University of Otago
ŌTĀKOU WHAKAIHU WAKA

SCHOOL OF COMPUTING

AIML402

Advanced Artificial Intelligence

Semester 2

Paper coordinator

Dr. Lech Szymanski

Office: Owheo 2.49
Email: lech.szymanski@otago.ac.nz

Lectures

2×1h per week

Labs

1×2h per week

Tutorials

3×2h (Week 9-11)

Course Description

In this paper we will look at different definitions of intelligence and at the concept of intelligent agents, concentrating on the issue of how to get information about the world and how to make use of it. We will consider techniques for machine learning and probabilistic reasoning. Almost every human ability results from learning from experience: we will look at how these learning processes can be modelled computationally.

Topics covered include:

- search and optimisation algorithms (including genetic algorithms);
- probabilistic reasoning methods (including Bayesian methods);
- machine learning algorithms (with a focus on neural networks).
- elements of deep learning

Learning Outcomes

By completion of this paper students are expected to:

- Understand fundamental concepts relating to classical Artificial Intelligence methods and recent advances in deep learning.
- Develop appreciation for challenges of developing robust and reliable AI acting in the real world
- Have fundamental understanding of a wide range of AI methods and their applicability for different situations.
- Implement a basic algorithm from each of the different approaches listed above on a toy example, including but not limited to: simple games, basic spam filters, identifying objects in an image.
- Become adept with usage of basic software libraries for general machine learning (i.e. Python + sklearn)
- Implement advanced deep learning models and develop understanding of the fundamental difficulties with analysis of what they learn when they learn a task.
- Become adept with the use of advanced software libraries for deep learning (i.e. Python + Tensorflow)
- Develop awareness of ethical issues related to increasing pervasiveness of AI in everyday life along with the impact and repercussions of use of self-learning/adaptable systems.

Lecture/Lab/Tutorial Schedule

Week	Lecture	Lab	Tutorial
1	Introduction Agents and environments	Robot navigation	
2	Uninformed search Informed search	8-puzzle	
3	Adversarial search Stochastic search	TicTacToe	
4	Uncertainty Bayesian reasoning I		
5	Bayesian reasoning II Learning	Spam filtering	
6	Reinforcement learning I Reinforcement learning II	Grid-world	
7	Classification I Classification II	Digits	
8	Classification III Regression		
9	Optimisation Multilayer perceptron I	Simple regression and classification	Tutorial 1
10	Multilayer perceptron II Multilayer perceptron III	Advanced regression and classification	Tutorial 2
11	Simple recurrent network Unsupervised learning I	Text generation	Tutorial 3
12	Unsupervised learning II Ethics	Data analysis	
13	Review		

Labs

Laboratory exercises provide opportunities for developing practical understanding and implementation of methods that complement the theory covered in the lectures. All the programming will be done in Python. The development environment consist of VSCode running Anaconda virtual environment Python interpreter.

Tutorials

Three tutorials scheduled for Weeks 9,10 and 11 will cover advanced machine learning concepts in the field of deep learning.

Course Readings

Prescribed Course Textbook:

Stuart J. Russel, Peter Norvig, *Artificial intelligence : a modern approach* (2020), 4th edition, Pearson, Boston.

The book provides extended reading for almost every lecture, as well as comprehensive coverage of related AI topics. It is strongly recommended that you purchase or borrow this book for the semester.

Course administration

The paper will be administered through Blackboard.

Course Workload and Expectations

This is a 20-point paper. For your guidance, we offer the following breakdown of hours:

Lectures	26 hours	(2h per week for 13 weeks)
Labs	26 hours	(2h per week for 13 weeks)
Tutorials	6 hours	(2h per week for 3 weeks)
Studies/Reading	60 hours	(5h per week for 12 weeks)
Assignments	122 hours	(for three assignments over 12 weeks)
Total	180 hours	

Students are expected to attend lectures. While handouts summarising the lecture slides will be provided, keeping good notes is recommended. The lectures will be recorded and available on Echo360 afterward, but these recordings will capture limited aspects of the lecture, such as the slides and the voice-over. The main purpose of these recordings is to serve as a supplement and a review of the lectures, and they should be viewed as a replacement only if necessary for those who cannot attend sporadically (due to sickness, etc.).

Students should prepare for labs by studying the class notes and reviewing the lab script. A demonstrator will be present in the lab to assist everyone in completing the exercise.

Assessment

Internal assessment	**Due**	
Assignment 1	Week 5 of the course	**10%**
Assignment 2	Week 9 of the course	**10%**
Assignment 3	Week 12 of the course	**20%**
Exam		**60**%
Total assessment		**100%**

Assignment 1 (10%)

Students will write a program in Python to implement an agent employing a search-based strategy for accomplishing a task, such as playing a simple game. The framework for the environment will be provided. Students will evaluate their agent and assess its effectiveness.

Students will be assessed on:

- the effectiveness of their agent in solving a given problem;
- the quality of the written code;
- the clarity of the technical report explaining methodology and the results.

Assignment 2 (10%)

Students will write a program in Python to implement an agent employing a reinforcement learning strategy for accomplishing a task, such as playing a simple game. The framework for the environment will be provided. Students will evaluate their agent and assess its effectiveness.

Students will be assessed on:

- the effectiveness of their agent in solving a given problem;
- the quality of the written code;
- the clarity of the technical report explaining methodology and the results.

Assignment 3 (20%)

An implementation of deep learning models and/or interpretability methods using deep learning libraries.

Submitting Assignments

All assignments must be submitted electronically via Blackboard.

Assignment Returns

It is University policy that assignments must be marked, graded and returned to students within three weeks from the date the assignment was due.

Students will be informed in lectures when assignments will be ready and when they will be available for viewing on Blackboard. Please note that late assignments and those with approved extensions will be returned three weeks from the date that they were submitted.

Extensions, Late Work and Deductions

In order to obtain an extension, you must see the course coordinator before the due date of the assignment. Except in cases of ill health supported by a medical certificate, students must submit evidence that they have already done some work on the assignment. Any work handed in after the closing date or after the last date of an extension is considered late.

Late work will have 5% (out of 100%) deducted each day (including weekends) after the due date. No assignments without an extension will be accepted more than 5 days (including weekends) after the due date, except in cases of illness or other exceptional circumstances beyond the student's control. In these cases, you must contact your course coordinator as soon as possible with appropriate supporting evidence.

Academic Integrity and Academic Misconduct

Academic integrity means being honest in your studying and assessments. It is the basis for ethical decision-making and behaviour in an academic context. Academic integrity is informed by the values

of honesty, trust, responsibility, fairness, respect and courage. Students are expected to be aware of, and act in accordance with, the University's Academic Integrity Policy.

Academic Misconduct, such as plagiarism or cheating, is a breach of Academic Integrity and is taken very seriously by the University. Types of misconduct include plagiarism, copying, unauthorised collaboration, taking unauthorised material into a test or exam, impersonation, and assisting someone else's misconduct. A more extensive list of the types of academic misconduct and associated processes and penalties is available in the University's Student Academic Misconduct Procedures.

Use of generative software such as ChatGPT is also considered academic misconduct unless specifically specified otherwise in the assignment handout. If generative software is used, students must specify how it was used and on what aspects of the assignment.

It is your responsibility to be aware of and use acceptable academic practices when completing your assessments. To access the information in the Academic Integrity Policy and learn more, please visit the University's Academic Integrity website or ask at the Student Learning Centre or Library. If you have any questions, ask your lecturer.

- Academic Integrity Policy
- Student Academic Misconduct Procedures

Support Services

Class representatives. Volunteers for class representative will be registered at the beginning of this course. The class representative is intended to help facilitate staff-student communication, enabling liaison, consultation and passage of information between teaching staff and the student body. Representatives usually meet with representatives of the academic staff twice a semester for about one hour. A full job description for the class representative is available from the Otago University Student's Association.

University Library. The Library website, https://www.otago.ac.nz/library, gives online access to Library services and resources including databases, past exam papers, referencing guides and the booking system for group study rooms.

Support for international students: https://www.otago.ac.nz/international/support-services.

Disabilities information and support: https://www.otago.ac.nz/disabilities.