

Northwestern | Kellogg

Human and Machine Intelligence (MORS 950)

Professor Matt Groh | he/him/his | matthew.groh@kellogg.northwestern.edu

Section 31 | Mondays and Thursdays | 10:30 am - 12:00 pm | Location TBD, Evanston
Section 81 | Thursdays | 6:00 pm - 9:00 pm | Location TBD, Chicago

Office Hours: Schedule available on Canvas and by appointment

Contact info: Please add MORS950 to the subject of any email you send me

Course overview

Artificial intelligence (AI) is a transformative technology and a modern day equivalent to fire in the early stages of human civilization. It is a tool that can be used to solve complex problems, make predictions, automate tasks, and enhance productivity. But like fire, it has a dual nature and has potential for both good and bad outcomes. This course requires no prior technical knowledge and is designed for people who want to lead the deployment of AI systems in the real-world, manage data science and design teams, and build and invest in AI companies. The goal of the course is to build intuition for what AI can do, how machine learning works, where these tools tend to succeed and fail, and how to navigate their ethical implications. We will explore a wide range of business applications, examine tools including ChatGPT, Midjourney, DeepBlue, Watson, AlphaZero, the recommendation systems behind Twitter and TikTok, and many more, and discuss best practices for managing teams of humans assisted by these tools. This course is a lecture-based course with case-based discussions, individual assignments, a midterm, and a final group project. By the end, you should be an expert at identifying promising use-cases, evaluating current limitations, and recognizing potential pitfalls such that you are capable of applying human and machine thought partnerships to grow new businesses and disrupt Grand Masters in any field.

Attendance

Students are expected to attend every class in-person unless they're feeling sick.

Laptops and cell phones

No laptops or cell phones are allowed to be used in class unless they're required for classwork.

Grading

Assignments (40%) Three individual assignments will give you first-hand experience in applying AI towards solving problems in business. The first assignment is worth 10% of your grade and the second and third are each worth 15% of your grade.

Midterm (25%) A mid-term will give you feedback on your grasp of the course work.

Northwestern | Kellogg

Final Project (25%) A group project will offer experience identifying, thinking through, and artfully pitching an AI-based startup idea, product, business plan, or policy recommendation.

Participation (10%) You are expected to attend every class and actively engage in group discussions. Your participation grade is based on your attendance and participation in class. In addition, for the group project, we will collect peer evaluations to encourage all group members to make active contributions. These peer evaluations will be directly used to help determine an individual's participation grade.

Honor code

At the Kellogg School of Management, we are committed to the highest levels of integrity, professionalism, and respect for others, both inside and outside the classroom. We have a responsibility to behave ethically in all situations, to honestly represent our own contributions and those of others in all the work we perform, and to acknowledge and respect diversity in individuals and cultures. Any work that you submit should be entirely your own work. If you have any questions regarding how the [honor code](#) applies to this course, please ask.

AI Policy

I expect you to think critically and use AI as a tool for assignments in this class. Assignments will require using large language models (LLMs) like ChatGPT and image generation tools. Do not blindly trust AI outputs. LLMs are prone to making up facts, and it's your responsibility to think critically and double check purported facts with credible sources. You are responsible for errors and omissions in the assignments you submit. If you copy and paste a response from an LLM, then you must cite the response accordingly. Failure to do so is in violation of academic honesty policies.

Student safety and wellness

Students can find useful resources for safety and security, academic support, and mental and physical health and well-being on the [NU help website](#). **If you are struggling, please know that you can reach out to me personally or the Office of Student Life. We will work to get you connected to the right resources.**

Students with disabilities

Northwestern University is committed to providing the most accessible learning environment as possible for students with disabilities (and I personally share this commitment). Should you anticipate or experience disability-related barriers in the academic setting, please contact AccessibleNU to move forward with the university's established accommodation process (email: accessiblenu@northwestern.edu; phone: 847-467- 5530).

If you already have established accommodations with AccessibleNU, please let me know as soon as possible, preferably within the first two weeks of the term, so we can work together to implement your disability accommodations. Disability information, including academic accommodations, is confidential under the Family Educational Rights and Privacy Act.

Northwestern | Kellogg

Prohibition of recording of class sessions

Unauthorized student recording of classroom or other academic activities (including advising sessions or office hours) is prohibited. Unauthorized recording is unethical and may also be a violation of University policy and state law. Students requesting the use of assistive technology as an accommodation should contact [AccessibleNU](#). Unauthorized use of classroom recordings – including distributing or posting them – is also prohibited. Under the University’s [Copyright Policy](#), faculty own the copyright to instructional materials – including those resources created specifically for the purposes of instruction, such as syllabi, lectures and lecture notes, and presentations. Students cannot copy, reproduce, display, or distribute these materials. Students who engage in unauthorized recording, unauthorized use of a recording, or unauthorized distribution of instructional materials will be referred to the appropriate University office for follow-up.

Required Books

The vast majority of the readings are available online, but you will need to purchase two books.

Unmasking AI by Joy Buolamwini (2023)
The Alignment Problem by Brian Christian (2020)

Further Background Reading (Optional)

Computer Power and Human Reason by Joseph Weizenbaum (1976)
Artificial Intelligence, A Guide for Thinking Humans by Melanie Mitchell (2019)
Race After Technology by Ruha Benjamin (2019)
The Second Machine Age by Andrew McAfee and Eric Brynjolfsson (2014)
Rebooting AI: Building AI We Can Trust by Gary Marcus, Ernest Davis (2021)
Power and Prediction by Ajay Agrawal, Avi Goldfarb, Joshua Gans (2022)
Genius Makers by Cade Metz (2022)

MORS 950 Winter 2024 Schedule (Jan 3 to Mar 7)

1. Artificial Intelligence: What, Why, and Where?

January 3rd in Evanston and January 4th in Chicago

What is artificial intelligence (AI)? What is AI capable of today? How have the capabilities changed over time? What is AI’s relationship to machine learning, deep learning, algorithms, and analytics? Why is AI relevant to business? When and where is AI well-suited and ill-suited for solving business problems?

Before Class Reading:

- [Computing Machinery and Intelligence](#) by Alan Turing (1950)
- [Why AI is Harder than We Think](#) by Melanie Mitchell (2021)

MORS 950 Human and Machine Intelligence (Professor Matt Grob)

Northwestern | Kellogg

Additional Optional Reading:

- [GPT-4 Technical Report](#) by Open AI (2023)
- [The Dawn of LLMs: Preliminary Explorations with GPT-4V\(ision\)](#) by Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, Lijuan Wang (2023)
- [New Navy Device Learns by Doing](#) in New York Times (1958)

2. Machine Learning: How?

January 4th in Evanston and January 4th in Chicago

How does a machine learn? How can machine learning be used to make predictions in business? How does the machine learning business flywheel work and where can it break down?

Before Class Reading:

- [Algorithms need managers too](#) by Michael Luca, Jon Kleinberg, and Sendhil Mullainathan (2017)
- “Introduction” (pages ix to xxi) in *Unmasking AI* by Joy Buolamwini
- [The Definitive Guide to Machine Learning for Business Leaders](#) by Hugo Bowne-Anderson (2020)

Additional Optional Reading:

- [But, what is a neural network? | Chapter 1 Deep Learning](#) by Grant Sanderson on the 3Blue1Brown Youtube channel (2017)
- [Prediction and explanation in social systems](#) by Jake Hofman, Amit Sharma, and Duncan Watts (2017)
- [Computer Scientist Explains Machine Learning in 5 Levels of Difficulty](#) by Hilary Mason (2021)

3. Computational Power and Deep Learning

January 8th in Evanston and January 11th in Chicago

What is compute and why does it matter? What is deep learning? What’s a neural network in machine learning, and how related is it to a biological neural network in the human brain? How can computer vision and natural language processing be tools to address business problems?

Before Class Reading:

- [Computational Power and AI](#) by Jai Vipra and Sarah Myers West (2023)
- “A Parable of Three Entrepreneurs” (pages 3-24) in *Power and Prediction: The Disruptive Economics of AI* by Ajay Agrawal, Joshua Gans, and Avi Goldfarb (2022)

Additional Optional Reading:

- [What is Computer Vision and Why Does it Matter](#) by NVIDIA (2022)
- [How Computer Vision Works](#) by Sara Robinson (2018)

4. Evaluating Algorithmic Performance

Due: Please submit assignment 1 “Large Language Models for Solving Data Science Problems in Business”

January 11th in Evanston and January 11th in Chicago

Northwestern | Kellogg

How do you evaluate AI systems and ML models? When should we care about accuracy, precision, recall, area under the curve, and other metrics? What trade-offs emerge? When are AI systems likely to drift, err, or otherwise go awry? When can we compare machine predictions to human decisions?

Before Class Reading:

- “Shield Ready” (pages 28 to 38) in Unmasking AI by Joy Buolamwini
- [Reliance on Metrics is a Fundamental Challenge for AI](#) by Rachel Thomas and David Uminsky (2020)
- “Representation” (pages 17 to 50) in The Alignment Problem by Brian Christian (2020)

Additional Optional Reading:

- [On the Folly of Rewarding A, While Hoping for B](#) by Steven Kerr (1975)
- [Performance vs. Competence in Human-Machine Comparisons](#) by Chaz Firestone (2020)

5. Evaluating Algorithmic Performance in Practice

January 17th in Evanston and January 18th in Chicago

What are the systematic yet surprising errors that crop up in business applications of AI systems and ML models? What is algorithmic bias? Why did Microsoft’s chatbot Tay fail? Why can GPT-4 pass the bar exam and medical licensing exam but fail at other basic tasks? When is computer vision prone to errors?

Before Class Reading:

- **Case Study:** [Challenges in Commercial Deployment of AI: Insights from the Rise and Fall of IBM Watson’s AI Medical Program](#) by Quy Huy, Timo Vuori, Tero Ojanpera, Lisa Simone Duke (2023)
- “Defaults Are Not Neutral and Facial Recognition Technologies” (pages 41 to 67) in Unmasking AI by Joy Buolamwini

Additional Optional Reading:

- [The Mythos of Model Interpretability: In machine learning, the concept of interpretability is both important and slippery](#) by Zachary Lipton (2018)
- [Beware explanations from AI in health care](#) by Boris Babic, Sara Gerke, Theodoros Evgeniou, and I. Glenn Cohen (2021)
- [The Final 11 seconds of a fatal Tesla autopilot crash](#) by Trisha Thadani, Rachel Lerman, Imogen Piper, Faiz Siddiqui, and Irfan Uraizee

6. Generative AI

Due: Please submit team and initial ideas for final project

January 18th in Evanston and January 18th in Chicago

Before Class Reading:

- [How to Capitalize on Generative AI](#) by Andrew McAfee, Daniel Rock and Erik Brynjolfsson

Northwestern | Kellogg

- [Art and Science of Generative AI](#) by Ziv Epstein, Aaron Hertzman, Memo Atken, Hany Farid, Jessica Fjeld, Morgan Frank, Matthew Groh, Laura Herman, Neil Leach, Robert Mahari, Alex Pentland, Olga Russakovsky, Hope Schroeder, and Amy Smith (2023)
- [AI Prompt Engineering isn't the Future](#) by Oguz Acar (2023)

Additional Optional Reading:

- [DAIR Prompt Engineering Guide](#) by Elvis Saravia (2023)
- [How Generative AI Can Augment Human Creativity](#) (2023) by Tojin Eapen, Daniel Finkenstadt, Josh Folk, Lokesh Venkataswamy
- [Fact Sheet: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence](#) (2023) US Government

7. Language Models: Stochastic Parrots or Conscious Computers

January 22nd in Evanston and January 25th in Chicago

Before Class Reading:

- "Introduction" (pages 1-12) in *Computer Power and Human Reason* by Joseph Weizenbaum (1976)
- [On the Dangers of Stochastic Parrots: Can Language Models Be Too Big](#) by Emily Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell (2021)
- [Could a Large Language Model Be Conscious](#) by David Chalmers

Additional Optional Reading:

- [Faith and Fate: Limits of Transformers on Compositionality](#) by Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jiang, Bill Lin, Peter West, Chandra Bhagavatula, Ronan Le Bras, Jena Hwang, Soumya Sanyal, Sean Welleck, Xiang Ren, Allyson Ettinger, Zaid Hacchaoui, and Yejin Choi
- [Human-like systematic generalization through a meta-learning neural network](#) by Brenden Lake and Marco Baroni

8. Recommender Systems: Who Wants to See What When?

Due: Please submit assignment 2 "AI Generated Creative Campaign from Design to Marketing"

January 25nd in Evanston and January 25th in Chicago

Before Class Reading:

- [The Netflix Recommender Systems: Algorithms, Business Value, and Innovations](#) by Carlos Gomez-Uribe and Neil Hunt (2015)
- [How the New York Times Recipe Team Makes Personalized Cooking Recommendations](#) by Kyelee Fitts and Celia Eddy (2023)

Additional Optional Reading:

- [Twitter's Recommendation Algorithm](#) by Twitter
- [How TikTok Reads Your Mind](#) by Ben Smith (2021)
- [Recommender Systems](#) by Charu Aggarwal (2016)

9. Beating Human Grandmasters in Chess, Go, Poker, Diplomacy, and More

MORS 950 Human and Machine Intelligence (Professor Matt Groh)

Northwestern | Kellogg

January 29th in Evanston and February 1st in Chicago

Before Class Reading:

- “Curiosity” (pages 181 to 210) in *The Alignment Problem* by Brian Christian
- [How AI Conquered Poker](#) by Keith Romer (2022)
- [AI learns the art of Diplomacy](#) by Matthew Hutson (2022)

Additional Optional Reading:

- [Chess, a Drosophila of Reasoning](#) by Garry Kasparov (2018)
- [Acquisition of Chess Knowledge in AlphaZero](#) by Thomas McGrath, Andrei Kapishnikov, Nenad Tomasey, Adam Pearce, Martin Wattenberg, Demis Hassabis, Been Kim, Ulrich Paquet, Vladimir Kramnik (2022)

10. Midterm

February 1st in Evanston and February 1st in Chicago

Due: Please come prepared for the midterm

11. Identifying and Evaluating AI Opportunities

February 5th in Evanston and February 8th in Chicago

Before Class Reading:

- [A Simple Tool to Start Making Decisions with the Help of AI](#) by Ajay Agrawal, Joshua Gans, and Avi Goldfarb
- [What AI-Driven Decision Making Looks Like](#) by Eric Colson

Additional Optional Reading:

- [Opportunities in AI](#) by Andrew Ng (2023)
- [How to Spot a Machine Learning Opportunity, Even if You Aren't a Data Scientist](#) by Kathryn Hume (2017)

12. Guest Lecture

Due: Please submit AI Canvas and Business/Policy Model Canvas for final project

February 8th in Evanston and February 8th in Chicago

Before Class Reading:

- To be announced

Additional Optional Reading:

- To be announced

13. Fairness, Accountability, and Transparency in AI

February 12th in Evanston and February 15th in Chicago

Before Class Reading:

Northwestern | Kellogg

- “Guardians Assemble, Power Shadows, Gender Shades, and Deserted Desserts” (pages 68 to 85, 125 to 154) in *Unmasking AI* by Joy Buolamwini
- [The AI Wars Have Three Factions, and They All Crave Power](#) by Bruce Schneier and Nathan Sanders (2023)

Additional Optional Reading:

- [The Alignment Problem from a Deep Learning Perspective](#) by Ricard Ngo, Lawrence Chan, and Sören Mindermann
- [Engineering Knowledge: The Construction of Knowledge in Artificial Intelligence](#) by Diana Forsythe (1993)

14. Managing Data Science and Machine Learning Teams

February 15th in Evanston and February 15th in Chicago

Due: Please submit assignment 3 “Interview with a Data Scientist”

Before Class Reading:

- [4 Skills the Next Generation of Data Scientists Need to Develop](#) by Joel Shapiro (2023)
- [What Data Scientists Really Do, According to 35 Data Scientists](#) by Hugo Bowne-Anderson (2018)

Additional Optional Reading:

- [Maker Manager Schedule](#) by Paul Graham (2009)
- [Please don't hire a Chief AI officer](#) by Kristian Hammond (2017)

15. Managing Human and AI Teams

February 19th in Evanston and February 22nd in Chicago

Before Class Reading:

- [Will A.I. Become the New McKinsey](#) by Ted Chiang (2023)
- [Biased Algorithms are Easier to Fix than Biased Humans](#) by Sendhil Mullainathan (2019)
- [The Super Mario Effect](#) by Mark Rober (2018)

Additional Optional Reading:

- [The H-Metaphor as a Guideline for Vehicle Automation and Interaction](#) by NASA (2003)
- [Do Algorithms Beat Us at Complex Decision Making](#) by Shane Parrish
- [MBA Students vs. ChatGPT on Generating Innovative Ideas](#) by Christian Terwiesch and Karl Ulrich (2023)

16. Affective Computing

February 22nd in Evanston and February 22nd in Chicago

Before Class Reading:

- **Case Study:** [Feeling Machines: Emotion AI at Affectiva](#) by Shane Greenstein and John Masko (2019)

Northwestern | Kellogg

- [Affective Computing](#) by Rosalind Picard (1995)

Additional Optional Reading:

- [Can AI Do Empathy Even Better Than Humans? Companies Are Trying It.](#) by Lisa Bannon (2023)
- [The Media Equation](#) by Byron Reeves and Clifford Nass (1996)

17. Deepfakes and Synthetic Media

Due: Please submit story telling techniques for final presentation

February 26th in Evanston and February 29th in Chicago

Before Class Reading:

- [Creating, Using, Misusing, and Detecting Deepfakes](#) by Hany Farid (2022)

Additional Optional Reading:

- [Deepfake Detection by Human Crowds, Machines, and Machine-Informed Crowds](#) by Matt Groh, Ziv Epstein, Chaz Firestone, and Rosalind Picard (2022)

18. Being Human in the Age of AI

February 29th in Evanston and February 29th in Chicago

Before Class Reading:

- “Conclusion” (pages 310 to 334) in *The Alignment Problem* by Brian Christian
- “Seat at the Table” (pages 283 to 291) in *Unmasking AI* by Joy Buolamwini

Additional Optional Reading:

- [Leading Thinkers on AI and What it Means to Be Human](#) by Maria Popova (2015)
- [Meaning in the Age of AI](#) by Maria Popova (2023)
- [Meaning in the Age of AI before ChatGPT](#) by Maria Popova (2023)
- [The Turing Trap: The Promise and Peril of Human-Like AI](#) by Erik Brynjolfsson (2022)

19 and 20. Final Presentations

March 4th and 7th in Evanston and March 7th in Chicago

Due: Please submit materials for your final presentation