COMP9444: Neural Networks and Deep Learning

Week 1a. Overview

Alan Blair

School of Computer Science and Engineering

September 13, 2023

Course Web Pages

- https://www.cse.unsw.edu.au/~cs9444/23T3/
- https://webcms3.cse.unsw.edu.au/COMP9444/23T3/
- https://edstem.org/au/courses/13843/

Lecturers, Course Admins

- Alan Blair (Lecturer-in-Charge)
- Sonit Singh (Lecturer)
- Arun Kumar Marndi (Course Admin)
- Hugh Wilson (Course Admin)
- cs9444@cse.unsw.edu.au

Lecture / Tutorial / Lab Schedule

- Lectures (Weeks 1-5, 7-10)
  - Wednesday 4-6pm (Ainsworth G03)
  - Thursday 2-4pm (Keith Burrows Theatre)
- Tutorials (Weeks 2-5)
- Group Mentoring Sessions (Weeks 6-10)
- CSE Help Sessions (Weeks 2-5, optional, tentative)

Teaching Strategies

- Course materials will be delivered through the course Ed page — including text, images, online discussion forums, quizzes, coding exercises, and links to Tutorial Questions.
- You are encouraged to read through the materials on Ed before each Lecture.
- Lecture time will be used to summarize the material, discuss recent developments, and answer questions.
- Tutorials in Weeks 2 to 5, to discuss worked examples and develop a deeper understanding of fundamental topics.
- Mentoring Sessions (Weeks 6 to 10) to assist with Group Project.

Teaching Strategies

You must keep up with lectures, either by attending in person or watching the recordings. Students enrolled in the Web stream are welcome to attend in person if space is available.

You are expected to:

- review course materials before and after each lecture
- attempt tutorial questions beforehand and be ready to ask questions
- complete quizzes, coding exercises, assignment and group project
- discuss the material with your fellow students if possible
- consider further exploring topics of particular interest
- ask questions and contribute to discussion in online Ed forums

Textbook

The textbook for this course is:

Deep Learning
by Ian Goodfellow, Yoshua Bengio and Aaron Courville
MIT Press, 2014
`http://www.deeplearningbook.org`
`https://mitpress.mit.edu/books/deep-learning`

Another good reference book is:

Understanding Deep Learning
by Simon J.D. Prince
MIT Press, 2023
`https://udlbook.github.io/udlbook/`

Assumed Knowledge

The course will assume knowledge of the following mathematical topics:

- Linear Algebra (2.1-2.8)
- Probability (3.1-3.14)
- Calculus and Chain Rule (6.5.2)

Students should study the relevant sections of the textbook (shown in brackets) and, if necessary, try to revise these topics on their own during the first few weeks of the course.

Planned Topics (Weeks 1-5)

Week 1. a)	Neuroanatomy and Perceptrons	(1.2, 9.10)
Week 1. b)	Multi Layer Perceptrons and Backprop	(4.3)
Week 2. a)	Probability, Generalization & Overfitting	(3.1-14, 5.1-6, 7.11-12)
Week 2. b)	PyTorch	
Week 3. a)	Cross Entropy, Softmax, Weight Decay	(6.1-5)
Week 3. b)	Hidden Unit Dynamics	(8.2-3)
Week 4. a)	Convolutional Networks	(7.9, 9.1-5)
Week 4. b)	Image Processing	(7.4, 8.4, 8.7.1)
Week 5. a)	Recurrent Networks	(10.2)
Week 5. b)	Long Short Term Memory	(10.7, 10.10)

Planned Schedule (Weeks 6-10)

Week 6.	(Flexibility Week)	
Week 7. a)	Word Vectors	(12.4)
Week 7. b)	Language Processing	(10.4)
Week 8. a)	Reinforcement Learning, TD- & Q-Learning	(12.5.1.1)
Week 8. b)	Policy Learning and Deep RL	(18.1, 20.9)
Week 9. a)	Autoencoders and Adversarial Training	(14.1-5, 20.10.3, 10.10.4)
Week 9. b)	Vision and Language Reasoning	
Week 10. a)	Generative Models	
Week 10. b)	Review	

Assessment

Assessment will consist of:

Assignment	20%
Group Project	35%
Written Exam (in-person)	45%

Students are expected to form themselves into groups of 5 for the group project, by the end of Week 4. Each group will be assigned a Mentor.

Assignment

The assignment may involve, for example:

- using code written in pytorch
- writing your own code
- running experiments and analysing the results

Further details will be provided on the course Web site.

Pytorch

These are the versions of modules currently installed on the CSE lab machines. Please try to install equal or later versions on your own laptop.

python3	3.9.3
torch	1.7.0
numpy	1.19.5
sklearn	0.23.2

Plagiarism

- Plagiarism is taken seriously by UNSW/CSE and treated as Academic Misconduct. ALL work submitted for assessment must be your own work.
- For an individual assignment, collaborative work in the form of “think tanking” is encouraged, but students are not allowed to derive code together as a group during such discussions. In the case of a group assignment, code must not be obtained from outside the group.
- Plagiarism detection software may be used on submitted work.
- Academic Integrity and Plagiarism:
  `https://student.unsw.edu.au/plagiarism`

Related Courses

- COMP3411/9414 Artificial Intelligence
- COMP9417 Machine Learning and Data Mining
- COMP9418 Advanced Topics in Statistical Machine Learning
- COMP4418 Knowledge Representation and Reasoning
- COMP9491 Applied Artificial Intelligence
- COMP9517 Machine Vision
- COMP3431 Robotic Software Architecture
- COMP9727 Recommender Systems
- COMP6713 Natural Language Processing
- 4th Year Thesis topics