What Is Domain Of AI?

Solving problems that are hard for people?

- Math?

Performing tasks that people can do but monkeys can’t?

- Sort numbers
- Flexible behavior

Ambitious Goal Of AI

Building a complete, integrated, general-purpose, autonomous agent which can interact with the world and maximize its reward.

Practical Goal of AI

Solving really hard problems for which no feasible or known algorithm exists.

Not necessarily problems people can do

e.g., weather prediction

Often problems that people do effortlessly

perception, language, motor control

Does it stop being AI when a solution is found?

AI Problems: Vision

Where are the people in an image?

Who are the people in an image?

What images are similar to this one?

X-ray baggage screening

Fingerprint mathing

Perception of arbitrary scenes
(3d reconstruction, object identification)

captchas

AI Domains: Speech

Recognizing single command

Recognizing spoken phone numbers

Recognizing arbitrary phrases

Speaker identification

Emotion identification

AI Domains: Natural Language

Story understanding

John bought a new kite and went to the park. When he heard the clap of thunder, he sighed.

Answer arbitrary queries

Google search

Watson / Jeopardy

AI Domains: Search / Inference

Play chess

Get lunch

reasoning under uncertainty

Recommend a good movie

integration of a lot of information

AI Domains: Robotics

Robot that can map out floor plan of house

Robot that can assemble cars

Robot that can juggle

Robot that can move over unknown surfaces

Robot whose components can break apart and reassemble themselves

Origins of AI

1943 **“A logical calculus of the ideas immanent in nervous activity” (McCulloch and Pitts)**

Boolean circuit model of brain

laying foundation for neural networks

1948 ***Cybernetics*** **(Norbert Wiener)**

control theorist wrote about biological and mechanical control systems

introduced public to notion of AI machines

Origins of AI (continued)

1950 Turing “Computing Machinery and Intelligence”

Can machines think? Can machines behave intelligently?

Operational test for intelligent behavior: imitation game

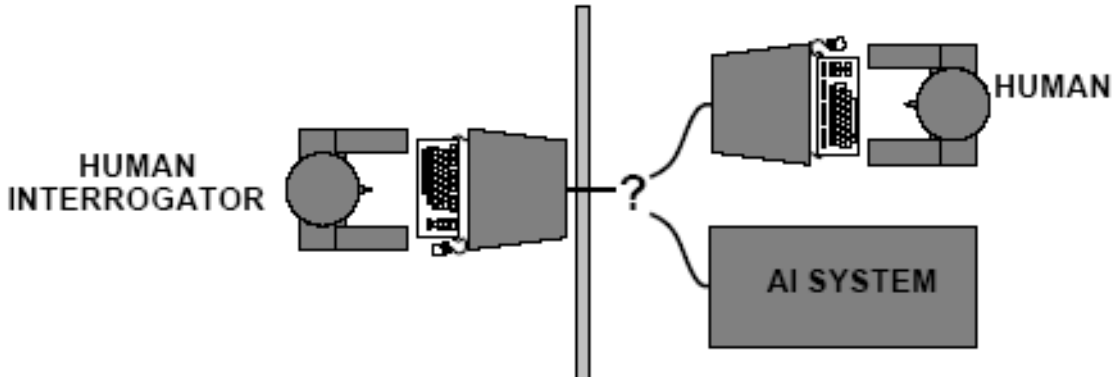

Predicted that by 2000, machine might have a 30% chance of fooling a lay person for 5 minutes

Suggested major components of AI: Knowledge representation, reasoning, language understanding, learning

1956 Dartmouth meeting

“Artificial Intelligence” adopted

Early Symbolic AI

1955 **Logic theorist (Newell, Shaw, & Simon)**

rederived theorems from *Principia Mathematica* (fundamental properties of numbers in terms of pure logic)

given: set of axioms of logic

derive implications (theorems) from axioms

heuristic search

1957 **General problem solver (Newell, Shaw, & Simon)**

domain independent

abstract rules of problem solving, e.g., “if you can, always try to bring your current state closer to the goal state”

did not scale

domain specific knowledge and heuristics were needed

Early Symbolic AI (continued)

1952-62 Samuels's checkers program

automatic tuning (learning) of heuristic evaluation function

strategy that followed in AI: microworlds, specialized domains

1958 Invention of LISP (McCarthy)

1963 ANALOGY (Evans)

1964 Algebra word problems (Bobrow)

early natural language understanding

1965 ELIZA (Weizenbaum)

interactive program, carries on a dialog on any topic

Rogerian psychotherapist version very popular

http://www.manifestation.com/neurotoys/eliza.php3

http://www.pandorabots.com/pandora/talk?botid=f5d922d97e345aa1

Early Symbolic AI (continued)

1965 **resolution theorem proving for first-order logic (Robinson)**

complete: guaranteed to draw conclusion if it follows from knowledge

limitation: scaling properties

1966 **semantic networks (Quillian)**

simple knowledge representation scheme

1966 **negative report on machine translation**

kills much work in NLP

1969 **Shakey (SRI)**

robot combining locomotion, perception, and problem solving to turn on lights, move from room to room

1971 **SHRDLU (Winograd)**

robot arm carries out instructions typed in English in blocks world

Early Symbolic AI (continued)

1972 **Development of Prolog (Coimerauer)**

theorem proving language

1979 **Nonmonotonic logic and truth maintenance systems (McDermott & Doyle; McCarthy)**

1980 **First AAAI (American Association for Artificial Intelligence) conference**

More Recent Symbolic AI

1969-79 Early development of knowledge-based systems

Realization that AI systems build on general-purpose search and inference mechanisms were not powerful enough to solve difficult problems.

Knowledge of specific domains is key to success.

DENDRAL: infers molecular structure from information provided by mass spectrometer

MYCIN: medical diagnosis (blood infections)

1980-88 Expert systems industry booms

R1: configures orders for new DEC computer systems

1988-93 Expert systems industry busts; “AI winter”

Other areas of AI pretty dead, esp. NLP

Early Subsymbolic AI

1949 ***The organization of behavior*** **(Hebb)**

proposed mechanism by which the brain adapts

neurobiological reality has since been verified

1962 **Perceptrons (Rosenblatt)**

A perceptron can learn to do any task it is programmed to do.

1969 **computational complexity analysis of perceptrons (Minsky and Papert)**

computational complexity analysis of perceptrons

origin of complexity analysis

1970-85 **neural networks abandoned**

1970s **computer vision (Marr)**

reconstructing 3D representation of world from camera images

Modern Subsymbolic AI

1982 **attractor networks (Hopfield)**

Neural network computing can be viewed as optimization

1986 **back propagation (Rumelhart, Hinton, & Williams)**

Neural network learning can be viewed as optimization

1988– **resurgence of probabilistic and statistical methods in AI**

Bayesian networks, support-vector machines, deep belief nets

Integration of ideas from symbolic and subsymbolic AI (e.g., Bayesian networks, artificial life, intelligent agents)

2005– **Machine learning and “big data” rule**

Issues

Scaling

Symbolic versus subsymbolic AI

Human versus rational intelligence

human intelligence: fidelity to human performance

rational intelligence: ideal performance; makes best possible inferences, performs best possible actions

For example...

Sam is a construction worker...
...who drives a pick up truck and goes out for a beer with his buddies after work on Friday.

Programmed knowledge vs. Learning

Bounded (limited) rationality

limitations on knowledge, time, computational power

“satisficing” (good enough) solutions