



ARTIFICIAL INTELLIGENCE

The Very Idea

Vasant G. Honavar

Dorothy Foehr Huck and J. Lloyd Huck Chair in Biomedical Data Sciences and Artificial Intelligence
Professor of Data Sciences, Informatics, Computer Science, Bioinformatics & Genomics and Neuroscience
Director, Artificial Intelligence Research Laboratory
Director, Center for Artificial Intelligence Foundations and Scientific Applications
Associate Director, Institute for Computational and Data Sciences
Pennsylvania State University

vhonavar@psu.edu
<http://faculty.ist.psu.edu/vhonavar>
<http://ailab.ist.psu.edu>

Introductions

- **Instructor**

- Dr. Vasant Honavar
- Professor with 30 years of AI research and teaching experience
- Director, Artificial Intelligence Research Lab
- Director, Center for Artificial Intelligence Foundations and Scientific Applications (CENSAI)
- E335 Westgate Building
- vhonavar@psu.edu
- <http://faculty.ist.psu.edu/vhonavar>



- **Learning Assistant**

- Apoorv Thite
- Junior, BS in Applied Data Science



- **Students?**

Why am I teaching this course?

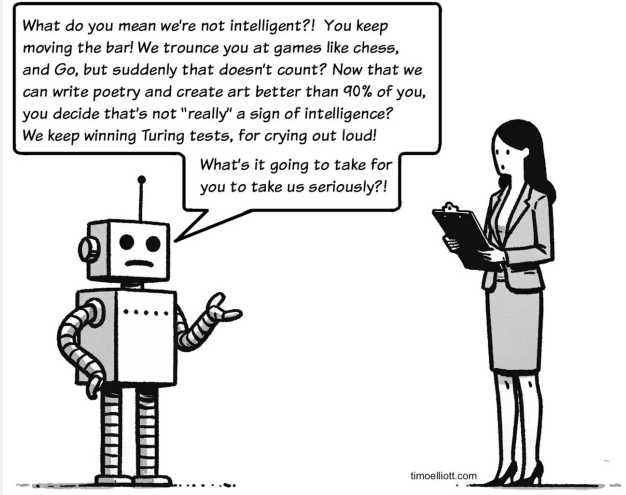
- Artificial Intelligence (AI) is increasingly powerful and ubiquitous
- AI is transforming all aspects of our lives, including
 - How we work
 - How we make sense of the world around us
 - We diagnose and treat diseases
 - How we learn
 - How we interact with others
 - How we create works of art
 - How we do science
- To thrive in a world being transformed by AI
- To shape highly consequential societal policies and regulations around AI
- Every citizen needs to
 - Have a basic understanding AI
 - Know AI's benefits and its potential for harm

This course is designed to provide such an accessible understanding of AI

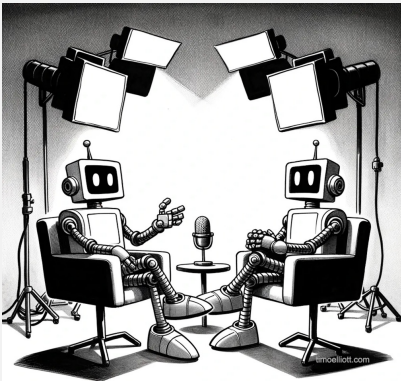
Ever-evolving state of artificial intelligence



Ever-evolving state of artificial intelligence

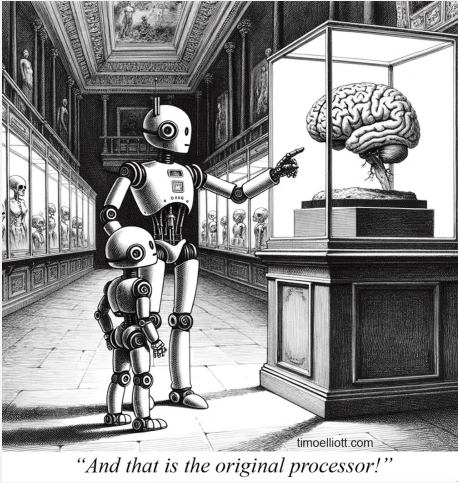


Ever-evolving state of artificial intelligence



And tonight's topic is "how can we tell if humans are actually intelligent?"

Ever-evolving state of artificial intelligence



About Artificial Intelligence – The Very Idea

- **What the course is about**
 - Key concepts needed to understand recent AI advances
 - Ethical and societal implications of AI technologies
- What you will learn
 - Intellectual roots of AI
 - Disparate goals of AI
 - Many types of AI and their applications
 - The risks and benefits of AI technologies
 - The societal impacts of AI technologies
 - Ethics of AI
 - How to tell AI science from science fiction

About Artificial Intelligence – The Very Idea

- This is a freshman introductory course
- It is open to all students regardless of major
- The only prerequisites are
 - Interest in learning about AI
 - Curiosity and willingness to ask questions
 - Ability and willingness to think
 - “I think, therefore I am” – Descartes
 - If you can’t think or don’t want to think, for all intents and purposes you are not in this course 😊
 - Ability to read critically and write coherently

About Artificial Intelligence – The Very Idea

This course is NOT going to teach you

- How to design and develop AI systems
- How to write AI programs
- All the math behind AI

There are other courses that do cover all these technical aspects

- AI courses at the 200, 300 and 400 levels
- A new BS in AI Methods and Applications is scheduled to launch next fall

Course materials

- Course web page
 - <https://faculty.ist.psu.edu/vhonavar/Courses/ai100/homepage.html>
 - Syllabus
 - Useful reference materials
 - Course policies
 - Study guide – assigned readings, lecture slides, videos
- Course page in Canvas
 - <https://psu.instructure.com/courses/2349092/>
 - Assignments
 - Grades
 - Announcements
 - Course-related email communications

Texts and references

- <https://faculty.ist.psu.edu/vhonavar/Courses/ai100/texts.html>
- **There is no required textbook**
- Primary references
 - Wooldridge, M. (2021). [A Brief History of Artificial Intelligence - What it is, where we are, and where we are going](#). MacMillan.
 - Lamb, H., Levy, J., and Quigley, C. (2023). [Simply Artificial Intelligence](#). Penguin Random House
 - Colombo, M. and Piccinini, G. (2023). [The Computational Theory of Mind](#). Cambridge University Press.
- The first two are inexpensive if you want to purchase them
- The third is accessible online if you are on Penn State network
- A long list of additional references available on the course web page

What to expect

- Lectures are designed to provide background and context for questions
- Discussions organized around questions posed
- Assigned readings etc. reinforce the material covered
- Assignments call for reading, reflection, critical thinking, writing
- To get the most out of this class, plan on each hour of time in class, expect to spend two to three hours working on the course outside the class

Course policies

<https://faculty.ist.psu.edu/vhonavar/Courses/ai100/policies.html>

- Attendance expected
- Distraction-free engagement and participation in class expected
- Use Canvas for all course-related emails
- Check study guide on course web page at least once a day
- Check canvas at least once a day

Course policies

- <https://faculty.ist.psu.edu/vhonavar/Courses/ai100/policies.html>
- Review academic integrity policies carefully
- Work on assigned work independently
- You may rely on reference books, authoritative online sources, with proper attribution
- **Do NOT** post course materials on online or share them with others
- **Do NOT** seek help from outside sources – Chegg, Course Hero, etc.
- **Do NOT** use ChatGPT or similar programs to do your assignments
- **Sanctions for cheating range from failing grade on the course to dismissal from the university**
- Use Canvas for all course-related email communications
- Check study guide on course web page at least once a day
- Check canvas at least once a day

Grading Policy

- Assignments: 50%
 - Quizzes: 10%
 - Exams: 30%
 - Class participation: 10%
- 95% - 100% A
 - 90% - 94.9% A-
 - 87% - 89.9% B+
 - 83% - 86.9% B
 - 80% - 82.9% B-
 - 77% - 79.9% C+
 - 70% - 76.9% C
 - 60% - 69.9% D
 - 0% - 59.9% F

Please consult course policies regarding late assignments

<https://faculty.ist.psu.edu/vhonavar/Courses/ai100/policies.html>

Other policies

- <https://faculty.ist.psu.edu/vhonavar/Courses/ai100/policies.html>
- Resources available
 - IST Tutoring services
 - Counseling and psychological services
- Disability accommodation
- Educational equity and non-discrimination
- Pandemic guidelines
- Emergency notifications

Artificial Intelligence – The very idea

I propose to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think'".



-- Alan Turing (1950)

Artificial Intelligence – The very Idea

Questions of concern to AI have tantalized humans for millenia

- What are minds?
- What is knowledge?
- What is a skill?
- How do we reason?
- How do we learn (acquire knowledge and skills)?
- What makes language possible?
- What is rational behavior?
- What is creativity?

Artificial Intelligence – The very Idea

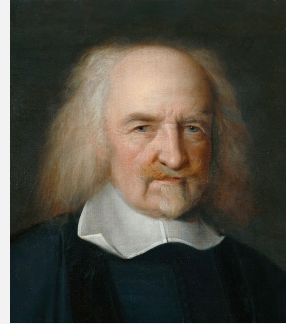
- Despite much philosophical speculation and debate, until recently, by scientific standards anyway, there was little progress in answering these questions
- But the past few decades have seen a sudden blossoming of ideas -- nothing short of an intellectual revolution in the sciences of the mind

Artificial Intelligence – The very Idea

- At the forefront of this revolution is Artificial Intelligence (AI) -- the exciting effort to build machines that think.
- The goals of those engaged in AI range from
 - building machines that behave as though they possess minds to
 - building machines with minds, in the full and literal sense
- This quest for artificial intelligence is based on
 - an idea that is as deep and imaginative as it is audacious
 - that thinking is, or at least can be modeled by, computation

Artificial Intelligence – The very Idea

- Back in the 17th century, Thomas Hobbes, the British philosopher claimed that reasoning or thinking – was a form of computation



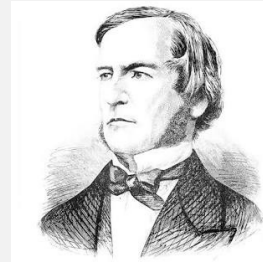
Artificial Intelligence – The very Idea

- Gottfried Leibniz (1646-1716), the German philosopher and mathematician, dreamed of a machine that could reason, or more precisely, reduce thinking to a kind of calculation.
- Leibniz's central thesis was that all human thoughts, no matter how complex, are made of basic and fundamental ideas, in much the same way that sentences are made of words, and words combinations of letters.
- He conjectured that if he could find a way to symbolically represent the fundamental ideas and develop a systematic method – logical calculus – by which to combine them, then he would be able to generate new ideas or thoughts on demand.



Artificial Intelligence – The very Idea

- Boole (1815-1864) develops algebraic approaches to logic and proposes logic and probability as the basis of laws of thought



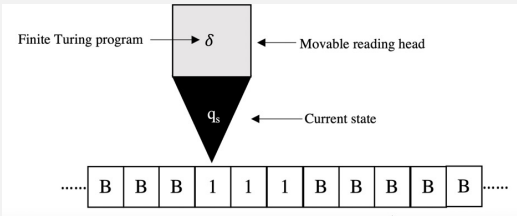
Artificial Intelligence – The very Idea

- In 1928, the German mathematician David Hilbert, posed the problem as to whether there exists
- an "effective procedure" that can be used
- to determine whether a given claim is provable from a given set of (non mutually inconsistent) assumptions.



Artificial Intelligence – The very idea

- In an attempt answer Hilbert’s decision problem, Turing (1912-1954) invents the **Turing Machine** to formalize the notion of an **effective procedure**
 - ✓ An effective procedure is a recipe for transforming a string of letters into another string of letters
 - ✓ A recipe (program) executed by the Turing machine



Artificial Intelligence – The very idea

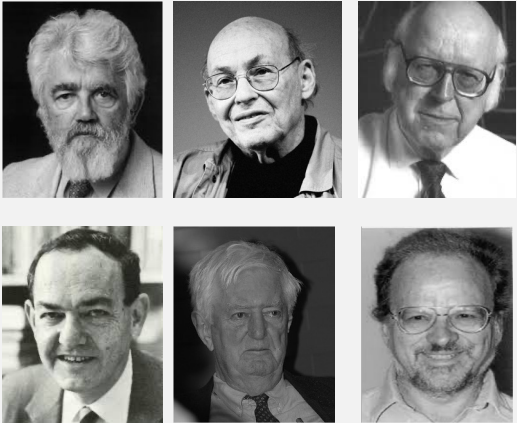
- Could practical realizations of computing machine of the sort conceptualized by Turing be used to realize Leibniz's dream of machines that could reason?
- Turing thought so



“I propose to consider the question, ‘Can machines think?’ This should begin with definitions of the meaning of the terms ‘machine’ and ‘think’”.

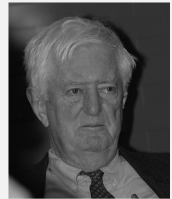
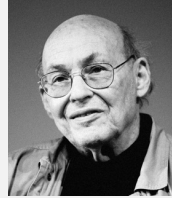
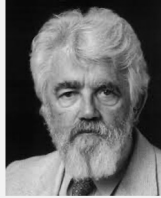
Artificial Intelligence – The very idea

- So did John McCarthy and a group of philosophers, psychologists, mathematicians and computer scientists.



Artificial Intelligence – The very idea

- John McCarthy organized a meeting of the group at Dartmouth in the summer of 1956
- “to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.”



Artificial Intelligence – The very idea

- Some find the quest for artificial intelligence quite preposterous.
- Others proclaim that computers with minds, if they are not already here, are just around the corner.
- The remarkable thing is how utterly confident each side is.

Artificial Intelligence – The very idea

- We will see that Artificial Intelligence is neither preposterous nor inevitable.
- Rather, it is based on a powerful idea, which very well might be right (or right in some respects) and just as well might not.
- But **attempts to build computers with minds**, regardless of our opinions about their success, **have already resulted in useful artifacts that to varying degrees behave as if they have minds**

Our goals in this course

- Demystify artificial intelligence
- Illuminate the rich intellectual and scientific underpinnings of artificial intelligence
- Discuss many types of artificial intelligence their capabilities, and their limitations, as well as their practical applications
- Discuss the risks, benefits, ethics, and societal impacts of artificial intelligence technologies

Artificial Intelligence – The very idea

Along the way, we will examine

- Computers, (computational) problems, algorithms, languages, and programs;
- Church-Turing thesis and its implications for theories of the mind and beyond;
- Discuss the disparate goals of those engaged in the quest for artificial intelligence and all that it entails
- Examine computational theories of minds and brains
- Discuss different kinds of minds and the corresponding kinds of AI
- Consider examples of
 - machines that reason,
 - machines that learn
 - machine that interact with humans in natural language
- Examine how we can maximize the societal benefits of AI while limiting its potential for harm.