A Gentle Introduction to Machine Learning PHIL 408F/PHPE 408J

Instructor:	Eric Pacuit (epacuit@umd.edu, pacuit.org)
Teaching Assistant:	Justin Helms (jhelms@umd.edu)
Semester:	Spring 2025
Course Website:	https://umd.instructure.com/courses/1380870
Class Times:	TuTh 12:30pm - 1:45pm
Location:	ASY 1213
Office Hours:	TBA
Office:	Skinner 1103A

Course Description

This course provides an introduction to the fundamental concepts and tools of machine learning. Students will explore both the theoretical foundations and practical techniques needed to implement machine learning algorithms. Key paradigms, including supervised learning, unsupervised learning, and deep learning, will be covered. These topics will be reinforced through hands-on projects and examples, enabling students to apply a wide range of machine learning methods with confidence.

A significant portion of the course focuses on deep learning. Using Keras as the primary framework, students will design, build, and experiment with various deep learning models. Applications such as image classification and text generation will be used to illustrate key concepts. Along the way, students will develop practical skills in coding, model optimization, and deployment. If time allows, the course will also address the strengths, limitations, and ethical implications of AI technologies.

Note for Computer Science or Machine Learning Majors: This course is intended for students with little to no background in computer science. Most of the material covered here will be explored in greater depth in computer science courses. Specifically, if you have taken or plan to take CMCS 422, I strongly recommend against enrolling in this course. Before registering, please consult your Academic Advisor to confirm whether this course fulfills your ULC requirements. (My recommendation is that this course should **not** count toward the ULC requirement.)

Prerequisites

This course is self-contained and does not require formal prerequisites in math or programming. However, it is important to be aware that it will involve programming exercises and discussions of several mathematical concepts. To ensure all students are prepared, the course will begin with a 2-3 week crash course in Python programming.

While prior experience with calculus (at the level of a high school course) is helpful, it is not required. Students should ideally have some familiarity with the following topics: the concept of a derivative; logarithms and exponents; basic trigonometry, particularly cosines; and graphs of functions. I will make every effort to explain the necessary mathematical concepts as they arise in the course. However, due to time constraints, we will not have the opportunity to explore these topics in depth.

Course Objectives

After successfully completing this course, students will be able to:

- Design, code, and train machine learning models using Keras.
- Identify and explain the fundamental concepts behind supervised learning, unsupervised learning, and deep learning.
- Implement and evaluate machine learning algorithms for tasks such as classification, regression, clustering, and pattern recognition, with an understanding of their strengths and limitations.

Required Resources

- Course Website: https://umd.instructure.com/courses/1380870
- We will **PollEverywhere** this semester. This will be used for short quizzes and surveys given during the lectures. You can sign-up for PollEverywhere for free by following this link

https://PollEv.com/epacuit/register?group_key=jKUzrkJWJcUsPeZ770aLs3mhg

• This term we will be using **Piazza** for class discussion:

https://umd.instructure.com/courses/1380870/external_tools/42711

The system is highly catered to getting you help fast and efficiently from both me and your classmates. Rather than emailing questions to me, I encourage you to post your questions on Piazza.

• **Readings**: Much of the course will discuss chapters from:

F. Chollet (2021). Deep Learning with Python: Second Edition, Manning Press.

Although I do recommend purchasing the book, any readings will be made available on the course website.

Supplemental optional reading:

- M. Hutson (2021). Robo-writers: the rise and risks of language-generating AI. Nature, 591, pp. 22 - 25.
- C. Stokel-Walker and R. Van Noorden (2023). What ChatGPT and generative AI mean for science, Nature, 614, pp. 214 - 216.
- E. van Dis, J. Bollen, W. Zuidema, R. van Rooij, and C. Bockting (2023). ChatGPT: Five priorities for research, Nature, 614, pp. 224 - 226.

Course Structure

This course will include weekly lectures on Tuesdays and Thursdays. During these sessions, we will introduce the week's material and work collaboratively on tutorials. Please bring your laptop to each lecture and be ready to engage with the tutorials.

Attending lectures is essential for success in this course. While attendance will not be formally recorded, you are responsible for all material covered during lectures, even if you are absent. It will be very easy to fall behind if you miss classes, so please reach out to me or the TA if you miss more than one or two sessions to ensure you stay on track.

Grading Policy

The course requirements are:

- Participation: The participation grade will be based on the following two components:
 - 1. Answers to PollEverywhere questions during lectures. Missed questions cannot be made up, but a percentage of the total questions will be dropped to accommodate occasional absences.
 - 2. Your contribution to Piazza. Students are expected to make a minimum of 10 substantive contributions on Piazza. Contributions may include asking questions, posting comments, or responding thoughtfully to another student's question or comment.
- **Tutorials**: There will be a number of tutorials assigned throughout the semester that will include many programming exercises. These will be assigned typically on Mondays and due within 1-2 weeks. Since the answers to the tutorials may be discussed during lectures, late assignments will not be accepted for full credit. If you hand in a tutorial late, we will give you 1/2 credit. In addition, the lowest grade of the tutorials will be dropped.
- **Exams**: There will be 2 exams given during the semester: A midterm and a final exam. Details about these exams will be made available later in the semester. The midterm and final exam may consist both an online portion and an in-class portion.

Grades will be assigned according to the following weights:

Participation	35%
Tutorials	45%
Midterm	10%
Final Exam	10%

Your final grade may be curved, but the final grade cutoffs are typically as follows:

Course Outline

The following is a *tentative* course outline. It may be adjusted based on the students' background, allowing us to either progress more quickly or take additional time as needed.

Week 1 (1-28, 1-30): Introduction; Crash Course in Python I *Reading*: Chollet, chapter 1; Online tutorials

Week 2 (2-4, 2-6): Crash Course in Python II *Reading*: Online tutorials

Week 3 (2-11, 2-13): First steps in machine learning: Classification and Regression *Reading*: Online tutorial

Week 4 (2-18, 2-20): Fundamentals of neural networks I: The forward pass *Reading*: Chollet, chapter 2

Week 5 (2-25, 2-27): Fundamentals of neural networks II: Backpropagation *Reading*: Chollet, chapter 2 *Note:* Thursday, 2-27 - No lecture (away at AAAI Conference)

Week 6 (3-4, 3-6): Introduction to Keras; Training, validation, and test sets *Reading*: Chollet, sections 3.2-3.3, 3.6, & 4.1-4.2; section 7.2.2

Week 7 (3-11, 3-13): Finish any outstanding topics *Reading*: Chollet, sections 5.2-5.4; chapter 6 (optional)

Week 8 (3-18, 3-20): No class - Spring Break

Week 9 (3-25, 3-27): Computer vision I: Introduction to convolutional neural networks *Reading*: Chollet, sections 8.1-8.2

Week 10 (4-1, 4-3): Computer vision II: classification; visualizing convnets *Reading*: Chollet, sections 8.3, 9.3

Week 11 (4-8, 4-10): Introduction to natural language processing: Representing words as numbers *Reading*: Jay Alammar, Illustrated Word2Vec; Chollet sections 11.1-11.3

Week 12 (4-8, 4-10): Transformers I: Introduction; Masked language models *Reading*: Jay Alammar, The Illustrated Transformer, Chollet chapter 11.4

Week 13 (4-15, 4-17): Transformers II: Masked language models continued; Generative models *Reading*: Chollet chapter 11.5.3, 12.1

Week 14 (4-22, 4-24): Transformers III: Generative models continued *Reading*: Chollet chapter 11.5.3, 12.1

Week 15 (4-29, 5-1): Putting it all together: Review and loose ends *Reading*: TBA

Week 16 (5-6, 5-8): Putting it all together: Review and loose ends *Reading*: TBA

Course Policies

A full list of course-related policies and relevant links to resources may be found at:

http://www.ugst.umd.edu/courserelatedpolicies.html.

Communication about this Course There are two ways to communicate about the course:

- 1. ELMS announcements: https://umd.instructure.com/courses/1380870/announcements. There will be weekly announcements (typically on Mondays) to convey important information about the course.
- 2. Piazza: https://umd.instructure.com/courses/1380870/external_tools/42711 This should be used for questions both about logistics and content. It is much faster to ask a question on Piazza than asking questions on email.

Students are responsible for keeping their email address up to date, and must ensure that forwarding to another email address functions properly. Failure to check email, errors in forwarding, and returned email are the responsibility of the student, and do not constitute an excuse for missing announcements or deadlines.

Names/Pronouns and Self-Identifications The University of Maryland recognizes the importance of a diverse student body, and we are committed to fostering inclusive and equitable classroom environments. I invite you, if you wish, to tell us how you want to be referred to both in terms of your name and your pronouns (he/him, she/her, they/them, etc.). The pronouns someone indicates are not necessarily indicative of their gender identity. Visit trans.umd.edu to learn more.

Additionally, how you identify in terms of your gender, race, class, sexuality, religion, and dis/ability, among all aspects of your identity, is your choice whether to disclose (e.g., should it come up in classroom conversation about our experiences and perspectives) and should be self-identified, not presumed or imposed. I will do my best to address and refer to all students accordingly, and I ask you to do the same for all of your fellow Terps.

Accessibility and Disability Services The University of Maryland is committed to creating and maintaining a welcoming and inclusive educational, working, and living environment for people of all abilities. The University of Maryland is also committed to the principle that no qualified individual with a disability shall, on the basis of disability, be excluded from participation in or be denied the benefits of the services, programs, or activities of the University, or be subjected to discrimination. The Accessibility & Disability Service (ADS) (https://www.counseling.umd.edu/ads/) provides reasonable accommodations to qualified individuals to provide equal access to services, programs and activities. ADS cannot assist retroactively, so it is generally best to request accommodations several weeks before the semester begins or as soon as a disability becomes known.

For assistance in obtaining an accommodation, contact Accessibility and Disability Service at 301-314-7682, or email them at adsfrontdesk@umd.edu.

Student Resources and Services

- Note taking assistance (https://counseling.umd.edu/ads/notetakers)
- Counseling Center (https://www.counseling.umd.edu/)
- UMD's Student Academic Support Services website (http://tutoring.umd.edu/)
- UMD's Writing Center (http://www.english.umd.edu/academics/writingcenter/schedule)
- UMD's Student Resources and Services website (https://sph.umd.edu/content/student-resourcesand-services)
- Basic Needs Security (https://studentaffairs.umd.edu/basic-needs-security)

Class Cancelations The University may be closed in the event of an emergency, in which case class will be cancelled. To find out if the University is closed you can check its main site (http://www.umd.edu), its emergency preparedness site (http://www.umd.edu/emergencypreparedness/), or call the "snow phone line" at 301-405-7669 (which covers more than just snow caused closings). If class is cancelled while the University remains open, then there will be an announcement posted on the course ELMS page.

Emergency Protocol: In the case of an extended closure to the University (e.g., because of inclement weather), consult the ELMS course page for announcements and changes to any due dates.

Attendance and Absences Students are expected to attend classes regularly. Consistent attendance offers students the most effective opportunity to gain command of course concepts and materials. Events that justify an excused absence include: religious observances; mandatory military obligation; illness of the student or illness of an immediate family member; participation in university activities at the request of university authorities; and compelling circumstances beyond the student's control (e.g., death in the family, required court appearance). Absences stemming from work duties other than military obligation (e.g., unexpected changes in shift assignments) and traffic/transit problems do not typically qualify for excused absence. Students claiming an excused absence must notify the course instructor in a timely manner and provide appropriate documentation. The notification should be provided either prior to the absence or as soon afterwards as possible. In the case of religious observances, athletic events, and planned absences known at the beginning of the semester, the student must inform the instructor during the schedule adjustment period. All other absences must be reported as soon as is practical. The student must provide appropriate documentation of the absence. The documentation must be provided in writing to the instructor by the means specified in this syllabus. The full university attendance/absence policy can be found here: http://www.ugst.umd.edu/courserelatedpolicies.html.

Academic Integrity The UMD Honor Code prohibits students from cheating on exams, plagiarizing papers, submitting the same paper for credit in two courses without authorization, buying papers, submitting fraudulent documents and forging signatures. On every examination, paper or other academic exercise not exempted by the instructor, students must write by hand and sign the following pledge: I pledge on my honor that I have not given or received any unauthorized assistance on this examination (or assignment).

Allegations of academic dishonesty will be reported directly to the Student Honor Council: https://studentconduct.umd.edu/

Copyright Notice Class lectures and other materials are copyrighted. They are the property of the instructor - do not sell them, do not post them on a website. They may not be reproduced for anything other than personal use without written permission from the instructor. Copyright infringements may be referred to the Office of Student Conduct.

Academic Accommodations for Students who Experience Sexual Misconduct The University of Maryland is committed to providing support and resources, including academic accommodations, for students who experience sexual or relationship violence as defined by the University's Sexual Misconduct Policy. To report an incident and/or obtain an academic accommodation, contact the Office of Civil Rights and Sexual Misconduct at 301-405-1142. If you wish to speak confidentially, contact Campus Advocates Respond and Educate (CARE) to Stop Violence at 301-741-3555. As 'responsible university employees' faculty are required to report any disclosure of sexual misconduct, i.e., they may not hold such disclosures in confidence. For more information: http://www.umd.edu/ocrsm/

Diversity The University of Maryland values the diversity of its student body. Along with the University, I am committed to providing a classroom atmosphere that encourages the equitable participation of all students regardless of age, disability, ethnicity, gender, national origin, race, religion, or sexual orientation. Potential devaluation of students in the classroom that can occur by reference to demeaning stereotypes of any group and/or overlooking the contributions of a particular group to the topic under discussion is inappropriate.

For information on elms, counseling, health, learning workshops, tutoring, writing help, student rights in undergrad courses, questions about graduation or add/drop/withdraw, please see http://www.ugst.umd.edu/courserelatedpolicies.html.