

Northwestern | Kellogg

AI Foundations for Managers (AIML 901)

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

Section 31 Tuesdays and Fridays	10:30 AM to 12:00 PM	L130, Global Hub
Section 81 Saturdays	9:00 AM to 12:00PM	350, Wieboldt

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 is a practical, fundamentals of AI in business course designed for students who want to lead and consult on the deployment of AI systems in organizations, manage data science and software teams, or build and invest in AI companies. You will learn the conceptual foundations for how AI uses machine learning with neural networks to help computers see (computer vision), talk and write (large language models), and learn by trial and error (reinforcement learning). Moreover, you will build intuition for what AI can (and cannot yet) do and discover where AI delivers value across industries but also where it fails spectacularly by examining a series of business applications, real-world examples, ethical considerations, and evaluation frameworks. This course is an interactive lecture-based course with four interrelated individual assignments and a final. The course is continually updated with the latest science and industry insights and is based on the first five weeks of Professor Groh's 2024 and 2025 Human and Machine Intelligence MORS950 course. The goal is to set you on your way to leading AI adoption in any organization and being equipped for AIML950, which dives deeper into how humans and AI interact, what that means for leadership and organizations, and how to successfully adopt AI in organizations.

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 except when we're using laptops for in-class exercises. This policy is designed to avoid distractions for you and your classmates. I highly encourage note taking with a pen and notebook because research shows that writing by hand is better for memory and learning than typing (see Flanigan et al 2024 "Typed versus Handwritten").

Northwestern | Kellogg

Grading

Final (20%) A final will give you feedback on your grasp of the course work.

Assignments (60%) Four individual assignments will give you first-hand experience in applying AI towards solving problems in business. Each assignment is worth 15% 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).

Participation (20%) You are expected to attend every class and actively engage in group discussions. Your participation grade is based on (a) your attendance and participation in class (b) contributing shared lecture notes or key takeaways for one class session, summarizing the main ideas, examples, and discussion points for your peers

AI Policy

I expect you to think critically and use AI as a tool for assignments in this class. Do not blindly trust AI outputs. Do not copy AI outputs – use your own words. You are responsible for errors and omissions in the assignments you submit.

Optional Background Reading

Computer Power and Human Reason by Joseph Weizenbaum (1976)*
Artificial Intelligence, A Guide for Thinking Humans by Melanie Mitchell (2019)*
The Alignment Problem by Brian Christian (2020)*
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)
AI 2041 by Kai-Fu Lee and Chen Qiufan (2024)
Why Machines Learn by Anil Ananthaswamy (2024)*
AI Snake Oil by Arvind Narayanan and Sayash Kapoor (2024)*
Co-Intelligence: Living and Working with AI by Ethan Mollick (2024)*
The AI Con by Emily Bender and Alex Hanna (2025)
These Strange New Minds by Christopher Summerfield (2025)

* indicated that at least one chapter of the book is assigned in this class

Podcasts, Youtubes Substacks, and Ted Talks (Optional)

[Complexity: Nature of Intelligence](#) (Podcast) by Sante Fe Institute
[One Useful Thing](#) Substack by Ethan Mollick
[Luiza's Policy Newsletter](#) Substack by Luiza Jarovsky
[Normal Technology](#) Substack by Arvind Narayanan and Sayash Kapoor
[Interconnects](#) by Nathan Lambert
[Dwarkesh Podcast](#) by Dwarkesh Patel
[Simon Willison's Weblog](#)
[Stratechery on AI](#) by Ben Thompson '11
[AI: A Guide for Thinking Humans](#) by Melanie Mitchell
[Andrei Karpathy Youtube Channel](#)
[3Blue1Brown](#) by Grant Sanderson
[Ted Talks on AI](#) and [More Ted Talks on AI](#)

Northwestern | Kellogg

AIML 901 Winter 2026 Schedule

1. Artificial Intelligence

What is AI? Why has AI progress been historically misunderstood? Why is AI today not just a technology trend, but a core management issue?

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

2. Machine Learning

How do machines learn? Why does machine learning need managers? How does the machine learning business flywheel work?

- [Computer Scientist Explains Machine Learning in 5 Levels of Difficulty](#) by Hilary Mason (2021)
- [*Algorithms need managers too](#) by Michael Luca, Jon Kleinberg, and Sendhil Mullainathan (2017)

3. Deep Learning

What's a neural network? How did deep learning trigger the current wave of AI breakthroughs? How should firms think about the build versus buy decision for deep learning solutions?

- "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)

4. Evaluating Algorithmic Performance

When should we care about accuracy, precision, recall, area under the curve, and other metrics? How does algorithmic bias emerge? How does Goodhart's Law apply to managing algorithms?

- [*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](#) (2024)

5. Computer Vision

How do computers see? What makes vision problems uniquely valuable and uniquely risky for businesses? How do organizations build advantage from data that "sees"?

- Chapter 11 "The Eyes of a Machine" (pages 346-381) in [Why Machines Learn](#) (2024)
- Optional: Read through [Colab Notebook](#)

6. Large Language Models

How does natural language processing work? How should executives cut through hype when evaluating LLMs for their organizations? What new business risks do LLMs introduce?

- “Creating Alien Minds” (pages 3 to 26) in **Co-Intelligence** by Ethan Mollick (2024)
- “Introduction” (pages 1-12) in **Computer Power and Human Reason** by Joseph Weizenbaum (1976)

7. Reinforcement Learning

How does AI learn to beat human grandmasters at complex games? Where does RL create business value? Why are RL experiments harder to manage and govern than other AI approaches?

- Chapter 6 “Curiosity” (pages 181 to 210) in **The Alignment Problem** (2019)
- “Game On” (pages 179 to 208) in **Artificial Intelligence A Guide for Thinking Humans** by Melanie Mitchell (2019)

8. LLMs++

How can reinforcement learning create even better large language models? Why do we need to align LLMs with human goals and values? What risks emerge when AI systems optimize too much for what we say we want rather than what we actually need?

- [Generative AI’s Act o1](#) by Sonya Huang and Pat Grady (2024)
- “The Long Road to Generative AI” (pages 99 to 149) in **AI Snake Oil**

9. Agents and Multimodal LLMs

How do multimodal models handle text, images, and audio? What’s the difference between an AI model and an AI agent? How should leaders think about orchestration where AI systems work with other AI systems?

- Chapter 2 “From Paradox to Payoff: How agents can scale AI” (pages 11 to 22) in **Seizing the Agentic AI Advantage** by McKinsey’s QuantumBlack
- “Why Do Myths about AI Persist” (pages 227 to 257) in **AI Snake Oil**

10. Artificial General Intelligence and Artificial Super Intelligence

How do we precisely define AGI and ASI? What scenarios should managers consider if AI rivals human cognition? How should organizations prepare today for long-term uncertainty and opportunity?

- [Debates on the nature of artificial general intelligence](#) by Melanie Mitchell (2024)
- Chapter 9 “AI As Our Future” (pages 193-212) in **Co-Intelligence**

Northwestern | Kellogg

Session	Module	Reading + Assignments
1	Artificial Intelligence	Computing Machinery and Intelligence (pages 1-6) by Alan Turing (1950) Why AI is Harder than We Think by Melanie Mitchell (2021)
2	Machine Learning	Computer Scientist Explains Machine Learning in 5 Levels of Difficulty by Hilary Mason (2021) Algorithms need managers too by Michael Luca, Jon Kleinberg, and Sendhil Mullainathan (2017)
3	Deep Learning	Assignment 1 Due – Evaluating LLMs on Business Tasks “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)
4	Evaluation	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 (2024)
5	Computer Vision	Assignment 2 Due – Designing Business Tasks for LLMs Chapter 11 “The Eyes of a Machine” (pages 346-381) in Why Machines Learn (2024)
6	Large Language Models	“Creating Alien Minds” (pages 3 to 26) in Co-Intelligence by Ethan Mollick (2024) “Introduction” (pages 1-12) in Computer Power and Human Reason by Joseph Weizenbaum (1976)
7	Reinforcement Learning	Assignment 3 Due – Human Solutions to Peer Challenges Chapter 6 “Curiosity” (pages 181 to 210) in The Alignment Problem (2019) “Game On” (pages 179 to 208) in Artificial Intelligence A Guide for Thinking Humans by Melanie Mitchell (2019)
8	LLMs++	Generative AI's Act o1 by Sonya Huang and Pat Grady (2024) “The Long Road to Generative AI” (pages 99 to 149) in AI Snake Oil
9	Agents	Assignment 4 Due – Evaluation of Human vs. Machine Chapter 2 “From Paradox to Payoff: How agents can scaleAI” (pages 11 to 22) in Seizing the Agentic AI Advantage by McKinsey's QuantumBlack “Why Do Myths about AI Persist” (pages 227 to 257) in AI Snake Oil
10	Artificial Super Intelligence	Debates on the nature of artificial general intelligence by Melanie Mitchell (2024) Chapter 9 “AI As Our Future” (pages 193-212) in Co-Intelligence
11	Final	Take Home Final