

Deep Learning

Yingzheng Li and Bernhard Kainz

Delivery team 2023



Bernhard Kainz, part I



Yingzheng Li, part II

CSLs:

- CW1: Dr Qingjie Meng
- CW2: Harry Coppock
- retiring: Alex Spies



GTAs:

- Alex Spies (afs219)
- Avinash Kori (agk21)
- Edward Stevinson (es1615)
- Konstantinos Barmpas (kb1716)
- Matthew Baugh (mb4617)
- Melanie Bernhardt (mb121)
- Myles Foley (mlf20)
- Osama Hassan (oh1117)
- Weitong Zhang (wz1820)
- Hadrien Reynaud (hjr119)

Learning outcomes

- After this course you will know a little bit more about:
 - Feature extraction, convolutions and CNNs
 - Common Network architectures
 - Automatic parameter optimisation
 - RNNs, LSTMs, GRUs
 - VAEs and GANs
 - GNNs
 - Deep learning programming frameworks
 - Applications of deep learning

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

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 – theory and main concepts: videos. Experimental new format. Feedback welcome but be lenient please.
- Lecture – watch videos, weekly recap
 - 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

Coursework

Lecture

Panopto

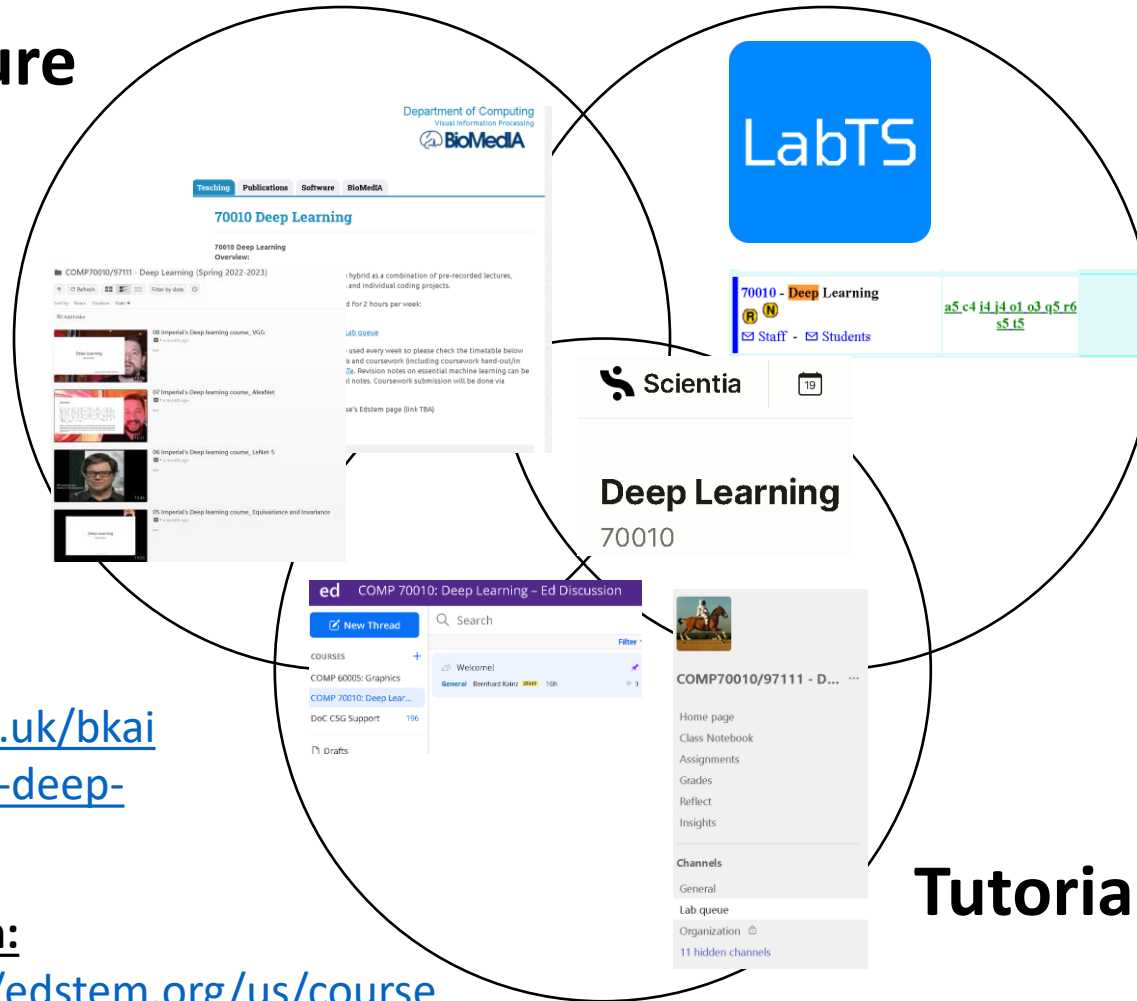
<https://imperial.cloud.panopto.eu/Panopto/Pages/Sessions/List.aspx#folderID=%228631beb3-a0b0-4b5d-b85e-aedd01736de2%22>

Course website:

<http://wp.doc.ic.ac.uk/bkainz/teaching/70010-deep-learning/>

EdStem:

<https://edstem.org/us/course/s/29422/discussion/2358988>



Scientia:

<https://scientia.doc.ic.ac.uk/2223/modules/70010/materials>

MS Teams lab queue

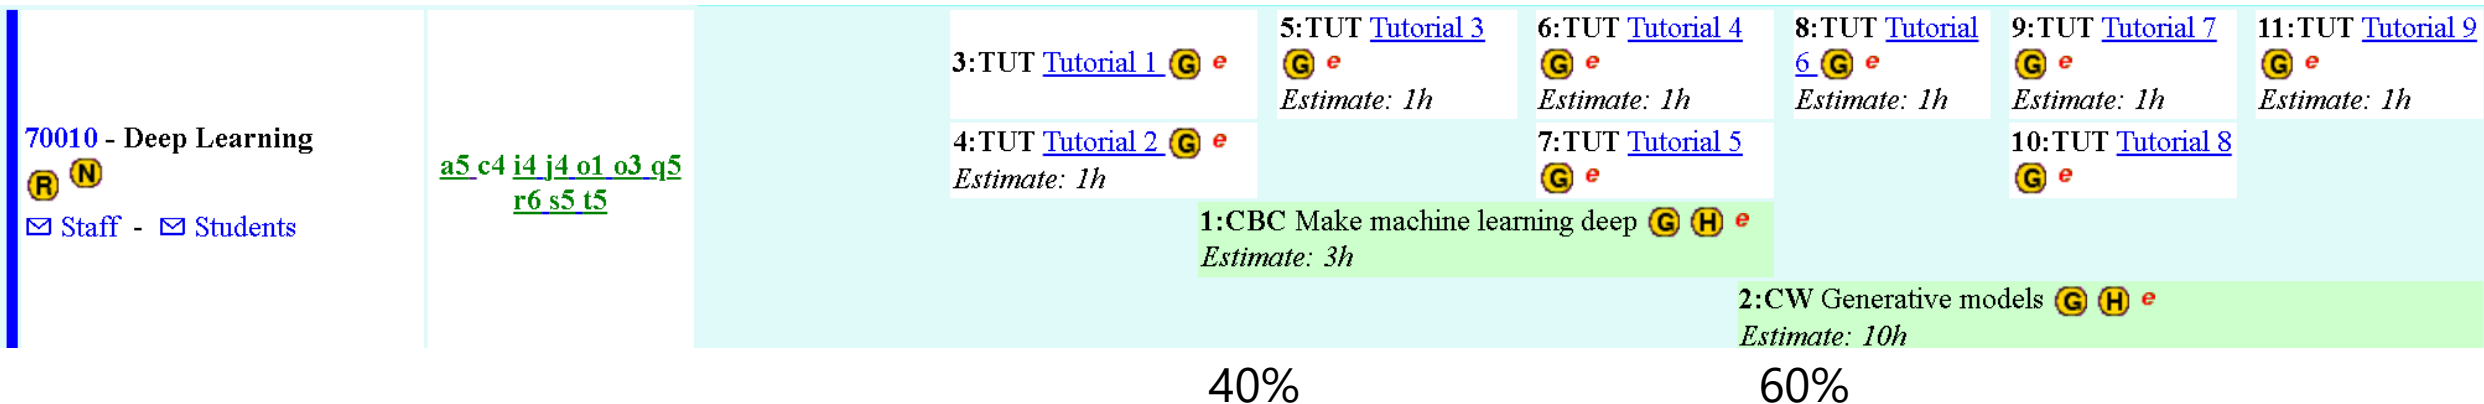
<https://teams.microsoft.com/l/channel/19%3ac0df226af9114ec7b8b9ba6916d4e32a%40thread.tacv2/Lab%2520queue?groupId=646a1ded-f8bf-4e89-8c02-74fc0fa3b9e5&tenantId=2b897507-ee8c-4575-830b-4f8267c3d307>

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



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 sufficient GPU time for you, you will be added to our Team – Alex intro
 - <https://colab.research.google.com/>
 - Activate GPU support: Edit -> hardware accelerator -> GPU (only if you need one, e.g. CW2 and CW3)
- Submission on CATe (and via LabTS for coursework 1)

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.
- 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 by 20-25%
- Lecture is structured on purpose into two 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.
- 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