2022 Cambridge AI+ Academic Programme (online)

Machine Learning and Neural Networks

Course Outline

Course Dates
September, 2022

Course Duration
6 Weeks

Course Hours
48 hours in total
It includes: 6 hours of pre-learning, 16 hours of professor live teaching, 8 hours of professor supervisions/office hour, 2 hours of professor evaluation, and 16 hours of group workshop and self-study.

Prerequisites
A background in basic computer science knowledge is required for the course. Programming experience is helpful but not necessary.

Assessment
Assessed individually and in groups through group projects.

Skills Trained
Problem Solving, Teamwork, Presentation, Communication.

Materials Required
Internet connection and laptop for writing, researching and preparing presentations.

Course Description

Neural networks have enjoyed several waves of popularity over the past half century. Each time they become popular, they promise to provide a general purpose artificial intelligence--a computer that can learn to do any task that you could program it to do. The first wave of popularity, in the late 1950s, was crushed by theoreticians who proved serious limitations to the techniques of the time. These limitations were overcome by advances that allowed neural networks to discover internal representations, leading to another wave of enthusiasm in the late 1980s. The second wave died out as more elegant, mathematically principled algorithms were developed (e.g., support-vector machines, Bayesian models). Around 2010, neural nets had a third resurgence. What happened over the past 20 years? Basically, computers got much faster and data sets got much larger, and the algorithms from the 1980s with a few critical tweaks and improvements appear to once again be state of the art, consistently winning competitions in computer vision, speech recognition, and natural language processing. The many accomplishments of the field have helped move research from academic journals into systems that improve our daily lives: apps that identify our friends in photos, automated vision systems that match or outperform humans in large-scale object recognition, phones and home appliances that recognize continuous, natural speech, self-driving cars, and software that translates from any language to any other language.

Neural networks are mentioned regularly in the popular press.Today people tout neural nets as the hot new thing, except that now the field has been rechristened deep learning to emphasize the architecture of neural nets that leads to discovery of task-relevant representations.

Goals & Objectives

By the end of this course, participants will gain detailed knowledge of:

1. A principled statistical foundation to machine learning focusing on what machine learning can do and cannot do, a narrative of the key components that makes us able to learn using data.
2. Current understanding of generalization in neural networks vs classical statistical models.
3. Optimization procedures for neural network models such as stochastic gradient descent and ADAM.
4. Automatic differentiation and at least one software framework (PyTorch, TensorFolow) as well as an overview of other software approaches.
5. Architectures that are deployed to deal with different data types such as images or sequences including a. convolutional networks b. recurrent networks.

In addition participant will gain knowledge of more advanced topics reflecting recent research in machine learning chosen from the following list.

1. Approaches to unsupervised learning including auto-encoders and generative adversarial networks.
2. Techniques for deploying models in low data regimes such as transfer learning and meta-learning.
3. Techniques for propagating uncertainty such as Bayesian neural networks.
4. Deployment of neural network models in hardware systems.

Course Structure

Modules	Date	Contents	Hours
Pre-learning	One Week before	Pre-learning Materials	6
Module 1	Week 1-3	Professors Live Lectures	6
		Professors Supervision	2
Module 2	Week 3-6	Professors Live Lectures	10
		Professors Office Hour	6
Evaluation	Week 6	Presentations	2

Course Topics

- Introduction to Statistical Learning
- Probability Theory
- Linear Models
- Non-Parametric Models
- Computation in a neural net
- Data Augmentation Method
- Reinforcement Learning
- Interpretability
- Assessment

Assessment:

Group presentation and paper assignments

Final Project:

Final project will be a group presentation.

Format

Zoom will be the primary tool to conduct the lectures and supervisions.

Instructor

Module1: Carl Henrik Ek
Associate Professor, Computer Laboratory, University of Cambridge, UK.
Docent in Machine Learning, Royal Institute of Technology, Sweden.

Module2: Pietro Lio'
Professor, Computer Laboratory, University of Cambridge, UK.
Member of the Artificial Intelligence group, Cambridge Centre for AI in Medicine, University of Cambridge, UK.

Reading List

A.Blum, J. Hopcroft and R. Kannan
Foundations of Data Science
https://www.cs.cornell.edu/jeh/book.pdf

G. James, D. Witten, T. Hastie, and R. Tibshirani.
An Introduction to Statistical Learning.
Springer, 2014.
http://www-bcf.usc.edu/~gareth/ISL/

J. Leskovec, A. Rajaraman, J.~D. Ullmann
Mining of Massive Datasets.
3rd edition, Cambridge University Press, 2020.

S. Shalev-Shwartz and S. Ben-David
Understanding Machine Learning: From Theory to Algorithms
Cambridge University Press, 2014.

https://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/understanding-machine-learning-theory-algorithms.pdf