**UC BERKELEY SCHOOL OF PUBLIC HEALTH**

**PH 290**

# Reliable Generative Artificial Intelligence for Health

Course Syllabus | Fall 2026 | 4 units | Updated October 15, 2026

## Course Information

| | |
|---|---|
| **Lecture** | Mondays, 1-3 pm.; August 31-November 30, 2026 |
| **Lab** | Thursdays, 3-4 pm. |
| **Locations** | Lecture: BWWB1213; Lab: ETCH3109 |

## Instructor Information

| | |
|---|---|
| **Instructor** | Jingshen Wang, Associate Professor of Biostatistics |
| **Phone** | (510) 643-8083 |
| **Email** | jingshenwang@berkeley.edu |
| **Office hours** | By appointment |
| **GSI** | Olivia Zhang, yueshanzhang@berkeley.edu |

## Course Description

Reliable Generative Artificial Intelligence for Public Health (PH 290, 4 units) gives graduate students across campus an overview and hands-on practice of the engineering foundation of modern AI systems. The course connects machine-learning and LLM foundations with health data, clinical workflows, prompt engineering, vibe coding, clinical phenotyping with LLMs, model selection, and post-training. Later sessions examine retrieval-augmented generation, agents and tools, multimodal and longitudinal applications, AI-generated health data, and evaluation in real-world clinical workflows. Lectures, labs, and guest lectures connect technical choices to public-health and clinical practice, with attention to validity, uncertainty, privacy, bias, human oversight, and practical usefulness.

**Course unit value:** 4 units

### Prerequisites

- Familiarity with Python.
- Calculus, linear algebra, and introductory probability or statistics.
- An introductory machine-learning or data science course will be helpful.

## Course Learning Objectives

After successfully completing this course, students will be able to:

- Explain the progression from supervised machine learning to foundation models, including tokenization, embeddings, transformers, pretraining, generation, and the distinction between training and inference.
- Translate a health or public-health problem into an AI task with defined inputs, outputs, intended users, constraints, and evaluation criteria.
- Design and evaluate prompts and context; use vibe coding to build AI tools.
- Compare model selection (e.g., open/close models, lightweight models, JEV) and model adaptation strategies (including prompting, retrieval, supervised fine-tuning, parameter-efficient adaptation, distillation, and preference-based post-training).
- Describe how multimodal, longitudinal, and agentic systems combine models, data, tools, memory, permissions, and reusable skills.
- Develop evaluation pipeline and failure modes for AI systems used in clinical and public-health workflows.

## Course Format

The course includes a two-hour Monday lecture (1:00–3:00 p.m. in BWWB1213) and a separate one-hour Thursday lab (3:00–4:00 p.m. in ETCH3109). Lectures combine instruction, demonstrations, and guest discussions. Labs provide guided practice; the lab schedule includes Rita Lyu’s October 8 guest session on LLM fine-tuning.

## Course Schedule

The schedule lists twelve instructional meetings and the final-project presentation session. It reflects current course materials and guest arrangements as of October 15, 2026. Topics and guest-session details may be adjusted; updates, readings, and lab materials will be posted in bCourses.

| Meeting | Topic | Content and health perspective |
|---|---|---|
| **1 \| Aug. 31** | **From machine learning to generative AI** | Supervised learning, neural networks, foundation models, and LLMs. Tokens, embeddings, transformers, pretraining, and generation.<br>Health case studies introduce the capabilities and limits of generative AI in clinical workflow. |
| **— \| Sept. 7** | **No class — Labor Day** | University holiday. |
| **2 \| Sept. 14** | **From health data to AI tasks** | Model training and inference; multimodal models; inputs, outputs, intended users, and decisions within a clinical workflow.<br>Emergency department and triage examples connect model performance to task-specific evaluation. |
| **3 \| Sept. 21** | **Prompt engineering and vibe coding** | Task specification, relevant context, objective input, and iterative feedback. AI-assisted vide coding, web search and source selection, context limits, and AI as a thought partner.<br>Health case study examine wording sensitivity in medical tasks for AI to produce robust answer. |
| **4 \| Sept. 28** | **Clinical phenotyping and model selection** | Reliable methods to convert clinical notes into structured phenotype data with LLM; compare and introduce LLM models, size, context, modalities, adaptation, and deployment. |
| **5 \| Oct. 5** | **Model post-training, Part 1** | MedGemma; supervised fine-tuning, full fine-tuning, LoRA and QLoRA; training and validation data, memory, batching, and stopping.<br>Guest lecture: Marcell Borhi, Lotus Health AI. |
| **6 \| Oct. 12** | **Model post-training, Part 2** | Teacher–student distillation; off-policy and on-policy training; reinforcement learning, reward models, and direct preference optimization. Health case studies examine how training pairs in clinical data shape model behavior.<br>Guest lecture: Will Ma, Hope AI, 2:00–3:00 p.m. |
| **7 \| Oct. 19** | **Retrieval-augmented generation and grounding** | Document preparation, embeddings, retrieval, reranking, context assembly, and citations. Assess whether retrieved evidence supports an answer and where retrieval-based systems fail.<br>Guest lecture: Dan Zeltzer, 2:00–3:00 p.m. |
| **8 \| Oct. 26** | **Agents, tools, MCP, and reusable skills** | Planning, tool use, Model Context Protocol (MCP), memory, orchestration, permissions, and reusable skills. Compare prompt-only, tool-based, and agentic workflows. |

| Meeting | Topic | Content and health perspective |
|---|---|---|
| **9 \| Nov. 2** | **Multimodal and longitudinal applications** | Combine multimodal data (text, tables, images, signals, and information) collected over time. Connect model inputs and outputs to clinical and public-health workflows.<br>Guest lecture: Tomás León, California Department of Public Health; session details to be announced. |
| **10 \| Nov. 9** | **AI-generated health data** | Synthetic health data, digital twins, and AI-generated labels. |
| **11 \| Nov. 16** | **Evaluation: validity, calibration, and robustness** | Benchmarks, human labels, external validation, calibration, and subgroup performance. |
| **12 \| Nov. 23** | **Clinical deployment and human judgment** | Decision thresholds, uncertainty, abstention, escalation, automation bias, and human judgment.<br>Guest discussion: Sanjay Juneja, on clinical AI applications, limitations, and adoption. |
| **Final \| Nov. 30** | **Final-project presentations** | Individual or team presentations, 5–10 minutes each. Final report and code, if used, due November 30. |

## Course Grading

Grading is based on the following:

- Participation - 10%
- Homework assignments - 40% total (20% each)
- Final project - 50% total

# Course Materials

## bCourses and Nova

The bCourses site contains the syllabus, announcements, readings, assignment instructions, and course resources. Nova provides interactive lecture materials, note-taking tools, and project activities used in class. Links to Nova resources and PDF materials will be provided through bCourses. Students should check announcements and assignment pages regularly for updates.

**Required materials:** None

# Course Requirements

## Participation

Participation is measured through student engagement in class and lab activities.

## Group Homework

Two group homework assignments account for 40% of the course grade (20% each). Assignment 1 examines how medical terminology versus everyday language affects AI responses and is due October 22, 2026. Assignment 2 instructions and its deadline will be posted in bCourses.

## Final Project

The final project accounts for 50% of the course grade. Students may work individually or in groups of two or three on a project in any AI-related field. Projects may build on all or part of a course assignment, provided students make their own changes or extensions. Submit one final report per project and code, if used, by November 30, 2026, and give a 5–10-minute presentation that day. Complete the project sign-up sheet before the presentations. See the October 15 final-project announcement and Final Project slides in bCourses for details; the report submission location will be announced there.

# Course Communication

## Announcements

Announcements will be posted on the bCourses home page. Students should check regularly for updates.

## Course Mail

Course announcements will also be sent through Canvas notifications. Students should check Course Mail and their configured notification destination.

# Policies

## Due Dates

Assignment 1 is due October 22, 2026. The final-project report and code, if used, are due November 30, 2026; presentations take place that day. Assignment 2 instructions and its deadline will be posted in bCourses. Please communicate with the instructors through Canvas Course Mail before a deadline if you will be unable to meet it.

## Late Assignments

Late work is accepted for up to 72 hours after the deadline with a 10 percentage-point deduction for every 24-hour block. After 72 hours, the assignment receives no credit unless an extension was arranged before the deadline or an approved accommodation applies. Extensions are granted for documented emergencies or university-sanctioned obligations.

## Attendance

Attendance is encouraged.

## Technology

Students may use laptops or other technologies for note-taking and course activities.

## Correspondence

For email correspondence, please put PH 290 in the subject heading. We will try to respond within three working days.

# Berkeley Public Health Course Policies

Descriptions and campus links for Berkeley Public Health policies on disability support, religious accommodations, course evaluations, and academic integrity are available at: Berkeley Public Health course policies