Welcome to COMP741/841 Practical Artificial Intelligence

COMP 741/841 Week 1

Spring 2024

Agenda

- Introduction
- Course information
  - Syllabus, requirements, resources
  - Activities, schedule, assessment
  - Policies
- Due next week
  - Tools: communication, reading, writing, and development
  - First assignment: Assigned reading and reading notes (RN1)
- What is AI?

Introductions

Instructor, classroom assistant, and students

- What is your preferred given name?
- What do you like the most about computing?
- What do you dislike the most about computing?

Fill in the attendance sheet: first 5 minutes of each class

Discord Sign In

- Create a Discord account with your UNH email (given-name.family-name@unh.edu and UNH password)
- Join the course Discord server using your Discord account
  - Go to Canvas (mycourses.unh.edu) and get the Discord server link from Announcements
- Log in with your Discord credentials
  - Do NOT sign in as a guest
- Change your Discord Server name to your ***Family name, Given name** by editing your server profile

Course Information

- Weekly 1 hr 50 min class meetings: bring your laptop and charger!
- Resources
  - Canvas, Discord, GitHub, Zoom (if needed), Zotero
  - Readings:
    - Wikipedia's outline of AI: https://en.wikipedia.org/wiki/Outline_of_artificial_intelligence
    - Wikipedia AI https://en.wikipedia.org/wiki/Artificial_intelligence#
  - Zotero Practical AI group library: https://www.zotero.org/
    - See Zotero Guidelines under **Resources and Tools** menu on left-hand side bar.

Course Resources: Development Tools and Services

Local tools

- Bash shell
- Python 3.10
- Visual Studio Code
- git

  Cloud services
- AWS SageMaker Studio Lab
- Google Colab
- GitHub

Learning Resources

- Course instructor, classroom assistants, and your peers
  - In-class and Discord participation
  - Guided, collaborative, and independent study activities
- Group working sessions
  - With peers, facilitated by classroom assistants
- One-on-one check-ins with course instructor
- Tutoring in the Center for Academic Enrichment
  - Programming, writing, multilingual learning

Learning Activities, Assignments, and Assessment

- Assigned reading and reading notes (15%)
  - Assigned every other week, Week 1 to Week 8
  - Draft, revise (RN Feedback), and finalize (RN Final)
- Lab projects (25%)
  - Assigned every week, Week 2 to Week 7
- Project Work (10%), starts Week 9
  - Proposal, Design, First Presentation
- Project Final (50%)
  - Codebase, Report, Final Presenataion
- In-class participation (10%)

Learning Goals

Balancing the **science of AI** with its **engineering applications**, you will learn:

- The AI foundations and principles to build intelligent computational systems.
- How AI systems solve real-world problems that require:
  - Reasoning, planning, learning, explaining, and acting with certainty and uncertainty
- The impact of AI on our society and related ethical, privacy, security, and safety implications.

Competencies that the course helps you develop

- Explore and critique AI applications and their impacts on individuals, communities, society, and humankind.
- Read and analyze relevant AI literature disseminated through journal articles, conference proceedings papers, and popular media. Present, discuss, and evaluate AI approaches and technologies.
- Examine, annotate, and evaluate the theoretical basis, design decisions, and implementation of open-source AI applications.
- Participate in and bring your own contribution to the development of a team project that demonstrates the applicability of neural and symbolic AI approaches.
- Practice with and develop personal qualities and behavioral patterns that are highly regarded in the workplace, such as being responsible, persistent, adaptable, and self-reflective.

Due: Platforms and Tools

Deadline: Wednesday night, before class

Communication: Discord and Canvas

Have Discord and Canvas apps on your phone and your laptop to be notified of all learning activities in this course

Readings: Zotero (see Zotero Guidelines on the website)

Create Zotero account, download Zotero app, and add Zotero connector to the browsers you use frequently (Chrome, Firefox)

Writing: OneDrive

Create directory COMP841 (or COMP741) in your OneDrive account. Add **Readings** subdirectory where you'll write your reading assignments.

Development

bash shell, Python 3, VS Code, git, AWS SageMaker Studio Lab, GitHub, Google Colab

Due: Reading Assignment

Read the following articles:

Roscoe, Jules. 2024. "A ‘Shocking’ Amount of the Web Is Already AI-Translated Trash, Scientists Determine.” Vice (blog), January 17, 2024.

Bastian, Matthias. 2024. “Deepmind’s AlphaGeometry Solves Complex Math and Is a ‘Crucial Step’ toward AGI.” THE DECODER, January 17, 2024.

See more information in Reading Assignments, also referenced from Canvas Assignments.

Optional Research Articles

Original research articles that inspired the assigned reading:

Thompson, Brian, Mehak Preet Dhaliwal, Peter Frisch, Tobias Domhan, and Marcello Federico. 2024. A Shocking Amount of the Web Is Machine Translated: Insights from Multi-Way Parallelism. arXiv, January 11, 2024. https://doi.org/10.48550/arXiv.2401.05749.

Trinh, T.H., Wu, Y., Le, Q.V. et al. 2024. Solving olympiad geometry without human demonstrations. Nature 625, 476–482. https://doi.org/10.1038/s41586-023-06747-5

What is Artificial Intelligence?

- As an academic discipline
  - Study of computational systems that act intelligently (a.k.a. computational agents)
  - To act intelligently means
    - Know what to do: Exhibit complex behavior to reach goals
    - Know how to do it:
      - Sense, game play, reason, optimize, plan, process natural language, etc.
    - Do it: Perform (or act) to meet those goals
    - Learn: from past experience and prior knowledge

What is Artificial Intelligence?

- As a technology
  - Software and tools platforms, services made by humans and used by humans
- As a professional career
  - Research scientist
  - Software developer/engineer
  - UX (user experience) designer
  - Data scientist, data engineer
  - More ...

Goals of AI

- Scientific goal
  - Understand how artificial and natural systems can act intelligently
    - Analyze artificial and natural agents
    - Formulate and test hypotheses about how to construct intelligent agents (AI)
    - Design, build, experiment with computational systems that perform tasks that require intelligence
- Engineering goal
  - Design and make intelligent computational systems that are useful
- Ethical goal

Ethical Implications of AI and the Alignment Problem

How to make AI align complex, autonomous systems with human norms and values?

AI systems as **objects** and their implications

- Automation and unemployment
- Privacy and manipulation
- Human-robot interaction and autonomy

AI systems as **subjects**

- Machine ethics: ethics of AI systems themselves
- Artificial moral agency

Ethical Implications of AI and the Alignment Problem

AI as **superintelligence**

- singularity

Sources

Brian Christian, The Alignment Problem, 2020
Ethics of AI and Robotics, 2020, Stanford Encyclopedia of Philosophy, https://plato.stanford.edu/entries/ethics-ai/

Open and Closed Source Neural AI

What makes a neural AI system open or closed source?

- Model weights are publicly available
  - Model weights auditability versus classical programs
  - What parameters were used to train the model?
    - Can the results be replicated?
- Dataset is publicly available
  - Is the dataset documented?
    - Are limitations and biases explicitly defined?
- AI-as-a-service
  - ChatGPT, Bing AI, GitHub CoPilot
    - CoPilot is trained on GitHub repos
    - Would this be considered a publicly available dataset?

Human-Machine Teaming

- AI as an assistant
  - ChatGPT as a replacement for Google searches
  - Article summarization
- Human as the controller
  - Human has final say as user – AI augments capabilities
- Consequential versus non-consequential AI
  - Shopping recommendations versus military use
  - Human user has more control as potential consequences increase
- Building trust in AI
  - AI must produce accurate outputs, but also
  - Reliable, safe, transparent, and explainable decisions