

What AI Is, How It Works, and Why It Matters: A Primer for Political Science

Konrad Posch, PhD

University of San Francisco – Politics, Economics, Sociology, & Public Policy

kposch@usfca.edu

June 19, 2026

Draft 2.1 – Prepared for

*Teaching Political Science in an Age of Generative Artificial Intelligence (GAI) Symposium,
American Political Science Association (APSA) & the Political Studies Association (PSA),*

Virtual, June 22-26, 2026

Preliminary Draft for Discussion;

Please contact the author for the latest draft if you wish to share further.

AI Usage Statement (for APSA GenAI Virtual Conference)

Since we are a conference about GenAI, and in looking at the JPSE guidelines on AI usage in manuscripts, I figured this was an excellent moment to engage in the meta-discussion of how WE as scholars should/can/do use AI in our own work. I'm a big believer that you shouldn't use AI to replace our writing; indeed, I often tell students that "the narrative is the science" so that they realize that it's not the numbers, or the models, or the literature review, or the as-if randomness that really contributes to knowledge but instead the ability to use those tools to explain what's really going on that is the true scientific contribution. And yet, when faced with turning a series of lectures I've given to multiple classes over multiple years into the "article version," I couldn't help being seduced by the promise that I could delegate the gruntwork of transcription, rearranging, and prose construction to an AI model, at least for the initial pass. As you can read in the usage statement below, that's what I did. Then, I systematically worked through the whole thing and changed quite a bit (if you're interested, I'm happy to share the initial Claude Fable 5 draft for comparison with the changes marked up). I still need to do some reference housekeeping (mea culpa). Could this article have happened without AI? Sure. How many months of unpaid labor do you want to do this summer? Would this article have happened, at the level of production value I have here, for our discussion? Probably not. Papa's gotta pay rent, after all, and while it may be a publish or perish world, somehow no one ever wants to pay us to write the things that get published.

This manuscript was prepared with the assistance of an artificial intelligence tool: Claude (Anthropic PBC), model Claude Fable 5, accessed through the claude.ai web interface in June 2026. The tool was used for manuscript preparation, not as a research methodology; no study data were generated or analyzed by AI, and clinical AI reporting guidelines (e.g., CONSORT-AI, STARD-AI, TRIPOD-AI) are accordingly not applicable. The author used the tool to:

- (1) restructure and condense verbatim transcripts of the author's own classroom lectures, recorded in February 2026, into article form;
- (2) render that material in the author's prose style, which the tool was instructed to emulate from samples of the author's prior writing supplied by the author;
- (3) scaffold the footnote apparatus and flag claims requiring citation by inserting conspicuous editorial tags rather than generating references, with every reference subsequently checked by the author; and
- (4) format the manuscript.

All substantive arguments, examples, interpretations, and pedagogical content originated in the author's lectures, writing, and editorial decisions; the author reviewed, revised, and approved all AI-assisted text. The author confirms the originality and accuracy of the content. The author confirms having reviewed the terms of use of the AI tool and confirms its suitability for publication. The author takes full responsibility for the integrity of the whole content, including the accuracy of all references. A full record of the methods – including the complete, unedited prompt-and-response transcript of the drafting process, the lecture transcripts, and the writing samples provided to the tool – has been retained by the author and will be shared with the journal upon request. The drafting chat conversation is available at: <https://claude.ai/share/5ff10fde-bc0b-441d-ba8b-75b4c1aa16e5>

Abstract

Artificial intelligence (AI) is a crisis striking every institution political scientists study at once: a single technology simultaneously on the agenda of legislatures, regulators, employers, universities, and utilities. That simultaneity makes AI unavoidable in the political science classroom, but it also makes AI pedagogically golden – if students understand what the technology actually does. Most undergraduates arrive fluent in chatbot interfaces but innocent of the mechanism behind them, marinating in a public discourse prone to hyperbole, exaggeration, and misconstrued expertise. This article is a self-contained primer designed to be assigned directly in undergraduate political science courses. Part I explains what large language models (LLMs) are and how they work – transformers, embeddings, conditional probability, and the temperature feature – using John Searle’s Chinese Room thought experiment as an anchoring metaphor. Part II connects those technical features to core questions of public policy: algorithmic accountability, employment, economic and social regulation, and the physical realities of energy and water. Throughout, the primer develops one central claim: the policy-relevant change LLMs introduce is the removal of coding literacy as the price of admission to the power of computing. Developed across six course offerings to political science and economics students, the primer treats AI as analyzable – not unprecedented.

Keywords: artificial intelligence; large language models; public policy; pedagogy; technology regulation; undergraduate teaching