JOHNS HOPKINS
WHITING SCHOOL
of ENGINEERING

Computer Science EN.601.482/682
Machine Learning: Deep Learning
Spring, 2021 (4 credits, E)

Instructor
Mathias Unberath
Office hours: TBD
Office hours Zoom link: TBD

This is a public version of the syllabus.

If you are a JHU student taking this class, make sure to join the Piazza and download the complete Syllabus from there.

Teaching Assistant
Hao Ding
Office hours: TBD
Office hour Zoom link: TBD

Keith Harrigian
Office hours: TBD
Office hours Zoom link: TBD

Meetings
Asynchronous consumption of lecture recordings
Synchronous Q/A and flipped classroom session (*chose one session*)
Session 1: Mondays 7.00 to 8.15 am
Session 2: Mondays 12.00 to 1.15 pm
Class Zoom link:

Textbook
This course does not explicitly follow a specific textbook, though many good textbooks on the topic are available. A fairly exhaustive list of additional reading material (including textbooks, blog articles, tutorials, and scientific articles) will be made available in the "Resource" section of Piazza.

Online Resources
Please log in to Piazza for all materials related to this course.

Course Information
Synopsis: Deep learning (DL) has emerged as a powerful tool for solving data-intensive learning problems such as supervised learning for classification or regression, dimensionality reduction, and control. As such, it has a broad range of applications including language processing, computer vision, medical imaging, and perception-based robotics. The goal of this course is to introduce the basic concepts of DL. The course will include a brief introduction to the basic theoretical and methodological underpinnings of

machine learning, commonly used architectures for DL, current challenges including ethics and fairness, and specialized applications with a particular focus on computer vision. Students will be expected to solve several DL problems on standardized data sets and will be given the opportunity to pursue team projects on topics of their own interest.

Prerequisites: (AS.110.201 or AS.110.212 or EN.553.291) and (EN.553.310 EN.553.311 or EN.553.420 or EN.560.348) and (EN.601.475 or equivalent); Calc III and numerical optimization recommended. Recommended co-req: EN.601.382.

Communication Plan: In Spring 2021, this course will be held entirely remotely. It consists of both asynchronous and synchronous elements to enable remote learning in various time zones and schedules while preserving some of the valuable face to face and personalized interactions that distinguish this course from other online material. Because technology-mediated interactions are not quite comparable to in-person events just yet, we will use multiple channels to communicate as effectively as possible.

How does this work? You will find below a list of lectures and/or other material for you to consume during the specified week of the semester. Consuming this material is asynchronous and you can decide your schedule as long as you finish that week. Questions that arise during self-study of the material can be 1) posted on Piazza and/or 2) submitted anonymously via Google Forms as a request for recitation and clarification during our next synchronous session. During our synchronous sessions, we will then review and emphasize the most important learning points, revisit the most frequently requested topics, and/or discuss flipped classroom tasks.

Where do I find the material? The course's home is on Piazza and you should sign up as soon as possible using this link: piazza.com/jhu/spring2021/cs482682. On **Piazza** you will find links to recorded lectures, slides, assignments, and other relevant course material and resources. You are encouraged to post any questions and discussions on Piazza and contribute to answering questions your peers have posted. Please note that, while you can remain anonymous to peers, posting anonymously to instructors is disabled.

What about homework assignments? Questions about homework assignments should be posted to Piazza and can also be discussed in greater detail during office hours. Homework assignments are submitted through **Gradescope** with course code **XXXXX**.

Any other communication platforms? Yes – We will use **Google Forms** for quizzes and anonymous recitation requests. Links to the recitation requests are provided together with the semester schedule below. Links to Quizzes will be shared during the synchronous Zoom sessions. Finally, we have set up a **Discord server** for this class to provide a more real-time platform for peer-to-peer interaction including study groups, collaborations, and simply "hanging out". Information on how to join and basic rules are available on Piazza.

Course Goals

Specific Outcomes for this course are that

- Students will learn fundamental concepts of machine learning
- Students will learn the theoretical underpinnings of deep learning

- Students will learn contemporary architectures, applications, and challenges of deep learning
- Students will learn to design, implement, and validate deep learning-based solutions to machine learning problems

This course will address the following CSAB ABET Criterion 3 Student Outcomes

Graduates of the program will have an ability to:

1. Design, implement, and evaluate a computing-based solution to meet a given set of computing requirements in the context of the program's discipline.
2. Recognize professional responsibilities and make informed judgments in computing practice based on legal and ethical principles.
3. Function effectively as a member or leader of a team engaged in activities appropriate to the program's discipline.
4. Apply computer science theory and software development fundamentals to produce computing-based solutions.

Course Topics

Homework assignments are released / are due Fridays (latest submission 11.59pm).

Week 1 (Jan 25): Introduction and Basics

Synchronous: Welcome and Course Logistics
Asynchronous: Submit questions here

L-1) Overview
L-2) Basics Part I: Image Features, Regression, and Classification

Friday, Jan 29: Homework 1 released

Week 2 (Feb 01): Basics

Synchronous: Q/A Week 1 and intro to Python
Asynchronous: Submit questions here

L-3) Basics Part II: Regularization and Optimization
L-4) Computational Graphs and Backpropagation Part I

Friday, Feb 5: Homework 1 due, Homework 2 released

Week 3 (Feb 08): Convolutional Neural Networks

Synchronous: Q/A Week 2
Asynchronous: Submit questions here

L-5) History of and Introduction to Neural Networks
L-6) Convolutional Neural Networks

Week 4 (Feb 15): Training Neural Networks

Synchronous: Q/A Week 3
Asynchronous: Submit questions here

L-7) Training Part I: Activation, Initialization, Preproc., Dropout, Batch norm
L-8) Training Part II: Updates & Momentum, Augmentation, Transfer Learning

Friday, Feb 19: Homework 2 due, Homework 3 released

Week 5 (Feb 22): Architectures
Synchronous: Q/A Week 4
Asynchronous: Submit questions here
L-9) **Inverse Classroom**: It's not working! Help!
L-10) Network Architectures: AlexNet, VGG, ResNet, U-Net, …
Friday, Feb 26: Homework 3 due, Homework 4 released

Week 6 (Mar 01): Architectures continued – **Form project groups of 3**
Synchronous: Q/A Week 5 and **Inverse Classroom Discussion**
Asynchronous: Submit questions here
L-11) **Inverse Classroom**: What does this network do?

Week 7 (Mar 08): Sequence Modeling
Synchronous: Q/A Week 6 and **Inverse Classroom Discussion**
Asynchronous: Submit questions here
L-12) RNNs and LSTM
Friday, Mar 12: Homework 4 due, Homework 5 released

Week 8 (Mar 15): Unsupervised Learning
Synchronous: Q/A Week 7
Asynchronous: Submit questions here
L-13) Unsupervised and Self-supervised Learning
L-14) Autoencoders, Variational Autoencoders, and Disentanglement
Friday, Mar 19: Homework 5 due, Homework 6 released

Week 9 (Mar 22): Spring Break
Monday Mar 22 is designated a spring break day; we give the whole week off

Week 10 (Mar 29): Generative Models
Synchronous: Q/A Week 8
Asynchronous: Submit questions here
L-15) Generative Adversarial Networks
L-16) **Inverse Classroom**: Labeling? Ain't nobody got time for that.
Friday, Mar 26: Homework 6 due, Homework 7 released

Week 11 (Apr 05): Current Topics–– **Start project work**
Synchronous: Q/A Week 10 and **Inverse Classroom Discussion**
Asynchronous: Submit questions here
L-17) Generalization, domain gaps, and explainable AI
L-18) Domain gaps and black boxes
Friday, Apr 09: **Project proposals due**

Week 12 (Apr 12): Current Topics
Synchronous: Q/A Week 11
Asynchronous: Submit questions here

L-19) DL4Health
L-20) Character and word embedding

Week 13 (Apr 19): Current Topics

Synchronous: Q/A Week 12
Asynchronous: Submit questions here
L-21) Natural Language Processing and Transformers
L-22) Deep Reinforcement Learning

Week 14 (Apr 26): Wrap up

Synchronous: Q/A Week 13
Asynchronous: Submit questions on Piazza for clarifications, this is the last week!
L-23) Human-centered AI, ethics, etc.
L-24) Wrap Up
Friday, Apr 30: Homework 7 due

Wednesday, May 07: Final project reports due

Final Exam Slot (we will address time zone concerns): Synchronous project pitch and breakout rooms

Course Expectations & Grading

We will have short weekly quizzes to test your comprehension and recollection of the course material you prepared asynchronously. The link to participate in quizzes will be shared over the zoom link during the respective synchronous session. Further, the course has 7 assignments, 3 of which are written and the remaining 4 focusing on programming. The bulk of the homework assignment workload (1 to 6) is condensed into the first two-thirds in the first two thirds of the semester to free up time for the final project. Starting in Week 10, most time will be spent on the final project that will be completed in groups of four. Individual grades will be computed as a weighted combination of these factors:

1) Quizzes: 10%
2) Homework Assignments: 50%
3) Final Project: 40%

There is opportunity for bonus points. Bonus points are earned by completing additional assignments that will be described in homework assignment 7.

All quizzes after the "course add" deadline will count towards the grade. The **worst quiz score** of every student will be automatically **dropped** from scoring.
For late assignment submissions, you have a total of **3 late days** that you can use at your discretion. However, no smaller quantity than "day" can be used but you can use multiple days for the same homework. **Late days must be requested ahead of submission deadline** via private message to instructors on Piazza.

Assignments & Readings

See above schedule for assigned lectures and materials in every week.