



Board of Studies

Course Proposal Template

PROPOSED COURSE TITLE: Machine Learning and Pattern Recognition

PROPOSER(S): Amos Storkey, Iain Murray (and the machine learning staff).

DATE: 11 Jan 2016

SUMMARY

This template contains the following sections, which should be prepared roughly in the order in which they appear (to avoid spending too much time on preparation of proposals that are unlikely to be approved):

1. Case for Support

– To be supplied by the proposer and shown to the BoS Academic Secretary prior to preparation of an in-depth course description

1a. Overall contribution to teaching portfolio

1b. Target audience and expected demand

1c. Relation to existing curriculum

1d. Resources

2. Course descriptor

- This is the official course documentation that will be published if the course is approved, ITO and the BoS Academic Secretary can assist in its preparation

3. Course materials

- These should be prepared once the Board meeting at which the proposal will be discussed has been specified

3a. Sample exam question

3b. Sample coursework specification

3c. Sample tutorial/lab sheet question

3d. Any other relevant materials

4. Course management

- This information can be compiled in parallel to the elicitation of comments for section 5.

4a. Course information and publicity

4b. Feedback

4c. Management of teaching delivery

5. Comments

- To be collected by the proposer in good time before the actual BoS meeting and included as received

5a. Year Organiser Comments

5b. Degree Programme Co-Ordinators

5c. BoS Academic Secretary

[Guidance in square brackets below each item. Please also refer to the guidance for new course proposals at <http://www.inf.ed.ac.uk/student-services/committees/board-of-studies/course-proposal-guidelines>. Examples of previous course proposal submissions are available on the past meetings page <http://www.inf.ed.ac.uk/admin/committees/bos/meetings/>.]

SECTION 1 – CASE FOR SUPPORT

[This section should summarise why the new course is needed, how it fits with the existing course portfolio, the curricula of our Degree Programmes, and delivery of teaching for the different years it would affect.]

1a. Overall contribution to teaching portfolio

[Explain what motivates the course proposal, e.g. an emergent or maturing research area, a previous course having become outdated or inappropriate in other ways, novel research activity or newly acquired expertise in the School, offerings of our competitors.]

This is a change of the number of points for the course. We have made a case for this in the cover letter. But in summary, MLPR (along with PMR) is considered to be one of the hardest MSc courses we offer, with students routinely spending more than twice as much time on it than other courses, and *even given that* having failure rates between 25% and 40%. On reflection, students are often very positive about what they learn from the course, but emphasise the exorbitant time cost it takes.

The course content cannot reasonably be reduced – it would be difficult to go beyond IAML given the review that must be done due to the diversity of backgrounds (and some review is beneficial for our own students).

Changing it to 20 points will allow us to provide more contact time, Q&A sessions, review and recap sessions, as well as better reflecting the real time this course takes and the commitment involved.

1b. Target audience and expected demand

[Describe the type of student the course would appeal to in terms of background, level of ability, and interests, and the expected class size for the course based on anticipated demand. A good justification would include some evidence, e.g. by referring to projects in an area, class sizes in similar courses, employer demand for the skills taught in the course, etc.]

The target audience is the same as the current MLPR, but explicitly making it 20 points to ensure that those taking it are properly aware of the time commitment involved.

1c. Relation to existing curriculum

[This section should describe how the proposed course relates to existing courses, programmes, years of study, and specialisms. Every new course should make an important contribution to the delivery of our Degree Programmes, which are described at http://www.drps.ed.ac.uk/15-16/dpt/drps_inf.htm.

Please name the Programmes the course will contribute to, and justify its contribution in relation to courses already available within those programmes. For courses available to MSc students, describe which specialism(s) the course should be listed under (see <http://web.inf.ed.ac.uk/infweb/student-services/ito/students/taught-msc-2015/programme-guide/specialist-areas>), and what its significance for the specialism would be. Comment on the fit of the proposed course with the structure of academic years for which it should be offered. This is described in the Year Guides linked from <http://web.inf.ed.ac.uk/infweb/student-services/ito/students>.]

No change here. This is not a new course. However moving it to 20 points will affect selections. We have provided typical pathways for certain student interest groups to show how the change to this (and PMR) affects choice.

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:

Ex 1: MLPR+ANLP, PMR+{2 of MT,NLU,TNLP}

Ex 2: MLPR+TTDS+EX, PMR+{2 of MT,NLU,TNLP}

Of course there are other combinations (ML+comp neuro, ML+bioinf, etc) which I think we can work out too.

Note 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.

1d. Resources

[While course approvals do not anticipate the School's decision that a course will actually be taught in any given year, it is important to describe what resources would be required if it

were run. Please describe how much lecturing, tutoring, exam preparation and marking effort will be required in steady state, and any additional resources that will be required to set the course up for the first time. Please make sure that you provide estimates relative to class size if there are natural limits to its scalability (e.g. due to equipment or space requirements). Describe the profile of the course team, including lecturer, tutors, markers, and their required background. Where possible, identify a set of specific lecturers who have confirmed that they would either like to teach this course apart from the proposer, or who could teach the course in principle. It is useful to include ideas and suggestions for potential teaching duty re-allocation (e.g. through course sharing, discontinuation of an existing course, voluntary teaching over and above normal teaching duties) to be taken into account when resourcing decisions are made.]

This requires extra 10 points resource of lecturer time. Given the size of this course (~190 this year), that does not seem like a major extra commitment. Note that this reduces the demand on other courses (the total points doesn't change) so there is no actual real change in resourcing per student. No other specific additional resources are required over PMR.

SECTION 2 – COURSE DESCRIPTOR

[This is the official course descriptor that will be published by the University and serves as the authoritative source of information about the course for student via DRPS and PATH. Current course descriptions in the EUCLID Course Catalogue are available at www.euclid.ed.ac.uk under 'DPTs and Courses', searching for courses beginning 'INFR']

2a. Course Title [Name of the course.] :

Machine Learning and Pattern Recognition

2b. SCQF Credit Points:

[The Scottish Credit and Qualifications Framework specifies where each training component provided by educational institutions fits into the national education system. Credit points per course are normally 10 or 20, and a student normally enrolls for 60 credits per semester. For those familiar with the ECTS system, one ECTS credit is equivalent to 2 SCQF credits. See also <http://www.scqf.org.uk/The%20Framework/Credit%20Points.>]

20

SCQF Credit Level:

[These levels correspond to different levels of skills and outcomes, see http://www.sqa.org.uk/files_ccc/SCQF-LevelDescriptors.pdf. At University level, Year 1/2 courses are normally level 8, Year 3 can be level 9 or 10, Year 4 10 or 11, and Year 5/MSc have to be level 11. MSc programmes may permit a small number (up to 30 credits overall) of level 9 or 10 courses.]

11

Normal Year Taken: 1/2/3/4/5/MSc

[While a course may be available for more than one year, this should specify when it is normally taken by a student. "5" here indicates the fifth year of undergraduate Masters programmes such as MInf.]

4

Also available in years: 1/2/3/4/5/MSc

Different options are possible depending on the choice of SCQF Credit Level above: for level 9, you should specify if the course is for 3rd year undergraduates only, or also open to MSc students (default); for level 10, you should specify if the course is available to 3rd year and 4th year undergraduates (default), 4th year undergraduates only, and whether it should be open to MSc students; for level 11, a course can be available to 4th and 5th year undergraduates and MSc students (default), to 5th year undergraduates and MSc students, or to MSc students only]

4/5/MSc

2c. Subject Area and Specialism Classification:

[Any combination of Computer Science, Artificial Intelligence, Software Engineering and/or Cognitive Science as appropriate. For courses available to MSc students, please also specify the relevant MSc specialist area (to be found in the online MSc Year Guide at <http://web.inf.ed.ac.uk/infweb/student-services/ito/students/taught-msc-2015/programme-guide/specialist-areas>), distinguishing between whether the course should be considered as “core” or “optional” for the respective specialist area.]

Artificial Intelligence. Machine Learning Specialism Major. No change from current.

Appropriate/Important for the Following Degree Programmes:

[Please check against programmes from http://www.drps.ed.ac.uk/15-16/dpt/drps_inf.htm to determine any specific programmes for which the course would be relevant (in many cases, information about the Subject Area classification above will be sufficient and specific programmes do not have to be specified). Some courses may be specifically designed for non-Informatics students or with students with a specific profile as a potential audience, please describe this here if appropriate.]

As current MLPR

Timetabling Information:

[Provide details on the semester the course should be offered in, specifying any timetabling constraints to be considered (e.g. overlap of popular combinations, other specialism courses, external courses etc).]

Semester 1.

2d. Summary Course Description:

*[Provide a brief official description of the course, **around 100 words**. This should be worded in a student-friendly way, it is the part of the descriptor a student is most likely to read.]*

Machine learning is a field devoted to developing algorithms that adapt their behaviour to data, providing useful representations of the data and/or predictions. This course covers some fundamental theoretical concepts in machine learning, and common patterns for implementing methods in practice. The intended audience are those wanting the background required to begin research and development of machine learning methods.

This is an advanced course. Students should not choose this class without the required background (see "Other Requirements" box); experience suggests the course moves too fast to pick up the mathematical proficiency on the fly. Instead such students are strongly advised to take a more practical-based course, such as Introduction to Applied Machine Learning (INFR10063).

[Provide an academic description, an outline of the content covered by the course and a description of the learning experience students can expect to get. See guidance notes at: http://www.studentsystems.ed.ac.uk/Staff/Support/User_Guides/CCAM/CCAM_Information_Captured.html#AcademicDescription].

The precise set of methods and algorithms used to illustrate and explore the main concepts will change slightly from year to year. However, the main topic headings are expected to be fairly stable.

Classification and Regression:

Linear Regression, logistic regression, Bayes classifiers

Expanded feature representations:

Basis functions, neural networks, kernel methods

Generalization, regularization and inference:

Penalized cost functions, Bayesian prediction, learning theory

Model selection, pruning and combination:

Cross-validation, Bayesian methods, sparsifying regularizers, ensemble methods.

Representation and metric learning:

dimensionality reduction, clustering, feature learning

To support these topics we will also cover:

Optimization and Inference algorithms:

Stochastic gradient descent, simple Monte Carlo ideas, and more specialized methods as required.

Practical issues:

Formulating problems as machine learning, adapting methods to fit problems.

Numerical and programming issues important for machine learning.

Ethical issues, such as responsible application of methods and privacy concerns.

Pre-Requisite Courses:

[Specify any courses that a student must have taken to be permitted to take this course. Pre-requisites listed in this section can only be waived by special permission from the School's Curriculum Approval Officer, so they should be treated as "must-have". By default, you may assume that any student who will register for the course has taken those courses]

compulsory for the degree for which the course is listed in previous years. Please include the FULL course name and course code].

None

Co-Requisite Courses:

[Specify any courses that should be taken in parallel with the existing course. Note that this leads to a timetabling constraint that should be mentioned elsewhere in the proposal. Please include the FULL course name and course code].

None

[Specify any courses that should not be taken in combination with the proposed course. Please include the FULL course name and course code].

Existing MLPR course:

INFR11073 Machine Learning & Pattern Recognition

Other Requirements:

[Please list any further background students should have, including, for example, mathematical skills, programming ability, experimentation/lab experience, etc. It is important to consider that unless there are formal prerequisites for participation in a course, other Schools can register their students onto our courses, so it is important to be clear in this section. If you want to only permit this by special permission, a statement like "Successful completion of Year X of an Informatics Single or Combined Honours Degree, or equivalent by permission of the School." can be included.]

This course is open to all Informatics students including those on joint degrees. For external students where this course is not listed in your DPT, please seek special permission from the course organiser.

This course requires practical mathematical application of algebra, vectors and matrices, calculus, probability, and problem solving. For example, you will need to be able to differentiate linear algebra expressions with respect to vectors, interpret inner-products and quadratic forms geometrically, and compute expectations of linear algebra expressions under simple distributions. Some of the required details can be learned during the course. However, practical mathematical skills take time to accumulate and a strong mathematical background is essential.

Practical exercises usually require using a particular numerical language such as Matlab or Python+NumPy. We will assume and require sufficient past programming experience that a new package can be learned on the fly.

Available to Visiting Students: Yes

[Provide a justification if the answer is No.]

YES

2e. Summary of Intended Learning Outcomes (MAXIMUM OF 5):

[List the learning outcomes of the course, emphasising what the impact of the course will be on an individual who successfully completes it, rather than the activity that will lead to this outcome. Further guidance is available from

<https://canvas.instructure.com/courses/801386/files/24062695>]

On completion of this course, the student will be able to

- Frame an applied problem as a machine learning task, identifying appropriate methods.
- Critically compare and contrast alternative machine learning methods for a given task.
- Derive and motivate novel variants of machine learning methods.
- Create accessible and useful explanations of the workings and failure modes of machine learning methods.

Assessment Information

[Provide a description of all types of assessment that will be used in the course (e.g. written exam, oral presentation, essay, programming practical, etc) and how each of them will assess the intended learning outcomes listed above. Where coursework involves group work, it is important to remember that every student has to be assessed individually for their contribution to any jointly produced piece of work. Please include any minimum requirements for assessment components e.g. student must pass all individual pieces of assessment as well as course overall].

Coursework 20%, Exam 80%

Assessment Weightings:

Written Examination: 80%

Practical Examination: 0%

Coursework: 20%

Time spend on assignments:

[Weightings up to a 70/30 split between exam and coursework are considered standard, any higher coursework percentage requires a specific justification. The general expectation is that a 10-point course will have an 80/20 split and include the equivalent of one 20-hour coursework assignment (although this can be split into several smaller pieces of coursework. The Practical Examination category should be used for courses with programming exams. You should not expect that during term time a student will have more than 2-4 hours to spend on a single assignment for a course per week. Please note that it is possible, and in many cases desirable, to include formative assignments which are not formally assessed but submitted for feedback, often in combination with peer assessment.]

20 hours worth of assessed coursework, which is a mixture of practical (implementing and using machine learning algorithms) and tests of understanding.

It is expected that students will also spend around 20 hours on formative exercises in the same format.

In 2014/15 the assignment was in 2 parts. Is meant to take ~20 hrs, but realistically takes longer. One of these parts could be sufficient for assessment. With another part delivered less formally, with more formative feedback, and more supporting material.

Academic description:

[A more technical summary of the course aims and contents. May include terminology and technical content that might be more relevant to colleagues and administrators than to students.]

See the main summary and course outline.

*[Provide a more detailed description of the contents of the course, e.g. a list of bullet points roughly corresponding to the topics covered in each individual lecture/tutorial/coursework. The description should **not exceed 500 words** but should be detailed enough to allow a student to have a good idea of what material will be covered in the course. Please keep in mind that this needs to be flexible enough to allow for minor changes from year to year without requiring new course approval each time.]*

See the main summary and course outline.

Relevant QAA Computing Curriculum Sections:

[Please see <http://www.qaa.ac.uk/en/Publications/Documents/Subject-benchmark-statement-Computing.aspx.pdf> to check which section the course fits into.]

Artificial Intelligence

Graduate Attributes, Personal and Professional skills:

[This field should be used to describe the contribution made to the development of a student's personal and professional attributes and skills as a result of studying this course – i.e. the generic and transferable skills beyond the subject of study itself. Reference in particular should be made to SCQF learning characteristics at the correct level http://www.sqa.org.uk/files_ccc/SCQF-LevelDescriptors.pdf].

The student will be able to reason about how to make predictions from and interpret data, an important transferable skill.

In addition the student will be able to

Undertake critical evaluations of a wide range of numerical and graphical data.

Apply critical analysis, evaluation and synthesis to forefront issues, or issues that are informed by forefront developments in the subject/discipline/sector.

Identify, conceptualise and define new and abstract problems and issues.

Develop original and creative responses to problems and issues.

Critically review, consolidate and extend knowledge, skills, practices and thinking in a subject/discipline/sector.

Deal with complex issues and make informed judgements in situations in the absence of complete or consistent data/information.

the students; everything else should be accounted for in the Directed Learning and Independent Learning hours.

The total being 10 x course credits. Assume 10 weeks of lectures slots and 10 weeks of tutorials, though not all of these need to be filled with actual contact hours. As a guideline, if a 10-pt course has 20 lecture slots in principle, around 15 of these should be filled with examinable material; the rest should be used for guest lectures, revision sessions, introductions to assignments, etc. Additional categories of learning and teaching activities are available, a full list can be found at:

http://www.euclid.ed.ac.uk/Staff/Support/User_Guides/CCAM/Teaching_Learning.htm

Lecture Hours: 30 hours

Seminar/Tutorial Hours: 10 hours

Supervise practical/Workshop/Studio hours: hours

Summative assessment hours: 20 hours

Feedback/Feedforward hours: 2 hours

Directed Learning and Independent Learning hours: 138 hours

Total hours: 200 hours

You may also find the guidance on 'Total Contact Teaching Hours' and 'Examination & Assessment Information' at:

http://www.studentsystems.ed.ac.uk/Staff/Support/User_Guides/CCAM/CCAM_Information_Captured.html

Keywords:

[A list of searchable keywords.]

Machine Learning Supervised Learning Probabilistic prediction Data science	
---	--

SECTION 3 - COURSE MATERIALS

3a. Sample exam question(s)

[Sample exam questions with model answers to the individual questions are required for new courses. A justification of the exam format should be provided where the suggested format non-standard. The online list of past exam papers gives an idea of what exam formats are most commonly used and which alternative formats have been

[http://www.inf.ed.ac.uk/teaching/exam_papers/.](http://www.inf.ed.ac.uk/teaching/exam_papers/)]

See previous MLPR exams (not new course)

3b. Sample coursework specification

[Provide a description of a possible assignment with an estimate of effort against each sub-task and a description of marking criteria.]

See previous MLPR coursework (not new course), with split suggested above.

3c. Sample tutorial/lab sheet questions

[Provide a list of tutorial questions and answers and/or samples of lab sheets.]

See previous tutorial/lab sheets (not new course), although these may be extended.

specify a set of concrete readings for the official course descriptor, please list examples here.]

None

SECTION 4 - COURSE MANAGEMENT

4a. Course information and publicity

[Describe what information will be provided at the start of the academic year in which format, how and where the course will be advertised, what materials will be made available online and when they will be finalised. Please note that University and School policies require that all course information is available at the start of the academic year including all teaching materials and lecture slides.]

As currently.

4b. Feedback

[Provide details on feedback arrangements for the course. This includes when and how course feedback is solicited from the class and responded to, what feedback will be provided

on assessment (coursework and exams) within what timeframe, and what opportunities students will be given to respond to feedback.

The University is committed to a baseline of principles regarding feedback that we have to implement at every level, these are described at

http://www.docs.sasg.ed.ac.uk/AcademicServices/Policies/Feedback_Standards_Guiding_Principles.pdf.

Further guidance is available from <http://www.enhancingfeedback.ed.ac.uk/staff.html>.]

The extra time that a 20 point course provides will allow much more room for feedback, including online questions and feedback from those questions within core lecture time. This will enhance the current feedback arrangements. The marked assignment will provide additional feedback. Finally exam like questions will be provided for logical chunks of the course that help students see what they would need to do with what they have learnt in an exam setting. The coursework will be returned before the end of the lecture period allowing within lecture feedback.

ruit,

train, and supervise other course staff, what forms of communication with the class will be used, how required equipment will be procured and maintained. Include information about what support will be required for this from other parties, e.g. colleagues or the Informatics Teaching Organisation.]

As currently

SECTION 5 - COMMENTS

[This section summarises comments received from relevant individuals prior to proposing the course. If you have not discussed this proposal with others please note this].

This course and the complementary proposal for PMR has been discussed at length with the Machine Learning faculty and this proposal results from a regular review of the machine learning syllabus. This has also been discussed with some of the staff in other institutes who teach course most likely to be impacted by this change. It has been discussed with previous PMR/MLPR students from a number of years.

Given the encouragement for more 20 point courses, the clear cost of this course to the students as it currently stands, the need to incorporate some extra material, and the increasing number of students taking MLPR, this change is necessary.

5a. Year Organiser Comments

[Year Organisers are responsible for maintaining the official Year Guides for every year of study, which, among other things, provide guidance on available course choices and specialist areas. The Year Organisers of all years for which the course will be offered should be consulted on the appropriateness and relevance on the course. Issues to consider here include balance of course offerings across semesters, subject areas, and credit levels, timetabling implications, fit into the administrative structures used in delivering that year.]

5b. BoS Academic Secretary

[Any proposal has to be checked by the Secretary of the Board of Studies prior to discussion at the actual Board meeting. This is a placeholder for their comments, mainly on the formal quality of the content provided above.]

