Using Large Language Models in Political Science

Course Duration: 4 weeks (week 1 - 2 classes, weeks 2/3/4 - 3 classes, 150 minutes each class, 11 classes in total)

Format: Seminar discussions with integrated student presentations and practical tasks

Instructor: Gennadii IAKOVLEV

Contact: iakovlevg@ceu.edu

Course Description

Recent advancements in Natural Language Processing (NLP) and the emergence of powerful Large Language Models (LLMs) such as OpenAI's ChatGPT, DepSeek R1, Meta LLAMA, Mistral have opened new avenues for research in political science. LLMs made methods that were technically possible but time consuming, such as automated content analysis of political speeches and party manifestos, accessible. Unimaginable methods – such as synthetic surveys (silicone samples) and synthetic experiments - are now reality. Some integrations can even generate an entire research paper from scratch, provided only with a research question and hypothesis. This course focuses on applying modern LLM-based techniques in political science research and offers a structured overview. In 4 weeks, we will cover the foundations of NLP, explore applications, employ LLM-based analysis, and learn how to assess the reliability and validity of these methods critically. By the end of the course, students will be equipped with both the theoretical understanding and practical tools to incorporate the LLM-driven NLP into their own research on political phenomena.

Course Design

- **Target Audience**: Undergraduate-level students in social sciences with some quantitative and political science background
- **Structure**: Each week comprises three days, two 75-minute sessions each day. The first session is devoted to a seminar and discussion (120 min). We will discuss the week's core readings, explore conceptual frameworks, and examine use cases. The second sessions (30min) is devoted to students' presentations (except for the first week). It consists of a

10-minute presentation by one student on a hypothetical political science case or data project that employs LLMs, followed by a 15-minute Q&A.

Learning Outcomes

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

1. Understand fundamental concepts in NLP, the architecture of modern LLMs – know what large language models are, how they work at a high level, and how they differ from traditional text analysis methods in the social sciences, especially when applied to political science.
2. Distinguish among different NLP and LLM techniques (l.g., text classification, sentiment analysis, topic modeling, entity recognition) relevant to social science.
3. Apply LLM-based tools to real-world social science data (e.g., speeches, social media posts, legislative texts) for tasks like classification, summarization, and content analysis.
4. Critically engage with LLM-based research and outputs, including questions of ethics, reliability, validity, bias, and other misuse. Understand where these tools add value or risk.
5. Integrate AI with Theory: Connect LLM applications to social science concepts – for example, using model outputs to support arguments about public opinion, political communication, or policy processes.
6. Design a research project that leverages an LLM to answer a social science question.

Activities and Evaluation

1. **Applying Methodology from Weekly Readings (3 required)**
    - Each week, students give a maximum five-minute oral presentation applying that week's research design to another imaginary research setting, with a brief summary of the original reading.
2. **Weekly Practical Assignments (3 required)**
    - At the end of the first three weeks, all students will have to do a short, half-hour-long coding-free exercise, where they will be invited to use an LLM as a research tool.
3. **Mock Research Presentation**

Each student will give one 10–15 minute presentation on a specific imaginary case how LLM NLP methods could be applied to a specific political science topic, case, or dataset (e.g., election manifestos, parliamentary speeches, social media controversies).

Presenters should include:

- Research Question.
- Connection to Theory: How does this application tie back to political science theory?
- Data Source, Collection, and Cleaning: How were texts (or speech data) collected and pre-processed?
- Methods: Which NLP or LLM technique is proposed? Why is it suitable for the research question?
- Findings: What insights could be generated? Are there any methodological or ethical challenges?

  A 15-minute discussion follows each presentation, and the instructor will provide feedback.

  Attendance and Participation

  - Active participation in seminar discussions is expected. Engaging with weekly readings and offering thoughtful questions/comments during peer presentations will enhance the learning experience.

Weekly Topics, Literature, and Schedule

Each class runs from 10:00–13:00 ET and has two 75-minute sessions (Session 1: 10:00-11:15; Session 2: 11:25–12:40) with a 10-minute break in between. Each day will be based on approximately two required readings. Seminars will cover key concepts and discuss examples from political science. All materials will be made available in advance. Students are expected to come to class prepared to discuss the assigned readings. Below is an outline of the four weeks.

There will be capstone project presentations at the end of the second session of each day, apart from the first week. Each student (or group, depending on the number of participants) presents their mock research project applying course concepts to a social science research question using LLMs. Presentations should outline research questions, methods, proof of the validity of the method, expected findings, and reflections on reliability and ethics.

Week 1 (June 9 & 10). Introduction to NLP and LLMs in Political Science: Fundamentals of Text Representation and Core LLM Techniques

Topics

- Different approaches to NLP: dictionary analysis, supervised and unsupervised machine learning

- Understanding of how large language models like GPT are trained
- Survey of political science areas that LLMs can influence (e.g. speech analysis, automated content analysis, simulating public responses)
- Discussion of examples where AI has been applied in political analysis

Key Questions

- What is NLP and why is it useful for political science research?
- What generations of NLP exist in political science research?
- What are pre-trained large language models, and how have they changed the landscape of NLP?

Schedule

- Day 1 – June 9
- Session 1 (10:00-11:15) **Introduction**
    - Outline of the course structure, grading system, and assignments
    - Overview of the evolution of text analysis techniques in political science
    - Overview of the evolution of text analysis techniques in political science
    - How LLMs are transforming research compared to human-based content analysis
    - Readings:
        1. Jurafsky & Martin, *Speech and Language Processing* (3rd ed.), Chapters 1 & 10
        2. Ornstein et al., "How to Train Your Stochastic Parrot: LLMs for Political Texts" (selected sections)
    - Session 2 (11:25–12:40) – **Technical Foundations of LLMs**
        - Introduction to neural language models and the transformer architecture
        - Comparison with earlier approaches (dictionary analysis, n-grams, etc.)
        - Readings:
            1. Jurafsky & Martin, *Speech and Language Processing*, Chapters 9 & 10
            2. Bommasani et al. (2021), "On the Opportunities and Risks of Foundation Models" (selected sections)

- Day 2 – June 10
    - Session 1 (10:00-11:15) – **Transformers and the Evolution of Neural Language Models**

        - How transformers differ from earlier deep learning architectures (RNNs, LSTMs)
        - Attention mechanism and why transformers are state-of-the-art

        Video to watch – But what is a neural network? | Deep learning chapter 1

    - Session 2 (11:25–12:40) – **Transfer Learning & Fine-Tuning for Political Science**
        - Pre-trained "foundation" models and domain adaptation
        - Steps in fine-tuning LLMs for tasks like classification or stance detection
        - Examples of GPT-based models for legislative texts or social media content

    Readings:

        1. Visual Storytelling Team & Murgia (2023), "Generative AI exists because of the transformer."
        2. Vaswani et al. (2017), "Attention Is All You Need."

Weekly Practical Assignment (Due Week 2) Reflection on LLM Usage.

Experiment with any online LLM to analyze a piece of political text. Submit a short (0.5 page of your own text) reflection describing:

1. Your prompt and the model's output,
2. Observed strengths and weaknesses (overall quality, biases, and limitations) of the LLM's performance.

Week 2 (June 16, 17, 18). Applications in Political Science I. Sentiment Analysis, Named Entity Recognition

Topics

- Fundamentals of textual data representation (dictionary analysis, bag-of-words, embeddings)
- Introduction to transformer architectures, transfer learning, and fine-tuning for political science tasks
- Prompt engineering techniques for classification, sentiment analysis, and stance detection
- Evaluating model performance against human-coded "ground truth"
- Comparing LLM-based classification/sentiment analysis with traditional methods

Key Questions

- How do we represent language in ways machines can process (tokens, embeddings, n-grams)?
- What are the strengths and limitations of dictionary-based or machine-learning text analysis, and how do LLMs improve on them?
- How do transformers differ from earlier neural approaches (RNNs, LSTMs)?
- What does "fine-tuning" an LLM entail, and why does it matter for specialized political texts?
- How can we design prompts (zero-shot vs. few-shot) for classification, sentiment, or stance analysis?
- What methods can we use to evaluate LLM outputs for accuracy and bias, and how do we validate them against human-coded data?

Schedule

- Day 1 – June 16
    - Session 1 (10:00-11:15) – **Automating Text Classification in Political Science**
        - Review of classification tasks (party manifestos, legislative speech)
        - Traditional dictionary or ML approaches vs. prompt-based LLM classification
        - Readings:
            1. Hunger (2024), "Virtuous people and evil elites? The role of moralizing frames and normative distinctions in identifying populist discourse."
            2. Ornstein et al., section on text classification in "How to Train Your Stochastic Parrot."
            3. Excerpt from Laver et al. (2003), "Extracting Policy Positions from Political Texts."
    - Session 2 (11:25–12:40) – **Sentiment Analysis in Political Contexts**
        - Detecting sentiment/tone in political texts using LLMs
        - Comparison of dictionary-based, machine learning, and LLM-driven approaches
        - Mock Research Presentations
        - Readings:
            1. Cocco & Monechi (2024), "How Populist are Parties? Measuring Degrees of Populism in Party Manifestos Using Supervised Machine Learning"
            2. Young & Soroka (2012), "Affective News: The Automated Coding of Sentiment in Political Texts."
            3. Jurafsky & Martin, Chapter 22 (Lexicons for Sentiment)
- Day 2 – June 17

- Session 1 (10:00-11:15) – **Evaluating Classification – Accuracy and Validity**
  - Methods to evaluate LLM performance vs. human-coded data
  - Systematic misclassifications, validation strategies
  - Readings:
    1. Ornstein et al., section on evaluation of classification methods
    2. Benoit et al. (2016), "Crowd-Sourced Text Analysis."
    3. Le Mens & Gallego (2025), "Positioning Political Texts with LLMs by Asking and Averaging" (selected sections)
- Session 2 (11:25–12:40) – **Student Presentations & Discussion**
  - Presentations applying classification/sentiment analysis to selected texts
  - Debate on accuracy, biases, evaluation methods
  - Mock Research Presentations

- Day 3 – June 18
  - Session 1 (10:00-11:15) and Session 2 (11:25–12:40) **Prompting and In-Context Learning**
    - How LLMs are "programmed" via natural language prompts
    - Best practices for prompt design (zero-shot, few-shot)
    - In-class hands-on exercises
    - Mock Research Presentations
    - Readings:
      1. Ornstein et al., sections on prompt design in "How to Train Your Stochastic Parrot"
      2. Google, "Prompt engineering: overview and guide"
      3. Google, "Tips to enhance your prompt-engineering abilities"

Weekly Practical Assignment (Due Week 3) Content Analysis Comparison.

1. Apply a dictionary-based (manual or tool-based) analysis to a mock corpus of 10 political texts.
2. Use an LLM approach on the same texts, ask an LLM a question for each of the 10 political texts, and make it a mock dataset.
3. Write a brief report comparing outcomes and discussing possible reasons for differences.

Week 3 (June 23, 24, 25). Applications in Political Science II. Sentiment Analysis and **Introduction to Synthetic Surveys**

Topics

- Practical applications of LLMs in sentiment analysis, ideological scaling, and stance detection
- Using LLMs to simulate or augment political data ("synthetic surveys," "simulated" public opinion)
- Strengths and limitations of automated text analysis for measuring ideology, sentiment, or stance
- Ethical and validity concerns with synthetic data

Key Questions

- How can we measure sentiment towards leaders, policies, or parties using LLM tools?
- What methods exist for estimating ideology from text, and what are their limitations?
- How do we validate or triangulate automated text results against "ground truth"?
- What are synthetic surveys, and how can LLMs aid in generating them for political research?
- In what ways does synthetic data augmentation help overcome small-sample or sensitive-data constraints?
- What are the ethical and validity concerns when using synthetic respondents or AI-generated responses?

Schedule

- Day 1 June 23
- Session 1 (10:00-11:15) and Session 2 (13:45-15:00)
  LLMs for **Sentiment, Ideology, and Stance analysis, and Data Collection**
- Continuing a discussion of how LLMs/NLP techniques measure sentiment or ideological positions in political texts
    - Real-world examples of stance detection (in speeches, manifestos, social media)
    - Mock Research Presentations
    - Readings:
        1. Jurafsky & Martin, *Information Extraction: Relations, Events, and Time*, Chapter 20
        2. Yoruk (2021), "Global Contentious Policy Dataset"
        3. Arslan et al. "Political Events using RAG with LLMs"
- Day 2 June 24
- Session 1 (10:00-11:15)
  LLMs as Survey Respondents – "Simulated" Public Opinion
    - How LLMs can mimic human responses under specified conditions
    - Readings:

1. Argyle et al. (2023), “Out of One, Many: Using Language Models to Simulate Human Samples”
2. Gilardi (2023), “ChatGPT in Political Science Research: Applications and Challenges”

- Session 2 (13:45-15:00)

  Evaluating Simulated Data – Promise and Pitfalls
  - Assessing accuracy and limitations of LLM-generated survey data
  - Readings:
    1. Bisbee et al. (2024), “Synthetic Replacements for Human Survey Data? The Perils of LLMs”
    2. Kassaar & Curtice (2023), “Can ChatGPT Replace Survey Panels? An Empirical Assessment”
    3. Revisit Argyle et al. (2023) – Methods section
  - Student Presentations (Week 3)
- Day 3 June 25
- Session 1 (10:00-11:15)

  LLMs for Data Augmentation and Experimentation
  - Using LLMs to generate additional data for training or experimental purposes
  - Readings:
    1. Wu et al. (2023), “Large Language Models as Simulated Policy Makers”
    2. Yang & Torrance (2023), “Using AI to Generate Political Text for Experiments”
    3. Weidinger et al. (2021), “Ethical and Social Risks of Large Language Models” (selected sections)
- Session 2 (11:25–12:40)

  Mock Research Presentations & Discussion
  - Further exploration of stance detection, ideology measurement, or synthetic data experiments
  - Q&A and feedback on applying these methods

Weekly Practical Assignment (Due Week 4) A Survey Simulation Request.

1. Try to come up with a prompt that makes any web-based LLM answer survey questions from the point of view of a certain person (region, education, ethnicity, age).
2. Run this prompt 10 times
3. Compare it with the actual survey results.

Week 4 (June 30, July 1, July 2). LLMs for Data Augmentation. The Future of LLMs. Mock Research Projects

Topics

- Synthetic surveys revisited, full data augmentation techniques, and deeper ethical/IRB considerations
- Advanced discussions on applying LLM-driven approaches to real data
- Course wrap-up and final capstone presentations

Focus

- Students work with provided R scripts, mock text corpora, and research questions to practice LLM-driven analyses (via API queries).
- Ethical considerations of using AI-generated text or synthetic survey data in real-world political science research.
- Summarizing and integrating insights from the entire course.

Key Questions

- How can advanced LLM tools be integrated into mainstream research workflows?
- What are the limitations, biases, and methodological trade-offs associated with data augmentation?
- How do we benchmark augmented vs. real datasets, and what ethical concerns arise?
- In what directions might future large language models push the boundaries of political text analysis?

Schedule

- Day 1 – June 30
    - Session 1 (10:00-11:15) – **LLMs for Data Augmentation**
        - Techniques for generating or expanding textual datasets (e.g., policy briefs, legislative debates)
        - Balancing authenticity vs. synthetic generation
        - Readings:
            1. Wu et al. (2023), “Large Language Models as Simulated Policy Makers”
            2. Yang & Torrance (2023), “Using AI to Generate Political Text for Experiments”
    - Session 2 (11:25–12:40) **Combining Synthetic Surveys & Data Augmentation**

        - Designing research that uses both synthetic respondents and AI-generated text
        - Potential synergy or conflicts between these two methods
        - Privacy, consent, and potential for misuse
        - Handling data that may be too "synthetic" to be trusted in real-world policy discussions
- Day 2 – July 1
    - Session 1 (10:00-11:15) – **Evaluation & Validation of Augmented Data**
        - Benchmarking augmented vs. real datasets
        - Identifying biases or unrealistic patterns in synthetic text
        - Privacy, consent, and potential for misuse
        - Handling data that may be too "synthetic" to be trusted in real-world policy discussions
    - Session 2 (11:25–12:40) **– Ethical & IRB Considerations**
- Day 3 – July 2
    - Session 1 (10:00-11:15) – **Future Directions & Limitations of LLM development**
        - Possibilities for AGI and Artificial Superintelligence in the near future
        - Next frontiers for LLMs in political text analysis
        - Is there any impact on the economy at all?
        - Readings:
            1. Aschenbrenner (2024), "Situational Awareness: The Decade Ahead"

               against

            2. Acemoglu (2025), "The simple macroeconomics of AI"

    - Session 2 (11:25–12:40) – **Final Synthesis & Course Wrap-Up**
        - Reviewing major takeaways from all four weeks
        - Student Mock Presentations
        - Q&A and concluding remarks

Assignments and Evaluation

Grading Breakdown:

- **Applying Methodology from Weekly Readings – 25%:** Each week, every student will have to deliver a 5-minute oral presentation of an idea for a mock paper, where the key points from that week's readings' research design are applied to their research context. Also, it is advisable to summarize the original reading.
- *Grading:* A top-quality methodology goes as far from the original paper as possible. For instance, one paper from the readings offers a two-step dictionary analysis of populism in European speeches; a perfect application case would be – the same two-step dictionary-based approach is applied to depicting migrants as terrorists in parliamentary speeches; a good application case would be – applying the same analysis to the United States material with certain adaptation. Discuss how LLMs would differ from older text analysis methods in this case.

- **Weekly Practical Assignments – 25%:** Each week, every student will submit an assignment (approximately half a page, or ~200-250 words) outlined by the end of every week in a syllabus. These summaries should address the main questions of the tasks. Summaries are due at the start of the first class each week (typically Monday).
- **Mock Research Presentation – 35%:** In lieu of a written final paper, students will complete a final project that culminates in an in-class presentation. Each student will choose a topic related to LLMs in political science (to be approved by the instructor by the end of Week 1), conduct deeper research or analysis on that topic, and present their findings and arguments to the class. A visual aid (slides) should accompany the talk, and students are encouraged to be creative and analytical – this is your chance to explore an aspect of the course that interests you most.

  Grading: This final presentation is the capstone of the course and carries the highest weight. It will be graded on the depth of research, clarity of argument, originality of thought, and effectiveness of communication. Excellent projects will demonstrate a strong grasp of the chosen topic, connect it to broader themes from the course, and possibly even contribute a unique perspective or proposal. The Presentation Grading Rubric applies here, with expectations of a more comprehensive treatment given the extra time and research invested.
- **Participation and Engagement – 15%:** Active participation is essential in a seminar-style course, especially one as condensed as this. Participation includes attending each class, contributing to discussions with informed comments, listening attentively to others, and engaging respectfully with differing viewpoints. It also encompasses activities like providing feedback in the Week 5 workshop and asking questions or offering constructive comments during your peers' presentations.

  Grading: Quality matters more than quantity. You don't need to speak every five minutes, but

when you do contribute, it should be clear you've done the reading and are thinking critically. Positive indicators of participation include referencing the readings in discussion, building on a classmate's point, or posing thoughtful questions. Negative indicators would be not participating at all, or dominating discussion without listening to others. Attendance is the prerequisite for participation – more than one unexcused absence will reduce this portion of the grade.

Note: There are no separate midterm or final exams. The course uses continuous assessment (weekly summaries, presentations, etc.) to encourage steady engagement with the material. If circumstances arise that affect your ability to meet a deadline or attend class (e.g. illness or emergency), please communicate with the instructor as early as possible to make alternative arrangements.

Grading of Presentations

All in-class presentations – both the weekly discussion-leading presentations and the final project presentations – will be evaluated using a rubric focused on content mastery and communication skills. This rubric is adapted from a written assignment rubric to fit oral presentations:

- **Content & Understanding:** *Does the presentation demonstrate a strong grasp of the material or topic?* An excellent presentation provides an accurate and comprehensive explanation of key concepts or findings, and it uses evidence or examples appropriately. The presenter shows depth of understanding and can answer questions about the topic. A weaker presentation might omit important points, show confusion about the topic, or include inaccuracies.
- **Critical Analysis:** *Does the presenter go beyond summary to provide insight?* A high-quality presentation doesn't just recite what was in the readings or research – it offers some analysis or critical thought. This could be raising questions about the implications, comparing different perspectives, highlighting strengths/weaknesses in an argument, or relating the topic to broader themes from the course. An average presentation may accurately summarize material but not delve into analysis. An excellent one will introduce original insights or connections.
- **Organization & Clarity:** *Is the talk well-structured and clear?* The presentation should have a logical structure (introduction of topic, main points, conclusion or takeaway). Ideas should flow in a coherent order, making it easy for the audience to follow. Visual aids (slides) should be clear and not overburdened with text. Staying within the allotted time is also part of good organization. A top-notch presentation is clearly organized and stays focused; a weaker one might jump between points without a clear structure or try to cover too much without cohesion.
- **Delivery Skills:** *How effectively does the student communicate their points?* This criterion covers public-speaking skills: speaking at a clear and audible volume, at a measured pace (not too fast or slow), making eye contact or otherwise engaging the audience, and demonstrating

enthusiasm or confidence about the material. Excellent delivery means the audience can easily understand the speaker and remains engaged. Issues like mumbling, reading directly from notes/slides, or frequent loss of train of thought would detract from the delivery score.

- **Engagement & Q&A:** *Does the presenter engage the audience and handle questions well?* A great presenter involves the class, whether by asking rhetorical questions, pausing for discussion, or actively encouraging questions at the end. During the Q&A or discussion, the presenter listens to questions carefully and responds thoughtfully, showing they can think on their feet and further explain or defend their points. If a presenter cannot address basic questions about their topic, it may indicate insufficient preparation or understanding. Being open to feedback or alternative perspectives during discussion also reflects well on this aspect.

Each of these criteria will be taken into account for presentation grades. Generally, an "A" range presentation excels in most or all of the above areas: it delivers rich content with clear organization and engaging delivery, spurring a good discussion. A "B" range presentation is competent but might be average in one or two areas (for example, the content knowledge is solid but the delivery is a bit dry, or vice versa). A "C" range presentation may have significant shortcomings such as shaky understanding of the material and poor organization. The rubric is meant to guide you in preparing your presentations — aim for strength in each area.

Feedback: After your presentation, the instructor will provide brief feedback highlighting what you did well and suggestions for improvement. Peer feedback may also be incorporated (for instance, we might spend a few minutes after each student presentation where classmates share positive comments or ask additional questions). The goal is to help you improve your communication skills in a supportive environment.

Course Policies

Attendance: Attendance is crucial, given the fast-paced 4-week format. If you must miss a class due to illness or emergency, notify the instructor as soon as possible. More than one unexcused absence may significantly affect your participation grade. Remember that even if you miss a class, you are responsible for that session's material and should catch up via classmates' notes or talking with the instructor.

Preparation and Etiquette: Complete the required readings before each seminar session so you can actively participate. Bring your notes or reading summaries to refer to during discussion. During class, be respectful of others' contributions: listen actively, do not interrupt, and keep your comments civil even when debating ideas. We aim to create an inclusive environment where everyone feels comfortable sharing insights or questions.

Deadlines: Weekly reading summaries are due at the start of the week's first class. Student discussion leaders should also share any presentation slides or an outline with the instructor *before* their session begins (at least the night before, if possible). Missing a deadline for a summary or not being prepared on your presentation day not only affects your grade but also lets down your classmates who rely on your contribution. If you anticipate any difficulty meeting a deadline, discuss it with the instructor in advance.

Academic Integrity: All work you do for this course should be your own. Plagiarism (copying someone else's words or ideas without attribution) is a serious violation. This applies to both written work (summaries) and presented content. If you draw on external sources for your final project or presentation, cite them appropriately (a simple slide or footnote citation is fine). Cheating or plagiarism will result in a zero on the assignment and referral to academic affairs per university policy.

Use of AI Tools: You are encouraged to use AI tools (like LLMs) as part of your research or even to experiment with summary generation. There's educational value in that, but you must **disclose and credit** any use of AI in your work, and ideally, share with colleagues how to do it more effectively. For example, if you have ChatGPT summarize an article as a starting point, say so in your summary, and then critically evaluate its output against the article. Any direct text generated by an AI that you include in your submissions or presentations must be clearly indicated and your own analysis added. Failure to do so could be considered a misrepresentation of work.

Inclusivity and Support: We value a diversity of backgrounds and opinions. If you have any concerns about the class climate or your ability to participate, please reach out. If you require any accommodations for a disability, kindly coordinate with the disability services office and inform the instructor so we can support your learning needs. We will be discussing technology in a societal context, which might involve topics like bias, discrimination, or political conflict; if you feel uncomfortable at any point, you can always speak to the instructor privately about the situation.

Summary

This four-week course offers social science students a foundational understanding of NLP and LLMs while highlighting how these methods transform research on elections, public opinion, policy-making, and more. By combining theoretical readings, real-world case discussions, and student-led presentations, participants will develop both the conceptual framework and practical skills required to navigate the evolving landscape of LLM-based methods in political science research.