2023-2024 Course Syllabus

Table of Contents

What is LearnAI? 3

Course Structure 4

Background Information 5

Curriculum 6

What is LearnAI?

LearnAI is a free course run by UofT AI that introduces Artificial Intelligence and Machine Learning concepts to University of Toronto undergraduates. This course is perfect for all students, even those with limited experience working with AI!

Through this course, students will gain a broad set of skills – from regression to complex Computer Vision and Natural Language Processing concepts, allowing them to build an attractive employment profile.

Our refined curriculum is complemented by valuable resources such as TAs and mentors! The program consists of a series of lectures and workshops to teach important concepts and to help practice their implementation. Students also have the opportunity to complete a mentored final project to showcase their new skills!

Course Structure

Lecture Phase (12 lectures – 12 weeks)

- 3 hours weekly: A hour and a half lecture covering ML/AI concepts followed by an hour and a half tutorial focused on coding concepts.
- Weeks 1-7 will be in the Fall and weeks 8-12 will be in the Winter.
- Week 12 is an optional final lecture that will cover Q&A as well as further extension content if requested by students.
- Lectures and Tutorials may be online or in-person, depending on availability.
- There will be no lectures or tutorials during midterms and finals.

Project/Hackathon Phase (1 week)

- Form a team of up to 4 students and create a project in the Generative AI themed hackathon hosted in March!
- This is an open hackathon, but there will be a separate prize category and criteria for LearnAI students.
- Experienced mentors will guide you through the project.
- More details about this phase will be released in the Winter.

The tentative dates for this year's sessions can be found below with an overview of each weeks content. Lectures will be on Saturday and Tutorials will be on Sunday, with the timings yet to be decided. Any information regarding course dates, timings and locations will be communicated to students enrolled in the course.

Background Information

While the course does cater towards complete beginners in the field of ML/AI, some previous knowledge will enhance the quality of learning immensely. Particularly, recommended background information (and resources for learning them) includes;

- High school calculus, functions and linear algebra
    - Derivatives
    - Vectors
- Basic Python programming knowledge
    - CS50's Introduction to Programming with Python
- Object-Oriented Python (optional)
    - CS50's Object-Oriented programming in Python

Other than these concepts, this course focuses on an intuitive, diagram based learning approach that hopes to communicate complex concepts in a more visual manner while abstracting away the precise mathematical details involved.

Curriculum

Week 1

- Lecture 1 (September 30th)
    - Introduction to Machine Learning and it's applications
    - Linear Regression
- Tutorial 1 (October 1st)
    - Object-oriented Python programming
    - Introduction to PyTorch
    - Working with PyTorch Tensors

Week 2

- Lecture 2 (October 7th)
    - Understanding the PyTorch API
    - Multiple Linear Regression in PyTorch
- Tutorial 2 (October 8th)
    - Loading custom datasets with PyTorch

Week 3

- Lecture 3 (October 14th)
    - Representing image data
    - Neural Networks
- Tutorial 3 (October 15th)
    - Training a Neural Network for digit classification

Week 4

- Lecture 4 (October 21st)
    - Convolutional Neural Networks (CNNs)
    - Revisiting digit classification with a CNN
- Tutorial 4 (October 22nd)
    - Transfer Learning
    - Basic facial recognition with transfer learning

Week 5

- Lecture 5 (October 28th)
    - Upsizing methods
    - The encoder-decoder architecture and Autoencoders
- Tutorial 5 (October 29th)
    - Various tasks with autoencoders (image super resolution, image denoising, watermark removal, compression, etc.)

---

3 week break for midterms

Week 6

- Lecture 6 (November 25th)
    - Reading and implementing ML literature
    - Implementing the ResNet architecture
- Tutorial 6 (November 26th)
    - Implementing the VGG architecture

Week 7

- Lecture 7 (December 2nd)
    - Gradient descent for arbitrary tasks
    - Introduction to neural style transfer
- Tutorial 7 (December 3rd)
    - Implementing neural style transfer

Winter break

Week 8

- Lecture 8 (January 20th)
    - Introduction to NLP
    - Sequence modelling
    - Generating Shakespeare-like poetry
- Tutorial 8 (January 21st)
    - Recurrent Neural Networks
    - Revisiting sequence modelling and generating better Shakespeare-like poetry

Week 9

- Lecture 9 (January 27th)
    - More advanced recurrent architectures with Long Short Term Memory units (LSTMs)
- Tutorial 9 (January 28th)
    - More advanced recurrent architectures with Gated Recurrent Units (GRUs)

Week 10

- Lecture 10 (February 3rd)
    - LSTMs and GRUs with PyTorch
    - Introduction to many-to-one tasks and sentiment classification
- Tutorial 10 (February 4th)
    - Training a sentiment classification model

Week 11

- Lecture 11 (February 10th)
    - The encoder-decoder architecture in NLP
    - Introduction to Machine Translation
- Tutorial 11 (February 11th)
    - Cleaning the Machine Translation dataset
    - Training our Machine Translation model

3 week break for midterms

Week 12

- Lecture 12 (March 9th)
    - Final Q&A session
    - Extension material if students are interested