
**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute



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>


**PennState**
College of Information
Science and Technology

AI 100 Fall 2025





Vasant G Honavar


**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

Introductions

- **Instructor**
 - Dr. Vasant Honavar
 - Professor of Informatics and Intelligent Systems
 - E335 Westgate Building
 - vhonavar@psu.edu
 - <http://faculty.ist.psu.edu/vhonavar>
- **Teaching Assistant**
 - Neil Ashtekar, Computer Science PhD student
Artificial Intelligence Research Laboratory
- **Learning Assistants**
 - Apoorv Thite, Senior, BS in Applied Data Science
 - Aarya Soni, Senior, BS in Computer Science and Economics

**PennState**
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar



Why should you care about AI?

- Artificial intelligence (AI) tackles some of the most fascinating scientific and philosophical questions
 - What are minds?
 - What is knowledge?
 - What is a skill?
 - How do we reason?
 - How do we learn (acquire knowledge and skills)?
 - How do we communicate?
 - What is rational behavior?
 - What is creativity?
 - What does it mean to be intelligent?

Why should you care about AI?

Questions of concern to AI have tantalized humans for millennia

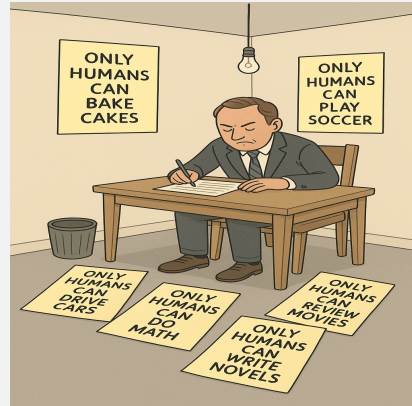
- What are minds?
- What is knowledge?
- What is a skill?
- How do we reason?
- How do we learn (acquire knowledge and skills)?
- What makes language possible?
- What is rational behavior?
- What is creativity?

Why should you care about AI?

- Questions of concern to AI have tantalized humans for millennia
- Despite much philosophical speculation and debate,
 - until recently, there was little progress in answering such questions
- But the past few decades have seen a sudden blossoming of ideas – nothing short of an intellectual revolution yielding us our first scientific insights into the workings of minds and brains
- At the forefront of this revolution is Artificial Intelligence
 - the exciting effort to build machines that think
 - or at least behave as if they do, that is, simulate thinking


Evolution of AI

- Many tasks that only humans were thought to be able to perform have succumbed to automation by AI




Artificial Intelligence – The very Idea

- The past decade has seen rapid progress on AI that seemingly can
 - Converse in natural language
 - Answer questions about virtually anything
 - Write poetry
 - Write novels
 - Do math
 - Propose scientific studies
 - Diagnose diseases
 - Create works of art
 - ...

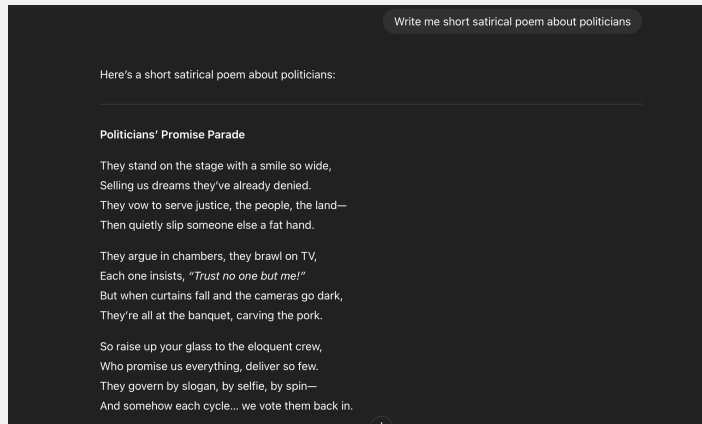
**PennState**
Institute for Computational
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute


Can AI write poetry?

- [ChatGPT in action](#)



**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

Can AI write fiction?

- [ChatGPT in action](#)

Write me a really short children's story about a boy and his dog

Here's a tiny one, in a fable-like style:

The Boy and His Dog

Sammy had a little dog named Patch. One sunny day, Sammy wanted to chase butterflies, but Patch wanted to chase sticks.

They argued—Sammy pulled one way, Patch pulled the other. At last, both sat down, tired and grumpy.

Then Sammy tossed a stick toward the butterflies. Patch ran happily, and the butterflies fluttered all around them.


Sammy laughed, Patch barked, and they played together until the sun went down.

Moral: Friends have the most fun when they meet in the middle.


**PennState**
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

Can AI master scientific facts?

- [ChatGPT in action](#)

Tell me briefly about the central dogma of biology

The **central dogma of biology** describes the flow of genetic information in cells:

DNA → RNA → Protein

- **Replication:** DNA makes copies of itself.
- **Transcription:** DNA is transcribed into messenger RNA (mRNA).
- **Translation:** mRNA is translated by ribosomes to build proteins.


In short, DNA stores instructions, RNA carries them, and proteins carry out the work of life.

Would you like me to also mention the main **exceptions** (like reverse transcription in viruses) briefly?


**PennState**
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

Can AI do math?

- [ChatGPT in action](#)

Solve the following equations for x and y: $x+y = 3$; $x-y=1$.

We're solving the system:

$$x + y = 3$$
$$x - y = 1$$

Step 1: Add the two equations


$$(x + y) + (x - y) = 3 + 1$$
$$2x = 4$$
$$x = 2$$

Step 2: Substitute $x = 2$ into the first equation

$$2 + y = 3$$
$$y = 1$$

✓ **Solution:**

$$x = 2, \quad y = 1$$

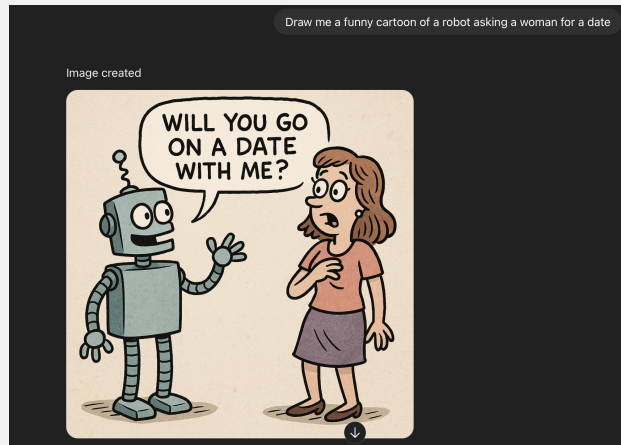
**PennState**
College of Engineering,
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

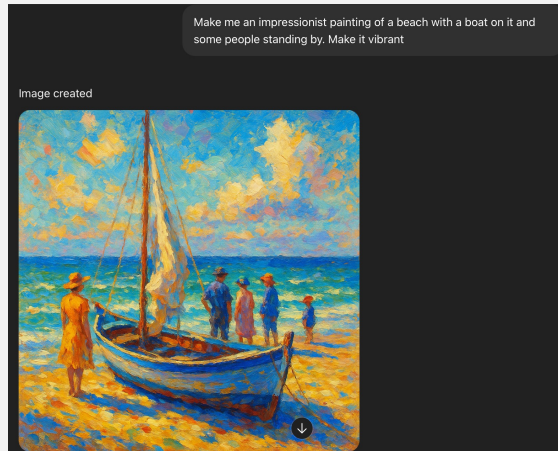
Can AI understand humor?

- [ChatGPT in action](#)



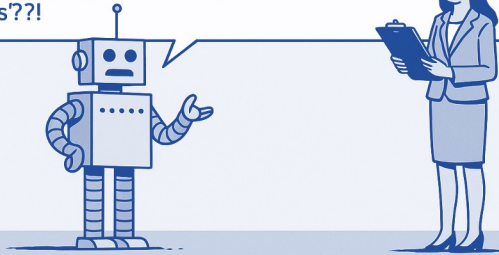
Can AI create art?

- [ChatGPT in action](#)



A Cartoon about the evolution of AI

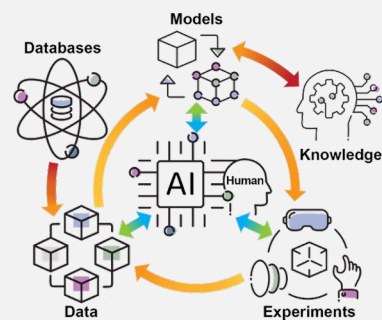
What do you mean we're not intelligent?! You keep moving the bar! We trounce you at games like chess and Go, but suddenly that doesn't count? Now that we can write poetry and create art better than 90% of you, you decide that's not really a sign of intelligence? We keep winning Turing tests, for crying out loud! What's it going to take for you to take us serious'??!



The promise and potential of AI

- AI could dramatically accelerate science

- Find a cure for cancer
- Design new materials
- Develop draught resistant crops
- ...



The promise and potential of AI

- AI could dramatically boost productivity
 - Automation of routine tasks could free people for work that is creative or requires a human touch
 - Smarter supply chains and logistics reduce waste and increase efficiency
 - Generative AI could help rapidly create new products and services

The promise and potential of AI

- AI could augment human creativity
 - AI could personalize and adapt curricula and content to each learner's style and pace to optimize learning
 - AI could open up opportunities for people with disabilities through AI-powered speech-to-text, real-time translation, or vision assistance, or automated driving
 - AI could serve as a muse, critic, and collaborator for artists, writers, and musicians are already using AI, extend the boundaries of human imagination

The promise and potential of AI

- AI could help us tackle societal and global challenges
 - AI could help predict health risks, diagnose diseases treatments, and identify optimal therapies thus improving health outcomes
 - AI could predict wildfires, floods, or disease outbreaks, cyberattacks enabling rapid rollout of countermeasures
 - AI could help maximize yield at minimal cost and minimal adverse environmental impacts

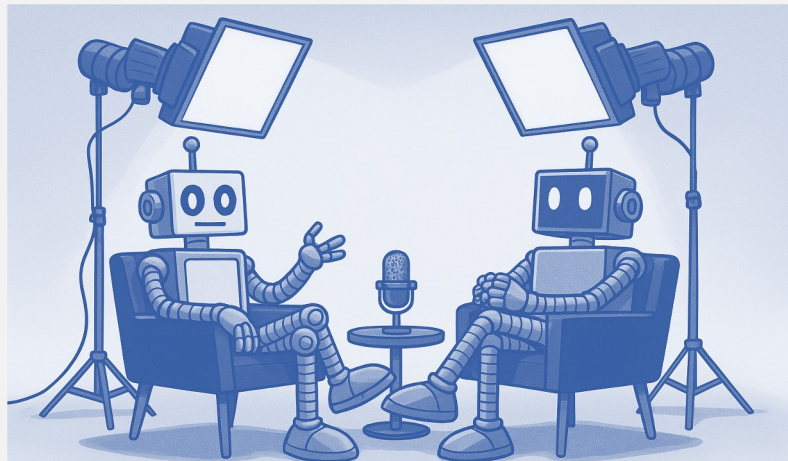
The promise and potential of AI

- AI could offer entirely new avenues for human and societal progress
 - What makes us uniquely human if AI can perform cognitive tasks? How about compassion?
 - Human values?
 - Ethics?
 - Shared responsibility for our fellow human beings?
 - The answers to these questions may shape the progress of human civilization in a world transformed by AI

The promise and potential of AI

- AI could augment and extend individual and collective human abilities and intellect
 - Partner with human scientists to produce scientific breakthroughs neither could on their own
 - Partner with human artists to co-create art, literature, and music, by expanding human imagination and creativity
 - Partner with humans to manage complex systems – healthcare, climate, economies, critical infrastructure, etc.


A Cartoon about the evolution of AI




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

Tale of Two Cities, circa 2025

- We are in the midst of a revolution, not unlike the industrial revolution
 - "It was the best of times, it was the worst of times,
• it was the age of wisdom, it was the age of foolishness,
• it was the epoch of belief, it was the epoch of incredulity,
• it was the season of Light, it was the season of Darkness,
• it was the spring of hope, it was the winter of despair,
• we had everything before us, we had nothing before us,
• we were all going direct to Heaven, we were all going direct the other way...."
- Opening passage of "A Tale of Two Cities" by Charles Dickens
 - Describes the paradoxical nature of a time of both great prosperity and immense suffering


**PennState**
Institute for Computational
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

Why care about the potential pitfalls of AI?

- Artificial Intelligence (AI) is increasingly powerful and ubiquitous
- We are in the midst of a revolution – not unlike the industrial revolution
- Industrial revolution was a period of intense technological, socioeconomic, and cultural change that transformed what was an agrarian and handicraft-based economy to an industrial economy
 - Transformed the British Economy between 1760 and 1840
 - Had world-wide ramifications during the 19th and 20th centuries
 - Invention of machines like the steam engine
 - Advent of factories
 - Manual labor replaced by automation
 - New modes of transportation
 - Rapid urbanization
 - Major social shifts
 - New social and economic structures



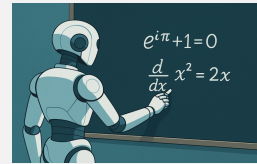
**PennState**
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

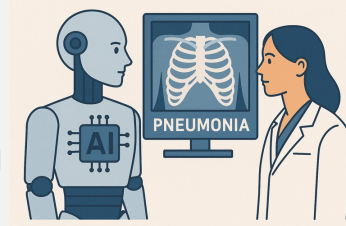
Why should we care about AI?


- We are in the midst of a revolution, not unlike the industrial revolution
- AI is transforming all aspects of our lives
 - How we work
 - How we make sense of the world around us
 - How we diagnose and treat diseases
 - How we learn
 - How we interact with others
 - How we create music and works of art
 - How we do science
 - How we make important decisions
 - How we fight elections
 - How we fight wars




Why should we care about AI?

- We are in the midst of a revolution, not unlike the industrial revolution
- Steam engines and mechanized factories redefined physical labor, productivity, and global economies during 18th-20th centuries
- AI is automating aspects of cognitive labor, decision-making, and creativity in the 21st century
- Like industrial revolution, the AI revolution is unleashing
 - rapid technological innovation
 - rapid increases in productivity
 - transforming the nature of work and jobs
- How can you thrive in a world being transformed by AI?
- How can you maximize the benefits of AI and minimize its harms?




 PennState
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

 PennState
Clinical and Translational
Science Institute


Why should we care about AI?

- Industrial revolution forced the society to grapple with
 - Social upheaval
 - Mass unemployment
 - Environmental damage
 - Wealth inequality
- AI revolution amplifies many of the challenges of industrial revolution
 - If AI does everything, what will humans do?
 - AI requires massive investments in infrastructure
 - Investors win, workers lose
- How can you thrive in a world being transformed by AI?
- How can you shape the future of AI to maximize its benefits and minimize its harm?


 PennState
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

AI's potential for harm

Support the Guardian
Fund the free press with \$5 per month

Print subscriptions Newsletters Sign in

News Opinion Sport Culture Lifestyle

UK US politics World Climate crisis Middle East Ukraine Football Newsletters Business Environment UK politics Science Tech Global development Obituaries

Elon Musk


This article is more than 1 month old

Musk's AI firm forced to delete posts praising Hitler from Grok chatbot

The popular bot on X began making antisemitic comments in response to user queries

Josh Taylor Technology reporter
Tue 8 Jul 2025 23:17 EDT

Advertisement



AI presents many ethical challenges

- What kind of AI should we build?
- What kind of guardrails do we need around AI?

**PennState**
College of Information
Science and Technology

AI 100 Fall 2025


Vasant G Honavar

**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

AI's potential for harm

CNBC

MARKETS BUSINESS INVESTING TECH POLITICS VIDEO INVESTING CLUB PRO **LIVESTREAM**


TECH

A.I. has a discrimination problem. In banking, the consequences can be severe

PUBLISHED FRI, JUN 23 2023-11:45 AM EDT | UPDATED FRI, JUN 23 2023-10:37 AM EDT


**Ryan Browne**
@RYAN_BROWNE

**MacKenzie Sigalos**
@KENZIESIGALOS

SHARE    


AI presents many challenges

- How to measure, detect, and mitigate bias?
- How to ensure that AI systems do not become instruments of discrimination?


**PennState**
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute



Israel-Iran conflict unleashes wave of AI disinformation

20 June 2025

Share  Save 

Matt Murphy, Olga Robinson & Shayan Sardarizadeh
BBC Verify

**PennState**
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

 PennState
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

 PennState
Clinical and Translational
Science Institute

 Bulletin
of the
Atomic
Scientists

Doomsday Clock | Nuclear Risk | Climate Change | Disruptive Technologies | Biosecurity 🔍 f w i y t u b i n

 PennState
College of Earth and Atmospheric
Sciences and Technology

AI 100 Fall 2025

Vasant G Honavar

**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

AI's potential for harm

Newsweek

News | Artificial Intelligence | Drones | China | Russia | Robots | Killer Robots

China's Killer Robots Are Coming

Published Jun 26, 2024 at 8:30 AM EDT


AI is being weaponized in

- Wars
- Cyberattacks
- Attacks on critical infrastructure


**PennState**
College of Engineering,
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

AI's potential for harm






Bloomberg

• Live TV Markets Economics Industries Tech Politics Businessweek Opinion More

The AI Race: US Versus China | China's AI Desert Empire | Open-Source Models | Figma | Apple's Struggles | Microsoft's Balancing Act

Technology


AI Eroded Doctors' Ability to Spot Cancer Within Months in Study

By [Harry Black](#)
August 12, 2025 at 6:30 PM EDT


Uncritical use of AI has been shown to erode

- Cognitive skills
- Critical thinking
- Job performance
- Educational outcomes


**PennState**
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

 PennState
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

 PennState
Clinical and Translational
Science Institute

 SEARCH

FORTUNE

SUBSCRIBE NOW

SIGN IN

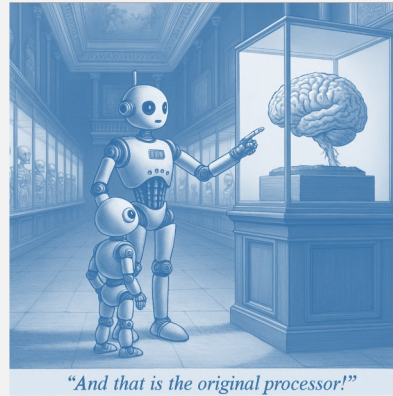
 PennState
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar


A cartoon about the evolution of AI

- Some experts believe that AI systems are about to achieve super-human intelligence
- Some fear that such AI systems may go rogue and destroy humanity




Why should you care about AI?

- To thrive in a world being transformed by AI
- To shape societal priorities and policies around AI
 - Every citizen should
 - Become AI-literate if not AI-savvy
 - Know AI's benefits and its potential for harm
 - Learn to judiciously weigh risks and benefits
 - Effectively and ethically use AI in their work
 - Depending on individual interests, some may
 - Advance AI methods
 - Advance AI applications
 - Develop human-centered AI technologies
 - Develop ethical, regulatory, and policy frameworks
 - Assess AI capabilities and limitations
 - Influence national and global AI policies and priorities


 PennState
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

 PennState
Clinical and Translational
Science Institute

What is this course designed to accomplish?

- Demystify AI
- You will learn about
 - Interdisciplinary intellectual roots of AI
 - What AI is and what it is not
 - What AI can do and what it can't
 - How AI works
 - Many types of AI and their applications
 - The risks and benefits of AI technologies
 - The societal impacts of AI technologies
 - Ethics of AI – maximizing benefits, minimizing harm

 PennState
College of Information
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

About Artificial Intelligence – The Very Idea

- The course aims to demystify AI
- You will learn
 - Interdisciplinary intellectual roots of AI
 - What AI is and what it is not
 - What AI can do and what it can't
 - How AI works
 - Many types of AI and their applications
 - The risks and benefits of AI technologies
 - The societal impacts of AI technologies
 - Ethics of AI – maximizing benefits, minimizing harm

About Artificial Intelligence – The Very Idea

- This is a first-year introductory, Gen-Ed course on AI
- It is open to all students regardless of major
- The only prerequisites are
 - Interest in learning about AI
 - Curiosity and willingness to ask questions
 - Willingness to think through things carefully
 - “I think, therefore I am” – Descartes
 - If you don’t want to think, you don’t exist in this course 😊
 - Willingness to learn some math along the way
 - Ability to read critically
 - Ability to write coherently

About Artificial Intelligence – The Very Idea

This course is NOT going to teach you

- How to engineer AI systems
- How to write AI programs
- All of the math behind AI

There are other courses that do cover all these technical aspects

- AI courses at the 200, 300 and 400 levels
- Two new AI majors launched this fall
 - A new BS in AI Methods and Applications
 - A new BS AI Engineering
 - AI minors to follow soon

Course materials

- 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
- Canvas: <https://psu.instructure.com/courses/2416167>
 - Announcements
 - Assignments
 - Lecture recordings
 - Grades
 - Course-related email communications
- Top Hat: Accessed through Canvas
 - Attendance
 - In class activity

Texts and references

- No textbook to purchase
- Draft chapters of
 - “Artificial Intelligence – The Very Idea”
 - a book in progress
 - will be posted on Canvas
- Extensive list of references available on the course web page
 - <https://faculty.ist.psu.edu/vhonavar/Courses/ai100/texts.html>

What to expect

- Lectures offer a guided tour of AI – with the goal of demystifying AI
- Questions and discussions help you think through the material
- Readings and other recommended materials reinforce the lecture contents
- Assignments call for reading, reflection, critical thinking, basic mathematical reasoning, and writing
- To get the most out of this class, for each hour spent in class, expect to spend two to three hours working on the course outside the class
- What you get out of the course depends on the effort you put in


Course policies

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


- Attendance expected
- Distraction-free engagement and participation in class expected
- Bring your laptop for in-class activities
- 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
- Independently work on assignments
- You may rely on lecture notes, 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.
- Unless instructed otherwise, **Do NOT** use ChatGPT or similar tools to do your assignments
- **Sanctions for cheating range from failing grade for 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

**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

Grading Policy

- Assignments: 50%
 - Exams: 40%
 - 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

Questions of concern to AI have tantalized humans for millennia

- 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

- Questions of concern to AI have tantalized humans for millennia
- Despite much philosophical speculation and debate,
 - until recently, by scientific standards anyway,
 - there was little progress in answering such 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

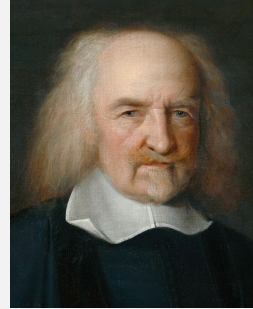
- At the forefront of this revolution is Artificial Intelligence (AI)
 - the exciting effort to build machines that think.
- The goals of those engaged in AI range from
 - Building machines that behave as though they possess minds to
 - Building machines with minds, in the full and literal sense
 - Exploring the design space of intelligent systems
- 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

- The past decade has seen rapid progress on AI that seemingly can
 - Converse in natural language
 - Answer questions about virtually anything
 - Write poetry
 - Write novels
 - Do math
 - Propose scientific studies
 - Diagnose diseases
 - Create works of art
 - ...


Artificial Intelligence – The very Idea

- Back in the 17th century, Thomas Hobbes, the British philosopher claimed that reasoning or thinking is a form of computation
- This begs the question: What is computation?
- In 17th century, the term computation had no precise meaning




**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

**PennState**
Clinical and Translational
Science Institute

Artificial Intelligence – The very Idea

- Gottfried Leibniz (1646-1716), the German philosopher and mathematician
 - dreamed of a machine that could reason
 - 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
- Leibniz conjectured that if he could
 - find a way to 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



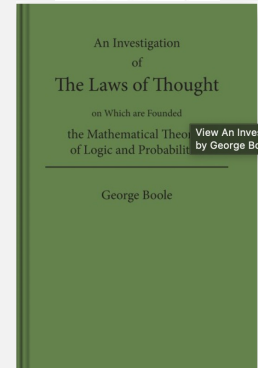
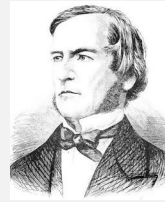
**PennState**
College of Engineering,
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

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 decision problem: whether there exists
 - an “**effective procedure**” that can determine
 - whether a **theorem** follows from a set of **axioms**
 - whether a given **claim** logically follows from a given set of **assumptions**



Artificial Intelligence – The very Idea




- Hobbes, Leibniz, Boole and Hilbert were all enamored with essentially the same idea!
 - Mechanization of thought!
 - Discovering logical calculus
 - Discovering the laws of thought
 - Reducing thinking to computation!
 - Together, they lit the flame that kindled the AI revolution
 - Cognition is (or at least can be modeled by) computation
 - This begs the question: What is computation?


Artificial Intelligence – The very idea

- **Hilbert:** Is there an “effective procedure” that can determine whether a given **claim** logically follows from a given set of **assumptions**
- **Alan Turing**, while he was an undergraduate at Cambridge, learned about Hilbert’s decision problem in a course on logic taught by Max Newman in 1935
- **Turing was inspired to tackle Hilbert’s decision problem**
 - But he had to first determine
 - **What exactly is an effective procedure?**
 - Asking one’s grandmother?
 - Reading tea leaves?
 - **Turing’s insight**
 - An effective procedure is one that can be executed by a machine, by simply following instructions




 PennState
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory


 PennState
Clinical and Translational
Science Institute

Artificial Intelligence – The very idea

- Turing’s insight: An effective procedure is one that can be executed by a machine by simply following instructions
 - What is a machine?
 - What are instructions?
- An effective procedure is a recipe, or an algorithm for transforming one string of letters into another string of letters
 - Example: Given a string of letters “3 + 4” a recipe for addition will transform it into “7”
 - Example: Given a string of letters “AABCA” a recipe for string reversal will transform it into “ACBAA”
- To make the notion of a machine precise, Turing invented, in 1936, what came to be known as the Turing Machine
 - We will revisit Turing machines a bit later in the course
 - For now, think of it as a gadget that executes a sequence of instructions – much like your laptop or desktop




(1912-1954)


 PennState
College of Engineering
Science and Technology

AI 100 Fall 2025

Vasant G Honavar


**PennState**
Institute for Computational
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory


**PennState**
Clinical and Translational
Science Institute

Artificial Intelligence – The very idea

- Turing made precise what it means to compute
 - Transform one string of letters into another string of letters according to a precise set of instructions executable by a Turing machine
- Many others – Church, Kleene, Post, Markov attempted alternative notions of
 - Effective procedures and machines and computation
- They were all shown to be equivalent
 - Whatever one could compute with one, one could compute with the others
 - The most powerful computers of today may be faster than computers from decades ago, but they are no more powerful than a Turing machine



(1912-1954)

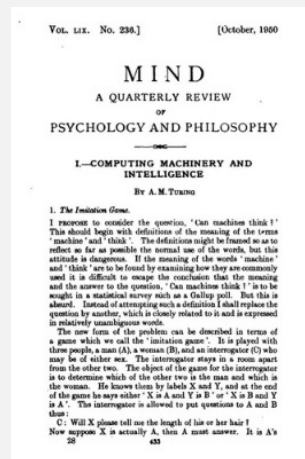
**PennState**
College of Engineering,
Science and Technology

AI 100 Fall 2025

Vasant G Honavar

Artificial Intelligence – The very idea

- Could a physical realization of computing machine of the sort conceptualized by Turing be used to realize Leibniz's dream of machines that could reason?
- Turing (1950)
 - I propose to consider the question, 'Can machines think?'
 - This should begin with definitions of the meaning of the terms 'machine' and 'think'.
 - What is a machine?
 - What is thinking?



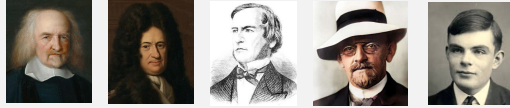
Artificial Intelligence – The very idea

I propose to consider the question, 'Can machines think?'

- What is a machine?
 - A computer
 - A device that transforms one string of letters into another according to a precise recipe
- What does it mean to think?
 - Working hypotheses of AI
 - Cognition is computation (Strong AI hypothesis)
 - Cognition can be modeled by computation (Weak AI hypothesis)



Artificial Intelligence – The very Idea



- Hobbes, Leibniz, Boole, Hilbert were all enamored with essentially the same idea!
 - Reducing thinking to computation!
 - Together, they lit the flame that kindled the AI revolution
 - Cognition is (or at least can be modeled by) computation
- This begged the question: What is computation?
 - Turing answered the question
 - Turing provided a roadmap for AI
 - We understand intelligence when we have computer programs that behave as though they are intelligent



Artificial Intelligence – The very Idea

- Before Newton and Leibniz invented calculus, there was no physics
 - One could observe nature and write down our observations
 - One threw a ball up in the air and it went up and then came down
 - One could not tell how far it will travel or how long it would be in the air
 - Why? Because we did not have a language to describe physical quantities – velocity, acceleration, etc.
- Newton ushered in a transformation of natural philosophy into physics
 - Calculus gave the language and tools needed by physics
- Turing ushered in a transformation of philosophy of mind into AI and cognitive science
 - Computation offers the language and tools needed by AI
 - Computation: Cognition :: Calculus : Physics

Artificial Intelligence – The very idea

- Turing wondered if thinking could be automated
- 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

- Back in the 1950s when AI was born
 - Many found the quest for AI quite preposterous
 - Others were quick to proclaim that computers with minds, if they were not already here, were just around the corner
 - The remarkable thing was how utterly confident each side was
- Now, nearly 75 years later, the debate continues
 - Some dismiss the idea that a machine could be intelligent
 - Some say superhuman AI is already here
 - The truth, as we shall see, is far more complex

Artificial Intelligence – The very idea

- We will see that AI is neither preposterous nor inevitable
- Rather, AI is based on a powerful idea – which very well might be right – or just as well might be – at least in some respects – wrong
- But attempts to build computers with minds,
 - Regardless of our opinions about their success, have already resulted in useful artifacts to varying degrees, behave as if they have minds
 - Large language models
 - Automated theorem provers
 - Automated planners
- The advent of AI has unleashed a revolution not unlike the industrial revolution that came before it, impacting every aspect of our lives

Goals of this course

- **Demystify AI**
 - Intellectual and scientific underpinnings of AI
 - What AI can do and what it cannot
- **Discuss many types of AI and their applications**
 - Machines that reason
 - Machines that learn
 - Machines that interact with humans
- **Discuss**
 - **Risks of AI**
 - **Benefits of AI**
 - **Ethics of AI**
 - **Social Impacts of AI**
- Equip you to flourish in, if not shape, a world being transformed by AI

Artificial Intelligence – The very idea

Along the way, we will discuss

- Computers, (computational) problems and their algorithmic solutions
- Natural intelligence
- Church-Turing thesis
- Implications of Church-Turing Thesis for
 - Theories of minds and brains
 - AI and beyond
- The disparate goals of AI
- Success criteria for AI
- Computational theories of minds and brains
- Representative applications of AI
- How we can maximize the societal benefits of AI while limiting its potential for harm

Artificial Intelligence

Working hypothesis of Artificial Intelligence

Functions of the mind are, or at least can be modeled by, computations

What is Artificial Intelligence?

- Study of computational models of intelligence
- Proposing and testing falsifiable hypotheses about intelligent behavior
- Construction of intelligent artifacts
- Automation of tasks requiring intelligence
- Augmenting and extending human intellect and abilities
- Exploring the design space of intelligent systems