

Frontiers of Computer Science

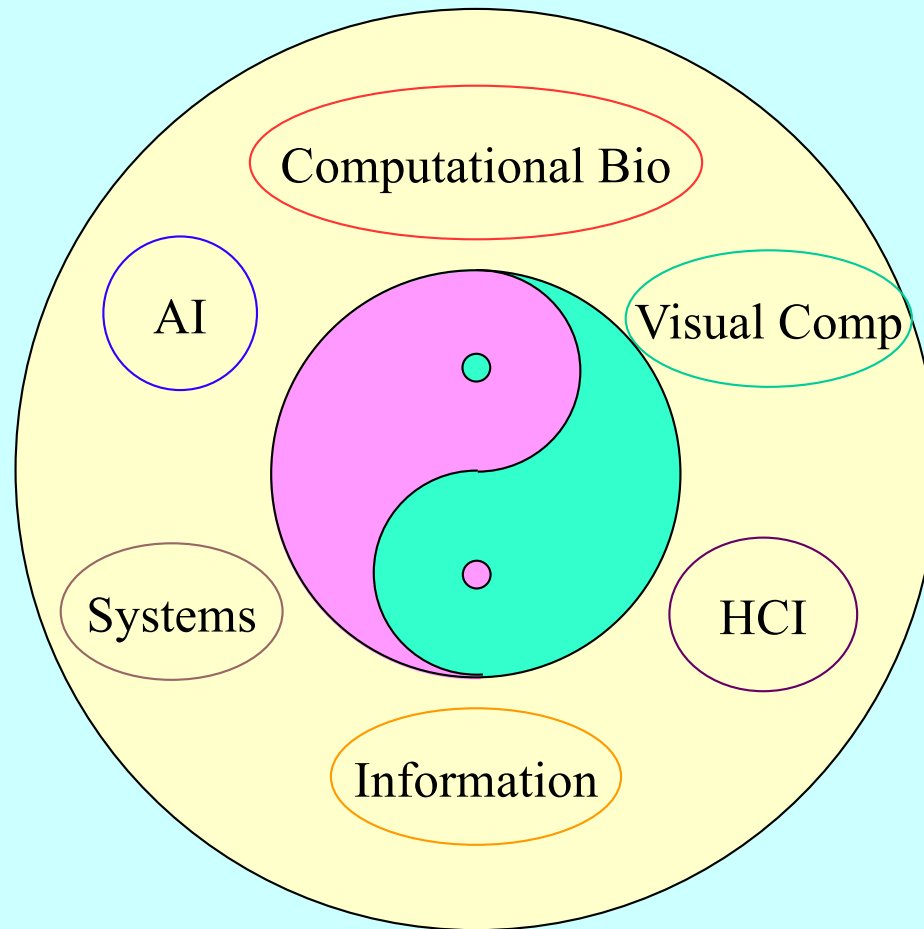
Jerry Cain

CS 106AX

December 1, 2023

slides courtesy of Eric Roberts

CS Major Tracks



Artificial Intelligence

Artificial Intelligence and the Turing Test

- In 1950, Alan Turing posed a thought experiment to address a question that underlies the field of Artificial Intelligence.

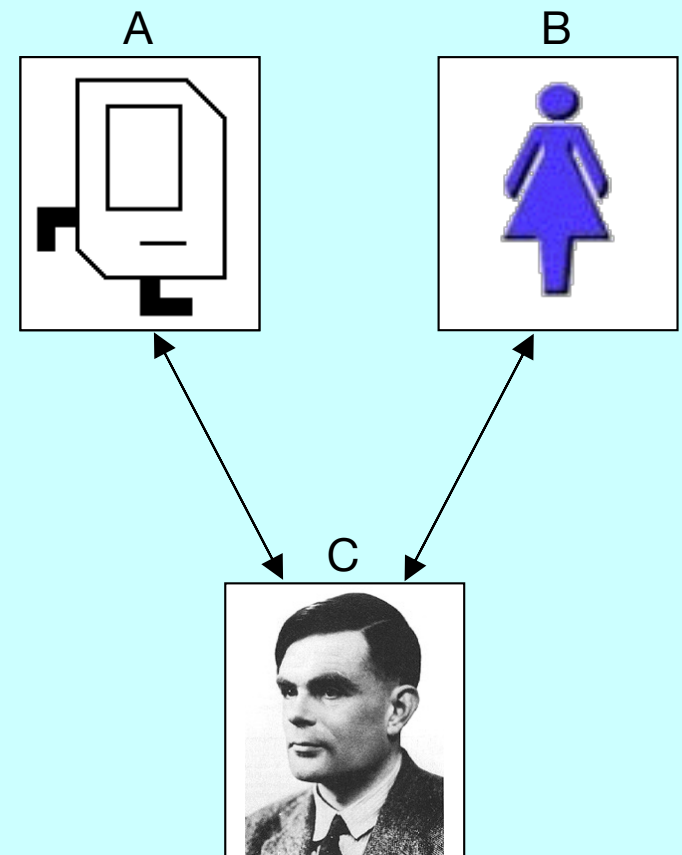
Computing Machinery and Intelligence A. M. Turing

1 The Imitation Game

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

[A revised form of the question] can be described in terms of a game which we call the 'imitation game'. It is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman. . . .

We now ask the question, 'What will happen when a machine takes the part of A in this game?'



A Dialogue in the Imitation Game

In his paper, Turing suggests the following as a possible dialogue between the interrogator and the unknown subject:




Turing notes that a computer trying to fool the interrogator might wait several seconds before answering the arithmetic problem, but does not point out that the answer is incorrect.

AI and Games

Chess as the Holy Grail for AI

- In 1950, Claude Shannon wrote an article for *Scientific American* in which he described how to write a chess-playing computer program.
- Shannon's strategy was to have the computer try every possible move for white, followed by all of black's responses, and then all of white's responses to those moves, and so on.
- Even with modern computers, it is impossible to use this strategy for an entire game, because there are too many possibilities.



*Positions evaluated: $\sim 10^{53}$
... millions of years later ...*

Deep Blue Beats Garry Kasparov

In 1997, IBM's Deep Blue program beat Garry Kasparov, who was then the world champion. In 1996, Kasparov had won in play that is in some ways more instructive.


Game 6

Kasparov

30. b6

Deep Blue

30. Bb8 ??



Watson Wins Jeopardy!



AlphaGo Beats the World Champion

The New York Times

Google's AlphaGo Defeats Chinese Go Master in Win for A.I.




Ke Jie, the world's top Go player, reacting during his match on Tuesday against AlphaGo, artificial intelligence software developed by a Google affiliate.

Machine Learning


Unsupervised Learning

- The goal of unsupervised learning is to find patterns even in the absence of a training set.
- The most common applications of unsupervised learning involve *clustering*, which is the process of dividing a data set into some number of independent, closely related clusters.
- As an example, if you have plotted a set of points as shown in the x - y grid at the bottom of the slide, unsupervised learning should be able to identify the three groups based on locality.

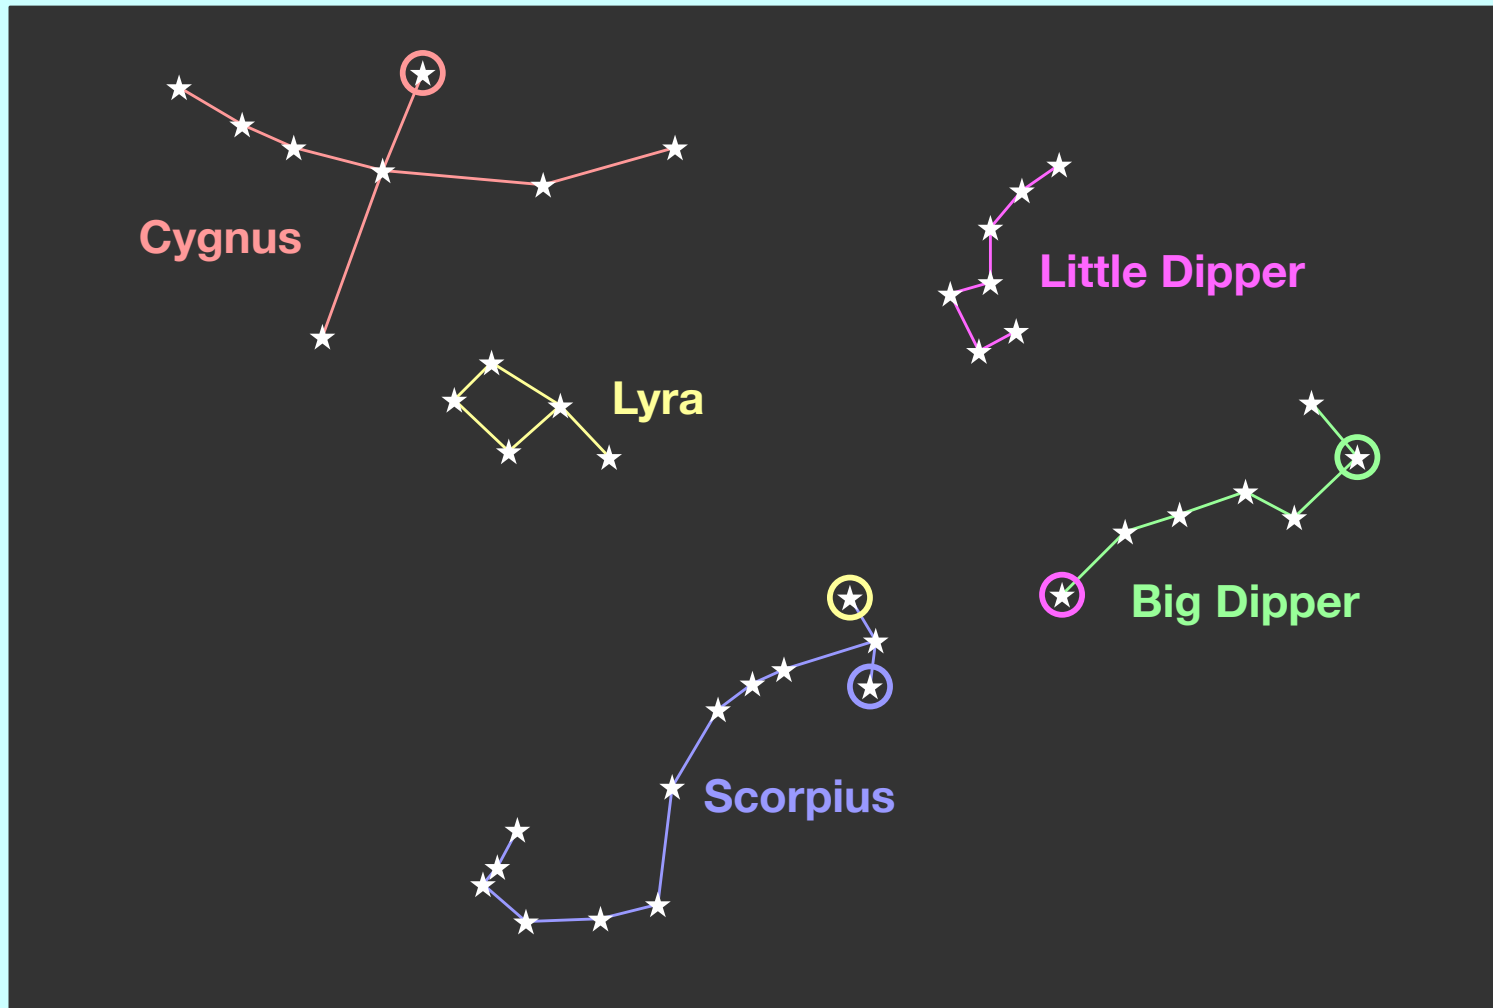


K-Means Clustering Algorithm

- One of the most common strategies for clustering data points is the *k-means algorithm*, which partitions a set into k clusters.
- The *k-means* algorithm requires the following steps:
 1. Choose k data points at random to serve as the initial centers of the clusters. These points are called *means*.
 2. Assign each point to the closest mean.
 3. Move each mean to the geometric center of its points.
 4. Repeat steps 2 and 3 until the means stop changing.




A Larger Clustering Example ($k = 5$)



1. Choose five random points as the initial means.
2. Assign each point to the closest mean.
3. Move each mean to the geometric center of its points.

Neural Networks

- In AI today, the hot topic is *neural networks*, which are collections of nodes that simulate neurons in the brain.
- Neural networks contain an *input layer* that provides a stimulus, an *output layer* that registers the results, and some number of *hidden layers* between the two.

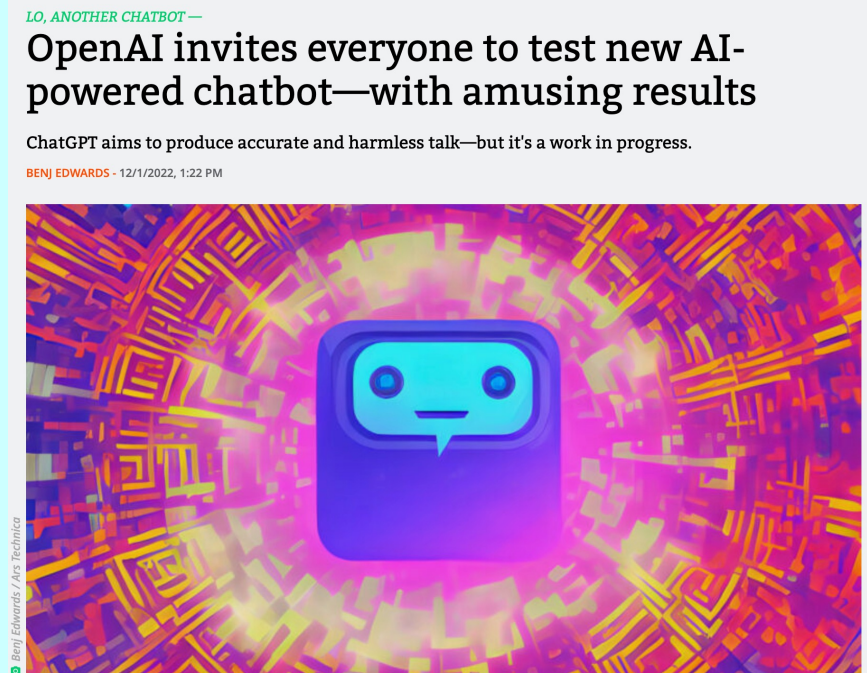


- Neural networks learn by updating the weights in the neurons through a process called *back propagation*.

Generative Pre-trained Transformer: Chat

"We've trained a model called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests. ChatGPT is a sibling model to InstructGPT, which is trained to follow an instruction in a prompt and provide a detailed response."

<https://openai.com/blog/chatgpt/>



Robotics

Stanley's Victory in the Desert

SFGATE

Sunday, October 9, 2005

NEVADA DESERT /
Computers, start your
engines / Stanford team
apparent winners in
robot car race

Tom Abate, Chronicle Staff Writer




Chronicle/Kurt Rogers

Primm, Nev. -- Stanford engineers steered the world toward a new era of driverless vehicles Saturday when their robotic Volkswagen SUV was the first to cross the finish line after a 132-mile race across the Nevada desert.

The Stanford car, nicknamed Stanley, unofficially edged out two robotic Hummers from Carnegie Mellon University. The three vehicles were competitors in a race sponsored by the Defense Advanced Research Projects Agency. DARPA, which backed the early research behind the Internet, drummed up interest in the race by offering \$2 million to the first team to complete the course in less than 10 hours.

OceanOne Explores the Ocean





Computational Biology

Computer-Assisted Pharmaceutical Design

—Professors Jean-Claude Latombe and Lydia Kavradi

- Designing new therapeutic drugs is an expensive, time-consuming process, in which computation can be of tremendous value.
- Most proteins are large molecules with a rigid, complex structure.
- Many drugs operate by blocking a reaction site in a protein. Such inhibitor drugs tend to be small and flexible.
- Understanding whether a drug molecule can fit into a reaction site is analogous to determining whether a robot arm can move in a particular way.




Visual Computing

Marc Levoy's Graphics Projects



<http://graphics.stanford.edu/projects/mich/>

The 3-D Fax Machine




1. Start with a statue of the Happy Buddha.
2. Use laser range scanning to produce a triangle mesh from a single perspective.
3. Merge scans from different perspectives.
4. Construct a digital model using shaded rendering.
5. Use stereolithography to construct a lucite copy.

<http://graphics.stanford.edu/projects/faxing/>

Scanning Michelangelo's David

In 1999, Professor Marc Levoy spent a year at Stanford's Overseas Studies campus in Florence, at which he and the approximately 30 students who spent at least part of their year there used the technology developed for the 3-D fax machine project to scan the Michelangelo sculptures open to the public.



The laser range-scanning technology allowed Marc to construct a model of each statue with sub-millimeter accuracy—accurate enough to tell when the sculptor changed chisels. This data makes it possible to view statues from any perspective and to analyze them mathematically.

The *Forma Urbis Romae* Project

While in Italy, Professor Levoy and his students also scanned fragments of the *Forma Urbis Romae*, a huge marble map from the 3rd century AD that is now a jigsaw puzzle with 1,186 pieces.



Scanning the fragments makes it possible to use computational techniques to reassemble the pieces of the puzzle. The first match found by Levoy's team is shown at the left. Since that time, the Stanford group has been able to determine the placement of missing pieces at a far faster rate than was previously possible.

<http://graphics.stanford.edu/projects/forma-urbis/>

Thank You!