

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 | TBD, Evanston
Section 81 | Saturdays | 9:00 am - 12:00 pm | TBD, Chicago

Office Hours: Schedule available on Canvas and by appointment.

Course overview

Artificial intelligence (AI) is a transformative technology and a modern day equivalent to fire in the early stages of human civilization. Organizations can use AI to solve complex problems, automate tasks, make predictions, and even boost human creativity. However, like fire, AI also carries significant risks when misused. This course requires no prior technical knowledge and is designed for students who want to lead and consult on the deployment of AI systems in organizations, manage data science and design teams, or 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 emerging from the capabilities and limitations of AI systems including ChatGPT, Copilot, Midjourney, Watson, AlphaZero, TikTok's recommender systems, and more, and we'll 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, quizzes, a midterm, and a final group project. By the end, you should be an expert at identifying promising business use-cases of frontier AI technologies and evaluating AI's current limitations such that you can apply human and machine thought partnerships to grow new businesses and disrupt Grand Masters in any field.

Required Books (please purchase these before the beginning of the quarter)

AI Snake Oil by Arvind Narayanan and Sayash Kapoor (2024)

The Alignment Problem by Brian Christian (2020)

Unmasking AI by Joy Buolamwini (2023)

Northwestern University Syllabus Standards

This course follows the [Northwestern University Syllabus Standards](#). Students are responsible for familiarizing themselves with this information. Students can find useful resources for safety and security, academic support, and mental and physical health and well-being on the [NU help website](#). Unauthorized copying or distribution of any course materials is strictly prohibited.

Class Attendance and Screen Policy

Students are expected to attend every class in-person unless they're feeling sick or have an approved absence. No laptops, tablets, or cell phones are allowed to be used in class unless they're required for classwork. I highly encourage note taking with a pen and notebook.

Grading

Assignments (25%) Three individual assignments will give you first-hand experience in applying AI towards solving problems in business. Each assignment is worth 10% of your grade. **All assignments are due one hour before the start of class.** Late assignments will incur a 10% penalty for the number of days the assignment is late (e.g. 1 minute late is a 10% deduction, 25 hours late is a 20% deduction, 3 days late is 30%, and so on).

Quizzes (10%) In class quizzes will give you feedback on your grasp of the course work.

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

Final Project (30%) A group project will offer experience identifying, thinking through, and artfully pitching an AI-based product and the organizational settings where such a product would be most successful: a new startup, a large corporation, a non-profit, a government, or something else. 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 grade on the final project.

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.

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. Do not blindly trust AI outputs. LLMs are prone to making up facts. 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 clearly mark the response as AI generated. Failure to do so is in violation of academic honesty policies.

Optional Background Reading (Optional)

Computer Power and Human Reason by Joseph Weizenbaum (1976)
The Second Machine Age by Andrew McAfee and Eric Brynjolfsson (2014)
Artificial Intelligence, A Guide for Thinking Humans by Melanie Mitchell (2019)
Race After Technology by Ruha Benjamin (2019)
Rebooting AI: Building AI We Can Trust by Gary Marcus, Ernest Davis (2021)
Power and Prediction by Ajay Agrawal, Avi Goldfarb, Joshua Gans (2022)
The Worlds I See by Fei-Fei Li (2023)
The Coming Wave by Mustafa Suleyman (2023)
Why Machines Learn by Anil Ananthaswamy (2024)

Podcasts, Substacks, and Ted Talks (Optional)

Complexity: Nature of Intelligence by Sante Fe Institute
The AI Daily Brief by Nathaniel Whittemore
The AI in Business Podcast by Daniel Faggella
[Luiza's Policy Newsletter](#) Substack
[AI Snake Oil](#) Substack
[Ted Talks on AI](#)

MORS 950 Winter 2025 Schedule (Jan 6th to Mar 15)

Before class begins, please purchase the three required books: **AI Snake Oil** by Arvind Narayanan and Sayash Kapoor, **Unmasking AI** by Joy Buolamwini and **The Alignment Problem** by Brian Christian. Please also purchase the course pack which has required reading for class.

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

January 6th in Evanston and January 11th in Chicago

What is artificial intelligence, and what are its most mind blowing capabilities? How have the capabilities and limitations changed over time? What is AI's relationship to machine learning, deep learning, algorithms, and analytics? When and where is AI well-suited and ill-suited for solving business problems?

Required Reading (Please read before the first class):

- [Computing Machinery and Intelligence](#) (pages 1-6) by Alan Turing (1950)
- [The Metaphors of AI](#) by Melanie Mitchell (2024)
- [Why AI is Harder than We Think](#) by Melanie Mitchell (2021)

Suggested Optional Reading:

- [World's Most Used AI Tools](#) (2024)
- [GPT-4 Technical Report](#) by Open AI (2023)
- [Gemini: A Family of Highly Capable Multimodal Models](#) by Google's Gemini Team (2024)
- [In the Age of AI, Major in Being Human](#) by David Brooks (2023)

2. Machine Learning: How?

January 9th in Evanston and January 11th in Chicago

How does machine learning differ from statistics? How does supervised learning work? How does unsupervised learning work? How can machine learning be used to make predictions in business? Why does machine learning need managers? How does the machine learning business flywheel work and where can it break down?

Required Reading:

- "Introduction (pages 1 to 35) in **AI Snake Oil**
- "Introduction" (pages ix to xxi) in **Unmasking AI**
- [Algorithms need managers too](#) by Michael Luca, Jon Kleinberg, and Sendhil Mullainathan (2017)
- [Computer Scientist Explains Machine Learning in 5 Levels of Difficulty](#) by Hilary Mason (2021)

Suggested Optional Reading:

- [Artificial Intelligence, Economics, and Industrial Organization](#) by Hal Varian (2018)
- [Prediction and explanation in social systems](#) by Jake Hofman, Amit Sharma, and Duncan Watts (2017)

3. Deep Learning, Computer Vision, and Natural Language Processing (Quiz Day!)

January 13th in Evanston and January 18th in Chicago

What is deep learning? What's a neural network? How is it related to a biological neural network in the human brain? How is deep learning useful for computer vision and NLP? Where can computer vision and NLP be applied in business? When is computer vision and NLP prone to errors and bias?

Required Reading:

- "Representation" (pages 17 to 29) in **The Alignment Problem**
- [But, what is a neural network? | Chapter 1 Deep Learning](#) by Grant Sanderson on the 3Blue1Brown Youtube channel (2017)

Suggested Optional Reading:

- [Global Deep Technology Startup Stories](#) by Federico Antoni, Stephen Ciesinski, Andrew Leon Hanna, and Howard Rozen (2023)
- [How Computer Vision Works](#) by Sara Robinson (2018)
- [The Illustrated Word2Vec](#) by Jay Alammar

4. Computational Power and Internet Scale Data

Due: Please submit assignment 1 "Large Language Models for People Analytics"

January 16th in Evanston and January 18th in Chicago

What is compute and why does it matter? What kind of data are useful for deep learning and large language models? How has compute, data, and algorithms enabled the deep learning revolution of 2010s and generative AI revolution of the 2020s? How is AI similar to and different from a 21st century version of electricity and what does this mean for managers?

Required Reading:

- [Training Data for the Price of a Sandwich](#) by Stefan Baack (2024)
- [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)

Suggested Optional Reading:

- [How People Are Really Using GenAI](#) by Marc Zeo-Sanders (2024)

- [What can machine learning do? Workforce implications](#) by Erik Brynjolfsson and Tom Mitchell (2017)
- [How AI Could Help Rebuild the Middle Class](#) by David Autor (2024)

5. Evaluating Algorithmic Performance: How and When?

January 22nd in Evanston and January 25th in Chicago

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? What are benchmarks and when are AI systems likely to drift, err, or otherwise go awry? When can we compare machine predictions to human decisions?

Required Reading:

- [Assessing Prediction Accuracy of Machine Learning Models](#) by Michael Toffel, Natalie Epstein, Kris Ferreira, and Yael Grushka-Cockayne (2020)
- “How Predictive AI Goes Wrong” (pages 36 to 59) in **AI Snake Oil**
- “Shield Ready” (pages 28 to 38) in **Unmasking AI**
- “Representation” (pages 29 to 50) in **The Alignment Problem**

Suggested Optional Reading:

- [Reliance on Metrics is a Fundamental Challenge for AI](#) by Rachel Thomas and David Uminsky (2020)
- [Performance vs. Competence in Human-Machine Comparisons](#) by Chaz Firestone (2020)
- [On the Folly of Rewarding A, While Hoping for B](#) by Steven Kerr (1975)

6. Evaluating Algorithmic Performance is a Leadership Problem

Due: Please submit team and three initial ideas per person for final project

January 24th in Evanston and January 25th in Chicago

What are the systematic yet surprising errors that crop up in business applications of AI systems and ML models? Why did IBM’s major investment in Watson fail? Why did Microsoft’s chatbot Tay fail? Why did Google’s 2024 launch of Gemini and its AI-generated summaries go awry? Why can GPT-4 pass the bar exam and medical licensing exam but fail at other basic tasks?

Required 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**

7. Generative AI (Quiz Day!)

January 27th in Evanston and February 1st in Chicago

What is generative AI, what are multimodal models, and how does self-supervised learning work? Specifically, how do text-to-image diffusion models? How might generative AI transform the future of creativity, labor, law, and media? What business frameworks can help us harness this technology to identify real-world business applications?

Required Reading:

- “The Long Road to Generative AI” and “Why Do Myths about AI Persist” (pages 99 to 149 and 227 to 257) in [AI Snake Oil](#)
- [The Intelligence Age](#) by Sam Altman (2024)
- [Hype, Sustainability, and the Price of Bigger is Better Paradigm in AI](#) by Gael Varoquaux, Alexandra Sasha Luccioni, and Meredith Whittaker (2024)

Suggested Reading:

- [How to Capitalize on Generative AI](#) by Andrew McAfee, Daniel Rock and Erik Brynjolfsson (2023)
- [How Generative AI Can Augment Human Creativity](#) (2023) by Tojin Eapen, Daniel Finkenstadt, Josh Folk, Lokesh Venkataswamy
- [AI Prompt Engineering isn't the Future](#) by Oguz Acar (2023)
- [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)

8. Large Language Models

Due: Please submit assignment 2 “Organization Onboarding with Generative AI”

January 30th in Evanston and February 1st in Chicago

How do LLMs work, where can LLMs offer practical value, and how are they transforming businesses? What are the capabilities and limitations of LLMs? How does evaluating LLMs differ from evaluating classification and regression? What is the Eliza effect, and how does it apply to LLMs?

Required Reading:

- [Generative AI's Act on](#) by Sonya Huang and Pat Grady (2024)
- “The Eyes of a Machine” (pages 346-381) in *Why Machines Learn* by Anil Ananthaswamy (2024)
- [What have language models learned?](#) by Google's People and AI Research Team
- “Introduction” (pages 1-12) in *Computer Power and Human Reason* by Joseph Weizenbaum (1976)
- [Debates on the nature of artificial general intelligence](#) by Melanie Mitchell (2024)

Suggested Reading:

- [Introduction to LLMs](#) by Andrej Karpathy (2023)

- [The magical AI assistants of the future – and the engineering behind them](#) by Harrison Chase (2023)
- [How Large Language Models Reflect Human Judgment](#) by Ajay Agrawal, Joshua Gans, and Avi Goldfarb
- [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 (2022)
- [Embers of autoregression show how large language models are shaped by the problem they are trained to solve](#) by Thomas McCoy, Shunyu Yao, Dan Friedman, Matthew Hardy, and Thomas Griffiths (2024)

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

February 3rd in Evanston and February 8th in Chicago

How do large companies connect expansive product offerings with diverse customers? What are the business decisions involved in designing and deploying recommender systems? What can go wrong with recommender systems?

Required Reading:

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

Suggested Reading:

- [Auditing YouTube’s recommendation system for ideologically congenial, extreme, and problematic recommendations](#) by Muhammad Haroon, Magdalena Wojcieszak, Anshuman Chhabra, Xin Liu, Prasant Mohapatra, and Zubair Shafiq (2023)
- [Twitter’s Recommendation Algorithm](#) by Twitter (2023)
- [How TikTok Reads Your Mind](#) by Ben Smith (2021)
- [Recommender Systems](#) by Charu Aggarwal (2016)

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

February 6th in Evanston and February 8th in Chicago

How does reinforcement learning work? How are AI systems built such that they can beat human grandmasters at complex games? When and why are AI systems’ performance on games relevant to the real-world? How does reinforcement learning with human feedback influence LLMs?

Required Reading:

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

Suggested 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)
- [The Necessity of Awe](#) (2020) by Helen De Cruz
- Chapter 9 “Game On” Artificial Intelligence A Guide for Thinking Humans by Melanie Mitchell

11. Midterm

February 10th in Evanston and February 15th in Chicago

Due: Please come prepared for the in-class midterm

12. Identifying and Evaluating AI Opportunities

February 13th in Evanston and February 15th in Chicago

What business models work for building AI-based startups? What is the “AI Canvas” and how do we apply it to AI products? From a managerial perspective, how can we integrate the “AI Canvas” with the “Business Model Canvas” and what do examples of “AI First Business Model Canvases” look like for AI-based companies?

Required Reading:

- [The New Business of AI \(and How It's Different from Traditional Software\)](#) by Martin Casado and Matt Bornstein (2020)
- [Opportunities in AI](#) by Andrew Ng (2023)
- [A Simple Tool to Start Making Decisions with the Help of AI](#) by Ajay Agrawal, Joshua Gans, and Avi Goldfarb

Suggested Reading:

- [What AI-Driven Decision Making Looks Like](#) by Eric Colson (2019)
- [How to Spot a Machine Learning Opportunity, Even if You Aren't a Data Scientist](#) by Kathryn Hume (2017)

13. AI Safety & Fairness

Guest Lecture

February 17th in Evanston and February 22nd in Chicago

What strategies do you pursue to limit risks while enabling general purpose technologies like LLMs, diffusion models, and AI agents? As a program manager at OpenAI, Anthropic, Google, Microsoft, Apple, Midjourney, Perplexity, or another generative AI company releasing a (a) diffusion model text-to-image platform (b) large language model, how do you navigate a successful product launch and avoid harmful, errant, and biased generations? What is “alignment”, and how do we identify human values and collective preferences? How do we precisely define fairness and what trade-offs emerge in

applications like credit scoring and people analytics? How do we fix word associations that match the historical data but do not match our social values? What is the role of algorithmic and evocative audits in transforming the use of AI in business?

Required Reading:

- “Is Advanced AI an Existential Threat” (pages 150 to 178) in **AI Snake Oil**
- “Guardians Assemble, Power Shadows, Gender Shades, and Deserted Desserts” (pages 68 to 85, 125 to 154) in *Unmasking AI* by Joy Buolamwini
- [Biased Algorithms are Easier to Fix than Biased Humans](#) by Sendhil Mullainathan (2019)

Suggested Reading:

- [The AI Wars Have Three Factions, and They All Crave Power](#) by Bruce Schneier and Nathan Sanders (2023)
- [OpenAI’s Approach to External Red Teaming for AI Models and Systems](#) by Lama Ahmad, Sandhini Agarwal, Michael Lampe, Pamela Mishkin

14. Leveraging AI for Persuasive Pitches

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

February 20th in Evanston and February 22nd in Chicago

What makes a business pitch persuasive? How well can humans and AI predict pitch success and failure? How could AI help us improve our pitches?

Required Reading:

- [The Secret to Successfully Pitching an Idea](#) by Mar Hershenson (2023)
- [The Secret Structure of Great Talks](#) by Nancy Duarte (2011)

Suggested Reading:

- [Entrepreneurial Pitching: A Critical Review and Integrative Framework](#) by Sai Gayathri Kalvapalle, Nelson Phillips, and Joep Cornelissen (2024)
- [Visuals Dominate Investors Decisions about Entrepreneurial Pitches](#) by Chia Jung-Tsay (2024)
- [Narrative Reversals and Story Success](#) by Samsun Knight, Matthew Rocklage, and Yakov Bart (2024)
- [A Fused Large Language Model for Predicting Startup Success](#) by Abdurahman Maarouf, Stefan Feuerriegel, Nicolas Prolochs (2024)

15. AI in the Music Industry

Guest Lecture

February 24th in Evanston and March 1st in Chicago

What capabilities emerge from self-supervised learning models trained on music? What tools and use cases are transforming the music industry? Where are the human elements in AI-generated music?

What is the difference between copying and inspiration by humans and machines in the production of music?

Required Reading:

- [The Future of Music: How Generative AI is Transforming the Music Industry](#) by Justine Moore and Anish Acharya (2023)
- [Jukebox](#) by OpenAI
- [MusicLM: Generating Music from Text](#) Andrea Agostinelli et al (2023)

Suggested Reading:

- [How generative AI could disrupt creative work](#) by David De Cremer, Nicola Morini Bianzino, and Ben Falk
- [5 ways AI has already changes the music industry](#) by Elias Light and Kristin Robinson

16. Managing Artificial Intelligence and Leading with Data

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

February 27th in Evanston and March 1st in Chicago

What do data scientists and machine learning engineers actually do, and how do we most effectively manage them? Why are managers’ perspectives on meetings different from the scientists’ and engineers’ perspectives? How can we lead with data based on the problem spotting, problem scoping, problem shepherding, and solution translating framework?

Required 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)
- [Maker Manager Schedule](#) by Paul Graham (2009)

Suggested Reading:

- [Please don’t hire a Chief AI officer](#) by Kristian Hammond (2017)
- [How Busy People Innovate](#) by Eric Athas (2024)

17. Affective Computing (Quiz Day!)

Due: Please submit Pitch Canvas for final presentation

March 3rd in Evanston and March 8th in Chicago

How well can AI systems recognize and respond to emotions? Where are the business applications of AI-based emotion recognition and empathic communication? How does the Affectiva case study reveal the promise build a business around emotion recognition technology, what was the promise of the business, and what technical factors limited business opportunities?

Required Reading:

- **Case Study:** [Feeling Machines: Emotion AI at Affectiva](#) by Shane Greenstein and John Masko (2019)
- [Can AI Do Empathy Even Better Than Humans? Companies Are Trying It.](#) by Lisa Bannon (2023)

Suggested Reading:

- [Affective Computing](#) by Rosalind Picard (1995)
- [The Media Equation](#) by Byron Reeves and Clifford Nass (1996)

18. Being Human in the Age of AI

March 6 in Evanston and March 8th in Chicago

How is human learning similar and different from machine learning? What are the implications of Alan Turing's imitation game for humans – what would it mean to be the “most human human”? How can you predict how AI will transform jobs and careers? What's changed in AI over the last 10 weeks, how can we make sense of the latest developments with this course, and how can you most effectively apply human and machine thought partnerships to grow new businesses and disrupt Grand Masters in any field?

Required 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
- [The Super Mario Effect](#) by Mark Rober (2018)

Suggested Reading:

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

19 and 20. Final Presentations

March 10th and 13th in Evanston and March 15th in Chicago

Due: Please submit slidedeck for your final presentation