Deep Learning

Yingzheng Li and Bernhard Kainz

This set will be updated shortly before the course

Delivery team 2024

Bernhard Kainz, part I

Yingzheng Li, part II

GTAs:

- TBA

CSLs:

- Harry Coppock

Learning outcomes

- After this course you will know a little bit more about:
  - Feature extraction, convolutions and CNNs
  - Common Network architectures, losses
  - Automatic parameter optimisation

  How it works

  - RNNs, LSTMs, GRUs
  - VAEs and GANs
  - GNNs

  Why it works

  - Deep learning programming frameworks
  - Applications of deep learning

  How to do it

Good to know

70015 Mathematics for ML (recommended)
60012 Introduction to ML (soft prerequisite, please read the basic ML notes if you haven't done this course)
60006 Computer vision
zz70014 ML for imaging
70016 Natural language processing
70028 Reinforcement learning

What we expect you to know already

- Perceptrons, Sigmoid neurons
- Feed-forward neural networks
- Computational graph, Chain rule and backpropagation
- Gradient descent and its variants
- Linear Algebra
- Familiarity with Python programming
- Prior exposure to libraries like TensorFlow or PyTorch would be beneficial

Reference

- Dive into Deep Learning https://d2l.ai/
- I. Goodfellow, Y. Bengio, A. Courville, *Deep learning*. MIT Press, 2016 www.deeplearningbook.org
- Some lectures have been heavily influenced by Material from Michael Bronstein, Kilian Weinberger, Stefanos Zafeiriou, Andreas Maier, Alex Smola, Serena Yeung, Fei-Fei Li

Structure

- Lecture – feedback form previous years integrated; videos + in-person lectures (also recorded)
- Lecture – watch videos, repetition and discussion in lecture
  - in Huxley 308 and Q&A on MS Teams Fri 9-10
  - Post and discuss questions in advance on EdStem please
- Tutorials – Q&A sessions with TAs
  - on Teams and Huxley 308 Fri 10-11
  - Post question in the Lab Queue Channel
- Coursework – hands-on programming exercises: individual with lab queue on MS Teams

Support structure

Lecture

Panopto
https://imperial.cloud.panopto.eu/Panopto/Pages/Sessions/List.aspx#folderID=%2213e9170f-ef01-4d9c-84f6-b02101004540%22

Course website:
http://wp.doc.ic.ac.uk/bkainz/teaching/70010-deep-learning-from-2023/

EdStem:
https://edstem.org/us/courses/29422/discussion/2358988

Coursework

Scientia:
https://scientia.doc.ic.ac.uk/2324/modules/70010/exercises

Tutorials

MS Teams lab queue
https://teams.microsoft.com/l/channel/19%3a59594c3f5432490e88331b3eef942db8%40thread.tacv2/Lab%2520queue?groupId=f5d4c850-1a83-4969-a72f-32bd6767aff8&tenantId=2b897507-ee8c-4575-830b-4f8267c3d307

Lab queue for coursework and tutorial

- Fri 10-11 UK time
- Post your question in the lab queue channel.
- Be specific, we will ignore you if you only post ‘I have a question’
- Reason: you will be picked up by one of the GTAs and invited into their break out room. There you can discuss in detail.
- Having a specific initial question allows other students with the same question to join and we can have a more engaging discussion.

Tutorial

- solve a problem sheet per week (provided on CATe and Scientia)
- not assessed
- help available on the lab queue or in Huxley 308 Fri 10-11 UK time
- solutions provided each following Monday on CATe

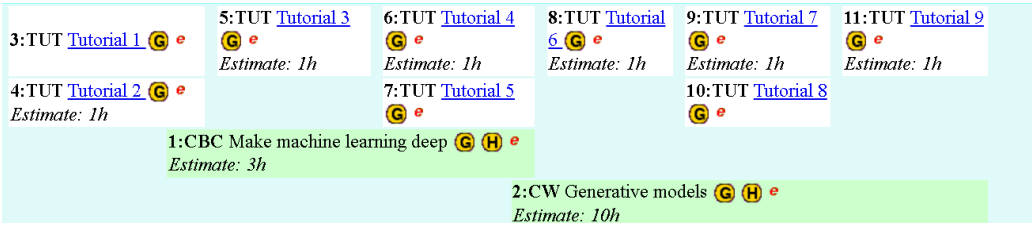

40% 60%

Coursework

- 2 Tasks, both assessed, Task 1: 40%, Task 2: 60%
- solution two weeks after deadline (to accommodate late submissions)
- results as soon as possible but you are many, which means we'll need some time for marking
- Jupyter notebooks
- Recommendation: use
  - https://www.paperspace.com/ -- We have purchased GPU time for you, you will be added to our Team – Harry intro
  - https://colab.research.google.com/
  - Activate GPU support: Edit -> hardware accelerator -> GPU (only if you need one, e.g. CW2 and CW3) – not there are GPU time limits for free tier
  - Computing lab GPUs: ssh into machine, if free train (risky re hard reset)
  - Imperial DoC GPU cluster https://www.imperial.ac.uk/computing/people/csg/guides/hpcomputing/gpucluster/
- Submission on CATe (and via LabTS for coursework 1)

Biggest pain of deep learning is finding the required resources. This is part of the learning and you will experience some frustration using different platforms. This forces you to keep your code universal and to learn a lot of platform engineering

changes to address student feedback from previous years

- more available GPUs – made a deal with paperspace in 2021, google colab has a time-limit on GPU use and does not want to let us pay for your use in bulk, more GPUs in local Imperial infrastructure.
- too much coursework: dropped coursework 3 (RNNs) this term, also too much overlap with NLP course
- more time for coursework 2, which is tempting to play around with it
- overall reduced coursework workload
- Lecture is structured on purpose into to points-of-view: One Engineering-focused view covering fundamental techniques and approaches and one mathematics-focused part, covering theory that can be approximated with deep networks.
- Reworked lecture part with new pre-recordings and structure
- all material available from the start of the course (but might be updated during the course)

Grading

- Assignments (2 assignments): 50%
- Exam 50% (2 questions)

github.com rabbit holes

- https://github.com/alievk/avatarify
- https://github.com/CompVis/stable-diffusion
- https://github.com/deepfakes/faceswap
- https://github.com/Avik-Jain/100-Days-Of-ML-Code
- https://github.com/facebookresearch/Detectron
- https://github.com/fastai/fastai
- https://github.com/CMU-Perceptual-Computing-Lab/openpose
- https://github.com/matterport/Mask_RCNN

There are lots of useful repositories on github.
Of course, deep fakes like the one I used in the beginning have also a lot of potential for mischief
And bear security issues. Don't trust everything you see on the media anymore!
Setups like this require a substantial setup and just loading the model already requires 3 GB of GPU memory. If you want to run this in real-time, then you'll also need a very fast GPU.