



Machine Learning: Artificial Intelligence System Design and Development
Computer Science 601.486/686
Spring 2020 (3 credits)
Syllabus

Instructors: Mathias Unberath (unberath@jhu.edu)
TA/CA: Aaron Mueller (amueller@jhu.edu)
Cong Gao (cgao11@jhu.edu)

Office hours: TBD

Time: Tuesday/Thursday 1.30pm to 2.45pm
Location: Croft Hall B32

Synopsis

The field of artificial intelligence (AI) has recently seen a substantial increase in popularity, largely fueled by the successes of training deep neural networks (deep learning, DL) that achieve state-of-the-art performance in a large variety of problems. These successes are not limited to academic benchmarks but have started to impact our everyday lives in the form of products such as [Google Lens](#), [Amazon Alexa](#), and [Tesla Autopilot](#). For such AI systems to succeed, several aspects are of key importance: 1) the impact on everyday life, 2) the capabilities of the underlying machine learning algorithm, and 3) and the effectiveness of the human-AI interaction. The importance of harmonic interplay between all these components is dramatically highlighted by recent catastrophic events in road transport and aviation. In this project-based course you will work in teams of 3-5 students to

- 1) Identify a need with high-impact implications on everyday life (e.g. entertainment, autonomous vehicles, healthcare), where an AI system can contribute substantially
- 2) Conceptualize and design an AI system (including a rudimentary concept for user interface to be demo-able, and focusing on the machine learning backend) that successfully improves the targeted scenario. This includes an investigation of ethical biases or other discriminatory challenges the envisioned system may face.
- 3) Develop the AI system by iteratively refining the demo-able prototype based on feedback received during course presentations

The course culminates in a public poster and demo session, where groups will present their AI systems to colleagues. While the course has a light human-computer interface component, the primary focus lies in the design, development, and training of a state-of-the-art machine learning

algorithm using techniques such as unsupervised or self-supervised learning, recurrent neural networks and LSTM, and reinforcement learning. This includes precise specification of the question to be answered, identification of suitable datasets, and training of generalizable models that strive to achieve state-of-the-art performance. Example projects could be:

- A smart movie recommender system: “Give me a movie like this one, but funnier!”
- A calendar plugin that suggests best time for meetings based on previous activity
- A Piazza plugin that automatically answers similar/duplicate questions
- A health monitoring system that interprets electronic health record data to suggest more intense monitoring
- A system that interprets medical images with state-of-the-art performance, but focuses on communicating the decision-making progress to users (saliency, explainability, etc.)

Enrollment

Due to the individualized nature of this course, enrollment is capped at 30 students maximum with 15 students enrolled in the 486 and 686 section, respectively.

Prerequisites

Students must have a solid training in theory and practice of machine learning, evidenced by successful completion of at least one of the following courses: EN 601.475/675 Machine Learning, EN 601.464/664 Artificial Intelligence, or EN 601.482/682 Machine Learning: Deep Learning.

Students must have a solid coding skills preferably in Python.

Experience with human computer interface design, e.g. EN 601.454/654 Augmented Reality, EN 601.290 User Interfaces and Mobile Applications, EN 601.490/690 Introduction to Human Computer Interaction, or EN 601.491/691 Human- Robot Interaction is desirable but not required.

Instructional Material

The learning experience in this course will be dominated by the guided hands-on experience in designing and developing an AI-based system, complemented with status reports in the form of presentation/demo, discussion sessions, and open feedback rounds. The AI system design process is ignited by lectures during the first weeks of the course. The material will be based on various papers, online resources, case studies, and slides (will be made available online through Piazza).

Online Resources

Assignments, course material, and links to relevant resources will be posted on Piazza. Students are encouraged to post any questions and discussions on Piazza as well.

Grading

The course has 1) 6 homework assignments that will be handed in and presented in class in form of bi-weekly update presentations (more information on what to prepare for these presentations will be provided in the respective assignment), 2) a presentation describing a case study of an AI system that is currently available on the market, and 3) a final poster and demo session. Since the course's success depends on everyone's engagement by providing mutual feedback and participating in discussions, participation and attendance during presentation and feedback rounds is mandatory. Grades will be computed as a weighted combination of those factors:

- 1) Individual: 20% (HW1, Task 2 -- 10% and Attendance/Participation -- 10%)
- 2) Group homework assignments: 60%
- 3) Case study presentation: 15%
- 4) Poster and demo: 5%

Topics and Schedule (subject to change)

Week 1 (Jan 27): Introduction and Basics

Monday, Jan 27: HW 1 (Group and Problem Finding, Refresher) released

- L-1) Overview, Course Logistics, and Brainstorming
- L-2) Human-centered AI: Biases, Ethics, and AI for Good

Week 2 (Feb 3): Refreshers

- L-3) Design Thinking and Human-machine Interaction
- L-4) Machine Learning Reminder: Important Concepts for Human-centered AI

Week 3 (Feb 10): Update Week – Initial Concept

Monday, Feb 10: HW 1 Task 1 due, HW 2 (Empathy, Definition & Ideation) released

- L-5) Presentation and Discussions: Part I
- L-6) Presentation and Discussions: Part II

Week 4 (Feb 17):

- L-7) Case Study + User Evaluations / Recitation / Open Question Session
- L-8) Case Study + User Evaluations / Recitation / Open Question Session

Week 5 (Feb 24): Update Week – Empathy, Definition & Ideation

Monday, Feb 24: HW 1 Task 2 & HW 2 due, HW 3 (Implementation and Testing) released

- L-9) Presentation and Discussions: Part I
- L-10) Presentation and Discussions: Part II

Week 6 (Mar 2):

- L-11) Case Study + User Evaluations / Recitation / Open Question Session
- L-12) Case Study + User Evaluations / Recitation / Open Question Session

Week 7 (Mar 9): Update Week – Implementation and Testing

Monday, Feb 24: HW 3 due, HW 4 (Finalizing Frontend) released

- L-13) Presentation and Discussions: Part I
- L-14) Presentation and Discussions: Part II

Spring Break (Mar 16)

Week 8 (Mar 23):

- L-15) Case Study + User Evaluations / Recitation / Open Question Session
- L-16) Case Study + User Evaluations / Recitation / Open Question Session

Week 9 (Mar 30): Update Week – Finalizing the Frontend

Monday, Mar 30: HW 4 due, HW 5 (Tuning and Analyzing the Backend) released

- L-17) Presentation and Discussions: Part I
- L-18) Presentation and Discussions: Part II

Week 10 (Apr 6):

- L-19) Case Study + User Evaluations / Recitation / Open Question Session
- L-20) Case Study + User Evaluations / Recitation / Open Question Session

Week 11 (Apr 13): Update Week – Tuning and Analyzing the Backend

Monday, Apr 13: HW 5 due, HW 6 (Finishing Touches and User Evaluation) released

- L-21) Presentation and Discussions: Part I
- L-22) Presentation and Discussions: Part II

Week 12 (Apr 20):

- L-23) Case Study + User Evaluations / Recitation / Open Question Session
- L-24) Case Study + User Evaluations / Recitation / Open Question Session

Week 13 (Apr 27): Update Week – Finishing Touches

- L-25) Presentation and Discussions: Part I
- L-26) Presentation and Discussions: Part II

Monday, May 11: HW 6 due

Wednesday, May 13 from 6.00 – 9.00pm (Exam slot): Demo and Poster Session

Ethics

The strength of the university depends on academic and personal integrity. In this course, you must be honest and truthful. Ethical violations include cheating on exams, plagiarism, reuse of assignments, improper use of the Internet and electronic devices, unauthorized collaboration, alteration of graded assignments, forgery and falsification, lying, facilitating academic dishonesty, and unfair competition. Report any suspected violations to the instructor. You can find more information about university misconduct policies on the web at these sites:

- Undergraduates: <http://e-catalog.jhu.edu/undergrad-students/student-life-policies/>
- Graduate students: <http://e-catalog.jhu.edu/grad-students/graduate-specific-policies/>

Personal Wellbeing

- If you are sick, in particular with an illness that may be contagious, and may require accommodation regarding homework notify us by email but do not come to class. Visit the Health and Wellness Center: 1 East 31 Street, (410) 516-8270. Also refer to: <http://studentaffairs.jhu.edu/student-life/support-and-assistance/absences-from-class/illness-note-policy/>
- All students with disabilities who require accommodations for this course should contact me at their earliest convenience to discuss their specific needs. If you have a documented disability, you must be registered with the JHU Office for Student Disability Services (385 Garland Hall; (410) 516-4720; <http://web.jhu.edu/disabilities/>) to receive accommodations.
- If you are struggling with anxiety, stress, depression or other mental health related concerns, please consider visiting the JHU Counseling Center. If you are concerned about a friend, please encourage that person to seek out these services. The Counseling Center is located at 3003 North Charles Street S-200 and can be reached at (410) 516-8278 and online at <http://studentaffairs.jhu.edu/counselingcenter/>.

Classroom Climate

We are committed to creating a classroom environment that values the diversity of experiences and perspectives that all students bring. Everyone here has the right to be treated with dignity and respect. I believe fostering an inclusive climate is important because research and my experience show that students who interact with peers who are different from themselves learn new things and experience tangible educational outcomes. Please join me in creating a welcoming and vibrant classroom climate. Note that you should expect to be challenged

intellectually by me, the TAs, and your peers, and at times this may feel uncomfortable. Indeed, it can be helpful to be pushed sometimes in order to learn and grow. But at no time in this learning process should someone be singled out or treated unequally based on any seen or unseen part of their identity.

If you ever have concerns in this course about harassment, discrimination, or any unequal treatment, or if you seek accommodations or resources, I invite you to share directly with me or the TAs. I promise that we will take your communication seriously and to seek mutually acceptable resolutions and accommodations. Reporting will never impact your course grade. You may also share concerns with the department chair (Randal Burns, chairadm@cs.jhu.edu), the Director of Undergraduate Studies (Joanne Selinski, joanne@cs.jhu.edu), the Assistant Dean for Diversity and Inclusion (Darlene Saporu, dsaporu@jhu.edu), or the Office of Institutional Equity (oie@jhu.edu). In handling reports, people will protect your privacy as much as possible, but faculty and staff are required to officially report information for some cases (e.g. sexual harassment).