Deep Learning for Data Science (DL4DS) / Spring 2024

Announcements

Dec 3, 2023:

This course web site is under active construction. Check back regularly for updates.

The Schedule, Lectures, Assignments, Projects and Materials pages are being updated and will be posted soon.

Course Abstract

In this course we will dive into Deep Learning. We'll balance important theoretical concepts with hands on network training and applications using modern deep learning python frameworks. We'll explore numerous network architectures like CNNs, transformers, and the rapidly developing state-of-the-art of large pre-trained foundation models. You'll have the chance to apply what you've learned in a final project.

Lectures: Tuesdays and Thursdays, 3:30pm – 4:15pm

Location: CAS 208

Discussion: Wednesdays, 11:15am – 12:05pm

Location: CDS 164

Instructor: Thomas Gardos

Office: CDS 1623

Office Hours: TBD

email: tgardos <at> bu <dot> com

Course Description

In this course students will gain an understanding of the fundamentals in deep learning and then apply those concepts in exercises and applications in python. We'll start with the origins of artificial neural networks, learn about loss functions, understand gradient descent, back propagation and various training optimization techniques. Students will be familiar with canonical network architecture such as multi-layer perceptions, convolutional neural networks, recursive neural networks, LSTMs and GRU, attention and transformers. Through explanations, examples and exercises students will build intuition on how deep learning algorithms work and how they are implemented in popular deep learning frameworks such as PyTorch. Students will be able to define, train and evaluate deep learning models as well as adapt deep learning frameworks to new functionality. Students will gain exposure to pre-trained large language models and other foundation models and the concepts of few-shot learning and reasoning. Finally, students will be able to apply many of the techniques they learned in a final class project.

Learning Outcomes

Upon successful completion of this course, students will be able to:

Grasp foundational theories and practices in the deep learning arena.

Design and implement various neural network architectures using Python and PyTorch.

Employ regularization, optimization, and advanced training techniques to enhance model performance.

Analyze real-world datasets, applying suitable deep learning techniques to derive actionable insights.

Understand the benefits and drawbacks of various neural architectures in specific contexts.

Understand the pros and cons of pre-trained large language and other foundation models and how best to employ them
Complete a data-centric project, showcasing end-to-end deep learning implementation.
Prerequisites
Python Programming – Should be proficient in python and associated data science packages, or studiously working towards proficiency. See for example Scientific Python Lectures for lessons on python language and relevant packages, or The Python Tutorial for a tutorial on the core language.
Packages such as NumPy NumPy - Learn and SciPy (SciPy User Guide have tutorials and documentation as well.
The more proficient you are, the more effective you will be at the assignments and projects. We will dedicate some discussion sessions to ensure your environment is setup correctly and review some of the basics as well as answer any questions.
Math Proficiency – In order to understand the foundational concepts, it is important to have proficiency in a number of areas of mathematics. These include linear algebra, first year calculus and trigonometry as well as some concepts from Real Analysis. We will cover these concepts in the class and some recitation sessions, but refreshing or building your foundation will help.
Reference Material
The primary textbook for this course will be Understanding Deep Learning, by Simon Prince . The book is available online as a preprint and should be available in print from MIT Press in early 2024.
We'll also reference Introduction to Linear Algebra, Sixth Edition (2023), by Gilbert Strang.
Lecture notes, consisting primarily of Jupyter notebooks will be posted online as well.
Given the fast moving nature of this area, we'll also be citing many articles available online as well as other online reference materials in each lecture. As part of the class, we will guide the students in constructing their own bibliography and give tips on how to efficiently and effectively read research papers.
Computing Environment
Students are of course free to use their own personal computer, but you will also have access to Boston University's Shared Computing Cluster and GPUs. Access instructions will be provided.
For more information, see:
Shared Computing Cluster : TechWeb : Boston University
GPU Computing : TechWeb : Boston University
Learning Management Software
To be added.
Piazza
Gradescope
Blackboard
Course Requirements
Approximately bi-weekly assignments based on lecture content
A mid-term take home exam
A final project where students apply image and video data analysis techniques to a real-world dataset

Course Assessment
Homework Assignments: tbd
Final Project: tbd
Participation: tbd
Final Project
TBD: more details will be provided here.
Student Code of Conduct
All students are expected to abide by University conduct policies as detailed in the following links:
Boston University Student Codes of Conduct:
https://www.bu.edu/policies/policy-category/student-codes-condcut/
College of Arts & Sciences Codes of Conduct:
https://www.bu.edu/cas/academics/undergraduate-education/academic-conduct-code/resources-for-students/
Boston University Student Responsibilities:
https://www.bu.edu/dos/policies/student-responsibilities/
Academic Honesty
You may discuss homework assignments with classmates, but you are solely responsible for what you turn in. Collaboration in the form of discussion is allowed, but all forms of cheating (copying parts of a classmate's assignment, plagiarism from books or old posted solutions) are NOT allowed. We – both teaching staff and students – are expected to abide by the guidelines and rules of the Academic Code of Conduct (which is at http://www.bu.edu/academics/policies/academic-conduct-code/ ).
Graduate students must also be aware of and abide by the GRS Academic Conduct code at http://www.bu.edu/cas/students/graduate/forms-policies-procedures/academic-discipline-procedures/ .
You can probably, if you try hard enough, find solutions for homework problems online. Given the nature of the Internet, this is inevitable. Let me make a couple of comments about that:
If you are looking online for an answer because you don't know how to start thinking about a problem, talk to the TA or instructor, who may be able to give you pointers to get you started. Piazza is great for this – you can usually get an answer in an hour if not a few minutes.
If you are looking online for an answer because you want to see if your solution is correct, ask yourself if there is some way to verify the solution yourself. Usually, there is. You will understand what you have done much better if you do that. So … it would be better to simply submit what you have at the deadline (without going online to cheat) and plan to allocate more time for homeworks in the future.
Accommodations for Students with Disabilities
If you have a disability and have an accommodations letter from the Disability & Access Services office, I encourage you to discuss your accommodations and needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with BU Disability & Access Services, I encourage you to find more information at https://www.bu.edu/disability/.

This syllabus provides a general plan for the course; deviations may be necessary depending on the progress of the class.

Instructor

Thomas Gardos
Teaching Assistants

Xavier Thomas

The AI Terrier Tutor