

# Proposal: Machine Learning refresh and 20 credits

Iain Murray and Amos Storkey

January 2016

As a result of a comprehensive review of the machine learning specialism courses by the machine learning faculty, and after discussion with other institutes, we propose to overhaul the DRPS entry for Machine Learning and Pattern Recognition (MLPR), move some current content in MLPR into PMR where it fits better, and make both MLPR and Probabilistic Modelling and Reasoning (PMR) be 20 credit courses.

## MLPR: case for making 20 credits

This class has grown very large: ~190 students in 2015/16. Half the MSc year took MLPR, making up 75% of the class. Only 15% were UoE BSc/BEng students. The rest were VUGs and integrated masters. Thus the cohort has diverse background knowledge.

While a minority of students find the course easy, many others struggle with the fast pace of the course. More than 40% of taught MSc students (excluding CDT) taking MLPR failed (mark <50%) in 2014/15! We are explicit to students about the background required, providing example materials and questions. We suggest they take IAML instead and provide quotes from former students telling them not to take the class if they can't do the maths. However, many MSc students come to do machine learning or datascience, and they will take the main MSc machine learning class no matter what we say.

The course is currently fast-paced. Lectures this year attempted to cover variational inference in about one lecture, and Gaussian processes in two lectures (both are advanced topics that could be courses in their own right). The students with weaker maths are obviously left totally behind. However, for the strong students this pace is too fast to put the topics in context, demonstrate applications, or to explore theoretical and practical subtleties.

We surveyed the current MLPR class: ~40% say they already spend 2× or more time on the class than other 10 credit courses. A clear majority (62%) would prefer it to be 20 credits (23% strongly so), either because it's too intensive, or they would like more depth. In contrast only 21% would rather keep the course at 10 credits (7% strongly so). Some value a large course choice: they are free to not do machine learning, but like robotics or NLP to do it in depth if they pick it. Others are worried about the risk of failing 20 credits instead of 10, which is a real concern. While more students should get to a minimum standard with more time, there will still be some that simply shouldn't do this course. It doesn't seem right to base curriculum choices on students who should be taking other options however.

Advantages of a 20 credit course:

- Time to cover topics in better context, with examples. It will then be more reasonable to test students' ability to choose and critically compare methods in the exam.
- Time for practical details, and to have formative practical exercises with feedback before assessment. MLP is great, but more narrow in focus and not everyone takes it. All our machine learning students should see how to run the algorithms they study. DME isn't currently running, where broader applications were previously discussed.
- Machine learning is becoming widely used commercially. The societal effects can be large, and it would be responsible to spend some time considering them. Discussing implications of methods isn't just nice-to-have gilding.

An example case study could be the Netflix prize challenge that was held for movie recommendations. A series of models from the course can be applied (does that solve the problem?) and ensembled (is that commercially worth it?). It turned out that releasing anonymized movie ratings was a privacy disaster: <http://www.wired.com/2009/12/netflix-privacy-lawsuit/>. At the time that was a surprise to me, but shouldn't be to students going into industry now. However, we don't currently have time to have this sort of discussion.

### **PMR: case for making 20 credits**

Probabilistic Modelling and Reasoning (PMR) is intended as an advanced course suitable as preparing for research in probabilistic modelling. It is just as technically demanding as MLPR (if not more so), and is just as fast-paced. It currently has 131 students enrolled for 2015/16. Most of the arguments above for MLPR being a 20 point course also apply to PMR. Students on this course also spend more time on it than many other courses combines. They view it as a critical course but also many struggle with it. The failure rates are comparable to PMR. In addition the review of courses recognised the approximate inference offering does not fit well in MLPR but is ideally suited to be in PMR where it is needed anyway. This additional content is not compatible with a 10 point PMR already at breaking point.

Making this course 20 credits would better reflect the importance and load of this course within the machine learning specialism. Those not wishing to 'spend' 40 credits on machine learning, could take only MLPR, or even only PMR is that was particularly appropriate.

### **Possible Pathways**

We envisage the following possible pathways for a 'full machine learning specialism' and various joint specialisms with other subject. A minor/major split for other specialisms may involve IAML as the machine learning course, IAML and MLP, or MLPR.

**Core ML** MLPR+{2 of MLP, IAML, EXC, TTDS, NC, IQC, IJP}; PMR+DME+RL

**ML+Vison** MLPR+MLP+{IAML, EXC, CG}, PMR+IVR+AV

**ML+Robotics** MLPR+RSS;PMR+IVR+RLSC

**ML+Language, ex1** MLPR+ANLP, PMR+{2 of MT,NLU,TNLP}

**ML+Language, ex 2** MLPR+TTDS+EX, PMR+{2 of MT,NLU,TNLP}

Note that none of the people in these areas have to choose both MLPR and PMR. But apart from the NLP people who might currently choose all of PMR,MT,NLU, and TNLP (which seems quite a heavy load for anyone), most subject combinations are catered for. At the end of the day we are only asking people with a ML focus to dedicate 50% of their course choice to ML. That is not too tall an ask. It allows for concepts of joint-major, major-minor, and single subject specialisms.

### **MLPR: changes to the descriptor**

The current DRPS descriptor is: <http://www.drps.ed.ac.uk/15-16/dpt/cxinfr11073.htm>

The new versions of the course summary and learning outcomes in the companion document don't change the original spirit of the descriptor, but are more concise and reflect an existing change in emphasis. The course description lists high-level topics, within which the details may change as the field evolves. We have a slightly broader list of machine learning principles than before. To compensate, the emphasis on probabilistic reasoning has been reduced, allowing lecturers the flexibility to reduce the amount we cover approximate Bayesian inference and more general probabilistic modelling. Allowing 20 credits for PMR would make it easier to ensure these topics are fully-covered in our MSc.