Course Outline

The outline from last year is below. This is subject to significant changes. An updated outline will be published soon. In the mean time, please note that:

- No pre-reqs waiver will be granted this year. So, if you haven't done COMP9417, you will not be able to enrol in COMP9418.
- For any admin question regarding enrolments, etc, you will need to email CSE teaching staff. I will not reply to these queries.

Contents

- Course Details
- Course Summary
- Assumed Knowledge
- Student Learning Outcomes
- Teaching Strategies and Rationale
- Student Conduct
- Assessment
- Special Consideration
- Course Schedule
- Resources for Students
- Course Evaluation and Development

Course Details

Course Code	COMP9418
Course Title	Advanced Topics In Statistical Machine Learning
Convenor	Edwin Bonilla (/users/z3503119)
Admin	Edwin Bonilla (/users/z3503119)
Classes	*Lectures:* Wed 12:00-15:00, OMB 230 Timetable for all classes (/COMP9418/17s2/timetable)
Consultations	Thursday 11:00-12:00, K17 Consultation room 403 (level 4)
Units of Credit	6
Course Website	http://cse.unsw.edu.au/~cs9418/index.html (http://cse.unsw.edu.au/~cs9418/index.html)
Handbook Entry	http://www.handbook.unsw.edu.au/postgraduate/courses/current/COMP9418.html (http://www.handbook.unsw.edu.au/postgraduate/courses/current/COMP9418.html)

Course Summary

This course provides an in-depth study of statistical machine learning approaches. The focus will be on methods for learning and inference in structured probabilistic models, with a healthy balance of theory and practice. This course is aimed at students who are willing to be go beyond basic understanding of machine learning. The course provides fundamental support for those willing to intensify their knowledge in the area of big data analytics.

It will cover topics on exact and approximate inference in probabilistic graphical models; learning in structured latent variable models; posterior inference in non-parametric models based on Gaussian processes; and relational learning.

Assumed Knowledge

Official Pre-requisites

Knowledge of machine learning at the level of COMP9417. This is a pre-requisite but, since this is the first time the course is offered, this can be waived subject to LiC's approval. *However, students are not entitled to any consideration if they discover that they do not have sufficient background* .

Mathematics

Solid mathematical background including linear algebra, basic probability theory, multivariate calculus. These courses are a good indication of the knowledge required (note that they are not official pre-requisites):

- Linear Algebra - MATH2501
- Several Variable Calculus - MATH2011
- Theory of Statistics - MATH2801

Programming

We will use Python in the practical part of our bring-your-own-device tutorials. However, it is expected that students unfamiliar with Python can get up to speed if they are able to construct (i.e., design, implement and test) working software in a general-purpose language such as C/C++, Java or Perl, at least to the level of a first-year computing course (e.g., COMP1927 Computing 2 or equivalent). Having a good working knowledge of and be able to construct working software in standard data analysis languages such as Matlab/Octave, or R can also be helpful.

Software Tools

Since an important part of practical machine learning is "data wrangling" (i.e. pre-processing, filtering, cleaning, etc.) of data files, students are expected to master Unix tools such as those taught in COMP2041 Software Construction, or equivalents such as those in R, Matlab/Octave or R.

Student Learning Outcomes

After completing this course, students will:

1. Derive statistical independence assumptions from a given graphical representation of a probabilistic model
2. Understand and implement exact inference methods in graphical models including variable elimination and the junction tree algorithm
3. Derive and implement maximum likelihood learning approaches to latent variable probabilistic models
4. Understand and implement approximate inference algorithms in graphical models including sampling and variational inference
5. Understand and apply basic methods for structured prediction
6. Understand and apply posterior inference and hyperparameter learning in models based on Gaussian process priors
7. Understand and apply basic methods for relational learning

This course contributes to the development of the following graduate capabilities:

Graduate Capability	**Acquired in**
Scholars capable of independent and collaborative enquiry, rigorous in their analysis, critique and reflection, and able to innovate by applying their knowledge and skills to the solution of novel as well as routine problems	Lectures, tutorials, assignments, quizzes and exam
Entrepreneurial leaders capable of initiating and embracing innovation and change, as well as engaging and enabling others to contribute to change	Assignments
Professionals capable of ethical, self-directed practice and independent lifelong learning	Lectures, tutorials and assignments
Global citizens who are culturally adept and capable of respecting diversity and acting in a socially just and responsible way	Lectures, tutorials, assignments and student interactions

Teaching Strategies and Rationale

Machine learning is at the intersection of Artificial Intelligence, Computer Science and Statistics. While the main goal of this course is to go beyond the basics of machine learning as provided by COMP9417 (focused on probabilistic modelling and inference), we will adopt a similar teaching rationale, where theory, algorithms and empirical analysis are all important components of the course. Therefore, the lectures, tutorials and assessments are design to address these components jointly.

The course involves lectures and practical work.

- *Lectures:* Aim to summarise the concepts and present case studies.
- *Tutorials:* Aim to reinforce the topics covered in lectures and will cover theoretical and practical exercises. The practical part of the tutorials will be based on a bring-your-own-device approach, where students will be introduced to the technology required for the assignments and follow a series of programming and data analysis questions. There will be no formal assessment of the tutorials.
- *Quizzes and assignments:* Aim the same as the tutorials at a higher degree of difficulty and will be assessed.
- *Final exam:* There will be a written final exam. This will be centrally timetabled and appear in your UNSW exam timetable.

Engagement Tools and Blended Learning

- All lectures (slides/recordings) will be on the Web

- All tutorial and lab materials (questions before, solutions after) will be on the Web
- All assignments and quizzes will have specifications on the Web and online submission
- Only the final exam will require attendance in person
- We will use Kahoot (https://kahoot.it/#/) as an engagement tool in the lectures
- Forum (https://webcms3.cse.unsw.edu.au/COMP9418/17s2/forums/course) for answering questions using WebCMS3

Student Conduct

*** New and Important ***

The **Student Code of Conduct** (Information (https://student.unsw.edu.au/conduct) , Policy (https://www.gs.unsw.edu.au/policy/documents/studentcodepolicy.pdf)) sets out what the University expects from students as members of the UNSW community. As well as the learning, teaching and research environment, the University aims to provide an environment that enables students to achieve their full potential and to provide an experience consistent with the University's values and guiding principles. A condition of enrolment is that students *inform themselves* of the University's rules and policies affecting them, and conduct themselves accordingly.

In particular, students have the responsibility to observe standards of equity and respect in dealing with every member of the University community. This applies to all activities on UNSW premises and all external activities related to study and research. This includes behaviour in person as well as behaviour on social media, for example Facebook groups set up for the purpose of discussing UNSW courses or course work. Behaviour that is considered in breach of the Student Code Policy as discriminatory, sexually inappropriate, bullying, harassing, invading another's privacy or causing any person to fear for their personal safety is serious misconduct and can lead to severe penalties, including suspension or exclusion from UNSW.

If you have any concerns, you may raise them with your lecturer, or approach the School Ethics Officer (mailto:ethics-officer@cse.unsw.edu.au) , Grievance Officer (mailto:grievance-officer@cse.unsw.edu.au) , or one of the student representatives.

Plagiarism is defined as (https://student.unsw.edu.au/plagiarism) using the words or ideas of others and presenting them as your own. UNSW and CSE treat plagiarism as academic misconduct, which means that it carries penalties as severe as being excluded from further study at UNSW. There are several on-line sources to help you understand what plagiarism is and how it is dealt with at UNSW:

- Plagiarism and Academic Integrity (https://student.unsw.edu.au/plagiarism)
- UNSW Plagiarism Procedure (https://www.gs.unsw.edu.au/policy/documents/plagiarismprocedure.pdf)

Make sure that you read and understand these. Ignorance is not accepted as an excuse for plagiarism. In particular, you are also responsible that your assignment files are not accessible by anyone but you by setting the correct permissions in your CSE directory and code repository, if using. Note also that plagiarism includes paying or asking another person to do a piece of work for you and then submitting it as your own work.

UNSW has an ongoing commitment to fostering a culture of learning informed by academic integrity. All UNSW staff and students have a responsibility to adhere to this principle of academic integrity. Plagiarism undermines academic integrity and is not tolerated at UNSW. Plagiarism at UNSW is defined as using the words or ideas of others and passing them off as your own.

If you haven't done so yet, please take the time to read the full text of

- UNSW's policy regarding academic honesty and plagiarism (https://student.unsw.edu.au/plagiarism)

The pages below describe the policies and procedures in more detail:

- Student Code Policy (https://www.gs.unsw.edu.au/policy/documents/studentcodepolicy.pdf)
- Student Misconduct Procedure (https://www.gs.unsw.edu.au/policy/documents/studentmisconductprocedures.pdf)
- Plagiarism Policy Statement (https://www.gs.unsw.edu.au/policy/documents/plagiarismpolicy.pdf)
- Plagiarism Procedure (https://www.gs.unsw.edu.au/policy/documents/plagiarismprocedure.pdf)

You should also read the following page which describes your rights and responsibilities in the CSE context:

- Essential Advice for CSE Students (https://www.engineering.unsw.edu.au/computer-science-engineering/about-us/organisational-structure/student-services/policies/essential-advice-for-cse-students)

Assessment

Item	Topics	Due	Marks	Contributes to (learning outcomes)
Take-home quizzes				
Q1	Weeks 1-2	Week 3	5%	1-2
Q2	Weeks 3-5	Week 6	5%	3-4
Q3	Weeks 6-8	Week 9	5%	3,5
Q4	Weeks 9-11	Week 12	5%	6
Assignment 1	Weeks 1-3	Week 5	10%	1-3
Assignment 2	Weeks 3-8	Week 11	20%	3-5
Final Exam	All topics	Exam period	50%	1-7

The overall course mark will be the weighted average of the individual components in the table above.

Special Consideration

If your work in this course is affected by unforeseen adverse circumstances, you can apply for Special Consideration (https://student.unsw.edu.au/special-consideration) through MyUNSW (https://iaro.online.unsw.edu.au/special_consideration/home.login#_blank) , including documentation on how your work has been affected. If your request is reasonable and your work has clearly been impacted, then

- for an assignment, you may be granted an extension
- for a take-home quiz, you may be granted an extension
- for the Final Exam, you may be offered a Supplementary Exam

Note the use of the word "may". None of the above is guaranteed. It depends on you making a convincing case that the circumstances have clearly impacted your ability to work.

Course Schedule

Week	Lecture	Tute/Lab	Assignments	Quizzes
1	Intro to probabilistic modelling		-	-
2	Exact Inference in graphical models	Week 1		Q1 released
3	Learning in graphical models	Week 2	A1 released	Q1 due
4	Approximate inference: Variational inference	Week 3	-	-
5	Sampling methods	Week 4	A1 due	Q2 released
6	Continuous latent variables	Week 5	-	Q2 due
7	Markov and Hidden Markov Models	Week 6	A2 released	-
8	Undirected graphical models	Week 7	-	Q3 released
9	Gaussian processes (GP) for regression	Week 8	-	Q3 due
10	GP classification and approximations	Week 9	-	-
11	Variational learning of GP models	Week 10	A2 due	Q4 released
12	Relational learning (guest lecture)	Week 11	-	Q4 due
13	Revision (optional, if enough demand)	Week 12	-	-

Resources for Students

Texts and recommended readings:

Prescribed Resources

- [Book] Machine Learning: A Probabilistic Perspective. (https://www.cs.ubc.ca/~murphyk/MLbook/) Kevin P. Murphy. MIT Press. 2012.

Recommended Resources

- [Book] Bayesian Reasoning and Machine Learning (http://web4.cs.ucl.ac.uk/staff/D.Barber/pmwiki/pmwiki.php?n=Brml.HomePage) . David Barber. Cambridge University Press. 2012.
  - References to online version (2 Feb 2017).
- [Book] Gaussian Processes for Machine Learning (http://www.gaussianprocess.org/gpml/) . Carl Edward Rasmussen and Christopher K. I. Williams. The MIT Press. 2006.
- [Book] Pattern recognition and machine learning (https://www.microsoft.com/en-us/research/people/cmbishop/#prml-book) . Christopher M. Bishop. Springer, 2006.
- [Book] Information Theory, Inference, and Learning Algorithms (http://www.inference.org.uk/mackay/itila/) . David J. C. Mackay. Cambridge University Press. 2003.
- [Book] Probabilistic Graphical Models: Principles and Techniques (http://pgm.stanford.edu/) . Daphne Koller and Nir Friedman. MIT Press. 2009.

Other resources

Other resources will be posted on a weekly basis under Course Work on the course website (http://cse.unsw.edu.au/~cs9418/index.html) .

Course Evaluation and Development

This course has never run before. It will be evaluated at the end of the semester using the myExperience system. However, you are encouraged to provide feedback during the semester so that we can address any problems ASAP.

Resource created 5 years ago (Wednesday 17 May 2017, 04:02:06 PM), last modified 3 years ago (Tuesday 29 May 2018, 01:03:52 PM).

Comments

(/COMP9418/17s2/forums/search?forum_choice=resource/9095) (/COMP9418/17s2/forums/resource/9095)



Omar Ghattas (/users/z5121260) 4 years ago (Wed Sep 27 2017 21:20:24 GMT+1000 (AEST))

we seem to be significantly behind schedule.. will we still have a second assignment and 3 more quizzes?



Edwin Bonilla (/users/z3503119) 4 years ago (Fri Sep 29 2017 00:09:01 GMT+1000 (AEST))

The amount and type of assessment remain. As explained in the first lecture, other details (including release date) have always been subject to changes. However, I do understand that students want to avoid assessment overload at the end of the semester. Just to clarify, there were only 2 quizzes remaining (at the time of your post), one of which has just been released. Assignment 2 will be released in the following few days. You will be given sufficient time to complete the assessments according to their difficulty.



Omar Ghattas (/users/z5121260) 4 years ago (Fri Sep 29 2017 00:30:15 GMT+1000 (AEST))

Thanks for the response, given that the due dates of these assessments will be very close to the final exam, it would be helpful for you to release solutions to these soon, for example, Assn1 theory was quite an involved question and it would be great to see the solution for that..