craffel Add quercus link 0e552ce · 3 days ago

94 lines (70 loc) · 7.55 KB

Preview Code Blame Raw

CSC413/2516: Neural Networks and Deep Learning

Instructor: Colin Raffel

TAs: Gül Sena Altıntaş (lead), Gyung Hyun Je, Jonah Strimas-Mackey, Malikeh Ehghaghi, Nikita Dhawan, Boglarka Ecsedi, Mustafa Ammous

Quercus: https://q.utoronto.ca/courses/396034/

Piazza (for all course communications): Link available on Quercus

Term: Fall 2025

Lectures: Thursdays 9-11am in BA 1170 (Bahen) and 1pm-3pm in AB107 (Astronomy and Astrophysics)

Tutorials: Tuesdays 10-11am in GB244 (Galbraith) and Wednesdays 11am-12pm in AB107 (Astronomy and Astrophysics) (additional office hours and rooms to be announced)

Deep learning is the branch of machine learning focused on training neural networks. Neural networks have proven to be powerful across a wide range of domains and tasks, including computer vision, natural language processing, speech recognition, and beyond. The success of these models is partially thanks to the fact that their performance tends to improve as more and more data is used to train them. Further, there have been many advances over the past few decades that have made it easier to attain good performance when using neural networks. In this course, we will provide a thorough introduction to the field of deep learning. We will cover the basics of building and optimizing neural networks in addition to specifics of different model architectures and training schemes. The course will cover portions of the "Dive into Deep Learning" textbook.

Prerequisites

Students should have taken courses on machine learning, linear algebra, and multivariate calculus. Further, it is recommended that students have some basic familiarity with statistical concepts. Finally, students must be proficient in reading and writing Python code. For a list of courses that serve as prerequisites for the undergraduate version of this course, see here.

Lecture and tutorial schedule

The following schedule is tentative; the content of each lecture may change depending on pacing. All readings refer to corresponding sections in Dive into Deep Learning. Because the book is occasionally updated, the sections listed may become out-of-date. If a reading seems incongruous with the topic of the lecture, please let me know and I will check if the sections changed. Homeworks assigned weekly and are nominally "due" on Thursdays at 11:59pm but there will be **no late penalty**, so homeworks can be handed as late as 11:59pm on the last day of the semester (December 4th, 2025). After 11:59pm on December 4th, 2025, no work will be accepted.

Date	Subject	Reading	Homework
9/4	Class introduction, linear & logistic regression	2.1-2.7 (optional), 3.1-3.5, 4.1-4.5	H1 assigned
9/11	Under/overfitting, regularization, multilayer perceptrons	3.6-3.7, 5.1-5.2, 5.6	H1 due, H2 assigned

Date	Subject	Reading	Homework
9/18	Backpropagation, numerical stability, automatic differentiation	5.3-5.4	H2 due, H3 assigned
9/25	Gradient descent, adaptive gradient methods	12.1-12.6, 12.10	H3 due, H4 assigned
10/2	Convolutional neural networks, batch/layer normalization, residual connections	7.1-7.5, 8.5-8.6	H4 due, H5 assigned
10/9	Sequences, recurrent neural networks, and sequence-to-sequence learning	9.1-9.7, 10.1-10.8	H5 due, H6 assigned
10/16	Attention	11.1-11.6	H6 due, H7 assigned
10/23	Transformers part 1	11.7, 15.8-15.10	H7 due, H8 assigned
10/30	No class - fall reading week	—	
11/6	Transformers part 2, large language models	11.8-11.9	H8 due, H9 assigned
11/13	Architecture grab bag part 1: GNNs, Autoencoders, and VAEs	—	H9 due, H10 assigned
11/20	Architecture grab bag part 2: Transposed convolution, UNet, and MoEs	14.1	H10 due, H11 assigned
11/27	Deep learning engineering; fairness, accountability, and transparency	13.5-13.6, 4.7	H11 due, H12 assigned

Grading

1. **Homework**, 78 points: There are 12 homeworks. Each is worth 6.5 points.
2. **Quizzes**, 22 points: There will be an in-class closed-book quiz during each lecture (which can be made-up at any of the tutorial or office hour slots). Each

will be worth 2 points, but your quiz with the worst grade will be dropped, resulting in 11 total quizzes counting toward your grade.

Late work, collaboration rules, and the honor code

There is no late penalty for handing in homework late, but all homeworks must be handed in by 11:59pm on the last day of the semester (December 4th, 2025). After that, no additional work will be accepted. You are welcome to work together with other students on the homework. You are also welcome to use any resources you find (online tutorials, textbooks, papers, chatbots, etc.) to help you complete the homework. However, **you must list any collaboration or resources you used to complete each homework on each assignment**. If you hand in homework that involved collaboration and/or makes use of content that you did not create and you do not disclose this, you will get a 0 for that homework. In addition, it is likely that you will be able to use some resource (be it another student, ChatGPT, or whatever) that can help you solve many of the homework problems. However, note that if you rely too much on such resources you will likely not learn the material and will do poorly on the in-class quizzes, during which such resources will not be available.

Conduct

I ask that we all follow the NeurIPS Code of Conduct and the Recurse Center Social Rules. I value the perspectives of individuals from all backgrounds reflecting the diversity of our students. I broadly define diversity to include race, gender identity, national origin, ethnicity, religion, social class, age, sexual orientation, political background, and physical and learning ability. I will strive to make this classroom an inclusive space for all students. Please let me know if there is anything I can do to improve.

Resources

Please be aware of the following resources at U of T which can help ensure your academic success and personal well-being:

- Writing at U of T
- University of Toronto Academic Integrity
- U of T Student Life
- Accessibility Services

- Health Services

Changes

I reserve the right to make changes to the syllabus, including project due dates and test dates. These changes will be announced as early as possible.