

Machine Learning - Course Outline

COMP 421: 2010 Trimester 1

This document sets out the workload and assessment requirements for COMP 421. It also provides contact information for staff involved in the course. If the contents of this document are altered during the course, you will be advised of the change by an announcement in lectures and/or on the course web site. A printed copy of this document is held in the School Office.

The course

This course looks at contemporary ideas and algorithms in machine learning. Broadly speaking machine learning is the study of how machines can learn from data and use that knowledge to help humans to make better decisions, or perhaps make those decisions themselves. Classic machine learning applications tend to be in robotic control, handwriting and speech recognition, spam filtering, DNA sequence classification, computer vision and so on. The course teaches the essential ideas of machine learning and makes reference to many of these application areas, but also looks at the larger question of how machines might become truly autonomous agents - what would it take, and how might we go about engineering them? The course makes much use of, and teaches, fundamental ideas such as probability calculus and dynamic programming: prior knowledge of these is not essential. It is primarily a lecture-based course, and ranges over a variety of topics based around adaptive behaviour, inference and information, learning to represent the world, and learning appropriate actions from reinforcement.

Preparation and expectations

COMP307 is a prerequisite for this course, but if you have other relevant background it can be waived: see Marcus if you don't have 307 but want to take the course. Simple "starter" code in python will be distributed for some of the algorithms. So previous familiarity with python (or Matlab) might be helpful, but is by no means essential. Python is very easily learned by 4th year computer science students.

Objectives

A solid pass in this course indicates that a student has:

1. a solid understanding of the core issues in machine learning, especially related to pattern recognition;
2. fluency with using probabilities for inference;
3. an understanding of the machine learning algorithms covered in the course, what their underlying assumptions are, and how they scale up;
4. the ability to understand new machine learning algorithms and problems encountered after the course.

Textbook

There is no *required* textbook for this course, but several are *recommended*.

- *Pattern Recognition and Machine Learning*, by Christopher M. Bishop, Springer publishing, 2006 - henceforth "PMLR". Unfortunately this is quite expensive (about \$200). It is an excellent book. It may be slightly cheaper at fishpond than at Vic Books. However it is (only) US\$ 60 from Amazon, so I suggest everyone who doesn't have a copy in week one gets together and we do a bulk order from Amazon.
 - There is one copy in the library and I have requested further copies to be put on closed reserve.
 - Chapter 8 (only) is available as a free PDF, and is attached to the lecture schedule page.
 - *I'm pretty sure I saw the whole book online somewhere that (oddly) didn't seem illegal. If someone else finds it in such a state, let me know.*
 - PRML goes into much more depth than we have time for on many issues, and has little to say about some others. It is not essential to buy it: you can take the course without recourse to the book. But I can guarantee it will enhance your understanding if you do.
- David MacKay's "Information Theory, Inference, and Learning Algorithms" (2003) - henceforth "DJCM". The entire book is freely available online - see under Readings. This is an excellent resource for many sections of the course. Some topics (especially Hopfield / Boltzmann machines for example) are in here but not in Bishop. I will recommend readings from here as we go along too.
- A third possibility is "Elements of Statistical Learning", by Hastie, Tibshirani and Friedman. I have the (legal as far as I can tell) PDF - see under Readings.

The lecture style for this course won't consist of exhaustive notes. Some handouts may contain detail, others may be just summary figures. It's up to students to take down notes and to go over these later re-working the material in their own way. Mostly, we're going to work things out directly on the board. While this is immediate and has a good side, it does mean you don't get to see nicely crafted takeaway descriptions, illustrative figures, lists of key points and so on all crafted into a wholesome whole. The most wholesome whole is the one by you, for you, built up by reflecting on the lecture material.

Lectures, Tutorials, Laboratories, and Practical work

The schedule of lecture topics, readings, and assignment due dates is available online. Lectures for COMP 421 are:

<u>day</u>	<u>time</u>	<u>room</u>
Monday	11am	Cotton 118
Tuesday	11am	Cotton 118
Friday	11am	Cotton 118

All lectures are standard 50 minutes. There are no tutorials scheduled but I am happy to do this on an ad hoc basis if students let me know.

Workload

In order to maintain satisfactory progress in COMP 421, you should plan to spend an average of 10-12 hours per week on this paper. A plausible and approximate breakdown for these hours would be:

- Lectures: 3
- Readings and reworking of the lecture material: 5
- Assignments: 3 (averaged)

School of Engineering and Computer Science

The School office is located on level three of the Cotton Building (Cotton 358). We share this with the School of Maths, Stats and Operations Research (MSOR).

Staff

The course organiser and lecturer for COMP 421 is Marcus Frean:

- Cotton 443
- +64 4 463 5672
- Marcus.Frean@ecs.vuw.ac.nz
- Office hours: Wednesday mornings

Announcements and Communication

Communication outside of lecture will be via the COMP 421 web area at http://ecs.victoria.ac.nz/Courses/COMP421_2010T1/.

Assessment

Your grade for COMP 421 will be determined based on the following assessment weightings:

<u>item</u>	<u>weighting</u>
assignment 1	10%
assignment 2	10%
assignment 3	10%
assignment 4	10%
Final Examination	60%

There will be four assignments, with details provided on the Assignments page. The assignments will delve further into the material covered in lectures, and play two roles: to enhance learning, and to assess the depth of understanding of that material.

One assignment will be a short presentation to the class. in the second half of the trimester. We will discuss topics in

class well before then.

Due dates for assignments will be decided soon after classes start.

Tests and Exams

There's no test - just the final exam. The exam will be similar in form and content to previous exams, apart from changes of emphasis and material as the course has changed over time.

The [\[\[http://www.victoria.ac.nz/timetables/exam-timetable.aspx\]\]](http://www.victoria.ac.nz/timetables/exam-timetable.aspx) [timetable for final examinations] will be available from the University web site and will be posted on a notice board outside the faculty office. The final examination will be three hours long. No computers, electronic calculators or similar device will be allowed in the final examination. Paper non-English to English dictionaries will be permitted. The study and examination period for trimester T1 is Mon the 7th to Wed the 30th June.

Practical Work

Details of assignments, including dates of submission are given on the assignments page.

Policies and penalties for late submission: it's important to get the assignments in on time, so late hand-in will be penalized at 20% per day unless a prior arrangement has been made with the lecturer.

Plagiarism

Working Together and Plagiarism

We encourage you to discuss the principles of the course and assignments with other students, to help and seek help with programming details, problems involving the lab machines. However, any work you hand in must be your own work.

The School policy on Plagiarism (claiming other people's work as your own) is available from the course home page. Please read it. We will penalise anyone we find plagiarising, whether from students currently doing the course, or from other sources. Students who knowingly allow other students to copy their work may also be penalised. If you have had help from someone else (other than a tutor), it is always safe to state the help that you got. For example, if you had help from someone else in writing a component of your code, it is not plagiarism as long as you state (eg, as a comment in the code) who helped you in writing the method.

Passing COMP 421

To pass COMP 421, a student must satisfy mandatory requirements and gain at least a **C** grade overall. The **mandatory requirements** are that you

1. hand in reasonable attempts at all assignments
2. at least 40% in the exam

Withdrawal

The last date for withdrawal from COMP 421 with entitlement to a refund of tuition fees is Friday 12th March, 2010. The last date for withdrawal without being regarded as having failed the course is Friday 14th May, 2010 -- though later withdrawals may be approved by the Dean in special circumstances.

Rules & Policies

Find key dates, explanations of grades and other useful information at <http://www.victoria.ac.nz/home/study>.

Find out about academic progress and restricted enrolment at <http://www.victoria.ac.nz/home/study/academic-progress>.

The University's statutes and policies are available at <http://www.victoria.ac.nz/home/about/policy>, except qualification statutes, which are available via the Calendar webpage at <http://www.victoria.ac.nz/home/study/calendar> (See Section C).

Further information about the University's academic processes can be found on the website of the Assistant Vice-Chancellor (Academic) at <http://www.victoria.ac.nz/home/about/avcacademic>

All students are expected to be familiar with the following regulations and policies, which are available from the school web site:

Grievances

Student and Staff Conduct

Meeting the Needs of Students with Disabilities

Student Support

Academic Integrity and Plagiarism

Dates and Deadlines including Withdrawal dates

School Laboratory Hours and Rules

Printing Allocations

Expectations of Students in ECS courses

The School of Engineering and Computer Science strives to anticipate all problems associated with its courses, laboratories and equipment. We hope you will find that your courses meet your expectations of a quality learning experience.

If you think we have overlooked something or would like to make a suggestion feel free to talk to your course organiser or lecturer.
