



Courses on Industrial Artificial Intelligence and Deep Learning

Prof. Xavier Bresson, NTU, Singapore

xavier.bresson@gmail.com

1. PROPOSED COURSES

- COURSE 1 : NON-TECH INTRODUCTION TO AI AND DISCUSSIONS – 1 DAY
- COURSE 2 : FUNDAMENTALS OF DEEP LEARNING – 2 DAYS
- COURSE 3 : ADVANCED DEEP LEARNING – 2 DAYS
- COURSE 4 : A LA CARTE

2. DESCRIPTION

These short courses provide the participants a clear understanding of the most successful AI/Deep Learning techniques. Course 1 offers a 1-day high-level understanding of modern AI and industrial benefits with ½ day lectures and ½ day brainstorming discussions with company staff to apply AI to their projects. Courses 2 and 3 present the fundamental concepts and theory behind Artificial Neural Networks. Practical exercises of Deep Learning applications focus on PyTorch, the library developed by Facebook AI. Python notebooks on the Cloud are offered to the participants to run the PyTorch exercises. The teaching approach provides a good balance of theory and practice. Course 4 is tailored to company's needs about this new technology.

3. ABOUT THE INSTRUCTOR

Prof. Xavier Bresson¹ (PhD, 2005, EPFL, Switzerland) is Professor in Artificial Intelligence at NTU, Singapore. He has an extensive experience in teaching courses in AI and Deep Learning to professionals worldwide. Since 2014, he has offered industrial courses on the newest AI algorithms to Fortune 500 companies including Deloitte, UnitedHealth, Nestle, UBS, SwissCom, Total, AMD, Apple, etc. He has also delivered trainings in Centers for Professional and Continuing Education at the world's best universities including the University of California, Los Angeles (UCLA)², the Swiss Federal Institute of Technology (EPFL) and the Nanyang

¹ LinkedIn : <https://www.linkedin.com/in/xavier-bresson-738585b>

Twitter : <https://twitter.com/xbresson>

² UCLA course: <http://www.ipam.ucla.edu/programs/special-events-and-conferences/an-industrial-short-course-on-deep-learning-and-the-latest-ai-algorithms-2019>

Technological University (NTU)³ to professionals with the aim of getting started with industrial Deep Learning technologies and applying them to their projects. Prof. Bresson always received the highest evaluations from the course participants. Prof. Bresson has also taught deep learning at universities to 500+ students at the Master and the PhD levels since 2014. He has supervised 100+ AI projects.

4. COURSE 1 : NON-TECH INTRODUCTION TO AI AND DISCUSSIONS – 1 DAY

OBJECTIVES

- Embrace the context of Deep Learning revolution – The Why?
- How to develop an AI project? Workflow and technical tools
- Why AI is inevitable? What AI can do and cannot do?
- Case studies with IT leading companies

PARTICIPANTS

This course is designed for non-technical industry professionals, that is for anyone who wants to understand AI opportunities and apply AI to industrial projects.

PREREQUISITE

No prerequisite necessary. No technical content.

OUTLINE

FIRST ½ DAY – Non-Tech Introduction to AI

Lecture 1: What is Deep Learning?

Lecture 2: Developing AI projects

Lecture 3: AI for Companies

Lecture 4: Case Studies

SECOND ½ DAY - Discussions

In-depth and brainstorming discussions with managers, engineers and data scientists to apply AI to the company projects.

5. COURSE 2 : FUNDAMENTALS OF DEEP LEARNING – 2 DAYS

OBJECTIVES

- Embrace the context of Deep Learning revolution – The Why?
- Learn the most important neural network architectures
- Assimilate in a clear, solid and intuitive manner Deep Learning algorithms
- Implement neural networks with Python notebooks
- Understand IT leading companies' industrial applications of Deep Learning

³ NTU course: <http://pace.ntu.edu.sg/LifelongLearning/PDP/Pages/CourseDetail.aspx?EventId=e1599a8e-19e3-44ec-97cf-b6bc7759d76a>

PARTICIPANTS

Anyone who wants to get started in Deep Learning, apply Deep Learning to their projects, learn how to code Deep Learning algorithms, and upgrade their skills to the newest AI algorithms.

PREREQUISITE

Basic knowledge of linear algebra (e.g. matrix multiplication) and script programming are needed (the coding will be done in Python). Participants must bring their own laptop.

OUTLINE

Lecture 1: Introduction to Deep Learning

AI/DL case studies, A brief history of AI, 2012 Breakthrough, DL is a universal learning technique, Terminology, Key ingredients, AI is the new electricity, Limitations

Lecture 2: Linear Algebra and PyTorch

Lecture 3: Vanilla Neural Networks - Inference and Learning

Image classification task, Neural networks, Forward pass for inference Understanding linear layers Understanding, The backward pass, Mini-batch learning

Lecture 4: Vanilla Neural Networks - Loss and Optimization

Quality of neural network, Maximizing quality by gradient ascent, Learning rate strategy, Neural network loss, Minimizing loss by gradient descent, Stochastic gradient descent

Lecture 5: Multi-Layer Perceptron - Inference and Learning

Multi-Layer Perceptron, Loss and gradient descent, Backpropagation, PyTorch implementation, Training with epochs, Monitoring loss, Test set

Lecture 6: Multi-Layer Perceptron - Depth

Three representations of neural networks, Importance of non-linearities, Importance of multiple layers

Lecture 7: Convolutional Neural Networks - Introduction

Data structures, Local reception fields, Modeling of hierarchical organization, One-layer convolutional neural network, Paradigm shift in computer vision

Lecture 8: Convolutional Neural Networks - Implementation

Padding, Stride, Multi-dimensional input, Convolution as linear operation, Depth, Convolutional layer with PyTorch, Pooling layer with PyTorch, LeNet5 architecture, Case studies

Lecture 9: Recurrent Neural Networks - Introduction

Introduction to natural language processing, PTB dataset, One-hot encoding, Vanilla RNNs, Training VRNNs, Word prediction, Text generation, Sentiment analysis

Lecture 10: Recurrent Neural Networks - Implementation

Word embedding, Vanilla RNN implementation, Batch of documents, Training loop of VRNNs, Application to NLP, LSTM, Deep RNNs and RNN variants

Lecture 11: Recurrent Neural Networks - Applications

Speech recognition , Machine Translation

6. COURSE 3 : ADVANCED DEEP LEARNING – 2 DAYS

OUTLINE

Lecture 1: Attention Neural Networks

Review of Neural Networks, Neural Networks for Sets, Memory Networks, Transformers, Sequence-To-Sequence Transformers, Language Model Transformers, Graph NNs vs Attention NNs

Note : Attention Neural Networks was a revolution in NLP in 2019

Lecture 2: Graph Neural Networks

Spectral Graph ConvNets, Graph Theory, Spectral Convolution, Graph Pooling
Spectral GCNs, Spatial Graph ConvNets, Spatial Architectures, Operations Research, DGL, Applications

Note : Graph Neural Networks has become the most popular technique in 2020

Lecture 3: Deep Reinforcement Learning

RL and Deep Learning, Agent, Environment and MDP Policy, Value Function, Policy Deep Q-Learning (DQN), Policy Networks, Actor-Critic Algorithm, Planning

Note : Deep Reinforcement Learning was a revolution in 2017

Lecture 4: Good Practices

Generalization, Regularization, Learning rate strategy, Loss and error metric, Hyperparameter search, Initialization and batch normalization, Deep networks, Multi-modal and ensemble technique, Layer placements, Interpretability, Libraries, GPUs