Applied Artificial Intelligence in Public Health Practice and Research HSMP 6635

Course details

Instructor:

Marcelo Coca Perraillon, PhD Department of Health Systems, Management & Policy Colorado School of Public Health Building 500, Rm. E3313 E-mail: marcelo.perraillon@cuanschutz.edu Office Hours: Available by appointment

Course Title:Artificial Intelligence in Public Health Practice and ResearchCredit Hours:3 credit hoursMeeting Time:Tuesday and Thursday at 5:30 pm. Online option availableMeeting Place:TBD

Course Description:

This course introduces artificial intelligence (AI) to public health students and researchers. The course is tailored for master's students. It explores the foundational concepts of AI focusing on Large Language Models (LLMs) and Generative AI, without requiring prior programming or advanced mathematical knowledge. The course emphasizes practical applications, equipping students to leverage existing AI tools for tasks relevant to public health research and practice, such as literature review, data summarization, qualitative and quantitative data analysis, communication, and tutoring. We will critically examine current AI capabilities, limitations, governance issues, data privacy concerns, and the significant economic and ethical implications of AI deployment in healthcare and public health, including alignment problems and impacts on the job market. The goal is to empower students to thoughtfully integrate AI into their future work while understanding its broader context and limitations. A comprehensive understanding of how generative AI models work will provide students with the necessary knowledge to assess current model capabilities, separating hype from reality. We will discuss examples of how to use LLMs on a local machine with open-source models (e.g., Llama 4, Gemma3). These options protect privacy since the data is not transmitted to AI companies. However, it requires well-equipped computers.

Student Audience and Prerequisites:

This course is designed for Master's-level students in Public Health. No background in computer programming or linear algebra is required. An interest in health policy, public health research, and the application of new technologies is assumed.

Learning Objectives (LO):

Upon successful completion of this course, students will be able to:

- 1. Explain core concepts, terminology, and the evolution of AI, particularly Large Language Models and Generative AI. (LO1)
- 2. Describe the approaches for "prompt engineering" (i.e., how to help generative AI modes understand instructions to get the desired output) and the different ways to interact with AI (e.g., chats, APIs, agents, RAGs, local LLMs) assessing privacy implications. (LO2)
- **3**. Evaluate current research assessing AI capabilities to have a more realistic view of models ' strengths and shortcomings. (LO3)
- 4. Critically evaluate the ethical considerations, governance challenges, data privacy risks, and societal implications (including job market impacts) of deploying AI in public health contexts. (LO4)
- 5. Demonstrate proficiency in using specific AI tools relevant to healthcare/research to perform tasks such as summarizing literature, assisting with qualitative data analysis, and generating text for research purposes. (LO5)
- 6. Analyze case studies of AI applications in public health, medicine, and research, identifying both the potential benefits and limitations (including challenges like causal inference with observational data). (LO6)

Course Competencies:

MPH 3. Analyze quantitative and qualitative data using biostatistics, informatics, computer-based programming and software, as appropriate.

MPH-HSMP 10. Develop a plan for engaging stakeholders and building coalitions in the development or implementation of health policy.

- 1. Accurately select, use, and interpret statistics commonly used in health services research
- 2. Apply and use appropriate study designs and methods to address research questions/hypotheses
- 3. Utilize healthcare databases and other information technologies used in research

Required Materials and Tools:

- Accounts for specific AI tools (details provided below). Many tools offer free tiers or trials, but some advanced features may require subscriptions (not mandatory for course completion). Copilot is available on campus Windows machines. We will prioritize freely accessible tools. Examples include:
 - General Purpose LLMs (e.g., ChatGPT, Gemini, Claude, Copilot)
 - Open-source LLMs in local machines (e.g., Gemma, Qwen, Llama)
 - Research Assistants (e.g., Elicit, Scite, Google and OpenAl's Deep Research)
 - Document Analysis Tools (e.g., NotebookLM, chats, RAGs)
- Weekly Readings: Key readings (articles, reports, book chapters) will be assigned each week via Canvas. See the schedule below.
- **Required Books**: Mollick, E. (2024). Co-intelligence: Living and working with Al. Penguin.

Optional books:

Selected chapters and materials from these books will be posted on Canvas. They are excellent sources for students who wish to gain a more technical understanding of AI tools (some require knowledge of Python).

- Raschka, S. (2024). Build a Large Language Model (From Scratch). Simon and Schuster.
- Huyen, C. (2024). AI Engineering. O'Reilly Media.
- Russell, S. and Norvig, P. (2021). Artificial Intelligence: A Modern Approach (4th ed.).
- Bender, E. and Hanna, A. (2025) *The AI Con: How to Fight Big Tech's Hype and Create the Future We Want*. Harper.

Assignments and Grading:

Understanding and application of the material will be assessed through bi-weekly homework assignments, a final project proposal, and a final project.

- Homework Assignments (7 total): 60% of final grade (approx. 8.57% each)
 - Assignments will involve using specified AI tools, analyzing case studies, evaluating AI outputs, using AI tools for public health and practice, or researching AI applications/implications in public health. Due dates in schedule.
- Final Project Proposal (Due Week 12): 10% of final grade
 - A 2-page proposal outlining a project where you will apply Generative AI to a specific public health or health policy task (e.g., a systematic approach for generating ideas for a policy proposal, designing a grant, creating patient education materials, summarizing public health legislation, qualitative analysis, data analysis, creating a customized tutor for a public health subject).
 - The proposal must include: Problem statement and goal, proposed AI tool(s), use strategy (including an outline of a detailed prompt or multi-step plan for interacting with the AI), expected output/deliverable, and potential challenges and ethical considerations, including details about any input data to ensure privacy.
- Final Project (Due end of Week 15): 30% of final grade [LO1-LO5]
 - Execution of the approved project proposal. Deliverables vary depending on the project.
 However, the final report should include a reflective summary (1-2 pages) on the process, the effectiveness of the AI tool, challenges encountered, and critical evaluation of the final output, including a process to ensure that the information generating by the AI is correct.

Grades. The following grading scale will be used for final grades:

А	94.00-100	B-	80.00-83.99	D+	67.00-69.99
A-	90.00-93.99	C+	77.00-79.99	D	64.00-66.99
B+	87.00-89.99	С	74.00-76.99	D-	60.00-63.99

No late assignments are allowed unless you have a valid justification, which needs to be communicated *before* the due date. 20 points deducted per day the assignment is late.

Summary of Homework and LOs

- Homework #1 Assigned (Due end of Week 2): Defining terms, identifying AI types, basic LLM interaction with prompts applied to public health. (Worth ~8.57%) [LO1]
- Homework #2 Assigned (Due end of Week 4): Prompting exercises, privacy policy analysis, evaluating AI text for accuracy/bias. (Worth ~8.57%) [LO2]
- Homework #3 Assigned (Due end of Week 6): Applying a framework for literature review and summarization using LLMs; critical reflection. (Worth ~8.57%) [LO5]
- Homework #4 Assigned (Due end of Week 8): Analyze AI in health case studies (benefits, risks, ethics). Exploring qualitative snippets with AI. (Worth ~8.57%) [LO3]
- Homework #5 Assigned (Due end of Week 10): Critically evaluate a study using AI for public health data analysis (focus on methods, claims, bias). (Worth ~8.57%) [LO4]
- Homework #6 Assigned (Due end of Week 12): Ethical analysis of complex AI-in-health scenario; discuss alignment risks. (Worth ~8.57%) [LO6]
- Homework #7 Assigned (Due end of Week 14): Critically evaluate the research on the impact of AI on jobs. (Worth ~8.57%) [LO6]

Accessibility Statement: My goal is for you to fully participate and have a great experience in this course. If you plan to use assistive technology or need any other accommodations, such as alternative formats to support your learning, please let me know so we can discuss and make sure you have the best possible support.

Workload: Students should expect to spend approximately 3 hours per week, per credit, outside of class, completing readings and assignments.

Inclusive Learning Environments:

In this class, we will work together to develop a learning community that is inclusive and respectful. Our diversity may be reflected by differences in race, age, sexual orientation, gender identity and expression, religion/spirituality, ability, socioeconomic background, and myriad other social identities and life experiences. In a diverse community, the goal of inclusiveness encourages and appreciates expressions of different ideas, opinions, and beliefs so that conversations and interactions are opportunities for intellectual and personal enrichment.

A dedication to inclusiveness requires respecting what others say, their right to say it, and the thoughtful consideration of others' communication. Both speaking up and listening are valuable tools for furthering thoughtful and enlightening dialogue. Respecting one another's individual differences is critical in transforming a collection of diverse individuals into an inclusive and collaborative learning community. We will hold ourselves and one another accountable, which includes bringing attention to times when microaggressions or macroaggressions happen in a classroom. Our core commitment shapes our core expectations for behavior inside and outside of the classroom. We encourage students to review the ColoradoSPH Equity, Diversity, and Inclusion Common Language and Commitment Statement

https://www1.ucdenver.edu/docs/librariesprovider151/default-document-library/edi-commitmentand-terms-5-11-21.pdf?sfvrsn=804479ba 0

Title IX: Non-Discrimination and Sexual Misconduct, Intimate Partner Violence, and Stalking

Non-Discrimination:

The ColoradoSPH and <u>University of Colorado Non-Discrimination Policy</u> prohibits discrimination on the basis of race, color, national origin, sex, age, disability, pregnancy, creed, religion, sexual orientation, veteran status, gender identity, gender expression, political philosophy or political affiliation in admission and access to, and treatment and employment in, its educational programs and activities.

Sexual Misconduct, Intimate Partner Violence, and Stalking:

The ColoradoSPH and University of Colorado <u>Sexual Misconduct</u>, <u>Intimate Partner Violence</u>, <u>and Stalking</u> <u>Policy</u> prohibits conduct including sexual assault, dating violence, domestic violence, Title IX stalking, stalking, sexual exploitation, Title IX harassment, hostile environment, Title IX quid pro quo sexual harassment, and quid pro quo sexual harassment.

ColoradoSPH Partner Campus Title IX Offices and Contact Information:

Incidents of discrimination, sexual misconduct, intimate partner violence, and stalking should be reported to the **Title IX office of the university where the incident occurred**. Incidents involving microaggressions or incidents that may not otherwise rise to the level of a policy violation, may also be reported to the appropriate university Title IX office listed below. Please refer to the CU Anschutz campus Office of Equity website for a self-learning guide about <u>microaggressions</u>.

If you have any questions on clarity related to the reporting of incidents, please contact the **ColoradoSPH Title IX Liaison**, Dr. Madiha Abdel-Maksoud, MD, PhD. You can reach Dr. Abdel-Maksoud at <u>Madiha.Abdel-Maksoud@cuanschutz.edu</u>.

CU Anschutz Campus: On the CU Anschutz campus, please contact the <u>Office of Equity</u>. The Office of Equity staff, including the University's Title IX Coordinator, may be reached at (303) 315-2567 or <u>equity@ucdenver.edu</u>.

Colorado State University Campus: On the Colorado State University campus, please contact the <u>Office</u> <u>of Title IX Programs and Gender Equity</u>. The Office of Title IX Programs and Gender Equity staff may be reached at (970) 491-1715 or through email by reaching out to one of the listed <u>Title IX coordinators</u>.

University of Northern Colorado Campus: On the University of Northern Colorado campus, please contact the <u>Office Institutional Equity and Compliance</u>. The Office Institutional Equity and Compliance staff may be reached at (970) 351-4899 or <u>titleix@unco.edu</u>.

Academic Conduct Policy:

All students are expected to abide the Honor Code of the Colorado School of Public Health. Unless otherwise instructed, all of your work in this course should represent completely independent work. Students are expected to familiarize themselves with the <u>Student Academic Honor and Conduct Code</u>. Any student found to have committed acts of misconduct (including, but not limited to cheating, plagiarism, misconduct of research, breach of confidentiality, or illegal or unlawful acts) will be subject to the procedures outlined in the ColoradoSPH Honor Code.

Academic Integrity Course: The ColoradoSPH requires students to complete an academic honesty online course. If you have not yet completed the online course, you must do so by [date]. Once the course is completed, upload your completion certificate on Canvas. Instructions for completing the course are available on Canvas.

Academic Integrity and Use of AI Tools Policy:

At the Colorado School of Public Health, we hold academic integrity as a cornerstone of your education. While artificial intelligence (AI) programs, such as ChatGPT and other generative AI tools, are rapidly evolving and can offer valuable assistance in various tasks, they can stifle your own independent thinking and creativity. Unauthorized use of AI tools undermines the academic standards of this institution and may constitute academic misconduct. Therefore, unless explicitly permitted by the instructor, the use of AI programs to complete, assist with, or generate content for any assignments, projects, or exams is not allowed. Violations will be investigated thoroughly and may result in serious consequences, including receiving a failing grade for the assignment or course, or further disciplinary actions in line with the university's academic integrity policies.

To ensure students are truly grasping course concepts, if unauthorized use of AI programs to complete an assignment is suspected, 1:1 consultation may be scheduled, where students will be asked to elaborate upon the work they have submitted in real-time.

- 1. First Instance: If a student is unable to adequately explain their own submission, they will receive a reminder of their responsibility to actively engage in their learning. While no institutional penalty will be imposed, no points will be awarded for the assignment in question.
- 2. Second Instance: A second inability to explain their work will lead to a loss of points for the assignment, a grade reduction for the course, and a referral to the Office of Academic and Student Affairs to assess for a pattern of bad faith actions.
- 3. Subsequent Instances: Additional violations of this policy will result in an automatic failure of the course.

This policy aims to encourage deep engagement with the material, ensuring that students are prepared for any future encounters with similar material in the field. Finally, this policy is evolving and subject to change as we continue to assess the most effective ways to integrate Al tools into our learning environment while upholding academic integrity.

Please consult your instructor if you are uncertain about whether an assignment permits the use of AI tools.

Accommodations for Disabilities: Virtual and In-Class:

The University of Colorado Anschutz Medical Campus is committed to providing equitable access to our programs for students with disabilities (e.g., psychological, attentional, learning, chronic health, sensory, and physical).

To engage in a confidential conversation about the process for requesting reasonable accommodations in the classroom and clinical settings please contact The Office of Disability, Access, and Inclusion at: <u>disabilityaccess@cuanschutz.edu</u> or begin the process via the <u>website</u> Accommodations are not provided retroactively, therefore, students are encouraged to begin this process early.

Mental Health Services:

Please visit the <u>Office of Student Affairs – Student Health Promotion</u> website OR the <u>Department of</u> <u>Psychiatry</u> website for information on mental health services.

Religious Observances:

The University of Colorado Denver, Anschutz Medical Campus has a legal obligation to accommodate students who must be absent from an educational activity in order to observe religious holidays or other observances. Students should speak to the faculty member to request accommodations for religious observances in advance during the first week of class. Requests received by faculty must be kept confidential and should be considered unless they create an undue hardship. If the student and faculty member cannot agree on an accommodation, the matter should be referred to the Associate Dean for Academic and Student Affairs for resolution.

CU Systemwide Lands Recognition Statement:

As we gather, we honor and acknowledge that the University of Colorado's four campuses are on the traditional territories and ancestral homelands of the Cheyenne, Arapaho, Ute, Apache, Comanche, Kiowa, Lakota, Pueblo and Shoshone Nations. Further, we acknowledge the 48 contemporary tribal nations historically tied to the lands that comprise what is now called Colorado.

Acknowledging that we live in the homelands of Indigenous peoples recognizes the original stewards of these lands and their legacies. With this land acknowledgment, we celebrate the many contributions of Native peoples to the fields of medicine, mathematics, government and military service, arts, literature, engineering and more. We also recognize the sophisticated and intricate knowledge systems Indigenous peoples have developed in relationship to their lands.

We recognize and affirm the ties these nations have to their traditional homelands and the many Indigenous people who thrive in this place, alive and strong. We also acknowledge the painful history of ill treatment and forced removal that has had a profoundly negative impact on Native nations.

We respect the many diverse Indigenous peoples still connected to this land. We honor them and thank the indigenous ancestors of this place. The University of Colorado pledges to provide educational opportunities for Native students, faculty and staff and advance our mission to understand the history and contemporary lives of Native peoples.

Course Schedule

The course is divided into 5 modules:

- Module 1: Overview of the class and AI foundations (Weeks 1-2)
- Module 2: Interacting with AI (Weeks 3-4)
- Module 3: AI for research and summarization (Weeks 5-6)
- Module 4: AI for Data Analysis (Weeks 7-10)
- Module 5: Future, Impact on Jobs, Summary (Weeks 11-15)

Module 1: Overview of the class and AI foundations (Weeks 1-2)

• Week 1, Session 1: What is AI?

- Topics: Course overview, AI definitions, history, type of tools, relevance to public health and healthcare policy.
- Readings:
 - Mollick, E. (2024), Chapter 1.
 - McCarthy, J. (2007). What is Artificial Intelligence? http://jmc.stanford.edu/articles/whatisai/whatisai.pdf

• Week 1, Session 2: Predictive models basics and uses in healthcare

- Topics: Overview of predictive models, from linear and logistic regression to neural networks; other approaches to uncovering patterns in data. Reducible and irreducible error. Understanding the information encoded in parameters.
- Readings:
 - Beam, A. L., & Kohane, I. S. (2018). Big Data and Machine Learning in Health Care. JAMA, 319(13), 1317–1318. (Short perspective)
 - Shivade, C., Raghavan, P., Fosler-Lussier, E., Embi, P. J., Elhadad, N., Johnson, S. B., & Lai, A. M. (2014). A review of approaches to identifying patient phenotype cohorts using electronic health records. *Journal of the American Medical Informatics Association*, 21(2), 221-230. (Focus on introduction and discussion)
 - Explore: OpenAI, Google AI, Anthropic, and Copilot models (more instructions during class).

• Week 2, Session 3: Large Language Models (LLMs)

- Topics: From natural language processing to LLMs. Conceptual workings (transformers, training data). Pre- and post-training. Models predicting the next word (token). Strengths and limitations. Examples from earlier models to show how the approach improved as a function of model size (scalability), resulting in the "emergence" of capabilities.
- Readings:

- Mollick, E. (2024), Chapter 2.
- Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., Gutierrez, L., Tan, T. F., & Ting, D. S. W. (2023). Large language models in medicine. *Nature Medicine*, 29(8), 1930-1940.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., ... & Amodei, D. (2020). Scaling laws for neural language models. arXiv preprint arXiv:2001.08361. (Introduction [1 and 1.1] and conclusion)
- Can AI Scaling Continue Through 2030? <u>https://epoch.ai/blog/can-ai-scaling-continue-through-2030</u> (Introduction)

• Week 2, Session 4: Introduction to Generative AI

- Topics: Defining generative AI (text, image, video) and more details on LLM models posttraining. Understanding the importance of post-training in going from text replication to useful tools: fine-tunning, reinforced learning, distillation, and other techniques.
- Readings:
 - Huyen C. (2024). Chapter 2 (Up to Sampling strategies section).
 - Juzek, T. S., & Ward, Z. B. (2024). Why Does ChatGPT "Delve" So Much? Exploring the Sources of Lexical Overrepresentation in Large Language Models. arXiv preprint arXiv:2412.11385. <u>https://aclanthology.org/2025.coling-main.426.pdf</u> (Introduction and discussion)
 - Cascella, M., Montomoli, J., Bellini, V., & Bignami, E. (2023). Evaluating the Feasibility of ChatGPT in Healthcare: An Analysis of Multiple Clinical and Research Scenarios. *Journal of Medical Systems*, 47(1), 33. (Focus on applications).
- **Homework #1 Assigned (Due end of Week 2):** Defining terms, identifying AI types, basic LLM interaction with prompts applied to public health. (Worth ~8.57%)

Module 2: Interacting with AI (Weeks 3-4)

- Week 3, Session 5: Prompt Engineering Fundamentals
 - Topics: Review of HW1. Effective prompting principles. Iteration. Examples for public health tasks. System versus user prompt and changing model hyper-parameters (e.g., temperature, sampling schemes). Extending and improving prompts with "reasoning" (chain-of-thought) and other advanced techniques, including models with "thinking" and "extended thinking."
 - Readings:
 - Mollick, E. (2024), Chapters 2 and 3
 - Boonstra L. (2025). Prompt Engineering. Google. (Prompt Engineering section starting on page 13)
 - White, J., Fu, Q., Hays, S., et al. (2023). A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT. *arXiv preprint arXiv:2302.11382*. (Skim patterns relevant to asking questions, summarization)
 - Detecting misbehavior in frontier reasoning models. <u>https://openai.com/index/chain-of-thought-monitoring/</u>

Week 3, Session 6: Modes of interaction and privacy

- Topics: Description of interaction methods beyond chats: APIs, agents, RAG, local LLMs conceptual. Data privacy implications (cloud vs. local, terms of service). HIPAA considerations and local LLMs
- *Readings and activity:*
 - Mehra, A. (2024). Retrieval Augmented Generation: Keeping LLMs Relevant and Fact-Based. Salesforce AI Research Blog. (Conceptual overview of RAG)
 - Introduction to NotebookLM as a type of "RAG" with a case study on health policy analysis
 - Examples from local LLMs in class and API calls to Claude, Gemini, and OpenAI

• Week 4, Session 7: AI Governance and risk Management

- Topics: Al governance concepts. Regulation challenges. Organizational policies. Risk identification/mitigation and the added considerations regarding healthcare data. Understanding consent and experimentation with LLMs.
- Readings:
 - Reddy, S., Allan, S., Coghlan, S., & Cooper, P. (2020). A governance model for the application of AI in health care. *Journal of the American Medical Informatics Association*, 27(3), 491-497.
 - Reddy, S. (2024). Generative AI in healthcare: an implementation science informed translational path on application, integration and governance. Implementation Science, 19(1), 27.

• Week 4, Session 8: Evaluating AI Capabilities and Limitations

- Topics: Model benchmarks, the emergent field of research evaluating LLMs. Hallucinations, bias, reasoning limits. Critical evaluation of outputs. Research on model capabilities/weaknesses. Human oversight ("human in the loop" approach). Developing your own strategy to evaluate model output.
- Readings:
 - Ji, Z., Lee, N., Frieske, R., et al. (2023). Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys*, *55*(12), 1–38. (Focus on Abstract, Intro, Types of Hallucinations)
 - Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? *FAccT* '21. (Focus on Intro, Risks/Harms sections)
 - Banerjee, S., Agarwal, A., & Singla, S. (2024). Llms will always hallucinate, and we need to live with this. arXiv preprint arXiv:2409.05746. (Introduction)
 - Circuit Tracing: Revealing Computational Graphs in Language Models. <u>https://transformer-circuits.pub/2025/attribution-graphs/methods.html</u> (Introduction and examples).
- Homework #2 Assigned (Due end of Week 4): Prompting exercises, privacy policy analysis, evaluating AI text for accuracy/bias. (Worth ~8.57%)

Module 3: AI for research and summarization (Weeks 5-6)

• Week 5, Session 9: AI for Literature Search and Review

- Topics: Review HW2. AI research assistants (Elicit, Scite, Deep Research) that anchor (aka, "ground") models to external data. Strengths and limitations and hallucinations/mistakes.
- Readings:
 - Birhane, A., Kalluri, P., Card, D., et al. (2022). The Values Encoded in Machine Learning Research. *FAccT* '22. (Consider potential biases even in research tools - skim relevant sections).
 - Buckley, T. A., Crowe, B., Abdulnour, R. E. E., Rodman, A., & Manrai, A. K. (2025, March). Comparison of frontier open-source and proprietary large language models for complex diagnoses. *JAMA Health Forum* (Vol. 6, No. 3, pp. e250040-e250040).

• Week 5, Session 10: Hands-on with AI Research Assistants

- Topics: Best practices. Formulating questions. Critically assessing outputs with examples on literature review and research synthesis. Anticipating mistakes based on task.
- Readings:
 - Mollick, E. (2024), Chapters 5 and 6.
 - Huyen, C. (2024). Chapter 2 (from Sampling strategies section).

• Week 6, Session 11: AI for Text Summarization and Document Analysis

- Topics: Summarization techniques (conceptual). Using LLMs for summaries. Documentgrounded tools (NotebookLM). The problem of missing the middle of the text (long prompts). Examples analyzing YouTube videos on the experience of autistic adults.
- Readings:
 - Goyal, T., Li, J. J., & Durrett, G. (2022). Training data augmentation for abstractive text summarization. *arXiv preprint arXiv:2204.00790*. (Technical, focus on Abstract/Intro for conceptual understanding of challenges)
 - Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., & Liang, P. (2023). Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*. (Introduction and conclusion)
 - Khalil, M., Mohamed, F., & Shoufan, A. (2025). Evaluating the quality of medical content on YouTube using large language models. *Scientific Reports*, 15(1), 9906. (Introduction and discussion)
- Week 6, Session 12: Applying Summarization
 - Topics: Use cases in public health (policy docs, interviews). Generating outputs. Ethical considerations and methods for evaluating output quality. Developing a strategy to assess output quality.
 - Readings:
 - Goodman, K. E., Paul, H. Y., & Morgan, D. J. (2024). AI-generated clinical summaries

require more than accuracy. JAMA, 331(8), 637-638.

- Gallifant, J., Afshar, M., Ameen, S., Aphinyanaphongs, Y., Chen, S., Cacciamani, G., ... & Bitterman, D. S. (2025). The TRIPOD-LLM reporting guideline for studies using large language models. Nature Medicine, 1-10.
- **Homework #3 Assigned (Due end of Week 6):** Applying a framework for literature review and summarization using LLMs; critical reflection. (Worth ~8.57%)

Module 4: AI for Data Analysis (Weeks 7-10)

• Week 7, Session 13: LLMs as Data Analysts

- Topics: Review HW3. Using LLMs as data analysis assistants, from Excel tools to coding in different statistical programs (R, SAS, Stata) and languages (SQL, Python). Guidelines for privacy precautions and examples using Excel. Using the models to generate simulated data (aka, synthetic data) to preserve privacy during future interactions.
- Readings:
 - Mollick, E. Chapter 7 and 8.
 - Xing, Y. (2024). Exploring the use of ChatGPT in learning and instructing statistics and data analytics. *Teaching Statistics*, 46(2), 95-104.

• Week 7, Session 14: AI as Data Analysis Tutors in Public Health Practice

- Topics: Examples of AI as tutors for data analysis, from design to interpretation of results. Understanding limitations and areas where models are likely to make mistakes.
- Readings:
 - Mollick, E. Chapter 7 and 8.
 - Bharel, M., Auerbach, J., Nguyen, V., & DeSalvo, K. B. (2024). Transforming Public Health Practice with Generative Artificial Intelligence. Health Affairs, 43(6), 776-782.

• Week 8, Session 15: AI and Causal Inference Challenges

- Topics: Correlation versus causation. Overview of the potential outcomes framework to understand when a statistical analysis provides causal answers.
- Readings:
 - Perraillon, M.C. and Hedeker, D. (2025). The potential outcomes framework (up to estimation section)
 - Hernán, M. A., Hsu, J., & Healy, B. (2019). A Second Chance to Get Causal Inference Right: A Classification of Data Science Tasks. CHANCE, 32(1), 42-49. (Focus on task classification and causal questions).

• Week 8, Session 16: AI Causal Inference Challenges and Conceptual frameworks

- Topics: The need for strong conceptual frameworks in research seeking causal explanations (i.e, "world models."). Limitations of current AI models in this area with an example on AI understanding of maps.
- Readings:
 - Evaluating the World Model Implicit in a Generative Model. K Vafa, J Y Chen, A

Rambachan, J Kleinberg, S Mullainathan. Neural Information Processing Systems (NeurIPS), 2024. <u>https://arxiv.org/abs/2406.03689</u>. (Introduction and discussion)

Prosperi, M., Min, J. S., Bian, J., & Modave, F. (2020). Big data hurdles in precision medicine and precision public health. *BMC Medical Informatics and Decision Making*, 20(1), 1-10. (Discusses challenges including causal inference).

• Week 9, Session 17: Introduction to AI in Qualitative Research

- Topics: Potential AI roles. Content and thematic analysis. Tools overview and examples from a project on Long COVID analyzing interview transcripts.
- Readings:
 - Paulus, T. M., & Marone, V. (2024). "In minutes instead of weeks": Discursive constructions of generative AI and qualitative data analysis. *Qualitative Inquiry*, 107780.
 - Guetterman, T. C., & Fetters, M. D. (2024). The Role of Artificial Intelligence in Qualitative Research. *Journal of Mixed Methods Research*, *18*(1), 3-11. (Editorial/overview)
 - Review features pages for AI capabilities in NVivo, Atlas.ti, MAXQDA. Explore Taguette websites.

• Week 9, Session 18: Using AI Tools for Qualitative Tasks

- Topics: Practical considerations (transcription check, code suggestions, theme exploration).
 Human oversight. Limitations with examples of systematic errors. Testing and refining prompt.
 Using the LLM model to verify model results.
- Readings:
 - Hamilton, A. B., Finley, E. P. (2019). Qualitative methods in implementation research: An introduction. *Psychiatry Research*, 280, 112516. (Contextualizes standard qualitative methods before considering AI).
- **Homework #4 Assigned (Due end of Week 8):** Analyze AI in health case studies (benefits, risks, ethics). Optional: Small exercise exploring qualitative snippets with AI. (Worth ~8.57%)

• Week 10, Session 19: AI for Data Visualization and Reporting

- Topics: Al assistance in visualization. Al for drafting quantitative summaries/reports. Word clouds, text clustering techniques. Common mistakes.
- Readings:
 - Explore features of tools claiming AI-assisted visualization or reporting (e.g., Tableau AI, Power BI Copilot features - conceptually).
- Week 10, Session 20: Evaluating AI Studies of Model Capabilities
 - Topics: Key questions for evaluating AI studies (data, bias, validation, interpretability, causality). Interpretability research.
 - Readings:
 - On the biology of large language models. Anthropic. https://transformercircuits.pub/2025/attribution-graphs/biology.html

- Stevens, L. M., Mortazavi, B. J., Deo, R. C., Curtis, L. H., & Saria, S. (2020). Recommendations for reporting machine learning analyses in clinical research. *Circulation: Cardiovascular Quality and Outcomes*, *13*(10), e006556. (Focus on reporting standards as a basis for appraisal).
- **Homework #5 Assigned (Due end of Week 10):** Critically evaluate a study using AI for public health data analysis (focus on methods, claims, bias). (Worth ~8.57%)

Module 5: Future, Impact on Jobs, Synthesis (Weeks 11-15)

- Week 11, Session 21: Diffusion of Innovation in Technology
 - Topics: Review HW5. Diffusion of innovation theory applied to computer technology changes. Innovation in healthcare and implications for the future of AI in healthcare and medicine.
 - Readings:
 - Mullick, E. (2024). Chapter 9
 - Autor, D. (2022). The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty (No. w30074). *National Bureau of Economic Research*. (Focus on introduction and conclusion)
 - Dearing, J. W., & Cox, J. G. (2018). Diffusion of innovations theory, principles, and practice. *Health affairs*, *37*(2), 183-190

• Week 11, Session 22: Potential Labor Market Impact

- Topics: Impact of AI technology on healthcare jobs, impact on disparities. Two views: AI as augmentation and AI as replacement. Two periods: short- and long-term changes.
- Readings:
 - Autor, D. Applying AI to Rebuild Middle Class Jobs. *National Bureau of Economic Research*. <u>https://www.nber.org/papers/w32140</u> (Focus on introduction and discussion)
 - Brynjolfsson, E., & McAfee, A. (2014). The second machine age: Work, progress, and prosperity in a time of brilliant technologies. WW Norton & company. (Chapter 1 on Canvas).
 - Artificial Intelligence and Its Potential Effects on the Economy and the Federal Budget. Congressional Budget Office, December 2024. <u>https://www.cbo.gov/publication/61147</u> (Focus on conclusions)

• Week 12, Session 23: AI Alignment Problem

- Topics: Defining alignment. How do we ensure that more autonomous, "intelligent" machines align to human values? Understanding trade-offs.
- Readings:
 - Askell, A., et al. (2021). A General Language Assistant as a Laboratory for Alignment. *arXiv* preprint arXiv:2112.00861. (Technical, focus on introduction and discussion)
 - Turpin, M., Michael, J., Perez, E., & Bowman, S. (2023). Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. Advances in Neural Information Processing Systems, 36, 74952-74965. https://arxiv.org/pdf/2305.04388 (Introduction and discussion)
 - Nong, P., & Platt, J. (2025). Patients' Trust in Health Systems to Use Artificial Intelligence.

JAMA Network Open, 8(2), e2460628-e2460628.

 Baker, B., Huizinga, J., Gao, L., Dou, Z., Guan, M. Y., Madry, A., ... & Farhi, D. (2025). Monitoring reasoning models for misbehavior and the risks of promoting obfuscation. arXiv preprint arXiv:2503.11926. (Introduction and discussion)

• Week 12, Session 24: Economic Implications for Public Health jobs

- Topics: The likely impact of AI in public health jobs: patient care versus healthcare administration, research, and policy analysis
- **Final Project Proposal Due.** Discussion of proposals
- Readings:
 - McKinsey Global Institute. (2023). Generative AI and the future of work in America. (Review sections related to healthcare, public policy).
- **Homework #6 Assigned (Due end of Week 12):** Ethical analysis of complex Al-in-health scenario; discuss alignment risks. (Worth ~8.57%)
- **Final Project Proposal Due.** (Worth 10%)

• Week 13, Session 25: AI as Interview Preparation Tools

- Topics: Review HW6. Prompts and models to leverage AI to find jobs, preparing for interviews, acquiring interview skills and updating skills.
- Readings:
 - Maloney, D. H., et al. (2023). Conversational Artificial Intelligence for Enhancing Medical Education: Opportunities and Challenges. *Academic Medicine*, *98*(11S), S48-S54.
 - Chamorro-Premusiz, T. (2025). Want to Use AI as a Career Coach? Use These Prompts. Harvard Business Review.

• Week 13, Session 26: Review of Project Proposals

- Topics: Peer-review of final projects, suggestions for improvement, ideas
- Readings:
 - Read a project proposal and provide suggestions.

• Week 14, Session 27: Review of Project Proposals (continuation)

- Topics: Peer-review of final projects, suggestions for improvement, ideas.
- Readings:
 - Read a project proposal and provide suggestions.
- Week 14, Session 28: Preparing for AI in Your Career
 - Topics: Staying current and a changing field. Communicating about AI. Final Project Q&A.
 - Readings:
 - (To be updated)
 - Homework #7 Assigned (Due end of Week 14): Critically evaluate the research on the impact of AI on jobs. (Worth ~8.57%)

• Week 15, Session 29: Course Review and Final Project Work Session

- Topics: Recap major themes, tools, ethics. Revisiting LOs. Open Q&A for final project.
- *Readings:*
 - Review course notes and previous readings, especially those related to final project.

• Week 15, Session 30: Final Q&A & Course Wrap-up

- Topics: Final opportunity for project questions. Ongoing learning resources. Evaluations. Conclusion.
- Activity: Open Q&A.
- Final Project Due. (Worth 30%)