University of Alberta
CMPUT 626 Machine Learning and Practical Privacy
Section Number A2
Fall 2024

Instructor: Bailey Kacsmar
Office: Athabasca Hall
E-mail: kacsmar@ualberta.ca

Office Hours: TBA
Lecture Room and Time: Thursday 11am-1:50pm

This document is for **reference only in advance of the course start date.** *This version of the syllabus is a guideline for the course and not a contract. As such, its terms may be altered when doing so is, in the opinion of the instructor(s), in the best interests of the students*

Course Description:

This is a research-seminar style course that focuses on current research in the space of machine learning and privacy.The course is structured around four core modules, each examining a distinct facet of privacy through different research methodologies:

- User Privacy and HCI: Understanding user experiences and perceptions of privacy in the digital age.
- Empirical Privacy: Investigating vulnerabilities in anonymous data and machine learning models.
- Semantic Privacy: Exploring the theoretical foundations and algorithmic implementations of differential privacy.
- Legal Privacy: Navigating the complex ethical, legal, and policy frameworks governing data use.

We will examine the literature as well as open questions in what is the multifaceted landscape of privacy within data analysis and machine learning. We explore the intricate interplay between privacy policies, regulations, and cutting-edge privacy-preserving technologies, emphasizing their practical application in user-centred design.

Students will be expected to read research papers, participate in their evaluation, and learn how to assess what makes meaningful work in this field. They will write reviews of published research papers from the field, present a paper to the class in the style of a research seminar presentation, write evaluations of the paper (in styles other than a review), and execute a novel research project that they will write up and present at the end of the term. *Prerequisites: This course has no prerequisites.*

Course Objectives and Expected Learning Outcomes:

By the end of this course students should be able to:

- Explain the notion of privacy within the machine learning space.
- Analyze security and privacy of machine learning protocols.

- Evaluate research on privacy of machine learning and articulate advantages and limitations.
- Design and conduct research
- Explain the significance of the different modules with respect to privacy and machine learning.

Grading Scheme

- 20% Seminar style presentations as discussion lead (1-2 per term)
- 5% Quality of feedback on peers
- 15% Paper reviews (15 papers across the term)
- 10% Class participation
- 10% Project proposal
- 10% Project presentation
- 30% Final project report

Class	Topics	
September 5	Intro + Usable Privacy + HCI + Exercises	
September 12	Empirical Privacy + Exercises	
September 19	Semantic Privacy + Exercises	
September 26	Legal Privacy + Exercises	
Oct 3	Reading (3)	
Oct 10	Reading (3)	
Oct 17	Projects	
Oct 24	Reading (3)	
Oct 31	Reading (3)	
Nov 7	Reading (3)	
Nov 14	Reading Week	Reading Week
Nov 21	Reading (3)	
Nov 28	Reading (3) / Projects	
Dec 5	Project Presentations	
December ?	Reports due	

Table 1: Class Schedule Overview