

Harnessing Artificial Intelligence

CS4000

SUMMER 2021

Peter Denning
Producer

Our Purpose

- EXPLORE
 - What is AI?
 - What is ML?
 - What are their benefits and risks?
 - What is their military significance?

AND ALSO


Establish a common language with precise definitions for discussing AI/ML.


Cut through the hype.

“Don’t mistake utility and speed for intelligence.”

AI field

- Founded 1956 to study how to build computing machines that perform intelligent functions
- Buoyed by enthusiasm, plagued by hype and over-promising





speech recognition
 board games (chess)
 language translation
 simple robots
 problem solving systems
 neural networks
 perceptrons

expert systems
 knowledge revolution
 multilevel neural networks
 statistical inference
 automatic classification
 sparse distributed memory
 Science Grand Challenges
 Fifth Generation Project

data analytics
 deep learning
 image recognition
 conversationalists (Siri, Alexis)
 grandmaster game players (Go, Poker)
 human machine symbiosis
 driverless vehicles
 automatic fire control
 swarms
 cyber defense and offense



What is AI?

Popular notions:

- Self-driving cars and auto-pilots
- Thinking machines
- Conversation machines
- Puzzle-solvers
- Games
- Speech recognizers
- Vision recognizers
- Expert systems

Hollywood notions:

- Terminator
- Skynet
- I, Robot
- Battlestar Gallactica
- Enders Game
- Commander Data
- R2D2

Things that worry us about AI

- Automatic weapon systems
- Drone swarms becoming WMD
- Mastering large battlespaces
- AI cannot explain its recommendations
- Mass surveillance and control
- Massive unemployment

What is AI?

- It's

Computers doing tasks previously considered human intelligent tasks

- But wait ...

What is intelligence?

What is intelligence?

- Problem: intelligence is ill-defined
- Some kinds of intelligence are not human – ants, bees, dogs, cats, chimpanzees, and others
- Elusive goalposts: “When a machine does it, we no longer think of it as intelligent”

- Intelligence is NOT speed
 - Machines do 1 Billion calculations per second
 - Humans do 1 calculation per second
- Your laptop and nearest supercomputer are not intelligent

Scientific Definition of Intelligence?

- Similar problem in other fields
- Biologists: is an organism living? Checklist to assess whether life is present:
 - Nutrition, respiration, movement, excretion, growth, reproduction, sensitivity
- Neuroscientists: when is someone conscious?
 - correlates such as eye movement or brainwave measures not definitive

Literal interpretations of AI and ML


- AI = artificial intelligence
 - Designing intelligent machines
 - Thinking, intuition, making conjectures, imagination, intention, making and keeping commitments, consciousness, cognition, self-awareness
- ML = machine learning
 - Designing machines that learn new functions
 - Speech recognition, image recognition, games, planning, deep learning, deep learning, conversations, driverless vehicles, automatic fire control, swarms

By these interpretations


- AI = artificial **intelligence**
 - No one has any idea how to do this
 - No progress since 1950s
- ML = machine learning
 - Almost all progress is actually ML
 - This is where the military applications are

The Turing Test (1950)

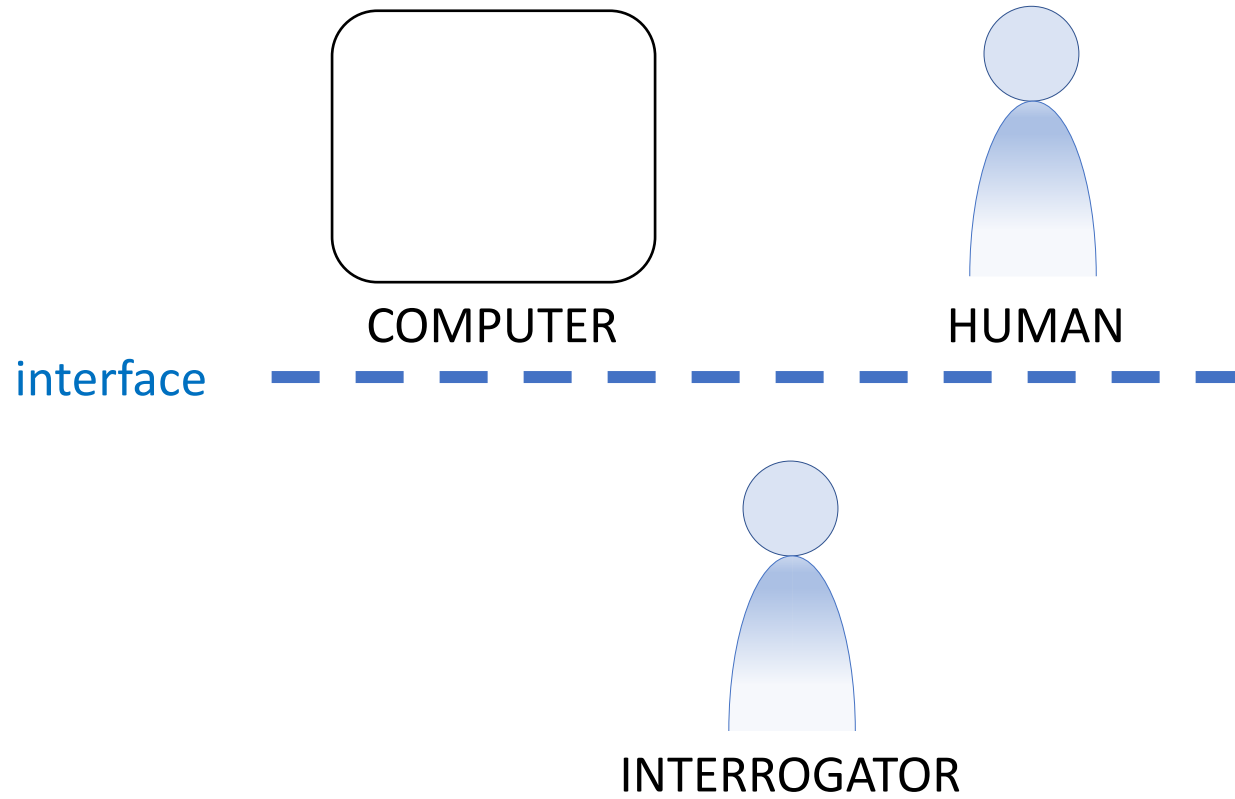
- Alan Turing sidestepped defining intelligence by specifying an "imitation game"
- Can a conversation machine fool a human interrogator?




COMPUTER




HUMAN





How long does it take
interrogator to distinguish
computer from human?



Even if indistinguishable,
is a simulation of
intelligence intelligent?

YOU: my relationship is in trouble
ELIZA: Is it important to you that your relationship is in trouble ?
YOU: yes
ELIZA: You are sure.
YOU: yes
ELIZA: I understand.
YOU: can we discuss my spouse
ELIZA: Why do you say your spouse ?
YOU: that's why my relationship is in trouble
ELIZA: Is it important to you that your relationship is in trouble ?
YOU: yes, i want to fix it
ELIZA: You seem to be quite positive.
YOU: that's what i want
ELIZA: What would getting mean to you ?
YOU: acquiring
ELIZA: Does that have anything to do with the fact that your relationship is in trouble ?
YOU: yes, my spouse nags me
ELIZA: You are sure.
YOU: all the time
ELIZA: But your spouse.
YOU: no one else
ELIZA: Are you thinking of a very special person ?
YOU: my spouse█

ELIZA: Please tell me what's been bothering you.
YOU: what's been bothering me
ELIZA: Why do you ask ?
YOU: what's been botherin me
ELIZA: Does that question interest you ?
YOU: what's been bothering me
ELIZA: What do you think ?
YOU: what's been bothering me
ELIZA: Does that question interest you ?
YOU: what's been bothering me
ELIZA: What comes to mind when you ask that ?
YOU: what's been bothering me
ELIZA: Are such questions much on your mind ?
YOU: what's been bothering me
ELIZA: What answer would please you most ?
YOU: what's been bothering me
ELIZA: What comes to mind when you ask that ?
YOU: what's been bothering me
ELIZA: Have you asked such questions before ?
YOU: what's been bothering me
ELIZA: What comes to mind when you ask that ?
YOU: what's been bothering me
ELIZA: Are such questions much on your mind ?
YOU: █

Turing test changed the (subjective) intelligence question to (objective) machine output question

But that did not help


We still don't know how to build a machine that passes the test

What is a machine?

- Apparatus of components that automates a process humans want done
- Grants an advantage such as amplified force or speed that humans do not naturally have
- Human partially or fully out of the loop
- Computer is machine with hardware components controlled by software – implements input-output functions described by algorithms

Hierarchy of Learning Machines

- Classify existing machines by learning power
- Learn: acquire a new capacity for action
 - Machine A is more powerful than machine B if A can learn to perform a function that B cannot
- Compare machine abilities, limitations, risks
- Precise language for evaluating learning machines



Machine A is more powerful than B:
A can learn functions B cannot

0	Basic Automation
1	Rule-based systems
2	Supervised learning
3	Unsupervised and reinforcement learning
4	Human-Machine Teaming
5	Aspirational AI

0	Basic Automation
1	Rule-based systems
2	Supervised learning
3	Unsupervised and reinforcement learning
4	Human-Machine Teaming
5	Aspirational AI



Different
degrees of
automation

Level 0 – Basic Automation

- Baseline – classical automation of unintelligent tasks
- Get a machine to perform a process, human (mostly) out of the loop
- May have feedback to ensure stable operation, but feedback does not change the function (no learning)
- ML may be used in some components
- Examples: autopilots, container shipping, robotic warehouse, voice response robots

Level 1 – Rule-based systems

- Programmed systems that make logical deductions from an input with help from a database of facts and rules of deduction
- Early examples: checker and chess
- Expert systems – aim to perform like an expert in a domain
- Examples: medical diagnosis, image interpretation, equipment repair, logistics

Level 2 – Supervised learning

- Most commonly neural networks
- Train the network to learn a particular input-output function given a large set of examples
 - Outputs are “labels” for inputs, network classifies
- Reliable training data a real problem
- Trained network approximates the desired function and may give inexact answers to some inputs
- What happens when network shown input not in the training set?
- Examples: success of image recognition; weird ways the network can go wrong

Level 3 – Unsupervised and reinforcement learning

- Learn a function from given data without being shown examples or guidance from a trainer
- Avoid the “sufficient data” problem
- Classification of unknown data into clusters
- Still need human interpretation of meanings of clusters
- Example: AUTOCLASS (NASA astronomy)

- Board games such as Chess and Go
- Use reinforcement learning, machine v. machine
- Reinforcement: feedback that rewards moves associated with wins and penalizes moves associated with losses
 - Complex algorithms not machine learning
- Example: AlphaZero (Go, Chess)
 - Learn grandmaster Chess in 4 hours
 - Learn grandmaster Go in 13 days

Level 4 – Human-Machine Teaming

- A system is designed with an interface that allows humans and machines to do what each is best at, producing a result that is better than any human or machine can do alone
- Examples: freestyle Chess, navigation-assisted cars

Level 5 – Aspirational AI

- These are all the intelligent machines AI researchers have dreamt about building ... the holy grails ... but to date do not exist

(and no one knows how to do them)

- Endless debates on whether such machines are possible
- No military AI depends on this level

The Context Problem

- To cross from learning machines to intelligent machines we will need a solution to the Context Problem
- Context: human ability to sense issues unstated in the background of our experience, bringing them forth to give meaning to what is around us

- Machines do not sense context
 - Their parts work “locally”: generate outputs (signals and symbols) and in response to inputs (signals and symbols)
 - Do not understand the meaning of any input or output
 - Source of their great speed
- Until we can formalize “bringing forth from the context” ... no way to get a machine to do it

CAUTIONS

1. Hype
2. Anthropomorphizing
3. Overbroad definitions
4. Master Algorithm claim

1. – Hype

Hype: a common bad habit of claiming more than machines can deliver

- Hype bandwagon: drop the term “AI” liberally into all presentations and proposals so as not to be left out of the AI game, even when AI has no known technology to do what you’re discussing

2. – Anthropomorphizing

Anthropomorphizing: attributing human attributes of intelligence to machines

- We love to project our ideas of intelligence into the machine:
 - “The computer is an **electronic brain**”
 - “The computer **knew** where I was”
 - “The computer **thought** the input was wrong”
 - “The computer **understood** my problem”
- But it’s not a brain
- It does not know, think, or understand

Humans -- Machines

- Social communities
- Empathy
- Compassion
- Commitments
- Judgments
- Invention
- Sensitive to context

- Calculations
- Logic
- Search
- Retrieval
- Comparisons
- Never bored
- Context free

3. – Overbroad definitions

- Example: “AI = machines with ability to perform functions normally considered intelligent”
 - No agreement on what functions make intelligence (if any)
 - Some people maintain that all human functions are intelligent and therefore all human-developed machines are AI, back to time immemorial
 - Some maintain that machine functions cannot be intelligent

4. – Master Algorithm

- Claim that deep learning is the “master algorithm” for all AI
 - Alan Turing: there can be no such thing as a master algorithm
- Deep learning is level 2 in hierarchy, many more powerful kinds of machines above it
- Related claim: automated systems prior to arrival of deep learning were “handcrafted” – i.e., built by ad hoc methods
 - Pre-AI systems are well engineered
 - Today’s deep learning systems are handcrafted

The Best Way Forward

- Make assessments whether AI can help a military system by looking at the requirements and evaluating the machines, algorithms, and data used to realize the requirements
- What machines of the hierarchy, and what other machines, must be combined to do the military job?