Mindluster Platform

About the course Deep Learning with PyTorch

Is PyTorch a deep learning? PyTorch is the premier open-source deep learning framework developed and maintained by Facebook. At its core, PyTorch is a mathematical library that allows you to perform efficient computation and automatic differentiation on graph-based models. Predictive modeling with deep learning is a skill that modern developers need to know. PyTorch is the premier open-source deep learning framework developed and maintained by Facebook. At its core, PyTorch is a mathematical library that allows you to perform efficient computation and automatic differentiation on graph-based models. Achieving this directly is challenging, although thankfully, the modern PyTorch API provides classes and idioms that allow you to easily develop a suite of deep learning models. In this tutorial, you will discover a step-by-step guide to developing deep learning models in PyTorch. After completing this tutorial, you will know: The difference between Torch and PyTorch and how to install and confirm PyTorch is working. The five-step life-cycle of PyTorch models and how to define, fit, and evaluate models. How to develop PyTorch deep learning models for regression, classification, and predictive modeling tasks. PyTorch Tutorial Overview The focus of this tutorial is on using the PyTorch API for common deep learning model development tasks; we will not be diving into the math and theory of deep learning. For that, I recommend starting with this excellent book. The best way to learn deep learning in python is by doing. Dive in. You can circle back for more theory later. I have designed each code example to use best practices and to be standalone so that you can copy and paste it directly into your project and adapt it to your specific needs. This will give you a massive head start over trying to figure out the API from official documentation alone. It is a large tutorial, and as such, it is divided into three parts; they are: How to Install PyTorch What Are Torch and PyTorch? How to Install PyTorch How to Confirm PyTorch Is Installed PyTorch Deep Learning Model Life-Cycle Step 1: Prepare the Data Step 2: Define the Model Step 3: Train the Model Step 4: Evaluate the Model Step 5: Make Predictions How to Develop PyTorch Deep Learning Models How to Develop an MLP for Binary Classification How to Develop an MLP for Multiclass Classification How to Develop an MLP for Regression How to Develop a CNN for Image Classification You Can Do Deep Learning in Python! Work through this tutorial. It will take you 60 minutes, max! You do not need to understand everything (at least not right now). Your goal is to run through the tutorial end-to-end and get a result. You do not need to understand everything on the first pass. List down your questions as you go. Make heavy use of the API documentation to learn about all of the functions that youâ€™re using. You do not need to know the math first. Math is a compact way of describing how algorithms work, specifically tools from linear algebra, probability, and calculus. These are not the only tools that you can use to learn how algorithms work. You can also use code and explore algorithm behavior with different inputs and outputs. Knowing the math will not tell you what algorithm to choose or how to best configure it. You can only discover that through carefully controlled experiments. You do not need to know how the algorithms work. It is important to know about the limitations and how to configure deep learning algorithms. But learning about algorithms can come later. You need to build up this algorithm knowledge slowly over a long period of time. Today, start by getting comfortable with the platform. You do not need to be a Python programmer. The syntax of the Python language can be intuitive if you are new to it. Just like other languages, focus on function calls (e.g. function()) and assignments (e.g. a = â€œbâ€). This will get you most of the way. You are a developer; you know how to pick up the basics of a language really fast. Just get started and dive into the details later. You do not need to be a deep learning expert. You can learn about the benefits and limitations of various algorithms later, and there are plenty of tutorials that you can read to brush up on the steps of a deep learning project. 1. How to Install PyTorch In this section, you will discover what PyTorch is, how to install it, and how to confirm that it is installed correctly. 1.1. What Are Torch and PyTorch? PyTorch is an open-source Python library for deep learning developed and maintained by Facebook. The project started in 2016 and quickly became a popular framework among developers and researchers. Torch (Torch7) is an open-source project for deep learning written in C and generally used via the Lua interface. It was a precursor project to PyTorch and is no longer actively developed. PyTorch includes â€œTorchâ€ in the name, acknowledging the prior torch library with the â€œPyâ€ prefix indicating the Python focus of the new project. The PyTorch API is simple and flexible, making it a favorite for academics and researchers in the development of new deep learning models and applications. The extensive use has led to many extensions for specific applications (such as text, computer vision, and audio data), and may pre-trained models that can be used directly. As such, it may be the most popular library used by academics. The flexibility of PyTorch comes at the cost of ease of use, especially for beginners, as compared to simpler interfaces like Keras. The choice to use PyTorch instead of Keras gives up some ease of use, a slightly steeper learning curve, and more code for more flexibility, and perhaps a more vibrant academic community. 1.2. How to Install PyTorch Before installing PyTorch, ensure that you have Python installed, such as Python 3.6 or higher. If you donâ€™t have Python installed, you can install it using Anaconda. This tutorial will show you how: How to Setup Your Python Environment for Machine Learning With Anaconda There are many ways to install the PyTorch open-source deep learning library. The most common, and perhaps simplest, way to install PyTorch on your workstation is by using pip. For example, on the command line, you can type: sudo pip install torch Perhaps the most popular application of deep learning is for computer vision, and the PyTorch computer vision package is called â€œtorchvision.â€ Installing torchvision is also highly recommended and it can be installed as follows: sudo pip install torchvision If you prefer to use an installation method more specific to your platform or package manager, you can see a complete list of installation instructions here: PyTorch Installation Guide There is no need to set up the GPU now. All examples in this tutorial will work just fine on a modern CPU. If you want to configure PyTorch for your GPU, you can do that after completing this tutorial. Donâ€™t get distracted!

Course Lesson(4)

Lesson 1 :
PyTorch Basics and Gradient Descent | Deep Learning with PyTorch Zero to GANs | Part 1 of 6

Lesson 2 :
PyTorch Images and Logistic Regression | Deep Learning with PyTorch Zero to GANs | Part 2 of 6

Lesson 3 :
Training Deep Neural Networks on a GPU | Deep Learning with PyTorch Zero to GANs | Part 3 of 6

Lesson 4 :
Image Classification with Convolutional Neural Networks | Deep Learning with PyTorch Zero to GANs |

Related courses

React Projects Workshops

Node js back end JavaScript

Python programming language

Social Network Theme UI With Sass

Full Stack React Django

Build a Node js App With Sequelize

for Business Contact
business@mindluster.com