Tinyml Machine Learning With Tensorflow On Arduin

LIVIA DONAVAN

Learning Deep Learning Packt Publishing Ltd
This practical book shows you how to employ machine learning models to extract information from images. ML engineers and data scientists will learn how to solve a variety of image problems including classification, object detection, autoencoders, image generation, counting, and captioning with proven ML techniques. This book provides a great introduction to end-to-end deep learning: dataset creation, data preprocessing, model design, model training, evaluation, deployment, and interpretability. Google engineers Valliappa Lakshmanan, Martin Görner, and Ryan Gillard show you how to develop accurate and explainable computer vision ML models and put them into large-scale production using robust ML architecture in a flexible and maintainable way. You'll learn how to design, train, evaluate, and predict with models written in TensorFlow or Keras. You'll learn how to: Design ML architecture for computer vision tasks Select a model (such as ResNet, SqueezeNet, or EfficientNet) appropriate to your task Create an end-to-end ML pipeline to train, evaluate, deploy, and explain your model Preprocess images for data augmentation and to support learnability Incorporate explainability and responsible AI best practices Deploy image models as web services or on edge devices Monitor and manage ML models
Hands-On Deep Learning for Images with TensorFlow Simon and Schuster
Gain expert guidance on how to successfully develop machine learning models in Python and build your own unique data platforms Key FeaturesGain a full understanding of the model production and deployment processBuild your first machine learning model in just five minutes and get a hands-on machine learning experienceUnderstand how to deal with common challenges in data science projectsBook Description Where there’s data, there’s insight. With so much data being generated, there is immense scope to extract meaningful information that’ll boost business productivity and profitability. By learning to convert raw data into game-changing insights, you’ll open new career paths and opportunities. The Data Science Workshop begins by introducing different types of projects and showing you how to incorporate machine learning algorithms in them. You’ll learn to select a relevant metric and even assess the performance of your model. To tune the hyperparameters of an algorithm and improve its accuracy, you’ll get hands-on with approaches such as grid search and random search. Next, you’ll learn dimensionality reduction techniques to easily handle many variables at once, before exploring how to use model ensembling techniques and create new features to enhance model performance. In a bid to help you automatically create new features that improve your model, the book demonstrates how to use the automated feature engineering tool. You’ll also understand how to use the orchestration and scheduling workflow to deploy machine learning models in batch. By the end of this book, you’ll have the skills to start working on data science projects confidently. By the end of this book, you’ll have the skills to start working on data science projects confidently. What you will learnExplore the key differences between supervised learning and unsupervised learningManipulate and analyze data using scikit-learn and pandas librariesUnderstand key concepts such as regression, classification, and clusteringDiscover advanced techniques to improve the accuracy of your modelUnderstand how to speed up the process of adding new featuresSimplify your machine learning workflow for productionWho this book is for This is one of the most useful data science books for aspiring data analysts, data scientists, database engineers, and business analysts. It is aimed at those who want to kick-start their careers in data science by quickly learning data science techniques without going through all the mathematics behind machine learning algorithms. Basic knowledge of the Python programming language will help you easily grasp the concepts explained in this book.
Making Embedded Systems "O'Reilly Media, Inc."
The design patterns in this book capture best practices and solutions to recurring problems in machine learning. The authors, three Google engineers, catalog proven methods to help data scientists tackle common problems throughout the ML process. These design patterns codify the experience of hundreds of experts into straightforward, approachable advice. In this book, you will find detailed explanations of 30 patterns for data and problem representation, operationalization, repeatability, reproducibility, flexibility, explainability, and fairness. Each pattern includes a description of the problem, a variety of potential solutions, and recommendations for choosing the best technique for your situation. You'll learn how to: Identify and mitigate common challenges when training, evaluating, and deploying ML models Represent data for different ML model types, including embeddings, feature crosses, and more Choose the right model type for specific problems Build a robust training loop that uses checkpoints, distribution strategy, and hyperparameter tuning Deploy scalable ML systems that you can retrain and update to reflect new data Interpret model predictions for stakeholders and ensure models are treating users fairly
40 Algorithms Every Programmer Should Know Lulu.com
Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively Make informed decisions by identifying the strengths and weaknesses of different tools Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity Understand the distributed systems research upon which modern databases are built Peek behind the scenes of major online services, and learn from their architectures
Python Image Processing Cookbook Packt Publishing Ltd
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
TinyML Packt Publishing Ltd
Explore TensorFlow's capabilities to perform efficient deep learning on images Key Features Discover image processing for machine vision Build an effective image classification system using the power of CNNs Leverage TensorFlow’s capabilities to perform efficient deep learning Book Description TensorFlow is Google’s popular offering for machine learning and deep learning, quickly becoming a favorite tool for performing fast, efficient, and accurate deep learning tasks. Hands-On Deep Learning for Images with TensorFlow shows you the practical implementations of real-world projects, teaching you how to leverage TensorFlow’s capabilities to perform efficient image processing using the power of deep learning. With the help of this book, you will get to grips with the

different paradigms of performing deep learning such as deep neural nets and convolutional neural networks, followed by understanding how they can be implemented using TensorFlow. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow and Keras. What you will learn Build machine learning models particularly focused on the MNIST digits Work with Docker and Keras to build an image classifier Understand natural language models to process text and images Prepare your dataset for machine learning Create classical, convolutional, and deep neural networks Create a RESTful image classification server Who this book is for Hands-On Deep Learning for Images with TensorFlow is for you if you are an application developer, data scientist, or machine learning practitioner looking to integrate machine learning into application software and master deep learning by implementing practical projects in TensorFlow. Knowledge of Python programming and basics of deep learning are required to get the best out of this book.

Learning Tensorflow. Js O'Reilly Media

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

Low-Power Computer Vision Packt Publishing Ltd

Build smarter systems by combining artificial intelligence and the Internet of Things—two of the most talked about topics today Key FeaturesLeverage the power of Python libraries such as TensorFlow and Keras to work with real-time IoT dataProcess IoT data and predict outcomes in real time to build smart IoT modelsCover practical case studies on industrial IoT, smart cities, and home automationBook Description There are many applications that use data science and analytics to gain insights from terabytes of data. These apps, however, do not address the challenge of continually discovering patterns for IoT data. In Hands-On Artificial Intelligence for IoT, we cover various aspects of artificial intelligence (AI) and its implementation to make your IoT solutions smarter. This book starts by covering the process of gathering and preprocessing IoT data gathered from distributed sources. You will learn different AI techniques such as machine learning, deep learning, reinforcement learning, and natural language processing to build smart IoT systems. You will also leverage the power of AI to handle real-time data coming from wearable devices. As you progress through the book, techniques for building models that work with different kinds of data generated and consumed by IoT devices such as time series, images, and audio will be covered. Useful case studies on four major application areas of IoT solutions are a key focal point of this book. In the concluding chapters, you will leverage the power of widely used Python libraries, TensorFlow and Keras, to build different kinds of smart AI models. By the end of this book, you will be able to build smart AI-powered IoT apps with confidence. What you will learnApply different AI techniques including machine learning and deep learning using TensorFlow and KerasAccess and process data from various distributed sourcesPerform supervised and unsupervised machine learning for IoT dataImplement distributed processing of IoT data over Apache Spark using the MLLib and H2O.ai platformsForecast time-series data using deep learning methodsImplementing AI from case studies in Personal IoT, Industrial IoT, and Smart CitiesGain unique insights from data obtained from wearable devices and smart devicesWho this book is for If you are a data science professional or a machine learning developer looking to build smart systems for IoT, Hands-On Artificial Intelligence for IoT is for you. If you want to learn how popular artificial intelligence (AI) techniques can be used in the Internet of Things domain, this book will also be of benefit. A basic understanding of machine learning concepts will be required to get the best out of this book.

Learning TensorFlow.js "O'Reilly Media, Inc."

Demonstrates the benefits of internet enabled embedded systems using real-world applications. This work examines the techniques required to achieve internet connectivity, starting with how to draw upon those TCP/IP implementations which already exist, right through to developing fresh ones. It also includes a CD-ROM with the TCP/IP stack.

Neural Computation in Hopfield Networks and Boltzmann Machines O'Reilly Media

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

Grokking Machine Learning University of Delaware Press

NVIDIA's Full-Color Guide to Deep Learning: All You Need to Get Started and Get Results "To enable everyone to be part of this historic revolution requires the democratization of AI knowledge and resources. This book is timely and relevant towards accomplishing these lofty goals." -- From the foreword by Dr. Anima Anandkumar, Bren Professor, Caltech, and Director of ML Research, NVIDIA "Ekman uses a learning technique that in our experience has proven pivotal to success—asking the reader to think about using DL techniques in practice. His straightforward approach is refreshing, and he permits the reader to dream, just a bit, about where DL may yet take us." -- From the foreword by Dr. Craig Clawson, Director, NVIDIA Deep Learning Institute Deep learning (DL) is a key component of today's exciting advances in machine learning and artificial intelligence. Learning Deep Learning is a complete guide to DL. Illuminating both the core concepts and the hands-on programming techniques needed to succeed, this book is ideal for developers, data scientists, analysts, and others--including those with no prior machine learning or statistics experience. After introducing the essential building blocks of deep neural networks, such as artificial neurons and fully connected, convolutional, and recurrent layers, Magnus Ekman shows how to use them to build advanced architectures, including the Transformer. He describes how these concepts are used to build modern networks for computer vision and natural language processing (NLP), including Mask R-CNN, GPT, and BERT. And he explains how a natural language translator and a system generating natural language descriptions of images. Throughout, Ekman provides concise, well-annotated code examples using TensorFlow with Keras. Corresponding PyTorch examples are provided online, and the book thereby covers the two dominating Python libraries for DL used in industry and academia. He concludes with an introduction to neural architecture search (NAS), exploring important ethical issues and providing resources for further learning. Explore and master core concepts: perceptrons, gradient-based learning, sigmoid neurons, and back propagation See how DL frameworks make it easier to develop more complicated and useful neural networks Discover how convolutional neural networks (CNNs) revolutionize image classification and analysis Apply recurrent neural networks (RNNs) and long short-term memory (LSTM) to text and other variable-length sequences Master NLP with sequence-to-sequence networks and the Transformer architecture Build applications for natural language translation and image captioning NVIDIA's invention of the GPU sparked the PC gaming market. The company's pioneering work in accelerated computing--a supercharged form of computing at the intersection of computer graphics, high-performance computing, and AI--is reshaping trillion-dollar industries, such as transportation, healthcare, and manufacturing, and fueling the growth of many others. Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

AI and Machine Learning for Coders O'Reilly Media

Develop and Deploy Powerful MSP432 Microcontroller Applications Bolster your electronics skills and learn to work with the cutting-edge MSP432 microcontroller using the practical information contained in this comprehensive guide. Programmable Microcontrollers: Applications on the MSP432 LaunchPad clearly explains each concept and features detailed illustrations, real-world examples, and DIY projects. Discover how to configure the MSP432, program custom functions, interface with external hardware, and communicate via WiFi. Ideal for practicing engineers and hobbyists alike, this hands-on guide empowers you to program all microcontrollers by thoroughly understanding the MSP432. Coverage includes: •MSP432 architecture •Code Composer Studio (CCS) •CCS Cloud and Energia •MSP432 programming with C and Assembly •Digital I/O •Exceptions and interrupts •Power management and timing operations •Mixed signal systems •Digital and wireless communication •Flash memory, RAM, and direct memory access •Real-time operating system •Advanced applications

Machine Learning with TensorFlow, Second Edition O'Reilly Media

Chapter 2. Introduction to Computer Vision -- Using Neurons for Vision -- Your First Classifier: Recognizing Clothing Items -- The Data: Fashion MNIST -- A Model Architecture to Parse Fashion MNIST -- Coding the Fashion MNIST Model -- Transfer Learning for Computer Vision -- Summary -- Chapter 3. Introduction to ML Kit -- Building a Face Detection App on Android -- Step 1: Create the App with Android Studio -- Step 2: Add and Configure ML Kit -- Step 3: Define the User Interface -- Step 4: Add the Images as Assets -- Step 5: Load the UI with a Default Picture.

IoT Streams for Data-Driven Predictive Maintenance and IoT, Edge, and Mobile for Embedded Machine Learning "O'Reilly Media, Inc."

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

The Embedded Internet "O'Reilly Media, Inc."

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

Programmable Microcontrollers: Applications on the MSP432 LaunchPad Pragmatic Bookshelf

Whether you’re a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley’s Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Deep Learning with PyTorch Simon and Schuster

Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better

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

The The Data Science Workshop Simon and Schuster

Build machine and deep learning systems with the newly released TensorFlow 2 and Keras for the lab, production, and mobile devices Key FeaturesIntroduces and then uses TensorFlow 2 and Keras right from the startTeaches key machine and deep learning techniquesUnderstand the fundamentals of deep learning and machine learning through clear explanations and extensive code samplesBook Description Deep Learning with TensorFlow 2 and Keras, Second Edition teaches neural networks and deep learning techniques alongside TensorFlow (TF) and Keras. You’ll learn how to write deep learning applications in the most powerful, popular, and scalable machine learning stack available. TensorFlow is the machine learning library of choice for professional applications, while Keras offers a simple and powerful Python API for accessing TensorFlow. TensorFlow 2 provides full Keras integration, making advanced machine learning easier and more convenient than ever before. This book also introduces neural networks with TensorFlow, runs through the main applications (regression, ConvNets (CNNs), GANs, RNNs, NLP), covers two working example apps, and then dives into TF in production, TF mobile, and using TensorFlow with AutoML. What you will learnBuild machine learning and deep learning systems with TensorFlow 2 and the Keras APIUse Regression analysis, the most popular approach to machine learningUnderstand ConvNets (convolutional neural networks) and how they are essential for deep learning systems such as image classifiersUse GANs (generative adversarial networks) to create new data that fits with existing patternsDiscover RNNs (recurrent neural networks) that can process sequences of input intelligently, using one part of a sequence to correctly interpret anotherApply deep learning to natural human language and interpret natural language texts to produce an appropriate responseTrain your models on the cloud and put TF to work in real environmentsExplore how Google tools can automate simple ML workflows without the need for complex modelingWho this book is for This book is for Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. This book gives you the theory and practice required to use Keras, TensorFlow 2, and AutoML to build machine learning systems. Some knowledge of machine learning is expected.

Interpretable Machine Learning Packt Publishing Ltd

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

Deep Learning with TensorFlow 2 and Keras O'Reilly Media

Another day without Test-Driven Development means more time wasted chasing bugs and watching your code deteriorate. You thought TDD was for someone else, but it's not! It's for you, the embedded C programmer. TDD helps you prevent defects and build software with a long useful life. This is the first book to teach the hows and whys of TDD for C programmers. TDD is a modern programming practice C developers need to know. It's a different way to program---unit tests are written in a tight feedback loop with the production code, assuring your code does what you think. You get valuable feedback every few minutes. You find mistakes before they become bugs. You get early warning of design problems. You get immediate notification of side effect defects. You get to spend more time adding valuable features to your product. James is one of the few experts in applying TDD to embedded C. With his 1.5 decades of training,coaching, and practicing TDD in C, C++, Java, and C# he will lead you from being a novice in TDD to using the techniques that few have mastered. This book is full of code written for embedded C programmers. You don't just see the end product, you see code and tests evolve. James leads you through the thought process and decisions made each step of the way. You'll learn techniques for test-driving code right nextto the hardware, and you'll learn design principles and how to apply them to C to keep your code clean and flexible. To run the examples in this book, you will need a C/C++ development environment on your machine, and the GNU GCC tool chain or Microsoft Visual Studio for C++ (some project conversion may be needed).