COMP9414 Artificial Intelligence

Term 2, 2019

Course Outline

Contents

- Course Staff
- Course Details
- Preamble
- Course Aims
- Student Learning Outcomes
- Assumed Knowledge
- Teaching Rationale
- Teaching Strategies
- Assessment
- Student Conduct
- Course Schedule
- Resources for Students
- Course Evaluation and Development

Course Staff

Contact Details

Name	Roles	E-Mail
Wayne Wobcke	Lecturer in charge	w.wobcke@unsw.edu.au
Alfred Krzywicki	Tutor, Admin	alfredk@unsw.edu.au
Armin Chitizadeh	Tutor	a.chitizadeh@unsw.edu.au
Jingying (Caroline) Gao	Tutor	jingying.gao@unsw.edu.au
Anna Trofimova	Tutor	a.trofimova@unsw.edu.au

Course Details

Units of Credit: 6

Web Site: http://www.cse.unsw.edu.au/~cs9414/

Timetable:

Lecture Time	Room	Lecturer
Mon 1.00-4.00	Rex Vowels	Wayne Wobcke

Tutorial Time(**)	Room	Tutor
Mon 4.00-5.00	Quad 1043	Alfred Krzywicki
Mon 5.00-6.00	Quad 1043	Alfred Krzywicki
Tue 12.00-1.00	Quad 1042	Armin Chitizadeh
Tue 1.00-2.00	Quad 1042	Armin Chitizadeh
Tue 3.00-4.00	Quad 1042	Armin Chitizadeh
Tue 4.00-5.00	Quad 1042	Anna Trofimova
Wed 12.00-1.00	Quad 1042	Anna Trofimova
Wed 2.00-3.00	Quad 1042	Anna Trofimova
Wed 3.00-4.00	Quad 1042	Alfred Krzywicki
Wed 4.00-5.00	Quad 1042	Alfred Krzywicki
Thu 2.00-3.00	Quad 1049	Jingying Gao
Thu 3.00-4.00	Quad 1049	Jingying Gao
Thu 4.00-5.00	Quad 1049	Jingying Gao

Note: (**) Tutorials may be cancelled depending on enrolments

Preamble

This course is an introductory postgraduate course on Artificial Intelligence, assuming only a facility with Python programming such as would be acquired in an introductory Python programming course, and some mathematical maturity. 2019 is a transition year in which the course includes a small number of undergraduate students enrolling under the code COMP3411 – lectures, tutorials and assessments are common to the two course codes. After 2019, COMP9021 Principles of Programming will be a prerequisite to COMP9414, and only postgraduate students will be able to enrol. Postgraduate students wishing to enrol in COMP9814 should enrol in Term 1, 2020, when COMP9814 will be the same course as COMP3411.

Course Aims

COMP9414 is an introductory course on Artificial Intelligence covering fundamental topics such as autonomous agents, problem solving, search, logic, knowledge representation, reasoning under uncertainty, natural language processing, machine learning and neural networks. The course provides a foundation for further study in AI such as COMP4418 Knowledge Representation and Reasoning, COMP9417 Machine Learning and Data Mining, COMP9517 Computer Vision, COMP9434 Robotic Software Architecture, and COMP9444 Neural Networks and Deep Learning.

The course is taught with an orientation towards data science and with a view to practical applications of Artificial Intelligence using Python, and focuses on the use of Python toolkits for AI such as NLTK (for natural language processing) and Scikit Learn (for machine learning). Some AI applications that make use of foundational concepts will be demonstrated in lectures.

Student Learning Outcomes

On successfully completing this course, students should:

- Gain in-depth understanding of a variety of fundamental AI techniques
- Understand the foundational issues of AI drawing on philosophical enquiry
- Be able to choose appropriate AI techniques to solve given problems
- Be able to implement programs based on standard AI algorithms
- Acquire experience in utlizing AI Python tookits in realistic contexts
- Be aware of a variety of application areas of current AI technology
- Appreciate the dangers of unthinkingly applying AI in business and industry

This course contributes to the following UNSW graduate attributes.

- *Skills involved in scholarly enquiry.* Students study a range of material that instills an appreciation of how AI draws on Philosophy, Psychology and Linguistics.
- *In-depth engagement with the relevant disciplinary knowledge in its interdisciplinary context.* The course introduces numerous foundational concepts and discusses applications across disciplines in the social sciences, law, business and medicine.
- *Capacity for analytical and critical thinking and for creative problem solving.* Assignments provide students with an opportunity to develop AI solutions.
- *Ability to engage in independent and reflective learning.* Students will be encouraged to experiment with AI concepts in tutorials and assignments.
- *Respect for ethical practice and social responsibility.* The course will consider ramifications of the increasing use of AI and data science in industry.

Assumed Knowledge

Students are assumed to have completed an introductory course in Python programming, such as COMP9021 Principles of Programming. Students are also assumed to have some mathematical ability, though any necessary mathematical concepts will be introduced in the course. Background knowledge of discrete mathematics would be useful.

Teaching Rationale

COMP9414 is an introductory course that provides fundamental knowledge and skills needed for further study in Artificial Intelligence and for the application of AI in industry, particularly machine learning and data science. Where possible, concepts are augmented with practical applications of AI techniques and systems based on widely-used paradigms, tools and platforms.

The Python programming language has been chosen for several reasons: (i) this follows on from the introductory postgraduate programming course, (ii) Python is used in the machine learning and neural networks/deep learning courses, (iii) there are many practical AI toolkits built using Python, (iv) these toolkits are widely used in industry, (v) this enables the course to adopt an emphasis on data analytics/data science, a major growth area in business.

The course will use the textbook by Poole and Mackworth. This book has been chosen because: (i) it is shorter and easier to read than other AI textbooks, (ii) it adopts an approach based on logical reasoning, (iii) the book is freely available online, (iv) the book has recently been updated, and (iv) the book comes with Python implementations of standard algorithms.

Teaching Strategies

The course has a 3 hour lecture and a 1 hour tutorial each week (lectures in Weeks 1–11, with the exception of Week 2 which is the Queen's Birthday holiday), and tutorials in Weeks 2–10). Lectures are **essential** for providing conceptual foundations for the course and preparation and understanding of tutorial exercises. Tutorials are **essential** for understanding specific techniques, programming experience, individual feedback from tutors, and preparation for the exam.

The instruction mode of the course is fairly traditional, however the style places a greater emphasis than usual on practical applications of AI technologies and existing tools (in Python) for AI, data analytics and data science. Lectures cover more conceptual material and give an overview of the main topics and their interconnections. Lectures also provide an essential foundation to apply in tutorials and assignments. Tutorials provide students an opportunity for more in-depth analysis of particular topics and practical help on assignments. As part of tutorials, students will be given self-guided programming activities to aid in understanding core concepts. While not assessed, these lab-style exercises will help with exam preparation. **Attendance is strongly encouraged to gain the maximum value from the course.**

There will also be a course forum where students can post (anonymously to other students if desired) course-related questions that can be answered by staff and other students and viewed by all students. The forum will be moderated by course staff and tutors. For matters of a general nature, students should use the forum in preference to e-mailing their tutor; for personal matters please e-mail the course administrator or lecturer in charge. Note that while forum questions will be endeavoured to be answered in a timely fashion, tutors work on a casual basis and cannot be guaranteed to work outside normal working hours. **So before posting to the forum, please use the search functionality to check if your question has already been asked.**

Assessment

The assessment for this course consists of the following weighted components.

- Assignment 1 (15%)
- Assignment 2 (25%)
- Exam (60%)

The two programming assignments are both auto-marked and marked manually for programming style and critical analysis. For correct auto-marking, it is the student's responsibility to ensure that the submitted code runs on the school environment using the version of Python installed in the labs. The final exam is a 2 hour multiple-choice examination worth 60% covering the major aspects of the course. The final mark for the course is determined by adding together these component marks according to the above weighting to give a result out of 100.

Late submission policy: For Assignments 1 and 2, late submissions are subject to the penalty that the mark obtainable reduces by 20% of the value of the assignment per day or part-day late, up to a maximum of 3 calendar days, after which a mark of 0 is received.

Student Conduct

The **Student Code of Conduct** (Information, Policy) 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, Grievance Officer, or one of the student representatives.

Plagiarism

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 is defined as "using the words or ideas of others and presenting them as your own" but also covers "students providing their work to another student". 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 online resources to help you understand what plagiarism is and how it is dealt with at UNSW:

> Plagiarism and Academic Integrity
> UNSW Plagiarism Procedure

Make sure that you read and understand these. Ignorance is not accepted as an excuse for plagiarism. In particular, you

are also responsible for your assignment files being inaccessible by anyone but you by setting the correct permissions in your CSE directory and code repository. 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. *Copying* (taking ideas and/or text from other students or the Internet and presenting them as your own) and *collusion* (working together on an assignment, or sharing parts of assignment solutions) are forms of plagiarism.

Take the time to read the full UNSW policy on academic honesty and plagiarism. The pages below describe the policies and procedures in more detail:

> Student Code Policy
> Student Misconduct Procedure
> Plagiarism Policy Statement
> Plagiarism Procedure

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

> Essential Advice for CSE Students

As applied to COMP9414, copying or sharing program code for individual assignments counts as plagiarism and is unacceptable. Copying or sharing code from the Internet is unacceptable in any circumstances. Assignment submissions are checked for similarity, both with one other and with code on the Internet. The school maintains a register of students with confirmed plagiarism offences. The penalties for plagiarism range from receiving 0 or negative marks for an assignment, through receiving a mark of 00 FL for the course, to expulsion from UNSW (for repeat offenders). **Students should be aware that the school takes plagiarism very seriously, and that these penalties are applied routinely.** Note that allowing someone else to copy your work counts as plagiarism, and makes you liable to a penalty, even if you can prove that the work was yours originally.

Course Schedule

The following is the rough sequence of lecture topics by week. Broadly speaking, there are four major topics, each with one or more application areas: (i) problem solving and search, (ii) knowledge representation and reasoning, (iii) natural language processing, and (iv) machine learning.

Week	Topic
1	Artificial Intelligence and Agents
2	Problem Solving and Search
3	Constraint Satisfaction Problems
4	Logic and Knowledge Representation
5	Reasoning with Uncertainty
6	Machine Learning
7	Natural Language Processing
8	Knowledge Based Systems
9	Neural Networks and Reinforcement Learning
10	Review

Resources for Students

Textbook

The textbook for this course is freely available online.

> Poole, D.L. & Mackworth, A.K. Artificial Intelligence: Foundations of Computational Agents. Second Edition. Cambridge University Press, Cambridge, 2017.

We will also draw on material from the COMP3411 textbook, which provides much more detail than is needed in this course, but covers some topics especially well, particularly search. It is recommended as a comprehensive reference.

> Russell, S.J. & Norvig, P. *Artificial Intelligence: A Modern Approach.* Third Edition, Pearson Education, Upper Saddle River, NJ, 2010.

References

The course draws on the use of two Python AI packages, namely NLTK (Natural Language Tookit) and Scikit Learn. An early edition of the NLTK book is available as a pdf file and in print; an up-to-date version compatible with Python 3 is available online. The Scikit Learn reference is a recently published text on the use of Python for machine learning. Scikit Learn will be used in COMP9417 Machine Learning and Data Mining.

> Bird, S., Klein, E. & Loper, E. Natural Language Processing with Python. O'Reilly, Sebastopol, CA, 2009.
> The latest version of the book is compatible with Python 3.

Müller, A.C. & Guido, S. Introduction to Machine Learning with Python. O'Reilly, Sebastopol, CA, 2017.

Lecture Recordings

Recordings of lectures will be distributed, however there is no guarantee that the recording system will always work satisfactorily. Note that, due to the general deficiencies of lecture recordings outlined below, recordings can **not** be recommended as a complete substitute for attendance, so students are encouraged to attend lectures in person. It is understood that students enrolling in the "Web Stream" of the course are fully aware of, and willingly accept, these limitations.

There are several deficiencies with lecture recordings that mean recordings cannot capture the full "experience" of lecture attendance, even if the system is working correctly.

- The output of only one projector (the primary projector) is recorded
- Audio and video quality may not be satisfactory at all times
- Use of the whiteboard/blackboard is not recorded
- Copyright protected material is not always able to be recorded (see below)
- Questions from students and class discussions may not be able to be recorded (see below)

The following elaborates on the issues concerning copyright and privacy. First, the copyright exceptions for educational institutions explain what is permissible during a course of research or study for lecturers and students. Note that images are covered by the Statutory Text and Artistic Licence defined in Part VB of the Copyright Act, which allows them to be recorded when presented in a lecture, but web pages, computer software and videos are not, so cannot be recorded if this is not permitted by the copyright holder (e.g. the YouTube standard licence does not allow recordings). **Note that students do not have permission to make a copy of a lecture recording, or to redistribute (send to anyone or post on the Internet) lecture recordings. In addition, students are not permitted to make their own copies (photographs or recordings) of any part of a lecture or tutorial.**

Second, the UNSW privacy policy prohibits disclosure, without consent, of personal information that may be used to identify an individual. This includes voice recordings of students that are captured by the lecture recording system. A notice warning that lectures could be recorded is displayed at the door to any lecture theatre where the recording system is enabled, but the issue of student consent is not mentioned. However, based on a trial of the system in 2017, it is extremely unlikely that the lecture recording system captures sufficient information to enable the identity of a speaker in the audience to be ascertained. Hence lecture recordings will *not* be stopped during class discussions. On the other hand, before or after the lecture or during breaks, when students are speaking to the lecturer close to the console, the recording system will be stopped, as in these situations, the system does record students' voices quite clearly. **If any student has any reservation about being recorded, they should notify the lecturer whenever this concern arises, and lecture recordings will be stopped while the student is speaking.**

Python Resources

The course uses two widely known Python toolkits for Artificial Intelligence: NLTK (Natural Language Toolkit) and Scikit Learn (for machine learning). Recent versions of both packages are included with Anaconda (see below) and the software packages and latest documentation are available online.

> NLTK
> Scikit Learn

This course involves Python programming. Students who have completed COMP9021 can use the resources provided in that course, however the following Python Tutorial should be helpful for reviewing the specifics of Python. Documentation on Python packages will be needed in addition. Note that Python 3 differs substantially from Python 2. **It is your responsibility to ensure that your (auto-marked) assignments run on lab machines under the latest version of Python 3 (currently Python 3.7.2) in the lab Linux environment, before submission.**

> Python Tutorial

An IDE (Interactive Development Environment) for Python may be useful, for example Spyder, which is included with Anaconda. This course will not make use of Jupyter notebooks. The Anaconda distribution also includes NLTK and Scikit Learn.

> Anaconda

Course Evaluation and Development

Computer Science and Engineering courses are evaluated by student survey each time they are taught. The survey includes standard questions asked of all comparable courses so that it is possible to compare a course with other relevant UNSW courses, and also includes space for free-form comments. Survey responses are anonymous. The completed survey forms are analysed statistically by someone independent of the course staff, and the results, including free-form comments, are made available to the lecturer in charge *after* grades have been reported and released.

Since 2013, Artificial Intelligence has been offered as a combined course for undergraduate and postgraduate students. However, the two cohorts have very different backgrounds, leading to difficulty in setting suitable assessments. The combined course has also grown extremely large. Thus the course has now been split into courses tailored for the two cohorts, COMP3411 for undergraduates and COMP9414 for postgraduates (COMP9814 students take COMP3411, which assumes a background in programming equivalent to CSE undergraduate students). With the introduction of Python in COMP9021 Principles of Programming, and the use of Python in COMP9417 Machine Learning and Data Mining, the new

COMP9414 course will be consistent with the postgraduate coursework programme. This year is the first time that this revised course has been offered.

Wayne Wobcke
Last updated: 19 November 2021
CRICOS Provider No. 00098G