Syllabus

CS598 Deep Learning for Healthcare

Instructor

Jimeng Sun (jimeng) (at illinois.edu), office hour: 9 PM - 10 PM CST, Wed

DL4H Text Book

We've updated our textbook for this year to include more up-to-date content. For any questions or concerns about the readings, please direct them to Piazza.

Generative AI Policy

Large Language Models (LLMs), including but not limited to ChatGPT, Claude, and Gemini, have demonstrated increasingly sophisticated programming capabilities. While independent completion of coding exercises without AI assistance remains optimal for developing fundamental programming skills and understanding core concepts, we acknowledge the growing integration of these tools in professional software development workflows.

Accordingly, this course **permits** the use of LLMs for assignments and actively encourages their integration in the final project. However, students must exercise critical judgment regarding the advantages and limitations of these tools. For the final project, students are required to implement and validate all LLM-generated outputs as part of their development methodology.

Course Description

Welcome to the Deep Learning for Healthcare course! In this course, we will explore deep learning (DL) methods and their applications in healthcare.

The course will include video lectures, self-guided labs, and homework assignments. During this time, you will learn about different DL and health applications topics, and develop practical experience in building deep learning models using healthcare data. Additionally, you will learn how to use "pyhealth," a package specifically designed for healthcare AI tasks. We encourage you to use pyhealth for your final project and, if possible, contribute to its improvement.

In the second half of the course will be a group project where you will work together to understand, replicate, and extend a recently published work in deep learning for healthcare. This project will give you hands-on experience in applying DL methods to real-world healthcare problems.

We hope you enjoy the course and learn a lot about deep learning and its applications in healthcare!.

Course Objectives

Upon completion of this course, you will be able to:

1. Understand and apply various deep learning models, including deep neural networks, convolutional neural networks, recurrent neural networks, autoencoders, attention models, graph neural networks, and deep generative learning.

2. Identify and implement different healthcare applications using DL methods, such as clinical predictive models, computational phenotyping, patient risk stratification, treatment recommendation, clinical natural language processing, and medical imaging analysis.

3. Develop practical experience in implementing various deep learning models on diverse medical data, using popular deep learning frameworks like PyTorch and data science software like Jupyter Notebook.

4. Gain hands-on experience in data preprocessing, data visualization, and model interpretation for healthcare applications.

5. Develop skills in critically evaluating and selecting appropriate DL models for healthcare applications.

By the end of this course, you will have a comprehensive understanding of deep learning methods and their applications in healthcare, as well as practical experience in implementing DL models using popular frameworks and tools. You will be well-prepared to work on real-world healthcare projects that involve DL and contribute to the development of innovative healthcare solutions.

Prerequisites

- Basic machine learning knowledge is helpful but not strictly required.*
- Strong programming skills in Python are required.
- A fair understanding of linear algebra and calculus is required.
- Sufficient system knowledge, such as using Linux and setting up programming environments on the cloud, is required.
- No knowledge in healthcare domain is required.

Textbook and Readings

Main reading:

Please find the PDF files in the "Textbook" section each week. This is a new textbook we are writing for the course.

Additional reading:

Introduction to Deep Learning for Healthcare by Cao Xiao and Jimeng Sun (link: https://link.springer.com/book/10.1007/978-3-030-82184-5#toc).

Machine Learning for Drug Discovery and Development by Tianfan Fu, Cao Xiao and Jimeng Sun (link:

https://ml4drug-book.github.io/)

Relevant research papers will be provided to you. You are also encouraged to conduct your own literature review, especially during the project phase.

Piazza

We will use Piazza as a private discussion board for this course, where you can ask questions and get answers from your classmates and the course staff. To join the forum, use (link: https://piazza.com/illinois/spring2025/cs598lho, code: hvkwqxfu0hb). Remember that all official course communication will happen through Piazza, make sure to check it often.

Before each office hour (OH), students can post their question or comment on other's questions on Piazza and choose a tag. The TA/instructor will choose to answer the questions with the corresponding tag and sort by popularity in descending order.

Elements of This Course

The following elements are in this course:

- **Video lectures:** They are recorded by the instructor to cover technical topics about deep learning and related healthcare applications.
- **Textbook:** Every week, we will release one chapter from the textbook.
- **Quizzes:** There are also quiz questions each week. Please do them as they help with your learning those topics.
- **Self-guided data science labs:** They are self-guided and should provide you with a step-by-step introduction to some basic concepts in deep learning. Though they are marked as optional, you are highly encouraged to finish the corresponding labs before attempting the homeworks.
- **self-guided pyhealth labs:** They are self-guided and should provide you with a step-by-step introduction to some basic concepts in pyhealth package modules. Students are encouraged to open the colab and as least read the input cells and the outputs to learn the capability of pyhealth. In the final project, these labs/tutorials can be very useful for data processing and module building references.
- **Homework assignments:** The homeworks help you gain practical deep learning programming skills. They are significantly harder than labs. Do not be surprised that homework assignments look very different from lectures and book chapters. Each homework will be released at least 2 weeks before the deadline. All homework is due at 11:59 PM Sunday Central Time. All homework assignments are automatically graded. You have limited attempts to submit each homework (see "Grading" section below). Submissions with the highest grade will be used.
- **Project:** a large group project at the second phase of the course.

- Each project team consists of **maximum of 2 students**. Individual projects are allowed though we encourage you to form a team due to the large amount of work.
- The project is about understanding, replicating, and extending some recently published work in deep learning for healthcare.
- **Detailed project requirements and paper pool will be released later.**
- Submissions will be done through gradescope.
- **Generative AI Form:** a Google form will ask you to submit relevant information related to the use of generative models, specifically that of LLMs like ChatGPT. This is to help us understand how it's being used and to help promote proper and safe LLM usage.
- **Project proposal:** a short clearly written description of your project that provides an overview of your project. It should demonstrate you have thought carefully about the paper you are planning to duplicate and communicate your understanding of the work and its importance to someone who most likely has not read the paper.
- **Project draft:** a complete write-up about the project including some preliminary experiment results. It should demonstrate you have fully understood your paper and realized all experimental setups and details. You are also expected to peer-review 3 other group's drafts.
- **Final submission**
  - **Project final report:** a well-written report that summarizes all the findings from your project. You should either reproduce the main experiments in a selected paper or report on your failed (but rigorous) attempt to reproduce a selected paper's main experiments.
  - **Project presentation:** a 5-min video presentation that clearly illustrates the main points in your work and the main results.
  - **Code:** A well-documented code repo with a README.md file describing the exact steps to run your code.

Grading

Component	**% of the total grade**
Homework assignments (1-5)	50% (10% each)
Project	50%

Homework will be graded by auto-graders. We allow **10 submissions** for homework 1, as many of you are new to Coursera's environment. For the rest of the homework, we allow up to **5 submissions**. Note that you will be able to do unlimited testing in your local environment. However, there are some hidden test cases, so you will need to submit your code to Coursera in order to see if you get full points. After each submission, you will receive feedback from the auto-grader and can modify your answers accordingly. Unlimited attempts for everything else. For anything staff-graded (mostly project-related works), we will take the last submission before the deadline.

Grade Cutoff (>=) (out of 100)	**Letter Grade**
99	A+
95	A
90	A-
85	B+
80	B
70	B-
60	C

Course Outline (all dates are tentative and on Coursera)

Week	**Topic**	**Activities**
1	Introduction	release HW1/2/3/4
2	Health data	release HW5
3	Machine learning basics	HW1 due
4	Deep Neural Networks (DNN)	
5	Embedding	HW2 due
6	Convolutional Neural Networks (CNN)	

Week	Topic	Activities
7	Recurrent Neural Networks (RNN)	HW3 due Project group formulation & paper selection due
8	Autoencoders	
9	Attention Models	
10	Graph Neural Networks	HW4 due Project Proposal Due
11	Memory Network (Transformers)	
12	Generative Models	HW5 due
13-14	Project activities	Project draft & peer-review due
15	Hand-in report and video presentation	Project final submission due

Late Policy

No late submission is allowed for project-related assignments. For all the other assignments, we apply a 10% late penalty per day.

Student Code and Policies

A student at the University of Illinois at the Urbana-Champaign campus is a member of a University community of which all members have at least the rights and responsibilities common to all citizens, free from institutional censorship; affiliation with the University as a student does not diminish the rights or responsibilities held by a student or any other community member as a citizen of larger communities of the state, the nation, and the world. See the University of Illinois Student Code for more information. Academic Integrity All students are expected to abide by the campus regulations on academic integrity found in the Student Code of Conduct. These standards will be enforced and infractions of these rules will not be tolerated in this course. Sharing, copying, or providing any part of a homework solution or code is an infraction of the University's rules on academic integrity. We will be actively looking for violations of this policy in homework and project submissions. **Any violation will be punished as severely as possible with sanctions and penalties typically ranging from a failing grade on this assignment up to a failing grade in the course, including a letter of the offending infraction kept in the student's**

permanent university record. Again, a good rule of thumb: Keep every typed word and piece of code your own. If you think you are operating in a gray area, you probably are. If you would like clarification on specifics, please contact the course staff.

Disability Accommodations

Students with learning, physical, or other disabilities requiring assistance should contact the instructor as soon as possible. If you're unsure if this applies to you or think it may, please contact the instructor and Disability Resources and Educational Services (DRES) as soon as possible. You can contact DRES at 1207 S. Oak Street, Champaign, via phone at (217) 333-1970, or via email at disability@illinois.edu.

Mental Health Service

Diminished mental health, including significant stress, mood changes, excessive worry, substance/alcohol abuse, or problems with eating and/or sleeping can interfere with optimal academic performance, social development, and emotional wellbeing. The University of Illinois offers a variety of confidential services including individual and group counseling, crisis intervention, psychiatric services, and specialized screenings at no additional cost. If you or someone you know experiences any of the above mental health concerns, it is strongly encouraged to contact or visit any of the University's resources provided below. Getting help is a smart and courageous thing to do -- for yourself and for those who care about you.

Counseling Center: 217-333-3704, 610 East John Street Champaign, IL 61820

McKinley Health Center:217-333-2700, 1109 South Lincoln Avenue, Urbana, Illinois 61801