

Course administration

Welcome to MLPR! I'm your lecturer, Iain Murray. Please just call me Iain, pronounced "ee-an" in emails and in person. Multiple people in the University share my name, so *please* don't select an "Iain Murray" you see in webmail. Use: i.murray@ed.ac.uk

The University of Edinburgh official course descriptor:
<http://www.drps.ed.ac.uk/16-17/dpt/cxinfo11130.htm>

The School administration page: <http://course.inf.ed.ac.uk/mlpr/>

Updates and materials all appear on the course page that I maintain:
<http://www.inf.ed.ac.uk/teaching/courses/mlpr/>
Check there for updates to my office hours, tutorials, and so on.

Machine Learning is growing in importance as a tool for other fields and in industry, and there's a lot of fun stuff in this course. I hope you'll enjoy it. However, as it's a 20 credit course, you're better not taking the course at all than doing badly. This course isn't necessary to apply machine learning, it's building up technical expertise towards being able to research new machine learning methods. If you don't have the mathematical background to do the course, you should (if possible) do the IAML course instead.

Course selection advice

If you haven't taken Introductory Applied Machine Learning (IAML), or a course like it, consider taking that instead of MLPR, providing you are allowed to take the 20 credits at level 10 in your degree. Every year some students take MLPR without the required background (often the maths background) and then fail it. Don't be one of these students! Take a look at the maths and programming self-test and notes on the course website, and ask yourself honestly whether this is material you understand.

If you are an Informatics undergrad student, this course reviews some of the same material as Inf2B learning and IAML, but will be more technical. If you didn't enjoy those courses, I would avoid MLPR! If you did like them, this course should reinforce and then extend that material.

Don't take both IAML and MLPR at the same time. MSc students should get more breadth out of a one-year programme. Undergraduates should space out the material to get full benefit.

The Machine Learning Practical (MLP) is a great course for spending more time on advanced practical skills. It is only a narrow part of machine learning though. Only take MLP if you are taking one of IAML and MLPR, or have already taken one of these or a similar broad machine learning course.

If you've already enrolled in MLPR, don't be afraid to change your course selection. Keep an open mind about whether you should really be taking the course, and don't be embarrassed to change if you find you don't have the required background. The MSc specialism in Machine Learning is *advisory*, no one has to take its core courses.

Notes

You should take your own notes in class. I will mainly work on the board, rather than show slides, so you should be able to keep up.

I provide notes to accompany the lectures. As well as reviewing the material covered, they contain questions to help check your understanding. The lecture notes are new this year. Sadly, despite my best efforts, they will contain some mistakes and unclear parts. *Please* make use of the web-based annotation tool. You can quickly highlight any part of the notes that need fixing or clarifying. Don't be afraid to be picky, I want to fix mistakes of any size (including typos). I am also more than happy to expand on the material where students have engaged with it.

I give pointers to textbooks where reviewing the material from another point of view may be useful. However, only the material we cover in class, the tutorials, and the assignments is examinable. Moreover, some material in the notes is marked either “non-examinable” or “for keen students”, which means that you don’t need to recall facts about this material in the exam. That said, you will be expected to be able to generalize your knowledge to models and machine learning problems that you haven’t seen. The more non-compulsory exercises and reading you have done, the easier the exam will be.

Assessment and study

A large majority of this course (80%) is assessed by the final exam. The aim is to test whether you, *as an individual*, can identify and compare appropriate machine learning methods, derive technical details and explain your choices clearly to others. Few people like exams, but a large fraction of exam assessment seems to be a necessary evil, so that your degree has external credibility.

I am on your side, and so are your fellow students. Talk with us! Most of the preparatory work during the course isn’t assessed, so we can discuss it freely in a non-adversarial way. Moreover, the exam is not “graded to a curve”: if you demonstrate excellent insight into the material, you will get an A, regardless of how everyone else does. As it’s not a competition, work with each other and you’ll all do better. The only exception is that the assessed assignment must be your own work.

It is a common mistake to neglect the “non-assessed” parts of the course: the review questions in the notes, the tutorial work, and formative (non-assessed) assignments. Work through this material as you go so you don’t get lost, so you enjoy the course, and so you pass the exam! If you don’t work through programming examples and theory as you go, the assignment will take longer than intended, and you’re likely to work unhealthy hours. Moreover, if you only ask questions related to the assessed assignment (only worth 20%), you are unlikely to do well in the course as a whole.