms - Design and build scalable data pipelines - Debug and troubleshoot AI services - Start fast with serverless templates About the reader For software developers familiar with cloud basics. About the author Peter Elger and Eóin Shanaghy are founders and CEO/CTO of fourTheorem, a software solutions company providing expertise on architecture, DevOps, and machine learning. Table of Contents PART 1 - FIRST STEPS 1 A tale of two technologies 2 Building a serverless image recognition system, part 1 3 Building a serverless image recognition system, part 2 PART 2 - TOOLS OF THE TRADE 4 Building and securing a web application the serverless way 5 Adding AI interfaces to a web application 6 How to be effective with AI as a Service 7 Applying AI to existing platforms PART 3 - BRINGING IT ALL TOGETHER 8 Gathering data at scale for real-world AI 9 Extracting value from large data sets with AI

You've decided to tackle machine learning - because you're job hunting, embarking on a new project, or just think self-driving cars are cool. But where to start? It's easy to be intimidated, even as a software developer. The good news is that it doesn't have to be that hard. Master machine learning by writing code one line at a time, from simple learning programs all the way to a true deep learning system. Tackle the hard topics by breaking them down so they're easier to understand, and build your confidence by getting your hands dirty. Peel away the obscurities of machine learning, starting from scratch and going all the way to deep learning. Machine learning can be intimidating, with its reliance on math and algorithms that most programmers don't encounter in their regular work. Take a hands-on approach, writing the Python code yourself, without any libraries to obscure what's really going on. Iterate on your design, and add layers of complexity as you go. Build an image recognition application from scratch with supervised learning. Predict the future with linear regression. Dive into gradient descent, a fundamental algorithm that drives most of machine learning. Create perceptrons to classify data. Build neural networks to tackle more complex and sophisticated data sets. Train and refine those networks with backpropagation and batching. Layer the neural networks, eliminate overfitting, and add convolution to transform your neural network into a true deep learning system. Start from the beginning and code your way to machine learning mastery. What You Need: The examples in this book are written in Python, but don't worry if you don't know this language: you'll pick up all the Python you need very quickly. Apart from that, you'll only need your computer, and your code-adept brain.

A hands-on, application-based introduction to machine learning and artificial intelligence (AI) that guides young readers through creating compelling AI-powered games and applications using the Scratch programming language. Machine learning (also known as ML) is one of the building blocks of AI, or artificial intelligence. AI is based on the idea that computers can learn on their own, with your help. Machine Learning for Kids will introduce you to machine learning, painlessly. With this book and its free, Scratch-based, award-winning companion website, you'll see how easy it is to add machine learning to your own projects. You don't even need to know how to code! As you work through the book you'll discover how machine learning systems can be taught to recognize text, images, numbers, and sounds, and how to train your models to improve their accuracy. You'll turn your models into fun computer games and apps, and see what happens when they get confused by bad data. You'll build 13 projects step-by-step from the ground up, including: • Rock, Paper, Scissors game that recognizes your hand shapes • An app that recommends movies based on other movies that you like • A computer character that reacts to insults and compliments • An interactive virtual assistant (like Siri or Alexa) that obeys commands • An AI version of Pac-Man, with a smart character that knows how to avoid ghosts NOTE: This book includes a Scratch tutorial for beginners, and step-by-step instructions for every project. Ages 12+

If you're an experienced programmer interested in crunching data, this book will get you started with machine learning—a toolkit of algorithms that enables computers to train themselves to automate useful tasks. Authors Drew Conway and John Myles White help you understand machine learning and statistics tools through a series of hands-on case studies, instead of a traditional math-heavy presentation. Each chapter focuses on a specific problem in machine learning, such as classification, prediction, optimization, and recommendation. Using the R programming language, you'll learn how to analyze sample datasets and write simple machine learning algorithms. Machine Learning for Hackers is ideal for programmers from any background, including business, government, and academic research. Develop a naïve Bayesian classifier to determine if an email is spam, based only on its text Use linear regression to predict the number of page views for the top 1,000 websites Learn optimization techniques by attempting to break a simple letter cipher Compare and contrast U.S. Senators statistically,

based on their voting records Build a “whom to follow” recommendation system from Twitter data

Artificial Intelligence: A Modern Approach offers the most comprehensive, up-to-date introduction to the theory and practice of artificial intelligence. Number one in its field, this textbook is ideal for one or two-semester, undergraduate or graduate-level courses in Artificial Intelligence.

Summary Data Wrangling with JavaScript is hands-on guide that will teach you how to create a JavaScript-based data processing pipeline, handle common and exotic data, and master practical troubleshooting strategies. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Why not handle your data analysis in JavaScript? Modern libraries and data handling techniques mean you can collect, clean, process, store, visualize, and present web application data while enjoying the efficiency of a single-language pipeline and data-centric web applications that stay in JavaScript end to end. About the Book Data Wrangling with JavaScript promotes JavaScript to the center of the data analysis stage! With this hands-on guide, you'll create a JavaScript-based data processing pipeline, handle common and exotic data, and master practical troubleshooting strategies. You'll also build interactive visualizations and deploy your apps to production. Each valuable chapter provides a new component for your reusable data wrangling toolkit. What's inside Establishing a data pipeline Acquisition, storage, and retrieval Handling unusual data sets Cleaning and preparing raw dataInteractive visualizations with D3 About the Reader Written for intermediate JavaScript developers. No data analysis experience required. About the Author Ashley Davis is a software developer, entrepreneur, author, and the creator of Data-Forge and Data-Forge Notebook, software for data transformation, analysis, and visualization in JavaScript. Table of Contents Getting started: establishing your data pipeline Getting started with Node.js Acquisition, storage, and retrieval Working with unusual data Exploratory coding Clean and prepare Dealing with huge data files Working with a mountain of data Practical data analysis Browser-based visualization Server-side visualization Live data Advanced visualization with D3 Getting to production

Describes ways to incorporate domain modeling into software development.

Why a new approach is needed in the quest for general artificial intelligence. Since the inception of artificial intelligence, we have been warned about the imminent arrival of computational systems that can replicate human thought processes. Before we know it, computers will become so intelligent that humans will be lucky to kept as pets. And yet, although artificial intelligence has become increasingly sophisticated—with such achievements as driverless cars and humanless chess-playing—computer science has not yet created general artificial intelligence. In Algorithms Are Not Enough, Herbert Roitblat explains how artificial general intelligence may be possible and why a robopocalypse is neither imminent, nor likely. Existing artificial intelligence, Roitblat shows, has been limited to solving path problems, in which the entire problem consists of navigating a path of choices—finding specific solutions to well-structured problems. Human problem-solving, on the other hand, includes problems that consist of ill-structured situations, including the design of problem-solving paths themselves. These are insight problems, and insight is an essential part of intelligence that has not been addressed by computer science. Roitblat draws on cognitive science, including psychology, philosophy, and history, to identify the essential features of intelligence needed to achieve general artificial intelligence. Roitblat describes current computational approaches to intelligence, including the Turing Test, machine learning, and neural networks. He identifies building blocks of natural intelligence, including perception, analogy, ambiguity, common sense, and creativity. General intelligence can create new representations to solve new problems, but current computational intelligence cannot. The human brain, like the computer, uses algorithms; but general intelligence, he argues, is more than algorithmic processes.

Summary Deep learning has transformed the fields of computer vision, image processing, and natural language applications. Thanks to TensorFlow.js, now JavaScript developers can build deep learning apps without relying on Python or R. Deep Learning with JavaScript shows developers how they can bring DL technology to the web. Written by the main authors of the TensorFlow library, this new book provides fascinating use cases and in-depth instruction for deep learning apps in JavaScript in your browser or on Node. Foreword by Nikhil Thorat and Daniel Smilkov. About the technology Running deep learning applications in the browser or on Node-based backends opens up exciting possibilities for smart web applications. With the TensorFlow.js library, you build and train deep learning models with JavaScript. Offering uncompromising production-quality scalability, modularity, and responsiveness, TensorFlow.js really shines for its portability. Its models run anywhere JavaScript runs, pushing ML farther up the application stack. About the book In Deep Learning with JavaScript, you’ll learn to use TensorFlow.js to build deep learning models that run directly in the browser. This fast-paced book, written by Google engineers, is practical, engaging, and easy to follow. Through diverse examples featuring text analysis, speech processing, image recognition, and self-learning game AI, you’ll master all the basics of deep learning and explore advanced concepts, like retraining existing models for transfer learning and image generation. What's inside - Image and language processing in the browser - Tuning ML models with client-side data - Text and image creation with generative deep learning - Source code samples to test and modify About the reader For JavaScript programmers interested in deep learning. About the author Shanging Cai, Stanley Bileschi and Eric D. Nielsen are software engineers with experience on the Google Brain team, and were crucial to the development of the high-level API of TensorFlow.js. This book is based in part on the classic, Deep Learning with Python by François Chollet. TOC: PART 1 - MOTIVATION AND BASIC CONCEPTS 1 • Deep learning and JavaScript PART 2 - A GENTLE INTRODUCTION TO TENSORFLOW.JS 2 • Getting started: Simple linear regression in TensorFlow.js 3 • Adding nonlinearity: Beyond weighted sums 4 • Recognizing images and sounds using convnets 5 • Transfer learning: Reusing pretrained neural networks PART 3 - ADVANCED DEEP LEARNING WITH TENSORFLOW.JS 6 • Working with data 7 • Visualizing data and models 8 • Underfitting, overfitting, and the universal workflow of machine learning 9 • Deep learning for sequences and text 10 • Generative deep learning 11 • Basics of deep reinforcement learning PART 4 - SUMMARY AND CLOSING WORDS 12 • Testing, optimizing, and deploying models 13 • Summary, conclusions, and beyond