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

On Artificial Intelligence

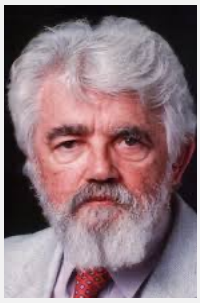

``[This] is or should be our main scientific activity
- studying the structure of information and
- the structure of problem-solving processes
- independently of applications and
- independently of its realization in animals or humans.''

- Implicit in this is the audacious idea is a functional or naturalist view of intelligence
- McCarthy envisioned artificial intelligence studying intelligence in a manner analogous to the way aerodynamics studies flight
  - independent of its applications and
  - independent of its realization in birds that fly

How is AI like aerodynamics?

Consider the following:

- Birds fly
- Ability to fly is natural (at least for birds)
- There are many types of bird flight (associated with different types of wings)

How is AI like aerodynamics?

- Albatrosses and some other seabirds have long, narrow wings and take advantage of winds over the oceans.
- Vultures and hawks have broad wings with slotted tips that permit more use of updrafts and winds deflected by hills.
- Birds such as ducks, pigeons, and falcons, which fly rapidly with continuous wing beats, tend to have moderately long, pointed wings.
- Many songbirds use their short, rounded wings to move with quick wing beats from perch to perch or from ground to perch.

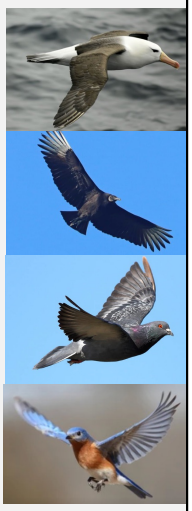

How is AI like aerodynamics?

Whether the flier is a bird or an airplane

- Flier must first become airborne
  - Must exert enough lift force to counteract its weight force of gravity
  - Lift is provided by air moving over and under the wings
- To move forward, the flier must exert thrust force to overcome the drag force or air resistance
  - Muscles provide the thrust in birds
  - Engines provide the thrust in airplanes

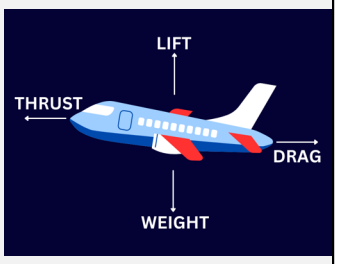

How is AI like aerodynamics?

- Aerodynamics, the science of flight explains both
  - How birds fly
  - How to design airplanes that fly
- Both birds and airplanes must obey aerodynamics
  - But they fly in very different ways
    - Airplanes don't flap their wings while in flight
    - The wings of airplanes are not made of feathers
  - Airplanes differ from birds more strikingly than birds differ from each other
  - But both have the ability to fly

How is AI like aerodynamics?

- Artificial intelligence is about understanding the essential principles and mechanisms that underlie intelligent behavior just as aerodynamics is about understanding the essential principles and mechanisms of flight.
- The resulting scientific understanding of intelligence should
  - enable us to build a a variety of machines that exhibit (different aspects of) intelligent behavior
  - help us understand our minds, brains, and ultimately, ourselves, in ways that would not otherwise be possible.

Working hypothesis of artificial intelligence

- Aerodynamics : : flight : : artificial intelligence : intelligence
- Working hypothesis of AI is that cognition is (or at least can be modeled by) computation
  - What computations underlie intelligent behavior?
    - We will have a theory of reasoning if we can devise algorithms that reason from assumptions to conclusions
    - We will have a theory of learning if we can devise algorithms for learning from experience
    - We will have a theory of linguistic communication if we can devise algorithms that effectively communicate using language
    - We will have a theory of cooperation if we have a computational model of multi-agent collaboration
    - We will have a theory of creativity if we can devise algorithms that exhibit creativity

Working hypothesis of AI

- Working hypothesis of AI is that cognition is (or at least can be modeled by) computation
- How do we know this hypothesis is true?
  - We really do not
  - Like any other working hypothesis in science, it is subject to revision or even outright replacement if scientific evidence so warrants
- The hypothesis has led to significant advances in
  - Our understanding of intelligence
  - Our ability to build AI systems
- Exactly what this hypothesis means is subject of debate
  - Strong versus weak AI

Computational theory of mind

- Cognition, or thought processes –
  - regardless of the physical substrates that realize them
    - be they living brains or silicon devices
    - can be modeled by computation

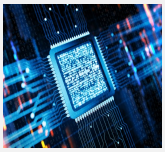

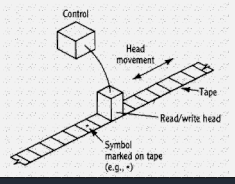

Unpacking the computational theory of mind

- When you speak to your mother on your cellphone,
  - A microphone turns your voice into electrical signals;
  - A microchip in the phone modulates (or varies) a radio wave using the electrical signal;
  - The radio wave travels through the air to a nearby cell tower;
  - The tower sends your voice to your mother and the process is reversed so that she hear your voice.

Unpacking the computational theory of mind

- Your message to your mother remains unaltered as it is carried by different media
  - from vibrations of air to
  - electrons in a wire to
  - electromagnetic waves and
  - back again in the reverse order
- Similarly, the message remains the same when
  - your mother conveys it to your father in another room,
  - after the vibrations of her eardrum turn into firings of neurons in her auditory cortex and
  - what she hears as a result get turned into airflow through her lungs,
  - up the windpipe (trachea) and through the voice box (larynx) which makes her vocal cords vibrate
  - creating sound which is shaped into words by the muscles controlling the soft palate, tongue and lips

Unpacking the computational theory of mind

- A given program can run on
  - computers made of vacuum tubes
  - electromagnetic switches
  - Transistors
  - integrated circuits, or
  - a group of humans passing notes to each other to accomplish the same functional behavior

Computational theory of mind implies a functional view of minds

- Functional view of minds
  - Is implicit in the attempts of Leibniz, Boole, and Turing to explain minds in computational terms
  - Is shared explicitly or implicitly by almost all of the work in AI.
  - Implies that intelligence is a functional capability independent of the specific physical substrates that support it

Computational theory of mind implies levels of analysis

- Functional view of minds implies levels of analyses

LEVELS			
Computation	1	What	Flight
Algorithm	2	How	Aerodynamics - lift, thrust
Implementation	3	How exactly	Birds, Airplanes, Drones ..

Computational theory of mind implies a functional view of minds

- Does this mean that the physical substrate -- whether it is the brain -- or the computer -- is irrelevant to understanding the minds?
- No!
  - Programs are executed by assemblies of simple information processing units -- of the sort that we say in the previous chapter -- that can match patterns, trigger some other circuit, or do other elementary logical and mathematical operations.
  - Circuits made from neurons may be a lot slower than circuits made from silicon in performing complex arithmetic calculations while neural circuits may be a lot faster than silicon in making sense of a video.
- It does mean that if our focus is on *what* is being computed and *how* it is computed by an algorithm, but not necessarily how exactly it is being computed, we need not worry about the physical substrate on which computation unfolds

Computational theory of mind and the mind-body problem

- The age-old mind-body problem has to do with
  - understanding the nature of the the relationship between the mind and the body
    - between the mental realm (the realm of thoughts, beliefs, sensations, pains, pleasures, emotions) and
    - the physical realm (matter, atoms, neurons, silicon circuits)

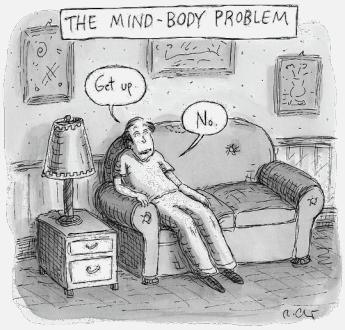

Discuss: Mind-body problem

- Are mind and body distinct entities?
- Is mind simply a part of the body?
- Is body a part of the mind?
- Are mind and body the same?

Computational theory of mind and the mind-body problem

- Why did Jill get on the plane?
- Because she wanted to visit her mother and knew the plane would take her to New York where her mother lives.
- If she hated her mother, or if she knew the plane would take her to San Francisco instead of New York, she would not be on the plane.

- Jill's desire to visit her mother and her knowledge that the plane would fly to New York are not physical, but mental entities.
- But they have the power to trigger or cause events in the physical world, such as getting Jill to get on the plane to New York.

Computational theory of mind and the mind-body problem

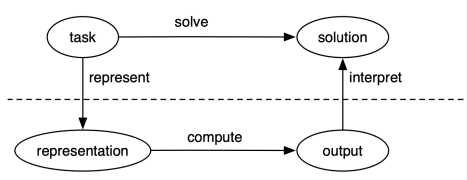

The computational theory of mind (CTM) offers a resolution of the mind-body problem.

- CTM says that contents of the mental realm, e.g., thoughts, are encoded by configurations of symbols.
- The symbols are the states of bits of matter, like chips in a computer or neurons in the brain.
- Symbols represent objects, events, and relationships in the physical world because they are triggered by them via our sense organs, and because of their physical effects once they are triggered.

Computational theory of mind and AI

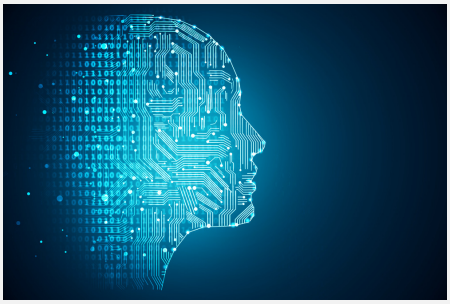

- Computational theory of mind is central to AI
- It is also central to modern psychology, cognitive and brain sciences

Artificial Intelligence – The very idea

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

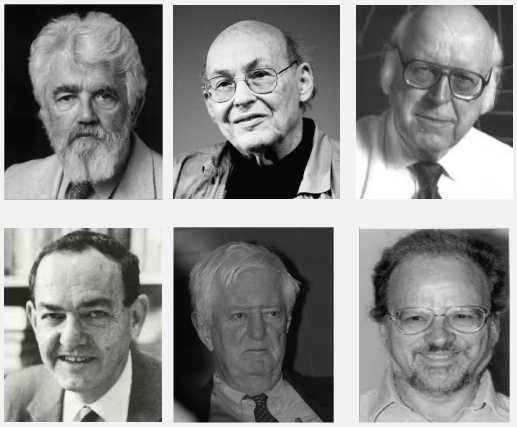

Artificial Intelligence – The very idea

- 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."

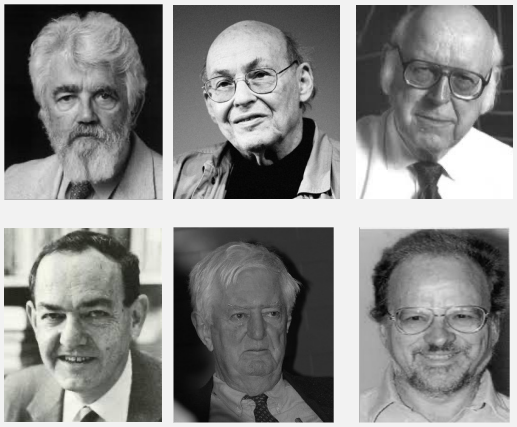

A brief (recent) history of AI

- Birth of artificial intelligence (1956)
- Early demonstrations of artificial intelligence and the publication of Computers and Thought (1959)
- 1960-1970
  - Un-tempered optimism fueled by early success on some problems thought to be hard (e.g., theorem proving) is tempered by slow progress on many problems thought to be easy (e.g., vision, language);
  - The field fragments into sub-areas focused on problem-solving, knowledge representation and inference, vision, planning, language processing, learning, etc.
- 1970-mid 1980s
  - investigation of knowledge representation and reasoning leads to many practical tools such as expert systems;
  - the difficulty task of knowledge engineering draws attention to the need for systems capable of learning from experience

A brief (recent) history of AI

- Mid 1980s-mid 1990s
  - Some of the failures of knowledge-based approaches to AI lead to renewed interest in biologically inspired neural network models leading to modest successes on some problems, e.g., vision.
  - Progress in machine learning begins to offer promising and practical alternatives to knowledge engineering.
  - Realization of the complementary strengths of knowledge-based and neural network approaches stimulates work on hybrid models.
  - Progress in the various sub-fields of AI refocuses attention on the design of intelligent agents that exhibit multiple aspects of intelligence.

A brief (recent) history of AI

- Mid 1990s-2010
  - The advent of the World-Wide-Web and advances in computing and storage technologies technologies make it possible for AI systems to be trained on massive amounts of data
  - Practical applications of AI in information retrieval, fault diagnosis, computer vision, information extraction from text, robotics, and related applications.
  - Major breakthroughs in learning theory offer insights that lead to practical advances, e.g., kernel machines, in machine learning.
  - There is growing interest in studying multi-agent systems, including inter-agent communication, coordination, and multi-agent organizations.

A brief (recent) history of AI

- 2010-2020
  - Increased availability of data and advanced hardware, e.g., graphical processing units, spur advances in deep learning and progress on computer vision, natural language processing, and related problems.
- 2020-present
  - Advances in powerful AI systems that could automate aspects of human intellectual work raise interest in socially responsible AI for maximizing the societal benefits of AI while minimizing its potential for harm.
  - There is a subtle shift in the goal of AI from automating intelligent behavior to augmenting and extending human intellect and abilities.