

# Course Syllabus

**Course Title:** Machine Learning in Practice

**Course Code:** CMU 17-691

**Instructor:** Isaac Faber

## Overview

**Course Description:** As Machine Learning and Artificial Intelligence methods have become common place in both academic and industry environments many resources have focused on methods and techniques for applications. However, there are other considerations that must be addressed when deploying such techniques into practice (or production). The purpose of this course is to cover topics relevant to building a machine learning system deployed into operations. Such systems have technical requirements including data management, model development, and deployment. However, business/organizational impacts must also be considered. Machine learning systems can be expensive to produce and operate. Students will learn about trade-offs in design, implementation, and expected value.

**Course Objectives:** After completing this course, students:

1. Can deploy products with machine learning and AI components
2. Understand how to implement data pipelines and data engineering systems
3. Calculate the approximate value provided by a machine learning system to an organization
4. Understand how to continually assess the value and quality of a deployed machine learning system

**Knowledge Prerequisites:** Introductory course in Machine Learning. Understanding of basic machine learning concepts (i.e., supervised/unsupervised learning, cost functions, confusion matrix, regression vs. classification). Working knowledge of Python 3.x, familiarity of Docker (not required be useful). This is a graduate level course.

**Course Prerequisites:** 17-634 Applied Machine Learning (or similar).

**Assessments:** Evaluation will be based on the following distribution: 60% Group Project (with peer grading), 20% in-class participation, 20% assignments.

**Waitlist:** The waitlist will be monitored and cleared periodically based on students' interest and instructor capacity.

## Schedule

Class	Section	Topic	Description
1	Foundations	What are we doing here?	Course overview and purpose. A discussion

			on the history of ML, AI, and data science.
2		What is a 'good' ML Product?	Concepts useful to define and scope a quality ML product
3		How an ML project works	The process and cadence of an ML project
4		Baselines: Do you even ML bro?	Building product baselines; metrics, context, and heuristics
5	MLOps	Tools and Infrastructure	A survey of the MLOps 'stack' and candidate technologies
6		Deployment and Monitoring	How to deploy ML components and keep track of their performance
7		Testing and Explain-ability	Evaluating testing Systems
8		Ethics and Governance	Topics about ethical ML development and governance considerations
9	Value	Understanding Decision Analysis	Core concepts in Decision Analysis
10		Applying ML concepts to DA context	How to use ML models to inform decision analysis problems
11	Guest Lecture Series		Guest Lecture in ML Value
12			Guest Lecture in MLOps
13		Product Demo's and Presentations	
14		Product Demo's and Presentations	

## Policies

Much of the following course policies are taken from Hthose used in 17-313, which we reuse almost directly (with minor modifications).

**Communication:** We make announcements through Canvas, including clarifying homework assignments. We will use the Canvas messaging system as the primary means of one-on-one communication.

**Course Syllabus and Policies:** The course is currently planned to be entirely live (in-person lecture with virtual option) for lectures, and recitation. Some lectures may be held over zoom as well. The course uses Canvas for homework submission, grading, discussion, questions, announcements, lecture recordings, and supplementary documents. We will also use Slack for communication and group work.

**Teamwork:** Teamwork is an essential part of this course. The group project is the majority factor in determining a student's grade. Most assignments are done in teams of 3-5 students. Teams will be assigned by the instructor and stay together for the entire course. Guidance on teamwork, reflection, and conflict resolution will be provided throughout the semester and are an essential component of the class.

**Textbook and Readings:** Various readings throughout the semester available online or through the library; we do not have a single textbook but rather assemble readings from different sources.

**Time management:** This is a 6-unit course, and it is our intention to manage it so that you spend close to 6 hours a week on the course, on average. In general, 3 hours/week will be spent in class and 3 hours on reading and assignments. Notice that most homework is done in groups, so please account for the overhead and decreased time flexibility that comes with groupwork. Please feel free to give the course staff feedback on how much time the course is taking for you.

**Late work policy:** Late work will receive feedback but no credit. Due to heavy reliance on teamwork in this course there are no late days. Exceptions to this policy will be made only in extraordinary circumstances, almost always involving a family or medical emergency—with your academic advisor or the Dean of Student Affairs requesting the exception on your behalf. Accommodations for travel (e.g., for interviews) are possible if requested at least 3 days in advance. Please communicate also with your team about timing issues.

**Writing:** Describing tradeoffs among decisions and communication with less technical stakeholders are key aspects of this class. Most homework assignments have a component that require discussing issues in written form or reflecting about experiences. To practice writing skills, the Global Communications Center (GCC) offers one-on-one help for students, along with workshops. The instructors are also happy to provide additional guidance if requested.

**Professionalism:** Your classmates are your colleagues. This is particularly true in this course, where we aim to provide you with principles, practices, tools, and paradigms that will enable you to be an effective, real-world Software Engineer. We ask that you treat one another like the professionals you are and that you are preparing to be.

To that end, we will not tolerate harassment in this class. We define harassment as unwelcome or hostile behavior of an ad hominem nature, i.e., that focuses not on ideas but on people and identity. This includes offensive verbal or written comments in reference to gender, sexual orientation, disability, physical appearance, race, or religion; sexual images in public spaces; deliberate intimidation, stalking, following, harassing photography or recording, sustained disruption of class meetings, inappropriate physical contact, and unwelcome sexual attention.

Harassment is against the law and we have no tolerance for it, and neither does the university. Even when behavior does not rise to the level of harassment (even if you think you're "just joking!"), it can still make people very uncomfortable, and harm their educational and professional career by forcing them to devote mental energy to something other than the material they are trying to learn or the professional successes they are trying to achieve. However, we expect that we do not need to threaten you to earn your respect on this matter: we simply ask that you treat one another like professionals, in the most positive sense.

This has two implications:

- 1) If you feel someone is violating these principles (for example, with a joke that could be interpreted as sexist, racist, or exclusionary), and you feel you have the standing to do so, speak up! Do not be a bystander to unprofessional behavior.
- 2) If you do not feel comfortable doing so, and/or if the behavior persists, send a private email to the course instructors or set up a meeting with us to discuss the matter. We will preserve your anonymity. We, the course staff, are committed to affording you the same respect we ask you to afford one another. If you feel that we are not doing so, we hope you will feel comfortable either telling us so directly or approaching another one of the course staff with your concerns.

**Academic honesty and collaboration:** The usual policies apply, especially the University Policy on Academic Integrity. Many of the assignments will be done in groups. We expect that group members collaborate with one another, but that groups work independently from one another, not exchanging results with other groups. Within groups, we expect that you are honest about your contribution to the group's work. This implies not taking credit for others' work and not covering for team members that have not contributed to the team. Otherwise, our expectations regarding academic honesty and collaboration for group work are the same as for individual work, substituting elevated to the level of "group."

The rest of this academic honesty and collaboration content is taken from the policy used in 17-214, which we reuse almost directly (with minor modifications, and attribution).

"You may not copy any part of a solution to a problem that was written by another student, or was developed together with another student, or was copied from another unauthorized source such as the Internet. You may not look at another student's solution, even if you have completed your own, nor may you knowingly give your solution to another student or leave your solution where another student can see it.

Here are some examples of behavior that are inappropriate:

Copying or retyping, or referring to, files or parts of files (such as source code, written text, or unit tests) from another person or source (whether in final or draft form, regardless of the permissions set on the associated files) while producing your own. This is true even if your version includes minor modifications such as style or variable name changes or minor logic modifications. Getting help that you do not fully understand, and from someone whom you do not acknowledge on your solution. Writing, using, or submitting a program that attempts to alter or erase grading information or otherwise compromise security of course resources. Lying to course staff. Giving copies of work to others, or allowing someone else to copy or refer to your code or written assignment to produce their own, either in draft or final form. This includes making your work publicly available in a way that other students (current or future) can access your solutions, even if others' access is accidental or incidental to your goals. Beware the privacy settings on your open source accounts! Coaching others step-by-step without them understanding your help. If any of your work contains any statement that was not written by you, you must put it in quotes and cite the source. If you are paraphrasing an idea you read elsewhere, you must acknowledge the source. Using existing material without proper citation is plagiarism, a form of cheating. If there is any question about whether the material is permitted, you must get permission in advance. We will be using automated systems to detect software plagiarism.

It is not considered cheating to clarify vague points in the assignments, lectures, lecture notes; to give help or receive help in using the computer systems, compilers, debuggers, profilers, or other facilities; or to discuss ideas at a very high level, without referring to or producing code.

Any violation of this policy is cheating. The minimum penalty for cheating (including plagiarism) will be a zero grade for the whole assignment. Cheating incidents will also be reported through University channels, with possible additional disciplinary action (see the above-linked University Policy on Academic Integrity).

If you have any question about how this policy applies in a particular situation, ask the instructors or TAs for clarification.”

Note that the instructors respect honesty in these (and indeed most!) situations.

**Accommodations:** If you wish to request an accommodation due to a documented disability, please inform the instructor as soon as possible and contact Disability Resources at 412.268.2013 or [access@andrew.cmu.edu](mailto:access@andrew.cmu.edu).

**A note on self care:** A note on self care. Please take care of yourself. Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress. All of us benefit from support during times of struggle. You are not alone. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is often helpful.

If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit their website at <http://www.cmu.edu/counseling/>. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.