

AI in Human Communication
Communication and Media 408-002
Josh Pasek, Ph.D.
Fall, 2023

Course Meetings:
TTh – 8:30-10:00AM
1110 NQ
jpasek@umich.edu

Drop-in Hours:
Tues – 10:00-11:00AM
5413 North Quad
Zoom (by appt. only)

Drop-in Office Hours

Tuesdays 10-11AM in 5413 North Quad and Zoom or by appointment (email: jpasek@umich.edu).

If you want to meet via Zoom, please sign up **in advance** at:

https://calendar.google.com/calendar/u/0/appointments/schedules/AcZssZ0PNCMvLfT9WPrdOjAfi0wB3YqK1JMC8kqNSLABn-Kx7k_mTK1-CUziVpt8ehEX4F5N6RHy_72n

Once you have signed up, the Zoom link is:

<https://umich.zoom.us/j/93710923108>

As we venture into the realm of *AI and Human Communication*, we find ourselves standing at the precipice of a new era. The ever-evolving landscape of artificial intelligence (AI) and particularly Generative Artificial Intelligence (GenAI or GAI) offers unprecedented opportunities and challenges that shape our societies and the way we interact with the world. In an era where AI tools are becoming increasingly integrated into our daily lives, it is crucial to explore their impact on human communication and understand the intricate dynamics that unfold.

Throughout this course, we will embark on an intellectual journey that transcends boundaries and delves into the multifaceted world of AI. Together, we will unravel the enigma of generative AI tools, unraveling their inner workings and uncovering their potential to revolutionize industries, reshape education, and influence various aspects of our lives.

Delving into this realm, we grapple with the ethical considerations that arise when AI intersects with human communication. We will critically examine issues of bias, fairness, and transparency in AI systems, and explore possibilities for mitigating these challenges. Through explorations, discussions, and debates, we examine the social and political

implications of AI, attempting to reach a nuanced understanding of its implications on society.

Drawing inspiration from the transformative nature of scientific discoveries, we will contemplate the future trajectory of AI and its potential impact. We will explore the possibilities of AI-driven advancements in the workplace, education, and scientific research, while considering the potential risks and societal implications that accompany these developments. By examining real-world case studies and engaging with expert guest speakers, we will gain insights into the latest AI innovations and their implications for human communication.

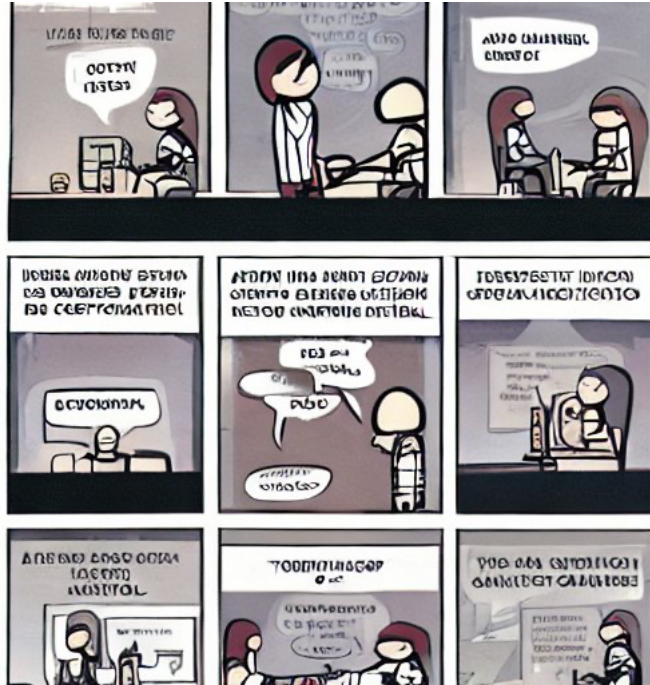


Image from DALL-E2

In this course, you will have the opportunity to actively engage with AI tools, exploring their capabilities and evaluating their potential impact. Through hands-on assignments and collaborative projects, you will develop critical thinking skills, enabling you to analyze and navigate the evolving landscape of AI with (some) confidence. And you will face the tricky social, ethical, and political questions posed by AI's growing suite of capabilities.

By the end of the course, you will emerge with as much of a comprehensive understanding of AI's impact on human communication and society as we can forecast at the moment. You will possess the skills to evaluate the ethical considerations and societal implications of AI applications, empowering you to contribute to the ongoing discourse surrounding AI in your future endeavors. And you will grapple with the ethical and social gravity of a future without a Turing test, where humans may not be the only or even the primary sources of intelligence.

Notably, given the rapidly changing nature of generative AI and the fact that little work has examined AI across the social sciences (let alone in the field of Communication and Media), this course inevitably serves as an experiment. There is no broadly accepted reading list for a class of this sort; nor is there a strong sense, even among experts, about what domains will experience AI's greatest impacts. While I will do my best to bring together a suite of materials that are both timely and relevant, we may find that some seminal works in the field no longer hold water after the developments of the last few years or that arguments that seemed salient in early 2023 read as hopelessly dated by the time we encounter them. Indeed, while the class description you just read was itself

heavily aided through the use of Generative AI tools (as are the weekly descriptions), the readings and activities mostly couldn't be, because the training data at the time only went through September of 2021, an eon ago in this rapidly evolving field.

Content Objectives

- Develop an understanding of how generative AI tools function and their limitations
 - Explore the societal implications of AI technologies in various domains
 - Analyze the ethical, legal, and social challenges posed by AI in human communication
 - Examine the potential risks and benefits of AI and its impact on employment and education
 - Foster critical thinking skills to evaluate and navigate AI technologies
 - Engage in discussions on policy implications and future developments in AI
-

Course Policies:

Special Accommodations:

Any student who has a need for accommodation based on the impact of a disability, religious practice, physical requirement, or medical need should contact me privately to discuss the specific situation as soon as possible.

Absences:

Sickness. If you are sick, please send me an email as soon as you are aware that you may miss class. In general, if you are sick for an extended period, it will be useful to provide a doctor's note in addition to emailing me.

Religious holidays. Within the first two weeks of the semester, please notify me of any religious holidays for which you will be absent.

Athletic and other university-related absences. If you are travelling to represent the University of Michigan, someone on your team will provide you with the appropriate paperwork to distribute to your instructors.

Other excusable absences. For family emergencies, funerals, and other such absences, please send me an email as soon as you are aware of the potential for missing class.

A Note on Academic Freedom and Controversial Subjects:

Some of the subjects we will be discussing in this class may be highly controversial and sometimes touch on matters of strongly held beliefs. It is both my responsibility as an instructor and your responsibility as students to respect the range of opinions held in the classroom and to recognize that aspects of every topic we will be discussing are open to debate. In addition, some of the materials that we will encounter are coupled with strongly held viewpoints. The viewpoints in these readings are not necessarily correct – indeed, you may see that they sometimes contradict one-another. A reading’s inclusion in the syllabus thus does not represent an endorsement of its content rather an assertion that a particular viewpoint is valuable to understand.

Academic Honesty:

A good student-teacher relationship operates on the basis of trust. From that basis, I trust that you will do your utmost to complete course assignments and to be honest with me if for any reason you are unable to fully meet a commitment to the class. I also trust your judgment that any collaboration with peers or use of online tools reflects an honest attempt to present your own work and contributions in the context of the class. That said, if I encounter evidence that you are taking sole credit for work that is not principally your own, you will fail the course and I will immediately report the incident to the Dean of Student Affairs.

Uses of Generative AI:

Over the course of the term, we will be proposing standards for what constitutes appropriate uses of generative AI in this class and what does not. These may vary over time and across assignments. Appropriate uses of generative AI are NOT the same as academically honest ones. A use of generative AI is appropriate if it is in concordance with standards (we have set as a class) for how and when the tools can and should be used. A use of generative AI is academically honest if any output clearly denotes (either from the terms of an assignment or through the use of citations) what contributions stem from various individuals and from GenAI tools. To be clear, uses of generative AI in this class could simultaneously be academically honest but still inappropriate (e.g., if they are out of bounds for the assignment but well documented) or academically dishonest but appropriate (e.g., material that is taken from GenAI at a time when it would be acceptable to do so, but that is passed off as work that does not use GenAI).

Resources to Know:

Counseling and Psychological Services (CAPS)

Counseling and Psychological Services offers counseling services, educational and preventive initiatives, training programs, outreach and consultation activities, and provide guidance on how to “do something” to fully contribute to a caring healthy community. CAPS can be contacted on their 24-hour crisis line, 734-996-4747 and online at caps.umich.edu.

Sweetland Center for Writing

The Sweetland Center for Writing is a comprehensive center for improving student writing at all levels. They offer one-to-one tutoring for undergraduate students. Find them at www.lsa.umich.edu/sweetland/.

MiTalk

MiTalk offers mental health resources including online screenings for depression and anxiety, skill-building tools to help you manage stress and academic life, and digitally recorded workshops, lectures, and some relaxation exercises. The site is completely free of charge to U-M Students. Find them at mitalk.umich.edu.

Sexual Assault Prevention and Awareness Center (SAPAC)

If you or someone you know has been harassed, assaulted, or stalked, you can receive confidential support and academic advocacy at the Sexual Assault Prevention and Awareness Center (SAPAC). SAPAC can be contacted on their 24-hour crisis line, 734-936-3333 and online at sapac.umich.edu. Alleged violations can be non-confidentially reported to the Office for Institutional Equity (OIE) at institutional.equity@umich.edu. Reports to law enforcement can be made to University of Michigan Police Department at 734-763-3434.

Grading

- Class Participation – 40%
- Weekly reflections – 30%
- Free-range final assignment – 30%

Assignments

- **Weekly Reflections (30%):** Each week (except the final week), write a 1-page (~250 word) reflection on the readings -OR- have a GenAI tool write a 1-page (~250 word) reflection and write a one paragraph (~150 word) evaluation of what the tool produced. Post this on the discussion on the class Canvas page and write at least 2 sentences of commentary on something someone else posted. These will be graded on a check/check-minus system, with 11 checks needed for full credit (i.e., you can miss 3 weeks). The reflections will be worth 20% of your grade, and the commentary/responses to other students' reflections will be worth 10% of your grade.
- **Free range assignment (30%):** At the end of the class, take one of the topics that we have explored, conduct some additional research on the area, and produce some kind of creative content that will help show (1) why it is an important issue and (2) what members of the public should know about it. You can complete this assignment in any way that you would like. You are welcome to work in groups on this assignment as long as you clearly denote the contributions of the various group members. You are also welcome to make prolific use of AI tools, again as long as you hew to the standards the class puts together on how to use AI sources appropriately and honestly. Some possibilities:
 - A traditional essay (boring, but acceptable)
 - A formal proposal to the university
 - An art piece or exhibit
 - A show or movie
 - An interactive website
 - A board game
 - A computer program
 - A graphic novel
 - Etc.

This should take approximately the same amount of effort as a 10-page paper, so plan accordingly and make sure to get started early! The assignment is due on December 5, our last class day.

Course Schedule

Part 1 – Making Sense of Generative AI: How AI Works

Week 1 (Aug 29 and 31) – Can “Intelligence” Be Artificial? An Introduction to AI and Human Communication. As we embark on our journey through the realm of artificial intelligence (AI), we begin by questioning the very nature of intelligence itself. Can machines genuinely think or feel? Or are they merely mimicking human behaviors without understanding? This week, we set the stage by diving deep into foundational readings and debates, probing the very essence of AI and its impact on human communication.

Readings (by Aug 31):

- Turing, A. (1950). Computing Machinery and Intelligence.
- Sejnowski, T.J. (2023). Large Language Models and the Reverse Turing Test.

Week 2 (Sept 5 and 7) – Understanding Generative AI Tools. Generative AI is a formidable force in the AI world, but what exactly is it, and how does it work? This week, we unravel the intricacies of these tools, shedding light on their underlying mechanics and vast capabilities. With insights from industry experts and groundbreaking research, we'll demystify the magic behind AI's ability to generate content.

Readings:

- Wolfram, S. (2023). What is ChatGPT doing and why does it work?
- Bommasani et al. (2021). On the Opportunities and Risks of Foundation Models
- Zhao et al. (2023). A Survey of Large Language Models

Guest Speaker: Dallas Card on 9/5

Week 3 (Sept 12 and 14) – What can AI do? Prompts and Outputs AI's capabilities are vast and ever-evolving, spanning from simple tasks like setting reminders to complex feats like predicting global trends. This week, we'll venture into the heart of AI's prowess, understanding its potential to transform various sectors as well as some of its surprising limitations. Through a range of readings and case studies, we'll illuminate the possibilities AI brings to the table and the art of harnessing its power effectively.

Readings:

- Wei et al. (2022) Emergent Abilities of Large Language Models
- Warszel, C. (2023) The Most Important Job Skill of This Century
- Liu & Chilton (2022) Design Guidelines for Prompt Engineering Text-to-Image Generative Models
- White et al. (2023) A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

Week 4 (Sept 19 and 21) – What is GAI Generating? Exploration and Exploitation

As Generative Artificial Intelligence (GAI) becomes more integrated into our digital landscapes, it's crucial to understand the content it produces. This week, we'll dive deep into the content that emerges from GAI, weighing its potential benefits against its associated risks. By examining the intricacies of prompt engineering and the consequences of the information produced, we'll garner a comprehensive understanding of GAI's footprint in today's digital age.

Readings:

- GPT-4 Technical Report & System Card (2023) [READ pp 1-14 & 41-70]
- Sisonal et al. (2023) ChatGPT: More than a “Weapon of Mass Deception.”
- Liu et al. (2023) Jailbreaking ChatGPT via Prompt Engineering: An Empirical Study

Week 5 (Sept 26 and 28) – When is AI Use Appropriate? – Round 1. With AI rapidly influencing diverse sectors, the lines between its appropriate and inappropriate uses blur. This week, we embark on a journey of ethical considerations, unraveling the complexities of deploying AI technologies. Through thoughtful discussions and readings, we'll confront the challenging questions surrounding the acceptable bounds of AI, particularly in our own classroom.

Readings:

- Generative Artificial Intelligence Advisory Committee Report
- Draft bylaws produced by GPT-4
- Anderson, R. (2023). Does Sam Altman Know What He's Creating?

Part 2 – Brace for Impact: AI Across Domains

Week 6 (Oct 3 and 5) – The Breadth of Applications. AI's influence extends far beyond a few niches—it has permeated or appears poised to permeate almost every sector. This week offers a panoramic view of the applications of AI, from healthcare to entertainment, logistics to finance. Grasping the breadth of AI's reach is essential to appreciate its transformative power.

Readings:

- West & Allen (2018). How artificial intelligence is transforming the world.
- HAI (2023). Generative AI: Perspectives from Stanford HAI.

Week 7 (Oct 10 and 12) – AI in Education. Investigation into the impact of AI on education, where it is currently being used as well as its potential to personalize learning,

encourage new forms of learning, and circumvent contemporary evaluation techniques (i.e., “cheat”). What are the opportunities and what are the risks?

Readings:

- Bogost, I. (2023). The First Year of AI College Ends in Ruin
- Cotton et al. (2023). Chatting and Cheating: Ensuring academic integrity in the era of ChatGPT
- Malesic, J. (2022). What ChatGPT Can't Teach My Writing Students
- Chronicle of Higher Education. (2023). How Will AI Change Higher Ed?
- Crompton et al. (2022). Affordances and challenges of artificial intelligence in K-12 education: a systematic review
-

Week 8 (Oct 19) - AI in the Workplace. The workplace is evolving, and AI plays a pivotal role in this transformation. This week, we'll delve into AI's impact on the future of work. How does AI reshape industries? What new opportunities arise, and what challenges do we need to confront? Through engaging readings and discussions, we'll forecast the future of the job market in the AI era.

[NOTE: NO CLASS ON OCT 17; FALL BREAK]

Readings:

- Eloundou et al. (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models
- White House (2022). The Impact of Artificial Intelligence on the Future of Workforces in the European Union and the United States of America
- Warzel, C. (2023). Here's How AI Will Come for Your Job
- Lowery, A. (2023). How ChatGPT Will Destabilize White Collar Work
- HAI (2023) Artificial Intelligence Index Report 2023. [Skim Chapter 4 The Economy pp. 168-233].

Week 9 (Oct 24) – AI in Creative Industries. Art and technology have always intersected, but the advent of AI brings novel dimensions to this relationship. As AI takes on roles in artistic creation and interpretation, what does it mean for human artists, the definition of creativity, and the value of originality? This week, we'll navigate the intersections of AI and art, discussing the controversies and the transformative potential of this union.

[NOTE: NO CLASS ON OCT 26; I WILL BE TRAVELING]

Readings:

- Roose, K. (2022). An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy.
- Science Friday (2022). How Will AI Image Generators Affect Artists? <https://www.sciencefriday.com/segments/ai-art/> [listen to the audio]

- Chayka, K. (2023). Is AI Art Stealing from Artists?
- Hutdon, M. (2023). AI learns to write computer code in 'stunning' advance.

Week 10 (Oct 31 and Nov 2) - When is AI Use Appropriate? – Round 2. Building on our earlier discussions about the ethics of AI, we'll focus to implications in educational and business contexts. Can we develop a standard that ensures AI is used responsibly and ethically across these domains? This week aims to confront these pressing questions head-on.

Readings:

- Ntoutsi et al. (2019). Bias in Data-Driven Artificial Intelligence Systems.
- Barbour, C. (2023). Can a Machine Be Racist?
- Eitel-Porter, R. (2021). Beyond the Promise: Implementing Ethical AI

Guest Speaker: Dean Perry Fittler on 10/31

Part 3 – An Intelligent Future: Where Is It Going?

Week 11 (Nov 7 and 9) – Futurecasting: Toward Artificial General Intelligence?

AI's journey doesn't end with chatbots and recommendation systems. Many believe we're on the brink of achieving Artificial General Intelligence (AGI)—a type of AI that could potentially outthink humans in every capacity. But what would such a future look like? Is it a promise of utopia or a warning of dystopia? Together, we'll look beyond the horizon of current AI capabilities and speculate about the road to AGI.

Readings:

- Bubeck et al. (2023). Sparks of Artificial General Intelligence
- Future of Life Institute (2023). Pause Giant AI Experiments: An Open Letter
- Marcus, G. (2023). Why Are We Letting the AI Crisis Just Happen?
- Warzel, C. (2023). What Have Humans Just Unleashed?
- Wong, M. (2023). AI Doomerism is a Decoy

Week 12 (Nov 14 and 16) – AI in Ethics and Law. AI doesn't exist in a vacuum; it plays a significant role in shaping global politics and public policies. This week, we'll explore how nations respond to the rise of AI, the geopolitical implications of AI supremacy, and the strategies governments might adopt in the AI era. Through diverse readings, we'll dissect the intricate dance between AI and global governance.

Readings:

- Zohny et al. (2023). Ethics of Generative AI.
- Liao and Vaughan. (2023). AI Transparency in the Age of LLMs
- Appel et al. (2023). Generative AI Has an Intellectual Property Problem

- HAI (2023) Artificial Intelligence Index Report 2023. [Read Chapter 3 Technical AI Ethics pp. 125-167].
- Alexandrei, D. (2023). Explainable AI, the key to open “black boxes”

Week 13 (Nov 21 and 23) – AI in public policy and geopolitics. AI doesn't exist in a vacuum; it plays a significant role in shaping global politics and public policies. This week, we'll explore how nations respond to the rise of AI, the geopolitical implications of AI supremacy, and the strategies governments might adopt in the AI era. Through diverse readings, we'll dissect the intricate dance between AI and global governance.

[NOTE: NO CLASS ON NOV 23; THANKSGIVING]

Readings:

- HAI (2023) Artificial Intelligence Index Report 2023. [Read Chapter 6 Policy and Governance, pp. 263-296].
- Hacker et al. (2023). Regulating ChatGPT and Other Large Generative AI Models.
- Lowrey, A. (2023). Before AI Takes Over, Make Plans to Give Everyone Money
- Anderson, R. (2023). Never Give Artificial Intelligence the Nuclear Codes

Week 14 (Nov 28 and 30) – When is AI use Appropriate? – Round 3. Our final dive into AI's ethical realm focuses on its societal and political implications. What are the potential repercussions of AI on our democracies, and how might it reshape political discourse? We'll grapple with the challenges and responsibilities of introducing AI into the broader societal fabric.

[What policies need to be in place to deal with future AI? How do we ensure AI is ethical and unbiased?]

Readings:

- Wong, M. (2023). The AI Crackdown is Coming
- Anderson et al. (2018). Artificial Intelligence and the Future of Humans [just skim it, I know it is too long]

Week 15 (Dec 5) – Summing it up – and a final project fair. As we culminate our exploration into the vast universe of AI, this week offers an opportunity to reflect, integrate, and showcase. Students will present their final projects, encapsulating their learnings and insights from the course. This collaborative event not only serves as a celebration of knowledge gained but also as a launchpad for future inquiries into the ever-evolving world of AI.