The Cultural Life of Deep Learning – History, Technics, and Politics

Level: 200-level

Class Size: 15-30

Course Description

Recent breakthroughs in Deep Learning (DL) have transformed automation and redefined the functional terms of artificial intelligence (AI). Alongside DL-based AI's widespread use comes a mix of high hopes and deep anxieties about an AI-steering future. This course debunks AI as a one-way ticket to either ultramodern prosperity or machine-dominated dystopia by introducing multiple perspectives – historical , theoretical, political, and technical – on the potential and constraints of DL-based AI. Together, we shall explore: What kind of "learning" is "deep learning"? How do its technical pipelines, which involve data sampling, algorithm design, model building, evaluation metrics, and the wider industrial landscapes, shape its interaction with the world? Are these technical designs value-neutral, or embedded with cultural-technical thinking and politico-economic drives? If the latter, how do these values influence contemporary AI-related social issues like cognitive bias, labor displacement, environmental impact, and geopolitical tensions?

A central emphasis of this course will be on reflecting on the nature of "meaning" in language models compared to human language. Along with weekly postings and a research essay, hands-on assignments will guide students to build their own language models using spaCy NLP, engineer prompts, and calculate word embeddings. These practical experiences will prompt students to consider: How do language models construct and interpret "meaning" differently from humans? By the end of the course, students will gain a nuanced understanding of DL-based AI, its capabilities and limitations, and its profound impact on modern life.

No prior technical knowledge is required, but basic programming (Python) engagement is expected.

Learning Goals: During the semester, students are expected to learn how:

1. To gain familiarity with the conventions of academic writing and the methods of research and citation.

2. To identify and describe some of the ways humanities and science studies can intersect research methods that inform more comprehensive critical perspectives.

3. To develop critical and technical skills in engaging, assessing, and intervening with important media objects and processes.

Learning Outcomes: By the end of the course, students will:

1. Gain a nuanced understanding of deep learning-based AI through interdisciplinary analysis, using approaches from both humanities and science to critically assess AI and its social implications.

2. Develop basic technical skills in natural language processing and apply these skills to explore the nature of "meaning" in language models compared to human language.

3. Familiarize themselves with academic writing, including research methods and citation practices, while developing the ability to critically engage with media objects through written analysis.

Course Requirements and Grading:

Canvas discussion posts – 15%: Students will take turns posting and replying to weekly reading responses on Canvas. Each discussion post should be approximately 250 words and concentrate on at least one of the class readings. These posts are designed to capture your insights and questions, foster interaction with your classmates, and spark class discussions. Posts will be evaluated on their thoughtfulness, originality, and critical engagement with the course materials. Submit your post by noon on Monday to give your peers time to respond. Respondents should submit their posts by noon on Wednesday.

Workshop on prompt engineering – 20%: Participate in a hands-on workshop to explore the potential and constraints of generative models like ChatGPT. First, design prompts to generate fictional stories about real historical figures, experimenting with different combinations of historical and fictional named entities and context information. Second, select one prompt to elicit three varied responses and compare their differences. These outputs will be used in a subsequent exercise on cosine similarity of word embeddings.

Workshop on custom language models - 20%: In a second workshop, follow a step-by-step tutorial on building a custom language model using spaCy NLP, working with your own corpus of raw language materials. Enrich the corpus through tokenization, lemmatization, part-of-speech tagging, dependency parsing and chunking, and named entity recognition. Finally, apply spaCy's nlp() function to analyze normalized word vector information from the three ChatGPT outputs. Compile a report documenting your process and findings, including a 1-2 page analysis comparing ChatGPT outputs and a 2-3 page report on building and enriching your custom language model. Submit the combined PDF by the specified deadline.

Final paper proposal - 10%: The paper proposal should be a 250-500 word summary of your argument. Your proposal must be approved by the instructor.

Final research paper - 25%: Write a research paper (6-8 pages) that explores a topic related to the evolution and impact of Deep Learning technologies as covered in this course. You may opt for either of the following ways to organize your essay. 1) Incorporate at least two technical texts and two critical texts from the course readings. Your task is to put them into dialogue to discuss

how the chosen texts reflect or critique the implications of Deep Learning-based technologies on society, culture, and politics. 2) Incorporate at least one cultural material (literary or visual), one technical text, and one critical text from the course readings. Your task is to draw connections between the fictional narrative and real-world advancements. Ensure your essay is well-structured, with a clear thesis, coherent arguments, and appropriate citations. Upon finishing your first draft, you will have the chance to revise your work, incorporating feedback from a peer-reviewed workshop before submitting the final version.

A list of suggested topics will be provided prior to the proposal due date, but you may also propose your own topic, subject to instructor approval.

Attendance/Participation - 10%: Attendance is mandatory and will be noted at the beginning of each class. Active and meaningful participation in discussions about the readings and films is expected. Throughout the semester, you will also be required to complete group assignments, participate in in-class activities, and take in-class quizzes.

Schedules

Week 1: Framing Critical AI Studies and Critical Deep Learning

Wizards, Data Science. 2023. "Machine Learning Pipeline: What It Is, Why It Matters, and Guide to Building It?" *Medium* (blog). May 23, 2023.

Tyler Reigeluth and Michael Castelle. "What Kind of Learning Is Machine Learning?" in *The Cultural Life of Machine Learning*. Cham: Palgrave Macmillan, 2021, pp. 79-116.

Gary Marcus. "Deep learning: A critical appraisal." *arXiv preprint arXiv*:1801.00631 (2018), pp. 1-27.

Rita Raley and Jennifer Rhee, eds. 2023. Introduction in *Critical AI: A Field in Formation*. *American Literature*, volume 95, number 2 (June 2023). Durham, NC: Duke University Press, pp. 185-204.

Week 2: Neural Networks

3Blue1Brown. "But What Is a Neural Network?" https://www.youtube.com/watch?v=aircAruvnKk.

Michael Wooldridge. "Deep Breakthroughs" in *A Brief History of Artificial Intelligence*, pp.95-118.

Pedro Domingos. "How Does Your Brain Learn" in *The Master Algorithm*. New York: Basic Books, 2015, pp. 104-129.

Stefanos Tsimenidis. “Limitations of Deep Neural Networks: a discussion of G. Marcus’ critical appraisal of deep learning.” *arXiv preprint arXiv*:2012.15754 (2020), pp.1-16.

Week 3: Architectures and Algorithm

Steve Brunton. “Neural Network Architectures & Deep Learning.” https://www.youtube.com/watch?v=oJNHXPs0XDk.

Yann LeCun, Yoshua Bengio, & Geoffery Hinton. (2015). “Deep learning.” *Nature*, 521(7553), pp. 436–444.

Matteo Pasquinelli. “The Material Tools of Algorithmic Thinking” in *The Eye of the Master: A Social History of Artificial Intelligence*. Verso Books, 2023, pp.28-51.

Pedro Domingos. “The Master Algorithm” in *The Master Algorithm*. New York: Basic Books, 2015, pp.38-68.

Week 4: Synthetic Data

TensorFlow. “Data Preprocessing for ML: Options and Recommendations.”

Sergey I. Nikolenko. Introduction and the Second Chapter in *Synthetic Data for Deep Learning*. Vol. 174. *Springer Nature*, 2021, pp.3-26.

B. N. Jacobsen (2023). “Machine learning and the politics of synthetic data.” *Big Data & Society*, 10 (1), pp.1-12.

James Steinhoff. “Toward a political economy of synthetic data: A data-intensive capitalism that is not a surveillance capitalism?” *New Media & Society* 26.6 (2024), pp. 3290-3306.

Week 5: Vector Space and the Calculation of Meaning

Jussi Karlgren and Pentti Kanerva. 2021. “Semantics in High-Dimensional Space.” *Frontiers in Artificial Intelligence* 4 (August), pp.1-6.

Fabian Offert. “Intuition and Epistemology of High-Dimensional Vector Space.” (2017), pp.1-7.

Lydia H. Liu. “Wittgenstein in the Machine.” *Critical Inquiry* 47.3 (2021): 425-455.

Week 6: Workshop I: Build your own language model and calculate meanings

Exercises with spaCy NLP to build your own language model and calculate cosine similarity of word embeddings.

Week 7: The Political Economy of Deep Learning

Kate Crawford and Vladan Joler, “Anatomy of an AI System.”

Matteo Pasquinelli. “Introduction: AI as Division of Labor” in *The Eye of the Master: A Social History of Artificial Intelligence*. Verso Books, 2023, pp.11-27.

Paolo Virno. “General Intellect.” *Historical Materialism* 15 (3), pp.3-8.

Simon Schaffer. “Babbage’s Intelligence: Calculating Engines and the Factory System.” *Critical Inquiry*, 21.1 (1994), pp. 203-227.

Week 8: Deep Learning and Cultural Production

K. Allado-McDowell. *Pharmako-AI*. 2020. UK: Ignota.

Martin Thoma (KIT Karlsruhe). “Creativity in Machine Learning.” 2016, pp.1-5.

Blaise Agüera y Arcas. “Art in the Age of Machine Intelligence.” *Arts*. Vol. 6. No. 4. MDPI, 2017, pp.1-9.

Week 9: Deep Learning and Image Culture

Tomasz Malisiewicz, “From feature descriptors to deep learning: 20 years of computer vision”, January 2015, computervisionblog.com.

Joanna Zylinska, *AI Art: Machine Vision and Warped Dreams* (excerpts). Open Humanities Press, 2020.

Lev Manovich. “Automating Aesthetics: Artificial Intelligence and Image Culture,” *Flash Art International*, n. 316, September–October 2017, pp.1-10.

Adrienne Lafrance. “When Robots Hallucinate: What do Google’s trippy neural network-generated images tell us about the human mind?” *The Atlantic*, 09/2015.

Week 10: Deep Learning and Science

Wang Hanchen, et al. “Scientific discovery in the age of artificial intelligence.” *Nature* 620.7972 (2023), pp. 47-60.

Raghu Maithra and Eric Schmidt, “A Survey of Deep Learning for Scientific Discovery,” *arXiv preprint* 2003.11755, 2020, pp.1-18.

Cade Metz, “A.I. Turns Its Artistry to Creating New Human Proteins,” *The New York Times*, January 9, 2023.

Fabian Offert, Paul Kim, Qiaoyu Cai. “Synthesizing Proteins on the Graphics Card: Protein Folding and the Limits of Critical AI Studies.” *AI & Society*, pp.1-19.

Week 11: Deep Learning and Philosophy

Tara Abraham. “(Physio)logical Circuits: The Intellectual Origins of the McCulloch-Pitts Neural Networks.” *Journal of the History of the Behavioral Sciences* 38, no. 1 (2002), pp.3-25.

Cameron Buckner. “Deep learning: A philosophical introduction.” *Philosophy Compass* 14.10 (2019), pp.1-19.

Hubert L. Dreyfus. “Why Heideggerian AI failed and how fixing it would require making it more Heideggerian.” *Philosophical psychology* 20.2 (2007), pp. 247-268.

Lydia H. Liu. “After Turing: How Philosophy Migrated to the AI Lab.” *Critical Inquiry* 50.1 (2023), pp. 2-30.

Week 12: Workshop II and Student Presentation I

Week 13: Student Presentation II