CS 1675: Introduction to Machine Learning

FALL, 2026

INSTRUCTOR: Weizhi Tang

OFFICE: N527

EMAIL: weizhi.tang@scai-ji.cn

OFFICE HOURS: Tuesday 11:10–11:55, 16:45–17:30, and 17:40–18:25 (N527), or by appointment.

LECTURES: Tuesday 13:50–16:25, Room S206

TEXTBOOKS:

- Zhi-Hua Zhou, *Machine Learning*, Springer, 2021, ISBN 978-981-15-1966-6 (for Part 1: Classical Machine Learning);
- Simon J. D. Prince, *Understanding Deep Learning*, MIT Press, 2023, free PDF at udlbook.github.io/udlbook (for Part 2: Deep Learning).

REFERENCES: scikit-learn Documentation (scikit-learn.org) and PyTorch Documentation (pytorch.org).

TEACHING ASSISTANT: Hang Chen

PREREQUISITE: CS 1501 Algorithm Implementation and STAT 1000 Applied Statistical Methods

DESCRIPTION: An introduction to machine learning, from classical methods to modern deep learning. Topics include machine learning basics (learning paradigms, inductive bias, and evaluation & model selection), regression and over-fitting, classification, ensemble methods, clustering and dimension reduction, neural networks and training, convolutional and recurrent neural networks, and modern architectures including Transformers, Diffusion Models, and Joint Embedding Predictive Architecture (JEPA).

COURSE OBJECTIVES: To learn to model data precisely, choose methods deliberately, and build machine learning models that actually work.

LEARNING OUTCOMES FOR THIS COURSE:

- **Foundations & Theory:** Students should understand classical machine learning techniques from both theoretical and practical perspectives, including the assumptions behind each method, so that they know why and when the methods work;
- **Deep Learning:** Students should understand the principles and capabilities of deep learning, including neural networks, CNNs, RNNs, Transformers, and JEPAs, along with their advantages and disadvantages;
- **Building ML Models:** Students should be able to design, implement, train, and evaluate machine learning models in practice, turning ML ideas into working models through the assignments and the course project.

GRADE DETERMINATION:

- Attendance 5%
- Assignments (4 × 5%) 20%

- Projects (10% + 20%) 30%
- Midterm Exam 20%
- Final Exam 25%

EXAMS: One in-class midterm and one comprehensive final. For each exam, students may bring one double-sided A4 cheat sheet. No AI tools may be used in any exam.

QUIZZES: N/A

GRADE REBUTTAL: Any request to have a grade reconsidered must be submitted in writing to me or to the teaching assistant by email within 7 days after the graded work is returned. The request must identify the specific problem in grading and explain why student believes the grade is incorrect. Requests received after that deadline will not be considered.

ASSIGNMENTS: There will be 4 assignments, each worth 5% of the final grade (20% in total). Each assignment is due 7 days after it is released. Late submissions receive a 5% penalty per day, up to 30%, and nothing is accepted more than 7 days late.

PROJECTS: The course project runs over the whole semester in teams of 3 to 4 students, and the topic is fully open, e.g. predicting on a dataset, reproducing a research result, or shipping an application. Every project must (1) train at least one model itself, and calling a ready-made API alone does not count, while using the scikit-learn or PyTorch API is fine; (2) evaluate in correct ways, following the train, validation, and test protocol; and (3) be feasible and reproducible, with everything runnable on a laptop or free Colab. There are four progress presentations: Week 4 (proposal, report of up to one page), Week 9 (milestone, report of up to one page, ideally with a code repo), Week 13 (progress, report of up to one page with a code repo), and Week 18 (final, report of up to two pages with a reproducible code repo). All four presentations are mandatory, but only those in Weeks 9 and 18 are graded, worth 10% and 20% respectively.

ATTENDANCE: Attendance counts 5% of final grade. Students have three absences free of penalty and without prior notification, with each additional absence reducing the attendance grade.

MAKE-UP POLICY: N/A

MATERIAL COVERED: The sequence of the topics covered in this class is:

Week	Readings	Descriptions
1 (09/01)	ML Ch 1	Syllabus, Overview & Introduction
2 (09/08)	ML Chs 1 & 2	Machine Learning Basics
3 (09/15)	ML Chs 2 & 3	Regression & Over-fitting, Assignment 1 released
4 (09/22)	Project	Project Presentation I (Proposal)
5 (09/29)	ML Chs 3, 4 & 6	Classification
6 (10/06)	ML Ch 8	Ensemble Methods
7 (10/13)	ML Chs 9 & 10	Clustering & Dimension Reduction, Assignment 2 released
8 (10/20)	Review	Midterm Exam on 10/20
9 (10/27)	Project	Project Presentation II (Milestone due, graded, 10%)
10 (11/03)	UDL Chs 2–5	Deep Learning Basics

Week	Readings	Descriptions
11 (11/10)	UDL Chs 6, 7 & 9	Neural Network Training, Assignment 3 released
12 (11/17)	UDL Ch 10	Convolutional Neural Networks
13 (11/24)	Project	Project Presentation III (Progress)
14 (12/01)	Lecture Notes	Recurrent Neural Networks
15 (12/08)	UDL Ch 12	Transformers, Assignment 4 released
16 (12/15)	UDL Chs 17 & 18	Diffusion Models
17 (12/22)	Lecture Notes	Joint Embedding Predictive Architecture
18 (12/29)	Project	Project Presentation IV (Final project due, graded, 20%)
19 (01/05)	TBA	TBA
20 (01/12)		Final Exam Week

For readings, "ML" refers to Machine Learning (Zhou) and "UDL" refers to Understanding Deep Learning (Prince). For TBA, depending on the progress of teaching and students' feedback, I may give reviews of previous knowledge or discuss additional topics.