

CS 4361

Introduction to Machine Learning

Course Website	Content is hosted on Google Classroom
Class Times	M/W: 6:10 - 7:30pm EST
Class Location	Google Meet (see calendar invites)
# of Credits	3 credits

1. General Course Description

1.1 Course Description

We will study how computers can learn from data. We will start with the basic machinery of training – *gradient descent* – and show how it applies for increasingly complex neural networks. Through practice with real datasets, students will learn to think critically about data and prediction tasks, choose appropriate baselines, experiment with training models, analyze results, and formulate ideas about how to improve models. Students will learn the basics of the powerful TensorFlow library as well as NumPy and Matplotlib, industry standards for machine learning. Finally, students will also be exposed to some specific applications of machine learning at Google and how to start thinking about the social implications for automated systems.

1.2 Prerequisites

Essential:

- Programming in Python
- Basics of vectors and matrices

Helpful:

- Basics of numpy and matplotlib libraries
- Experience reading documentation for Python libraries
- Basics of Probability and Statistics
- Basics of Linear Algebra
- Basics of Multivariate Calculus: what is a gradient?

2. Course Resources

2.1 Course Components

- Lecture

- 80 minutes, 2 times per week
- Introduce concepts with some code examples
- Some interactive activities
- **Lab**
 - 50 minutes 1 time per week run by TAs
 - Focus on applying concepts with code, reviewing homework, questions, etc.

2.3 Tools and Resources

The course utilizes several platforms to host material and exams:

- Google Classroom: An educational tool which will host all lecture slides, collaborative projects, quizzes, and surveys.
- CodeSignal: An technical interview platform that the course leverages to host all exams.

2.2 Components of course grade

% of grade	Component	Description
15%	Attendance and Participation	Students will be graded for their attendance and cameras on policy during class. Each student will receive 2 pts for attending/participating and 2 pts for having their cameras on. Engagement will be graded based on in-class activities during lecture and lab, as well as weekly check-ins. Please let us know if you will need to miss any lecture or lab sessions by submitting the Attendance Form .
40%	Homeworks	There will be 4 graded homework assignments. These will be colab notebooks with coding questions and analysis. While you may discuss these with classmates, everyone must write their own code.
20%	Midterm Exam	This will be an in-class exam. No coding.
25%	Final Project	Coding project in pairs.

2.3 Asking Questions

If students have questions outside of class or lab, there are 4 venues where they can ask them:

- Student group chat
 - When students start the course, they will be added to two Google Chat groups, one for their section, and one for the entire course: "TechX ML Chat (Everyone)". The course group chat will be monitored by TAs who will answer questions.
- Course email:

- If students have questions that they don't want to send publicly in the group chat (e.g. a question that includes a snippet of their code), they can email the course staff at tx24-ml-googler-leads@techexchange.in.
- Office hours
 - Instructors will host office hours throughout the semester via video chat. Instructions will be sent out for how to join those office hours later.
- Directly email your instructors or Course PgM
 - If you have questions specifically for your instructors or your Course PgM (see their emails below).

2.4 Required textbooks

Required linked readings will be posted every few weeks in preparation for upcoming topics. Students will access assignments and readings via Google Classroom.

2.5 Department resources

The course will be co-taught with instructors from Google and faculty from partner universities.

2.6 Background resources

- [Machine Learning Glossary](#) – Definitions for many ML concepts
- [Khan Academy](#) – Videos, summaries, practice questions for school math concepts
 - [Working with Matrices](#)
 - [Understanding Gradients](#)
- [3Blue1Brown](#) – Smooth animated videos of more advanced math
 - Visualizing Calculus [[intro](#)] [[derivatives](#)] [[derivative formulas](#)] [[chain rule](#)]
 - Neural Networks [[intro](#)] [[training](#)] [[backpropagation](#)] [[backprop math](#)]
- [Andrej Karpathy](#) – ML luminary explaining NN concepts
 - [Intro to neural networks and backpropagation](#)
- [StatQuest](#) – Lots of explainer videos of ML concepts (statistics perspective)

In addition, we will add more resources over the course of the semester to the following document: [\[TechX ML 2024\] Reading Material and Online Resources](#).

3. Course Aims and Learning Objectives

3.1 Course aims

Students who complete this course will be able to:

- apply machine learning framing and techniques to practical problems;
- train linear and non-linear models using gradient descent methods;
- analyze data and perform appropriate feature transformations to improve model performance;
- implement machine learning solutions using TensorFlow;
- evaluate the accuracy of a trained model both qualitatively and quantitatively;

- decide at a high level how to approach a task with the tools of machine learning and consider how the resulting system will be used

4. Staff & Course Times

Role	Name	Email
Instructor	Maxwell Chen	maxwellchen@techexchange.in
Instructor	Chunmei Liu	chuliu@techexchange.in
Instructor	Vojislav Stojkovic	vojislav.stojkovic@techexchange.in
Course PgM	Nerrisa Calagno-Mendoza	nerrisac@techexchange.in
Course PgM	Natalia Karauda	nkarauda@techexchange.in
TA	Achuith Bhandarkar	achuith@techexchange.in
TA	Allen Shen	allenshen@techexchange.in
TA	Ankush Garg	ankugarg@techexchange.in
TA	Arcad (Hsiao Chieh) Tseng	arcadtseng@techexchange.in
TA	Danie Theron	dtheron@techexchange.in
TA	Joseph Thomas	josephgec@techexchange.in
TA	Laz Karydas	lkary@techexchange.in
TA	Lief Esbenshade	liefesbenshade@techexchange.in
TA	Pavan Adharapurapu	apavan@techexchange.in
TA	Sruthi Bhavanam	sruthibhavanam@techexchange.in
TA	Wenxuan Cai	wenxuancai@techexchange.in
TA	Zhiwen Zhu	zhuzhiwen@techexchange.in

	TIME	TA
Lab 1	Friday 8-8:50 am PT / 11-11:50 am EST	TBD
Lab 2	Friday 8-8:50 am PT / 11-11:50 am EST	TBD
Lab 3	Friday 9-9:50 am PT / 12-12:50 pm EST	TBD
Lab 4	Friday 10-10:50 am PT / 1-1:50 pm EST	TBD
Lab 5	Friday 12-12:50 pm PT / 3-3:50 pm EST	TBD
Lab 6	Friday 2-2:50 pm PT / 5:00-5:50 pm EST	TBD

5. Weekly Topics

Course Schedule						
Module	Week	Date	Assigned	Due	Lab	Topic
Foundations	1	1/15			Lab 1: Python Libraries	[MLK Day] No Class
		1/17	HW 0			Intro to ML
	2	1/22	HW 1	HW 0	Lab 2: Linear Regression	Generalization and Overfitting
		1/24				Linear Regression
	3	1/29			Lab 3: Gradient Descent	Gradient Descent
		1/31				Features
	4	2/5	HW 2	HW 1	Lab 4: Logistic Regression	Logistic Regression
		2/7				Binary Classification
	5	2/12			Lab 5: Multiclass Classification	Multiclass Classification
		2/14				Multiclass Classification Cont.
Neural Networks	6	2/19	HW 3	HW 2	Lab 6: Neural Networks	[President's Day] No Class
		2/21				Nonlinear problems, XOR
	7	2/26			Lab 7: Midterm Review	Backpropagation
		2/28				Review Session
	8	3/4	HW 4	HW 3	No Lab	Midterm
		3/6				Backpropagation Cont.
9	3/11				[Spring Break] No Class	
	3/13				[Spring Break] No Class	
Natural Language Processing	10	3/18		Project Partners	Lab 10: Text Embeddings	Language Data
		3/20				Language Embeddings
Computer Vision	11	3/25		HW 4	Lab 11	Language Models, Transfer Learning
		3/27		Project Baselines		Convolution, Padding, Pooling
Generative AI and Fairness	12	4/1			Lab 12	Feature Maps, Non-Vision Applications
		4/3		Project Data Analysis		[TENTATIVE, WIP] Intro to GenAI: Transformers and Attention
	13	4/8			Lab 13	[TENTATIVE, WIP] LLMs and Prompt Engineering
4/10				ML Fairness, Bias, Privacy, Ethics		
ML at Google	14	4/15			Lab 14	Guest Speaker (MakerSuite, LLMs)
		4/17		Project Report		Guest Speaker (Responsible AI)
	15	4/22			No Lab	Guest Speaker Panel (RL, Robotics, ?)
4/24				The state of ML today, careers in ML, Grad School		

[Drive folder](#) of all lecture recordings (can only access with your techexchange.in account)

6. Final Project

[Link to Kaggle Project](#)

[Project description, due dates, grading rubric, etc.](#)

7. Assessment and Grades

7.1 Grading Scale

Grades will be determined using the rubric below. This course is *not* graded on a curve. We will map grades to include + and - for those schools that use these.

Percent of points earned	Letter Grade
≥ 90	A
[80, 90)	B
[70, 80)	C
[60, 70)	D
< 60	F

7.2 Late Work policy

Late assignments will be deducted 10% for each day late up to a maximum of 5 days; assignments not turned in within 5 days of the deadline will get no credit. We understand there may be extenuating circumstances, which we'll evaluate on a case by case basis.

7.3 Attendance policy

Our class and lab sessions will focus more on activities than a typical course. As a result, students are expected to attend classes and lab sessions regularly and on time, and should try to participate in class activities. If students must miss a class, they should notify their instructor in advance or as soon thereafter as possible. Unexcused absences or non-participation will result in a loss of participation credit; course staff will speak with students on a case-by-case basis about participation as necessary. Please fill out the [Attendance Form](#).

7.4 Plagiarism policy

Students are allowed to talk with other students about homework assignments. However, they are *not* allowed to share or look at other student's code for homework or exams. We should not see two homework submissions that have identical snippets of code. If that occurs, both homework submissions will receive an automatic **0**. Likewise, exams will receive a **0** that are flagged for identical code snippets. Students should also not copy code from online sources. That is, we should not see code snippets in homework solutions that also appear online. If you're worried that you'll be mistakenly flagged for plagiarism, include a comment in your homework submission explaining why. Students in general must be prepared to explain any code they submit.

In addition to this policy, all instances of plagiarism will be directed to the university administration, which will conduct the appropriate hearings.

8. ADA Policies and Procedures

If a student needs particular accommodations to be made, they must fill out [this form](#).