Advanced Computational Methods for Social Science Research

Graduate Seminar

Instructor: [TBD]
Email: [TBD]
Office Hours: [TBD]
Meeting Times: [TBD]
Location: [TBD]

Course Description

This seminar explores the integration of large language models and other AI tools into social science research methodology. The course emphasizes practical implementation skills, pipeline development, and rigorous validation frameworks, while maintaining close connections to fundamental research principles. Students will develop hands-on experience with current best practices in AI-assisted research, from annotation workflows to advanced prompting and chain-of-thought strategies. Throughout the course, we will also engage with the theoretical and substantive social science questions that these methods help address and critically evaluate the ethical implications, model governance, and bias in AI-driven research.

Prerequisites and Refresher Resources

Prerequisites:

- Familiarity with social science research design
- Basic programming experience (Python)
- Statistics through multiple regression

Refresher Resources (Optional):

- Grolemund, Garrett, and Hadley Wickham. *R for Data Science* (for data wrangling principles; Python equivalents available online)
- Wickham, Hadley, and Mine Çetinkaya-Rundel. *Data Science: A First Introduction* (for statistical inference review)
- Python Tutorials: `https://docs.python.org/3/tutorial/`
- Scikit-learn user guide: `https://scikit-learn.org/stable/user_guide.html` (for quick reference on machine learning basics)

Course Objectives

By the end of this seminar, students will be able to:

- Develop systematic frameworks for integrating AI tools (including large language models) into social science research pipelines.
- Master key technical skills in model deployment, prompt engineering, and validation.

- Design and implement robust, documented, and ethically sound research workflows that blend computational methods with substantive social science questions.
- Critically evaluate the strengths, limitations, biases, and ethical considerations of AI-driven research methodologies.
- Connect methodological approaches with core social science theories and empirical puzzles.

Assessment

- **Research Project (40%)**:
  - *Proposal (5%)*: A short proposal due Week 4 outlining the research question, theoretical motivation, and initial approach.
  - *Mid-Semester Update (5%)*: Due Week 8, detailing progress, adjustments, and initial findings.
  - *Final Implementation (20%)*: A fully documented pipeline, code, and associated data analysis.
  - *Final Paper (10%)*: A 10-12 page paper describing the research question, theoretical relevance, methods, results, validation steps, and interpretation.
- **Technical Exercises (30%)**: Periodic hands-on assignments to practice annotation, prompt design, model validation, and pipeline integration.
- **Research Design Document (20%)**: A detailed methodology specification (due Week 10) demonstrating how the chosen computational techniques address a specific social science question.
- **Class Participation (10%)**: Active engagement in discussions, peer feedback on projects, and participation in in-class activities.

Resources

- Computing resources and API access to major language models provided through the department.
- Technical support available through scheduled lab meetings.
- Sample code repositories and implementation guides posted on the course website.

Course Schedule and Learning Objectives

Below is the tentative schedule. Each segment includes key learning objectives to clarify expectations.

Weeks 1-2: Foundations and Current Landscape

Learning Objectives:

- Understand the current capabilities, architectures, and limitations of large language models.
- Identify ethical challenges and discuss model governance and bias in AI-driven research.
- Connect the use of LLMs to substantive social science questions (e.g., measuring policy positions, detecting ideological framing).

Topics:

- Capabilities of large language models
- Technical architecture and implementation considerations
- Ethical implications, model governance, and bias considerations

Readings:

- OpenAI. 2023. "GPT-4 Technical Report." arXiv:2303.08774.
- Krippendorff, Klaus. 2004. "Reliability in Content Analysis: Some Common Misconceptions and Recommendations." *Human Communication Research* 30(3): 411-433.
- Bommasani, Rishi, et al. 2021. "On the Opportunities and Risks of Foundation Models." arXiv:2108.07258.

Weeks 3-4: Research Design with AI Tools

Learning Objectives:

- Develop annotation schemes suitable for large-scale analysis.
- Understand how to scale annotation workflows efficiently.
- Assess reliability and validity of annotated data, linking these choices to substantive research questions (e.g., coding political speeches, media content).

Topics:

- Developing annotation schemes
- Scaling annotation workflows
- Reliability assessment

Readings:

- King, Gary, Robert O. Keohane, and Sidney Verba. 1994. *Designing Social Inquiry.* Princeton University Press. (Focus on measurement and inference)

- Radford, Alec, et al. 2019. "Language Models are Unsupervised Multitask Learners." OpenAI Blog.
- Wang, Alex, et al. 2022. "Self-Instruct: Aligning Language Models with Self-Generated Instructions." arXiv:2212.10560.

Project Proposal Due (Week 4)

Weeks 5-6: Implementation Strategies

Learning Objectives:

- Choose appropriate model architectures and deployment strategies.
- Manage cost and efficiency trade-offs in large-scale implementations.
- Begin drafting a pipeline that can process real-world social science data (e.g., legislative texts, news articles, survey responses).

Topics:

- Model selection and deployment
- Cost and efficiency trade-offs
- Pipeline development

Readings:

- Kaplan, Jared, et al. 2020. "Scaling Laws for Neural Language Models." arXiv:2001.08361.
- Brown, Tom B., et al. 2020. "Language Models are Few-Shot Learners." *NeurIPS* 33.
- Longpre, Shayne, et al. 2023. "The Flan Collection: Designing Data and Methods for Effective Instruction Tuning." arXiv:2301.13688.

Weeks 7-8: Multimodal Models

Learning Objectives:

- Extend NLP pipelines to multimodal contexts (images + text).
- Understand how integrating visual and textual data can answer social science questions (e.g., analyzing campaign ads that combine text and images).
- Incorporate multimodal models into a research pipeline.

Topics:

- Visual content analysis
- Multimodal research design
- Integration with text analysis

Readings:

- Joo, Jungseock, and Zachary C. Steinert-Threlkeld. 2022. "Image as Data: Automated Visual Content Analysis for Political Science." *Political Analysis* 30(4): 510-532.

- Zhang, Yuhao, et al. 2023. "GPT-4V(ision) System Card." arXiv:2311.11568.
- Torres, Michelle, and Francisco Cantú. 2022. "Learning to See: Convolutional Neural Networks for the Analysis of Social Science Data." *Political Analysis* 30(1): 113-131.

Mid-Semester Update Due (Week 8)

Weeks 9-10: Advanced Prompting and Chain-of-Thought

Learning Objectives:

- Apply prompt engineering principles to improve model performance.
- Implement chain-of-thought prompting to encourage step-by-step reasoning.
- Design prompts that help extract theoretically meaningful insights (e.g., identifying causal claims in policy documents).

Topics:

- Prompt engineering principles
- Chain-of-thought implementation
- Iterative refinement

Readings:

- Wei, Jason, et al. 2022. "Chain of Thought Prompting Elicits Reasoning in Large Language Models." *NeurIPS* 35.
- Kojima, Takeshi, et al. 2022. "Large Language Models are Zero-Shot Reasoners." *NeurIPS* 35.
- Zhou, Xuhui, et al. 2023. "SELF-INSTRUCT: Aligning Language Model with Self Generated Instructions." ACL.

Research Design Document Due (Week 10)

Weeks 11-12: Pipeline Development

Learning Objectives:

- Create a robust, end-to-end workflow integrating data sourcing, annotation, modeling, and validation.
- Implement error handling and documentation best practices.
- Ensure the pipeline supports reproducible social science research (e.g., can another researcher replicate the study on a different corpus?).

Topics:

- Robust workflow design
- Error handling
- Documentation standards

Readings:

- Mitchell, Margaret, et al. 2019. "Model Cards for Model Reporting." FAT*: 220-229.
- Bandy, Jack, and Nicholas Vincent. 2021. "Addressing 'Documentation Debt' in Machine Learning Research: A Retrospective Datasheet for BookCorpus." arXiv:2105.05241.
- Gebru, Timnit, et al. 2021. "Datasheets for Datasets." *Communications of the ACM* 64(12): 86-92.

Weeks 13-14: Validation and Research Applications

Learning Objectives:

- Establish validation frameworks and replication standards.
- Critically evaluate final results against social science theories.
- Identify future directions for integrating advanced AI methods into substantive political science and social research.

Topics:

- Validation frameworks
- Replication standards
- Future directions

Readings:

- Nosek, Brian A., and Timothy M. Errington. 2020. "What is Replicability?" *Annual Review of Psychology* 71: 167-193.
- Ribeiro, Marco Tulio, et al. 2020. "Beyond Accuracy: Behavioral Testing of NLP Models with CheckList." ACL: 4902-4912.
- Dodge, Jesse, et al. 2019. "Show Your Work: Improved Reporting of Experimental Results." EMNLP-IJCNLP: 2185-2194.

Final Implementation and Paper Due End of Week 14

Policies

Academic Integrity: All work must be original. Cite sources properly, including code snippets and external models.

Late Work: Late submissions lose 10% per day unless prior arrangements are made.

Accommodations: Students requiring accommodations should contact the instructor and the university's disability services office as early as possible.