



PWM-CS-HS2: Computational Social Science II

Artificial Intelligence and Democracy

Syllabus

Lecturer:

Prof. Dr. Andreas Jungherr

Contact:

Email: andreas.jungherr@uni-bamberg.de

Webseite: <https://www.uni-bamberg.de/politikdigital/>

FMA/01.12

Feldkirchenstraße 21

96052 Bamberg

Office hours:

By appointment. For appointments please reach out to Ms. Katharina Kachelmann (katharina.kachelmann@uni-bamberg.de)

Secretary:

Katharina Kachelmann

Email: katharina.kachelmann@uni-bamberg.de

Phone: +49 951 863-2738

FMA/01.11

Feldkirchenstraße 21

96052 Bamberg

1 Course description:

As artificial intelligence (AI) continues to advance, it becomes crucial to understand its economic, social, and political implications. The course provides students with a comprehensive and realistic understanding of AI's impact on democracy and the public arena.

We first establish foundational knowledge about the workings of AI and the foundations and context conditions for its success. We then discuss core issues in AI's impact on democracy. This includes:

- AI risks and safety,
- AI Governance,
- the future of work,
- AI and the public arena,
- AI in campaigning, and
- AI and autocracies.

The seminar will introduce the history, concepts, theories, and research methods at the intersection of AI and democracy. We will examine the interplay between digitalization, society, and politics, as well as the various investigative approaches used in this field. Relevant topics will be discussed in the context of current international case studies, offering students an in-depth understanding of the complex relationship between AI and democratic processes.

Please address your questions regarding the course to Ms. Katharina Kachelmann (katharina.kachelmann@uni-bamberg.de).

Learning objectives:

- Advanced understanding of concepts, theories, causal relationships and methods relevant to computational social science;
- Knowledge of the central paradigms in theory and research methods relevant to computational social science;
- Understanding of the applicability or transfer of theories and paradigms from different scientific areas in relation to computational social science.

2 Course requirements

2.1 Regular and active participation

The course features the discussion of the required readings. To benefit, students are expected to read the texts listed as required readings before each session and actively participate in the discussion for each session. You will find it useful to keep notes.

Background Readings:

- On taking notes: Ahrens, S. (2022). *How to take smart notes: One simple technique to boost writing, learning and thinking* (2nd ed.).

2.2 Presentation

Participants will be asked for a presentation on one of the topics discussed during the course. The task of the presenter will be to present:

- a problem associated with the topic of each session;
- a research design fit to create insight about the problem, underlying mechanisms, and potential solutions.

Make sure to answer the following questions:

- What problems, conflicts, challenges, or opportunities emerge from the use of AI in this area?
- What are the relevant actors? How and for what purpose are they (not) dealing with AI?
- What are the relevant mechanisms determining AI use and outcomes?
- How do we know?
- How is this phenomenon connected with other cases that we have a better understanding of? Or is it truly completely new?
- What is a research design that helps us understand this phenomenon and problem better?

To answer these questions and to develop your presentation start with the texts in the required and the background readings. But this is only meant to get you started. I expect that you build on these texts by looking at instructive media coverage and further academic articles.

During the first session, each student will be assigned a topic from the listed readings for presentations. Please keep the following considerations in mind in preparing your presentation:

- Please plan your presentation to take about 10-15 minutes;
- Please prepare a slide deck with a presentation program of your choice (except for Prezi);
- In preparing the presentation please follow the guidelines discussed in the first session;
- Discuss your plan for the presentation with the lecturer at least one week before your presentation date;
- Please prepare a handout of one to two pages for your fellow students, summarizing the main points of your presentation;
- Upload the handout on the day your presentation is due to the assignment folder on the course's VC repository. Use the following template for the filename "your_last_name-handout.pdf";

- The presentation will be graded and contribute 30% to your final grade.

If you do not follow these questions and guidelines this will be reflected in your grade.

Background Readings:

- Schwabish, J. (2017). *Better presentations: A guide for scholars, researchers, and wonks*. Columbia University Press.

2.3 Term paper

Participants will be asked to hand in a term paper. Please adhere to the following guidelines:

- Style requirements:
 - Font: Times New Roman, 12pt;
 - Line separation: 1.5;
 - Page borders: 2.5 cm left and right, 2cm above and below;
 - Page set: Block;
 - The first line of each paragraph is indented;
- Citation Style: Please follow the citation convention of the American Political Science Review (APSR) available at <http://www.apsanet.org/APSR-Submission-Guidelines-August-2016>, or you could simply use the reference style *APA* in the references manager of your choice;
- Cover page: University, department, course title, paper title, name, Matriknr., semester count, study program, and e-mail-address;
- Length: ca. 5000 words +/-10%
- Deadline: Please upload the paper up until September 30 electronically in the VC. The date is mandatory and can only be extended in case of officially certified illness;
- Use the following template for the filename "your_last_name-paper.pdf".
- The term paper will be graded and contribute 70% to your final grade.

Background Readings:

- Basbøll, T. (2018a). The paper. *Inframethodology*. https://blog.cbs.dk/inframethodology/?page_id=614
- Becker, H. S. (1998). *Tricks of the trade: How to think about your research while you're doing it*. The University of Chicago Press.
- Becker, H. S. (2020). *Writing for social scientists: How to start and finish your thesis, book, or article* (3rd ed.). University of Chicago Press.
- Howard, C. (2017). *Thinking like a political scientist: A practical guide to research methods*. The University of Chicago Press.

2.4 Policy on ChatGPT

By now, you have likely encountered accounts of ChatGPT's potential for assisting you in writing tasks. Perhaps you have even tried it out. This is excellent and highly recommended, as there is every reason to expect that your future life in work or research will involve working with AI-enabled assistants, whether for software development, data analysis, or managing mundane office tasks. Therefore, familiarizing yourself with these tools and learning about their strengths and weaknesses is crucial. However, as a student, certain uses may be more advisable than others.

Before you start using ChatGPT, consider what you might lose by relying on it. We assign research papers to help you practice specific tasks repeatedly throughout your studies, offering you the opportunity to learn and improve your skills. However, this will only happen if you actually *perform* the tasks and *do* the work. Relying on ChatGPT or other models too early in your education may prevent you from acquiring or refining these skills over time. At the same time, we can expect workflows in academia and industry to be shaped by collaboration between humans and AI-enabled systems, such as LLMs, sooner rather than later. Consequently, developing the necessary skills to use these models effectively is also essential.

One approach is to consider the skills or tasks you are expected to learn, perform, or improve with a given paper assignment. Challenge yourself to complete these tasks independently, write down your solutions, and then compare them with the output of your current AI-enabled model of choice or even competing models. By doing this, you can reflect on the accuracy of your work and the model's output, identify areas of improvement, and understand where the model's strengths and weaknesses lie. This process transforms LLMs into a supporting tool rather than a substitute, while also providing valuable insight into your own work.

For transparency reasons, we ask students at the Chair for Political Science, esp. Digital Transformation to include a short disclaimer in their papers, indicating if and which AI model they used and for what tasks.

Possible tasks include:

1. Exploring a phenomenon, mechanism, or literature;
2. Formulating a research question;
3. Developing theory-driven hypotheses;
4. Analyzing data;
5. Structuring the paper;
6. Writing;
7. Editing.

Please document

- If and for what tasks you have used ChatGPT or comparable large language models (LLMs);
- How your work built upon the results provided by the model;

- List prompts used by you and responses in an online appendix.

Remember, you are solely responsible for the text you submit. Undocumented use of AI-enabled models, plagiarism, flaws in reasoning or analysis, and fabricated sources may result in significant grade reductions or even failure of the class. It does not matter whether these issues originated from you or the model – as the author, you are accountable for the strengths and weaknesses of your submitted work.

Be aware that when grading papers, we may place greater emphasis on aspects where models perform poorly and discount tasks where models excel.

Background Readings:

- Jungherr, A. (2023e). Using ChatGPT and other large language model (LLM) applications for academic paper assignments. *SocArxiv*. <https://doi.org/10.31235/osf.io/d84q6>

3 Course plan

Class will meet in person (FMA/00.07) at the following dates and times:

Monday 14:00–16:00 c.t.

- 3.1 Week 1: Introduction and Housekeeping (April 15)**
 - 3.2 Week 2: AI and Democracy (April 22)**
 - 3.3 Week 3: No meeting (April 29)**
 - 3.4 Week 4: No meeting (May 6)**
 - 3.5 Week 5: What are we talking about, when we talk about AI? (May 13)**
 - 3.6 Week 6: No meeting (May 20)**
 - 3.7 Week 7: How does AI learn about the world? (May 27)**
 - 3.8 Week 8: AI Risk and AI Safety (June 3)**
 - 3.9 Week 9: AI Governance (June 10)**
 - 3.10 Week 10: AI and the Future of Work (June 17)**
 - 3.11 Week 11: AI and the Public Arena (June 24)**
 - 3.12 Week 12: AI and Campaigning (July 1)**
 - 3.13 Week 13: AI and Autocracies (July 8)**
 - 3.14 Week 14: Discussion and Open Questions (July 15)**
-

3.1 Week 1: Introduction and Housekeeping (April 15)

3.2 Week 2: AI and Democracy (April 22)

Session Prep:

- What are core features of democracy?
- How visible is power?
- What is organized uncertainty?
- What is AI's impact on democracy?

Required Readings:

- Jungherr, A. (2023a). Artificial intelligence and democracy: A conceptual framework. *Social Media + Society*, 9(3), 1–14. <https://doi.org/10.1177/20563051231186353>

Background Readings:

- Bobbio, N. (1987). *The future of democracy: A defence of the rules of the game* (R. Dellamy, Ed.; R. Griffin, Trans.). University of Minnesota Press. Trans. of *Il futuro della democrazia*. (1984). Einaudi.
- Dahl, R. A. (1998). *On democracy*. Yale University Press.
- Przeworski, A. (2018). *Why bother with elections?* Polity Press.
- Risse, M. (2023a). Artificial intelligence and the past, present, and future of democracy. In *Political theory of the digital age: Where artificial intelligence might take us* (47–72). Cambridge University Press. <https://doi.org/10.1017/9781009255189.004>

3.3 Week 3: No meeting (April 29)

3.4 Week 4: No meeting (May 6)

3.5 Week 5: What are we talking about, when we talk about AI? (May 13)

Session Prep:

- What is narrow AI?
- What is artificial general intelligence (AGI)?
- What are conditions for the successful application of AI?

Required Readings:

- Mitchell, M. (2019). *Artificial intelligence: A guide for thinking humans*. Farrar, Straus; Giroux. Chapters 1-3 (pp. 17-65).
- Wolfram, S. (2023b). What is ChatGPT doing ... and why does it work? *Stephen Wolfram Writings*. <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/>

Background Readings:

- Lee, K.-F., & Quifan, C. (2021). *AI 2041: Ten visions for our future*. Currency.
- Nilsson, N. J. (2010). *The quest for artificial intelligence: A history of ideas and achievements*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511819346>
- Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson Education. (Original work published 1995).
- Suleyman, M., & Bhaskar, M. (2023). *The coming wave: Technology, power, and the twenty-first century's greatest dilemma*. Bodley Head.

3.6 Week 6: No meeting (May 20)

3.7 Week 7: How does AI learn about the world? (May 27)

Session Prep:

- What are symbolic representations of the world?
- What is the relation between representation and world?
- What representations of the world does AI have access to?
- What is bias?

Required Readings:

- Shanahan, M. (2024). Talking about large language models. *Communications of the ACM*, 67(2), 68–79. <https://doi.org/10.1145/3624724>

Background Readings:

- Hand, D. J. (2004). *Measurement theory and practice: The world through quantification*. Wiley.
- Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale University Press.
- Smith, B. C. (2019). *The promise of artificial intelligence: Reckoning and judgment*. The MIT Press.

3.8 Week 8: AI Risk and AI Safety (June 3)

Session Prep:

- What are risks associated with AI?
- Which risks are mainly imaginative and which are pertinent?
- How do we ensure safety of AI models and products?

Required Readings:

- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Rismani, S., Shelby, R., Smart, A., Jatho, E., Kroll, J., Moon, A. J., & Rostamzadeh, N. (2023). From plane crashes to algorithmic harm: Applicability of safety engineering frameworks for responsible ML. In A. Schmidt, K. Väänänen, T. Goyal, P. O. Kristensson, A. Peters, S. Mueller, J. R. Williamson, & M. L. Wilson (Eds.), *Chi '23: Proceedings of the 2023 chi conference on human factors in computing systems* (pp. 1–18). ACM. <https://doi.org/10.1145/3544548.3581407>

Background Readings:

- Bahr, N. J. (2015). *System safety engineering and risk assessment: A practical approach* (2nd ed.). CRC Press. <https://doi.org/10.1201/b17854> (Original work published 1997).
- Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and machine learning: Limitations and opportunities*. The MIT Press.
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.
- Christian, B. (2020). *The alignment problem: Machine learning and human values*. W. W. Norton & Company.
- Raji, I. D. (2024). The anatomy of AI audits: Form, process, and consequences. In J. B. Bullock, Y.-C. Chen, J. Himmelreich, V. M. Hudson, A. Korinek, M. M. Young, & B. Zhang (Eds.), *The Oxford handbook of AI governance*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780197579329.013.28>

3.9 Week 9: AI Governance (June 10)

Session Prep:

- How should AI be governed?
- What is the role of the state in AI governance?
- What are alternative binding governance approaches besides the state?

- What are the regulatory approaches of the USA, EU, and China to AI?

Required Readings:

- Dafoe, A. (2018). *AI governance: A research agenda*. Centre for the Governance of AI Future of Humanity Institute University of Oxford. <https://www.fhi.ox.ac.uk/wp-content/uploads/GovAI-Agenda.pdf>

Background Readings:

- Anderljung, M., Barnhart, J., Korinek, A., Jade Leun and, C. O., Whittlestone, J., Avin, S., Brundage, M., Bullock, J., Cass-Beggs, D., Chang, B., Collins, T., Fist, T., Hadfield, G., Hayes, A., Ho, L., Hooker, S., Horvitz, E., Kolt, N., Schuett, J., ... Wolf, K. (2023). Frontier AI regulation: Managing emerging risks to public safety. *arXiv*. <https://doi.org/10.48550/arXiv.2307.03718>
- Bradford, A. (2023). *Digital empires: The global battle to regulate technology*. Oxford University Press. <https://doi.org/10.1093/oso/9780197649268.001.0001>
- Bullock, J. B., Chen, Y.-C., Himmelreich, J., Hudson, V. M., Korinek, A., Young, M. M., & Zhang, B. (Eds.). (2024). *The Oxford handbook of AI governance*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780197579329.001.0001>
- Harding, V. (2024). *AI needs you: How we can change ai's future and save our own*. Princeton University Press.
- Nowotny, H. (2021). *In AI we trust: Power, illusion and control of predictive algorithms*. Polity Press.

3.10 Week 10: AI and the Future of Work (June 17)

Session Prep:

- Will AI displace or create jobs?
- What types of jobs and tasks appear to be vulnerable?
- What are the democratic consequences of AI's impact on work?

Required Readings:

- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. *NBER Working Paper Series*, (31161), 1–65. <https://doi.org/10.3386/w31161>
- Free exchange. (2022). Economists are revising their views on robots and jobs. *The Economist*. <https://www.economist.com/finance-and-economics/2022/01/22/economists-are-revising-their-views-on-robots-and-jobs>

Background Readings:

- Autor, D. H., Mindell, D. A., & Reynolds, E. B. (2022). *The work of the future: Building better jobs in an age of intelligent machines*. The MIT Press.
- Dixon, J., Hong, B., & Wu, L. (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science*, 67(9), 5301–5967, iii–iv. <https://doi.org/10.1287/mnsc.2020.3812>
- Eggleston, K., Lee, Y. S., & Iizuka, T. (2001). Robots and labor in the service sector: Evidence from nursing homes. *NBER Working Paper Series*, (28322). <https://doi.org/10.3386/w28322>
- Frey, C. B. (2019). *Technology trap: Capital, labor, and power in the age of automation*. Princeton University Press.
- Koch, M., Manuylov, I., & Smolka, M. (2021). Robots and firms. *The Economic Journal*, 131(638), 2553–2584. <https://doi.org/10.1093/ej/ueab009>
- Mann, K., & Püttmann, L. (2023). Benign effects of automation: New evidence from patent texts. *The Review of Economics and Statistics*, 105(3), 562–579. https://doi.org/10.1162/rest_a_01083

3.11 Week 11: AI and the Public Arena (June 24)

Session Prep:

- What is the public arena and what are its functions for democracies?
- In which areas of the public arena does AI feature?
- How do we know?
- What is the impact of AI on the democratic functions of the public arena?

Required Readings:

- Jungherr, A., & Schroeder, R. (2023). Artificial intelligence and the public arena. *Communication Theory*, 33(2–3), 164–173. <https://doi.org/10.1093/ct/qtad006>

Background Readings:

- Diakopoulos, N. (2019). *Automating the news: How algorithms are rewriting the media*. Harvard University Press.
- Narayanan, A. (2023). Understanding social media recommendation algorithms. *Knight First Amendment Institute at Columbia University*. <https://knightcolumbia.org/content/understanding-social-media-recommendation-algorithms>
- Nielsen, R. K. (2024). How the news ecosystem might look like in the age of generative AI. *Reuters Institute for the Study of Journalism*. <https://reutersinstitute.politics.ox.ac.uk/news/how-news-ecosystem-might-look-age-generative-ai>

- Peters, B. (2007a). Der Sinn von Öffentlichkeit. In H. Wessler (Ed.), *Der Sinn von Öffentlichkeit* (pp. 55–102). Suhrkamp. (Original work published 1994).
- Peters, B. (2007c). Die Leistungsfähigkeit heutiger Öffentlichkeiten – einige theoretische Kontroversen. In H. Wessler (Ed.), *Der Sinn von Öffentlichkeit* (pp. 187–202). Suhrkamp. (Original work published 2002).
- Simon, F. M., Altay, S., & Mercier, H. (2023). Misinformation reloaded? Fears about the impact of generative AI on misinformation are overblown. *Harvard Kennedy School Misinformation Review*, 4(5), 1–11. <https://doi.org/10.37016/mr-2020-127>
- Simon, F. M. (2024). *Artificial intelligence in the news: How ai retools, rationalizes, and reshapes journalism and the public arena*. Tow Center for Digital Journalism, Columbia University. <https://academiccommons.columbia.edu/doi/10.7916/ncm5-3v06>

3.12 Week 12: AI and Campaigning (July 1)

Session Prep:

- What are the uses of AI for campaigns?
- What are the democratic functions of campaigns?
- Can AI strengthen campaigns' democratic functions?

Required Readings:

- Davenport, T. H., Guha, A., & Grewal, D. (2021). How to design an AI marketing strategy. *Harvard Business Review*, (4). <https://hbr.org/2021/07/how-to-design-an-ai-marketing-strategy>
- Jungherr, A., Rauchfleisch, A., & Wuttke, A. (2024b). Artificial intelligence in elections. *Working Paper*.
- Sifry, M. L. (2024). How AI is transforming the way political campaigns work. *The Nation*. <https://www.thenation.com/article/politics/how-ai-is-transforming-the-way-political-campaigns-work/>

Background Readings:

- Bai, H., Voelkel, J. G., Eichstaedt, J. C., & Willer, R. (2023). Artificial intelligence can persuade humans on political issues. *OSF Preprints*. <https://doi.org/10.31219/osf.io/stakv>
- Bueno de Mesquita, E., Canes-Wrone, B., Hall, A. B., Lum, K., Martin, G. J., & Velez, Y. R. (2023, November 1). *Preparing for generative AI in the 2024 election: Recommendations and best practices based on academic research*. The University of Chicago Harris School of Public Policy; Graduate School of Stanford Business. https://www.gsb.stanford.edu/sites/default/files/publication/pdfs/white-paper-2023-ai-and-elections-best-practices_0.pdf

- Cerina, R., & Duch, R. (2023). Artificially intelligent opinion polling. *arXiv*. <https://doi.org/10.48550/arXiv.2309.06029>
- Hersh, E. D. (2018, May 16). Cambridge Analytica and the future of data privacy: Written testimony of Eitan Hersh. In *Hearing before the United States Senate Committee on the Judiciary*. United States Senate. <https://www.judiciary.senate.gov/imo/media/doc/05-16-18%20Hersh%20Testimony1.pdf>
- Nickerson, D. W., & Rogers, T. (2014). Political campaigns and big data. *The Journal of Economic Perspectives*, 28(2), 51–74. <https://doi.org/10.1257/jep.28.2.51>
- Yarhi-Milo, K., Schiller, V., Clinton, H. R., Jourova, V., Ressa, M., Hajdu, D., Pallero, J., Tu, E., Schiffrin, A., Agranovich, D., Green, Y., Watts, C., Benson, J., Chertoff, M., Lindenbaum, D., Makanju, A., & Schmidt, E. (2024, March 28). *A.I.'s impact on 2024 global elections*. The Aspen Institute. <https://www.aspendigital.org/event/ai-2024-global-elections/>

3.13 Week 13: AI and Autocracies (July 8)

Session Prep:

- What are information challenges of autocracies compared to democracies?
- What are the conditions under which AI could help offset these challenges?
- Is there a Chinese AI advantage?

Required Readings:

- Beraja, M., Kao, A., Yang, D. Y., & Yuchtman, N. (2023a). AI-tocracy. *The Quarterly Journal of Economics*, 138(3), 1349–1402. <https://doi.org/10.1093/qje/qjad012>
- Yang, E. (2023). The limits of AI for authoritarian control. *Working Paper*.

Background Readings:

- Beraja, M., Kao, A., Yang, D. Y., & Yuchtman, N. (2023b). Exporting the surveillance state via trade in AI. *NBER Working Paper Series*. <https://doi.org/10.3386/w31676>
- Brussee, V. (2023). *Social credit: The warring states of China's emerging data empire*. Palgrave Macmillan. <https://doi.org/10.1007/978-981-99-2189-8>
- Creemers, R. (2018). China's social credit system: An evolving practice of control. *Social Science Research Network*. <https://doi.org/10.2139/ssrn.3175792>
- Lee, K.-F. (2018). *AI superpowers: China, Silicon Valley, and the new world order*. Houghton Mifflin Harcourt.

- Lindblom, C. E. (2001). *The market system: What it is, how it works, and what to make of it*. Yale University Press.
- Spufford, F. (2010). *Red plenty: Inside the fifties' Soviet dream*. Faber & Faber.
- Xu, X. (2021). To repress or to co-opt? Authoritarian control in the age of digital surveillance. *American Journal of Political Science*, 65(2), 309–325. <https://doi.org/10.1111/ajps.12514>

3.14 Week 14: Discussion and Open Questions (July 15)