

COURSE REPORT


Course	CS2109S - INTRODUCTION TO AI AND MACHINE LEARNING
Academic Year/Sem	2023/2024 - Sem 1
Department	COMPUTER SCIENCE
Faculty	SCHOOL OF COMPUTING

Note: Class Size = Invited; Response Size = Responded; Response Rate = Response Ratio

Raters	Student
Responded	251
Invited	391
Response Ratio	64%


1. Overall opinion of the course

Distribution of Responses




Rating Scores

Question	Dept Avg (COMPUTER SCIENCE)		Fac Avg (SCHOOL OF COMPUTING)		Dept Avg by Activity & Level (COMPUTER SCIENCE-LECTURE (Level 2000))		Fac Avg by Activity & Level (SCHOOL OF COMPUTING-LECTURE (Level 2000))		Course Average (2310-CS2109S-L)	
	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation
What is your overall opinion of the course?	3.9	0.9	3.9	0.9	3.9	0.9	3.9	0.9	3.7	0.9




2. Expected Grade

Distribution of Responses




Rating Scores

Question	Dept Avg (COMPUTER SCIENCE)		Fac Avg (SCHOOL OF COMPUTING)		Dept Avg by Activity & Level (COMPUTER SCIENCE-LECTURE (Level 2000))		Fac Avg by Activity & Level (SCHOOL OF COMPUTING-LECTURE (Level 2000))		Course Average (2310-CS2109S-L)	
	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation
The grade that I am most likely to get in the course is:	4.1	0.8	4.1	0.8	4.2	0.7	4.2	0.7	4.1	0.7




3. Difficulty Level of the course

Distribution of Responses



Rating Scores

Question	Dept Avg (COMPUTER SCIENCE)		Fac Avg (SCHOOL OF COMPUTING)		Dept Avg by Activity & Level (COMPUTER SCIENCE-LECTURE (Level 2000))		Fac Avg by Activity & Level (SCHOOL OF COMPUTING-LECTURE (Level 2000))		Course Average (2310-CS2109S-L)	
	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation
I rate this course as:	3.9	0.8	3.8	0.8	3.9	0.8	3.8	0.8	4.0	0.7



WHAT I LIKE / DISLIKE ABOUT THE COURSE

What I liked about the course:

Comments
a lot of work
This course provides a good opportunity for students to learn and understand how AI works in this era, in which AI is prevalent with the rise of ChatGPT and various generative AI.
NA
It introduces to me what AI and ML is all about, which is what I was trying to find out in the first place.
The content
I like all the content, especially GPT, very interesting

Comments

One of the standout features of CS2109S is the structure and content of the weekly problem sets. These assignments are exceptionally instructive, playing a crucial role in guiding us through the various concepts of machine learning. The problem sets are thoughtfully designed to reinforce the week's learning objectives, helping to solidify our understanding of each topic. This regular practice is not only beneficial for grasping the theoretical aspects but also invaluable in applying these concepts to practical scenarios.

Additionally, the course goes beyond just teaching the algorithms used in machine learning; it delves into the underlying deductions and concepts that form the basis of these algorithms. This deeper exploration into the 'why' and 'how' behind the algorithms provides a much more comprehensive understanding of machine learning. It's this emphasis on the foundational principles and logic that sets this course apart, as it ensures that we are not merely learning to apply algorithms but are also gaining a fundamental understanding of the principles that drive these machine learning techniques.

NA

Interesting to learn some AI/ML

- The ML and AI concepts!
- sufficient math justification

Problem sets are very practical

It provides a nice breadth in the field of AI/ML, spanning several different topics that are fascinating in their own way. I also like that problem sets, when introducing new libraries or concepts, often have a portion teaching the syntax or demonstrating an example first, so we are not entirely lost and stuck. This has helped in our learning a lot.

The implementation aspect (i.e, problem sets) give a good intro to utilising ML and AI packages

Concepts taught were interesting and ties closely with algorithms.

Teaches basic concepts about AI and ML well, while showing some of the mathematics that will be relevant in future modules

Makes me know more about AI and decided to pursue my focus area in AI.

Getting some practice with search algorithms helps with leetcode.

The problem sets were really fun.

Problem Sets do a good job of helping me understand the practical side of the coding.

I liked learning about machine learning.

Useful to learn about concepts.

fun psets

Interesting concepts

na

Great implementation of algorithm.

Made me think outside the box

NIL

Gamified system

I like that it delves into the practical aspects and answered many of the questions I had about machine learning. It is a good introductory course to ML. I also now have more hands on experience with the machine learning libraries.

none

First half of content

Its focus on implementation.

Interesting

Introducing to various machine learning and AI models

covers both AI and machine learning, a good introduction to this field

Intorudciton to AI

Insightful.

Problem sets are quite practical

nil

The topics are very useful in modern days, and the decision to introduce both Machine Learning and AI is a very good one.

it has thoughts me a lot

Very interesting introduction to AI

Comments
nil
It actually revealed the secrets behind how AI and ML works in general.
The knowledge is very useful
The prof.
Assignments walk us through implementing AI concepts from scratch – very useful, and they use real tools from industry like numpy, sklearn, PyTorch
The slides are very detailed, the lectures are good. The problem sets given help in fully understanding the lecture content.
Learning about different models.
nil
very informational
I learnt a lot of new AI ML algorithms
learn cool things
Interesting learning about AI and ML concepts
The basic content of the AI. and the challenge question in Coursemology
Interesting
Practical aspects of machine learning techniques were covered
Learn very relevant concepts
Content is interesting. The components of the course, especially the problem sets, are well–designed and well–structured.
Problem sets were quite challenging and interesting. First course that actually taught me "library abuse" which IMO is needed for SWE/other tech–related careers.
AI is interesting and is the boom right now.
Content is very interesting
Very interesting while challenging
One of the better intro modules for CS. Good inclusion of the different topics and concepts. Good structure of problem set.
Many different areas of AI and ML covered which have enlightened my previous confusions about them. Plenty of hands–on problem sets to try with powerful ML libraries which is a good introduction to this new coding experience for me.
Learning about the AI fundamentals is interesting.
I enjoy solving the problem sets. It helps to tie lecture concepts to "hands–on" knowledge.
Learn how Machine learning and Deep learning works
Implementing AI/ machine learning
fun, teaches key ai & ml concepts
Problem sets is rather hands–on.
The concepts we learned about AI like the search algorithms were quite interesting. Also appreciate that finals is a take home assignment.
Weekly prompts (problem sets) to keep up with our knowledge on the course.
NIL
I loved how the midterm was conducted, i.e. in–depth testing of key AI concepts without the need for me to "spot" which topics will/will not be tested. I hope this format remains and it doesn't become just MCQs.
The game–like experience is fun and interesting.
Weekly PSETs
Math not as central to the module as I was expecting, the concepts taught were presented with intuitive examples most of the time.
Got to learn many new things about AI/ML that help understand current technology.
learning about basics of AI stress–free CA component that promotes learning instead of worrying about grades
pacing was decent
i love ai!! the algorithms
A bit of introduction to AI and ML.
NA

Comments
Covers and introduces basic AI concepts extensively.
interesting
The hands-on practices were good, and contrived enough that they helped me to understand the course content while not being a "just throw some random model at the problem and pray" sort of thing that courses from elsewhere might do. The lectures were also fairly digestible, if not a little too heavy on the math side. But that's just a personal skill issue, and I believe that anyone who has stronger foundations in math would benefit greatly from understanding the inner workings of the system.
Have always been interested in ML, this course taught me a lot
Good introduction to machine learning with a wide variety of concepts taught. I think there's a good balance between math and practical concepts.
The mathematics for the second half of the semester is not exactly complicated, but is very tedious. I like that this isn't examined (so you don't feel the need to practice doing partial differentials) but we still get the chance to be exposed to it through tutorials.
Informative mod
Very interesting course in AI and machine learning. I liked that problem sets provide real-life examples.
It makes me start to step into the realm of AI and Machine Learning, which are my interest.

What I did not like about the course:

Comments
its AI
Math. Hunger games for finals. Math.
The final assessment is kinda brutal. You will potentially fail this course if your final assessment is terrible.
The workload was quite high for a 2K mod with submissions almost every week.
Due to my lack of interest in this field, the stuff that were taught are just not interesting to me. Furthermore, the concepts are way too difficult.
Workload
I dont like the lack of mathematical proofs for some parts. The amount of mistakes in the lectures are too much as well.
The format of the Finals is vey messed up. I am particularly angry because we should be given the freedom to clean data, build machine learning model in our own manner, instead of following a fixed given template. I didnt realise that we MUST use nn.Sequential to build our networks after DOZENS of failure to incorporate a class-based InternalModel(nn.Module) into the Model template. Either we should be notified that we MUST write our model in this manner, or Coursemology could specify some hints. But there is NONE. It is just rejecting the model, saying there is syntax error. This is very silly and very energy-consuming. I do not even want to mention that I wait for 5 hours until my code got run on Coursemology.
Also, the course is good that it covers a lot of underlying stuff, but would also be nice if it can cover some more practical stuff, such as the different Torch models (things like RNN, LSTM, Attention are just mentioned but never taught). Also, there is little teaching about data preprocessing, which could almost be equally important.
AI ML is a niche part of computer science, this module should not be a core module.
Workload.
Too much maths which is unavoidable in an AI/ML course, hard to grasp concepts without having a strong foundation in maths
Finals is take home assessment, which should have been a project instead. There should also be a finals paper instead of having midterms as a differentiating factor. Would be better if our midterm scripts were marked faster as well.
The prof
Anticipation (since the final practical test has not yet been released): The second part of the course (after the midterm) seems to be only loosely related to the final grading rubrics, which might undermine the effort put into learning all of those challenging concepts.
Math concepts can feel irrelevant sometimes
The final exam sounds like it will be quite random and luck-based. I don't see how it can be a good gauge of our competency. Also, after midterms, it feels like content difficulty has skyrocketed since things are no longer officially examinable in the traditional sense, so we don't really have to know all this theory.
Math heavy at some points
There was too little time for mid terms, especially for the first part of the exam paper.

Comments
Midterm weightage too high, doesn't really give room to get back after screwing up.
Too many overpowered peers.
The math covered in the lectures seem like an afterthought. I understand this is supposed to be an ML course but with attempts at not boring students with Math, but in the end I feel like I come out of ML lectures more confused than before.
Content can be better updated for more modern ML. First few topics felt pointless, and better suited in a algorithms course, for example.
There's too much math in the lectures that aren't really well-explained other than half the stuff being "magic", so it's really hard to follow along sometimes.
None so far
Nil.
hard
all the maths
Very Difficult
na
The lectures taught are really really vague. I do not understand the concept clearly sometimes.
Too tedious
NIL
Too content heavy! Too high workload! Competitive environment...
NIL
it being a core module
All the math heavy portions
The midterms had way too many questions. The course course tries to cover way too many things at once, but I guess that's to be expected for an introductory mod. The psets are great, but given the quantity of contents coming in week after week, they don't really give enough time to allow for acclimatization to the topics.
Very difficult
The errata made in almost all the slides, final is a take-home exam that is a time-sink
NIL
too much math
Challenging.
I feel the tutorials are not that helpful in learning as it tries to teach too many "unnecessary math" and some tutors don't care much about the tutorials.
too theoretical
Lecture topics are not interesting and have too many formulas and not so many visualizations.
it was difficult
Certain psets private test cases are not clear
nil
Finals format
The mathematics are very vague and the course has focused relatively more on the practical aspects.
the exams is too bad, midterm do not offer enough time, questions involves too many variation of different kind.
Cohort size too big taking very long to mark the midterms. Maybe can provide some help for the professor.
Too hard exam; Too hard homework
Midterm had questions that were too hand-compute heavy
The midterms test time was very short.
NIL. overall, pleasant experience
nil
na
Some difficult content in lectures isn't explained clearly. Could modify the content to be more organized and acceptable.
midterm exam had very subjective questions and answers (search formulation)

Comments
very tough :(
The amount of corrections at the start of each lecture
Too abstract, too much math
Too much workload.
A lot of content
There was more than necessary mathematical rigour for a 2k level module. Furthermore, the vast majority of it is not tested. To begin with, I feel that this module in its current form should not be a graduation requirement.
The course material (those involving math) usually lacks context and skips quite a bit of steps. For students who preferred more rigour, it's difficult to understand the slides. At the same time it is still too "mathy" for the others. Would be better if the materials included more math (including the notations and the intermediate steps) but indicates which parts are extra material.
The module administration can be better. Many things (e.g. problem sets / recorded lectures / practice exam / graded midterms) are released late without any signs. I would prefer receiving a realistic prediction on when something is ready, instead of delaying it bit by bit (e.g. the complexity of the midterm grading should be clear weeks ago).
The finals feel like another whole new course. It might be better if some resources on data science in Python can be provided during the semester (before reading week). It would be best if related optional exercises are available on Coursemology (on top of the CS1010 level practices).
MA2001/MA1522 Linear Algebra should be a prerequisite for this course (although most CS students should have taken it by this stage of their degree, it is not good to assume that all have taken it and/or are CS students). The linear algebra covered in the second half of the course is nontrivial and students should be expected to have at least a rudimentary knowledge of matrix operations, covariance, transposes, bases and diagonalisation before taking CS2109S.
Mathematical rigour.
Breadth >>>> depth
Too much problem sets
NIL
A lot of math. The lectures are too fast and shallow when they touch some difficult math calculations. Therefore, a lot of self-study is needed.
Many of the mathematical concepts covered in lecture are very in-depth, difficult, but covered within very short time. I think this creates unnecessary confusion, and would prefer if the lecture content can either allocate more time to explain the mathematical details clearly, or simply omit the introduction of these mathematical details and focus more on applications.
it is so heavy maths and algorithm content, and it is very difficult if i cannot visualise faster. but i hope i can explore more in this track
Very high weekly commitment
All the math content such as derivatives are incredibly boring and dry.
Mid Term Timing is too short, and make into multiple quizzes in lecture. Lecture could be 3 hours instead of 2 hours
Final assessment is unfortunately very "what is the mass of the sun"
Mathematical concepts are brushed over
Content of the modules seems to cater to those that are already interested in AI focus area. Tutorial are too math heavy at times , hard to really understand what is supposed to be learnt from the tutorial. For a core module. the content could be scaled down.
The last few weeks of ML lectures had too much math. While I agree it is good to know how differentiation makes gradient descent/back propagation work, at the end of the day, we are using libraries such as PyTorch to help us do the differentiation. I don't see the point in us spending 1 hour talking about how differentiation works, and then having to spend hours on tutorials to do differentiation ourselves. And even more so if finals do not require us to do a lot of (if any) differentiation.
N.A
Midterm format + Improve speed and accuracy of midterm. marking
The handwavy Math in the ML, and maybe how long it took for the midterms to be graded (but I understand that its complex and I value having such comprehensive midterms over trivia like MCQs)
None.
Lack of resources in terms of the practical aspects (e.g. more tutorials on Pytorch/Sklearn)
The differentiation and linear algebra (not taken MA1522 yet)
Feel a bit lost at times about course content because we go through many things, some of which are not important to the course (e.g. don't need to delve to deep into Automatic Differentiation, etc.)

Comments

There is a lot of content and it gets very confusing, especially the later chapters like SVM and backpropagation

finals being set on reading week when reading week is suppose to be for students to prepare for exams.

finals being set on such a weird timing (8pm to 11:59pm the next day), especially bad for people who have finals in the morning the next day

mid-terms weightage for entire score component really high (at 35%, usually its 20% for other modules).

content felt quite rushed, did not know what was relevant and what was trivial math.

very hard

The machine learning part of the course can be better structured. The content should either be further simplified to cover more algorithms or cover fewer algorithms but in greater mathematical detail.

How it is run

Heavy and high workload.

workload

This was somewhat acknowledged by Professor Rizki himself, but the whole chapter on Support Vector Machines was a little too confusing for a lot of students. But if it's an important concept, perhaps supplementary readings/resources can be provided and more time can be devoted to helping students understand the intuition or even rigour of the system to better grasp the concept. Also, if Professor Ben Leong can have Ben Bitdiddle as a name in the Midterm exams, I'm sure Professor Rizki can come up with a name for himself too, like Rizzmaster or something that would be pretty cool.

Bulk of skills required in finals was not taught prior

NIL

I feel an extra problem set (or replacement for some other problem set) on data processing would be very helpful to prepare for the finals.

Too hard the math part really hard to understand requires more explanation

The midterm is a bit difficult and the time is too short. Not sure if the final format is a fair way of assessing a student's knowledge.

The confusion i got from the lecture.

Some of the tutorial content seems unnecessary and some are too difficult.