UC BERKELEY SCHOOL OF PUBLIC HEALTH

PH 290

Reliable Generative Artificial Intelligence for Public Health

Course Syllabus | Fall 2026 | 4 units

Course Information

Lecture	Mondays, 1:00-3:00 p.m.; August 31-November 30, 2026
Lab	Thursdays, 3:00-4:00 p.m.
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 a comprehensive, practice-oriented foundation in designing, adapting, and evaluating generative-AI systems for population-health applications. The course moves from machine-learning and transformer foundations through prompt engineering, model adaptation and post-training, retrieval-augmented generation, multimodal and longitudinal modeling, human-AI decision-making, and agentic systems with tools and reusable skills. A special data-science component examines synthetic health data, digital twins, AI-generated labels, bias, provenance, causal claims, and feedback loops. Across the course, students evaluate validity, calibration, external performance, robustness, hallucinations, adversarial failures. Integrated lab sessions use public-health and clinical examples while emphasizing uncertainty, human oversight, patient autonomy, equity, traceability, and clinical or public-health utility.

Course unit value: 4 units

Prerequisites

- Familiarity with Python for processing data files, manipulating data frames, and using open-source machine-learning libraries.
- Calculus, linear algebra, and introductory probability or statistics.
- An introductory machine-learning or data science course covering supervised learning, model validation, and basic ethics.

Course Learning Objectives

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

- Explain the progression from supervised machine learning and neural networks to foundation models and modern generative AI systems, including tokenization, embeddings, transformers, pretraining, and generation.
- Translate a health or public-health problem into a well-defined AI task with appropriate data, labels, intended use, outputs, constraints, and evaluation criteria.
- Design and compare zero-shot, few-shot, retrieval-based, fine-tuned, and post-trained approaches, and justify model choices using performance, cost, latency, privacy, and reproducibility considerations.
- Understand multimodal, longitudinal, and agentic workflows that combine models, tools, MCPs, memory, permissions, guardrails, and reusable skills.
- Develop evaluation plans for LLM deployment in real-world workflow.
- Communicate model capabilities, limitations, uncertainty, and safety considerations to clinicians, policymakers, public-health practitioners, patients, and community partners.

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).

Course Schedule

The schedule below lists the twelve instructional meetings and the final-project presentation session. 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 the development of LLMs. Tokens, embeddings, transformers, pretraining, generation. Health examples introduce what current generative AI can do beyond conventional predictive machine learning.
- | Sept. 7	**No class - Labor Day**	University holiday.
2 | Sept. 14	**From health data to AI tasks**	Health data sources, missingness, labels, input and output definitions, ground truth, and evaluation criteria.

Meeting	Topic	Content and health perspective
3 | Sept. 21	**Prompt engineering and in-context learning**	Prompt; zero-shot, one-shot, and few-shot prompting; context construction; structured outputs; prompt templates, versioning, sensitivity, and reproducibility. Health examples examine biased or unrepresentative demonstrations, privacy, unsupported clinical claims, uncertainty, and escalation.
4 | Sept. 28	**Model selection and adaptation**	Open versus closed models, hosted versus local deployment, fine-tuning, parameter-efficient adaptation, cost, latency, and privacy.
5 | Oct. 5	**Changing model behavior**	Policy distillation, supervised fine-tuning, preference data, reward models, RLHF and RLAIF, DPO, and reinforcement learning. Health considerations include whose preferences define the reward, clinician-patient disagreement.
6 | Oct. 12	**Retrieval-augmented generation**	Embeddings, document ingestion, retrieval, reranking, prompt assembly, citations. Applications ground outputs in clinical guidance.
7 | Oct. 19	**Multimodal and longitudinal AI**	Combining text, tables, images, signals, and time-series data for cross-modal combinations.
8 | Oct. 26	**Human-AI decision-making in health**	Diagnosis, prognosis, triage, treatment recommendations, surveillance, and resource allocation; prediction versus causal claims; decision thresholds; uncertainty, abstention, and escalation; automation bias, alert fatigue, clinician judgment, and patient autonomy.
9 | Nov. 2	**Agents, tools, MCPs, and reusable skills**	Agents, skills, planning, memory, orchestration. Students compare the same task using a prompt alone, a prompt plus tools, and an agent with tools and a reusable skill.
10 | Nov. 9	**Data science with AI-generated health data**	Digital twins, AI-generated. LLM bias, provenance, and feedback loops.
11 | Nov. 16	**Evaluation I - validity and usefulness**	Benchmarks, calibration, human baselines, external validation.
12 | Nov. 23	**Evaluation II - failure and stress testing**	Hallucinations, robustness, prompt injection, data and benchmark leakage, adversarial testing, and red-teaming. Health cases examine unsafe recommendations, misinformation, missed populations, automation bias, acceptable risk, incident response, and recovery.
Final | Nov. 30	**Final-project presentations**	

Course Grading

Grading is based on the following:

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

Course Materials

bCourses Website

To access the course website, go to bCourses at bcourses.berkeley.edu. The course site will contain required and optional readings, the syllabus, assignment descriptions, starter materials, and additional course resources. Any changes will be reflected in the assignment section of the site.

Required materials: None

Course Requirements

Participation

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

Group Homework

Two group homework assignments (2 x 20%) translate lecture concepts into practice.

Final Project

The final project (50%) asks teams of three to build, document, evaluate, and present a generative-AI system addressing a real public-health need. Work begins with a one-page proposal, continues through an error-analysis and risk checkpoint, and culminates in a clean repository, model card, bias and safety report, and team presentation on November 30, 2026.

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

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