

# **Artificial Intelligence**

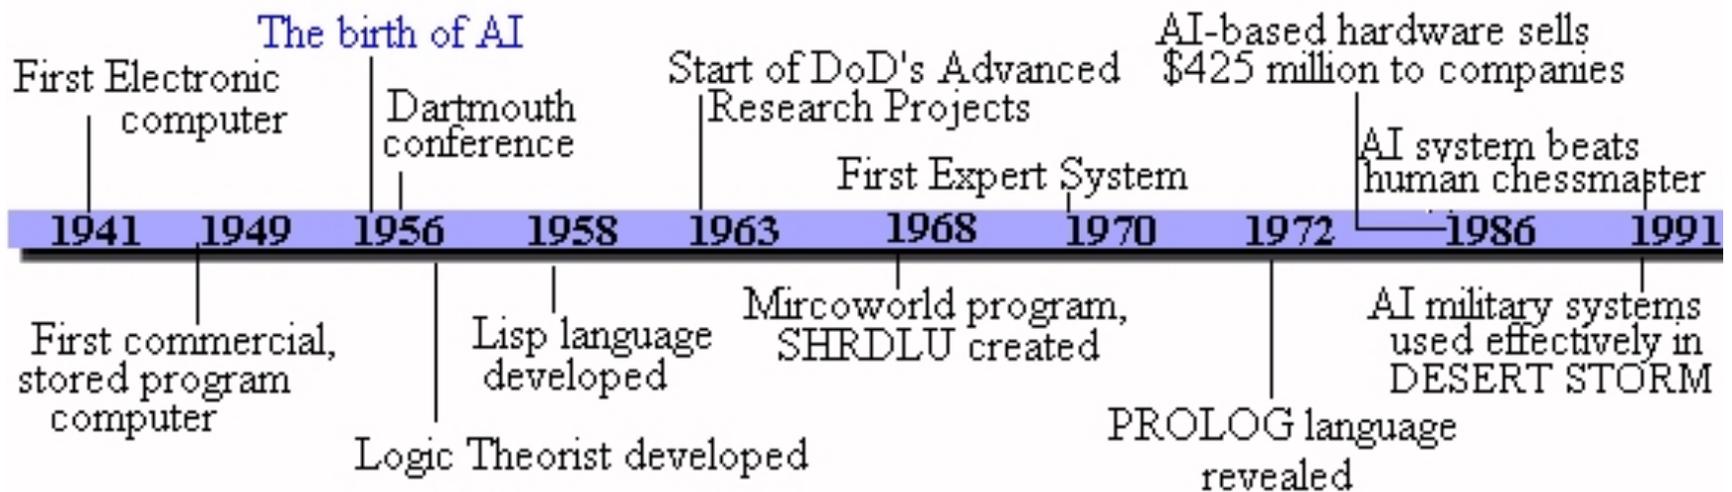
**Prof. Eric Eaton**

# Today's class

- Course overview
- Introduction
  - Brief history of AI
  - What *is* AI? (and why is it so cool?)
  - What's the state of AI now?

**What is AI??**

# History



It is not my aim to surprise or shock you – but the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until – in a visible future – the range of problems they can handle will be coextensive with the range to which the human mind has been applied. ~ Herbert Simon, 1957

# Eliza

- ELIZA: A program that simulated a psychotherapist interacting with a patient and successfully passed the Turing Test.
- Coded at MIT during 1964-1966 by Joel Weizenbaum.
- First script was DOCTOR.
  - The script was a simple collection of syntactic patterns not unlike regular expressions
  - Each pattern had an associated reply which might include bits of the input (after simple transformations (my → your))
- Weizenbaum was shocked at reactions:
  - Psychiatrists thought it had potential.
  - People unequivocally anthropomorphized.
  - Many thought it solved the NL problem.

# Eliza

- This is a transcript of a conversation with an Eliza-like program found in Emacs

- Try typing

*Meta-x  
doctor*

in your favorite Emacs.

- I am the psychotherapist. Please, describe your problems. Each time you are finished talking, type RET twice.
- i'm having a hard time with my students
- Why do you say that?
- well, they keep asking me for money
- Maybe your plans have something to do with this.
- what plans?
- Is it because of your life that you say what plans?
- No, it's because I didn't know what you were talking about
- Is it because you didn't know what I was talking about that you came to me?
- no, it's because i wanted to see if you could pass the turing test
- Does it bother you that you wanted to see if I could really pass the turing test?
- .....
- How do you reconcile problems at school?
- This is a waste of time. Good bye
- My secretary will send you a bill.

# Colby's PARRY

- Kenneth Colby modeled a paranoid using the same techniques circa 1968.
- PARRY has basic emotions. If it gets angry, its replies become more hostile.
- In a Turing test, expert psychiatrists were unable to distinguish PARRY's ramblings from those of real paranoids.

H: Good afternoon. My name is Doctor Jones.

P: Pleased to meet you.

H: How are you?

P: All right I guess.

H: Could you tell me your name?

P: Pat Smith.

H: I thought you might have some problems.

P: I've been upset lately.

H: By what?

P: People get on my nerves sometimes.

H: Tell me more.

P: Do you know anything about bookies?

...

# Parry meets Eliza

Please tell me *\*this\** isn't AI!

P: I don't understand your motives.

E: Why don't you understand my motives?

P: Shouldn't I be?

E: Why do you say that just now?

P: That's just the way it is.

E: Suppose that's the way it were not.

P: So?

E: I am not sure I understand you fully.

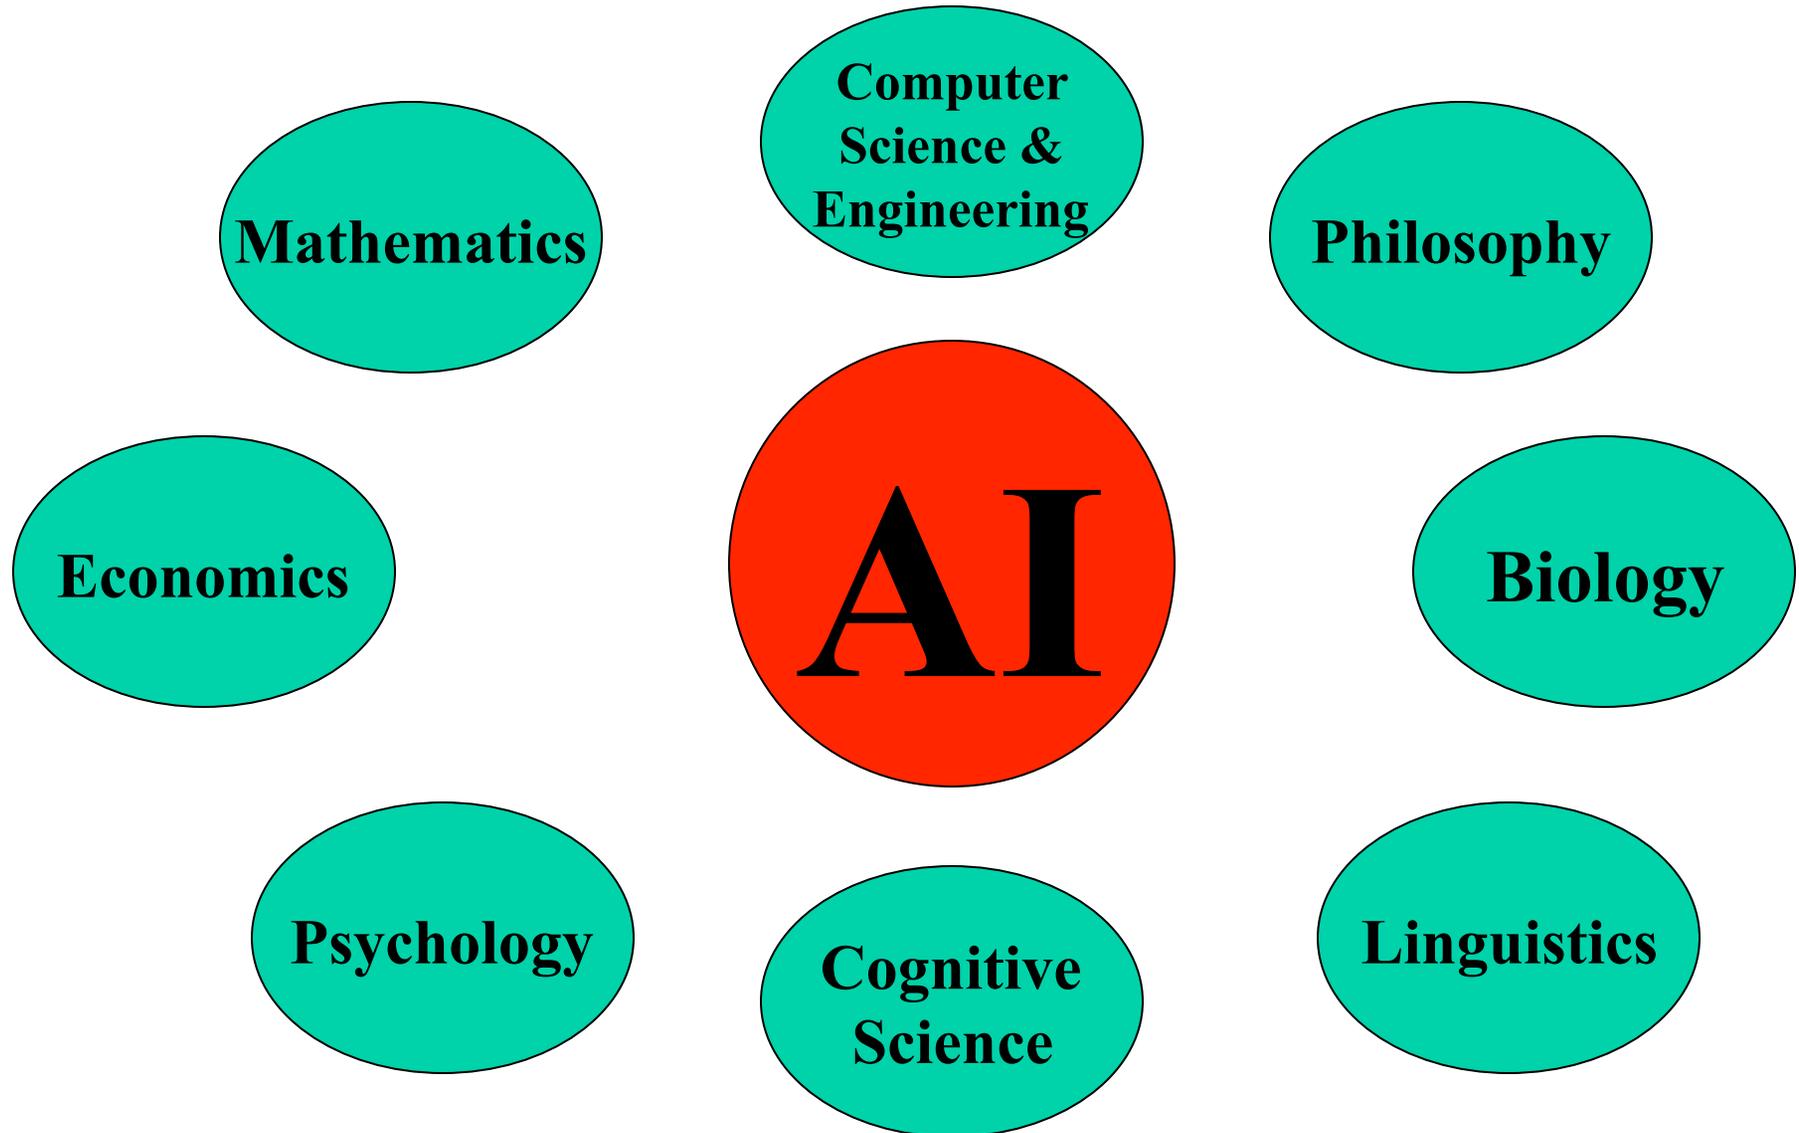
P: You harp on the same questions too much.

E: What does that similarity suggest to you?

P: I don't get you.

E: Do you wish that you did get I?

# Foundations of AI



# Big questions

- Can machines think?
- If so, how?
- If not, why not?
- What does this say about humans?
- What does this say about the mind?

# Why pursue AI?

- **Engineering:** To get machines to do a wider variety of useful things
  - e.g., understand spoken natural language, recognize individual people in visual scenes, find the best travel plan for your vacation, etc.
- **Cognitive Science:** As a way to understand how natural minds and mental phenomena work
  - e.g., visual perception, memory, learning, language, etc.
- **Philosophy:** As a way to explore some basic and interesting (and important) philosophical questions
  - e.g., the mind body problem, what is consciousness, etc.

# What's easy and what's hard for AI?

- It's been easier to mechanize many of the high-level tasks we usually associate with "intelligence" in people
  - e.g., symbolic integration, proving theorems, playing chess, medical diagnosis
- It's been very hard to mechanize tasks that lots of animals can do
  - walking around without running into things
  - catching prey and avoiding predators
  - interpreting complex sensory information (e.g., visual, aural, ...)
  - modeling the internal states of other animals from their behavior
  - working as a team (e.g., with pack animals)
- Is there a fundamental difference between the two categories?

# Turing Test

- Three rooms contain a person, a computer, and an interrogator.
- The interrogator can communicate with the other two by teleprinter.
- The interrogator tries to determine which is the person and which is the machine.
- The machine tries to fool the interrogator into believing that it is the person.
- If the machine succeeds, then we conclude that the machine can think.

# The Loebner contest

- A modern version of the Turing Test, held annually, with a \$100,000 cash prize.
- Named after Hugh Loebner
- <http://www.loebner.net/Prizetf/loebner-prize.html>
- Restricted topic (removed in 1995) and limited time.
- Participants include a set of humans and a set of computers and a set of judges.
- Scoring
  - Rank from least human to most human.
  - Highest median rank wins \$2000.
  - If better than a human, win \$100,000. (Nobody yet...)

# What can AI systems do?

Here are some example applications

- **Computer vision:** face recognition from a large set
- **Robotics:** autonomous (mostly) automobile
- **Natural language processing:** simple machine translation
- **Expert systems:** medical diagnosis in a narrow domain
- **Spoken language systems:** ~1000 word continuous speech
- **Planning and scheduling:** Hubble Telescope experiments
- **Learning:** text categorization into ~1000 topics
- **User modeling:** Bayesian reasoning in Windows help (the infamous paper clip...)
- **Games:** Grand Master level in chess (world champion), checkers, etc.

# What can't AI systems do yet?

- Understand natural language robustly (e.g., read and understand articles in a newspaper)
- Surf the web
- Interpret an arbitrary visual scene
- Learn a natural language
- Play Go well
- Construct plans in dynamic real-time domains
- Refocus attention in complex environments
- Perform life-long learning

**Exhibit true autonomy and intelligence!**



# Who does AI?



- Academic researchers (perhaps the most Ph.D.-generating area of computer science in recent years)
  - Some of the top AI schools: CMU, Stanford, Berkeley, MIT, UIUC, UMd, U Alberta, UT Austin, ...
- Government and private research labs
  - NASA, NRL, NIST, IBM, AT&T, SRI, ISI, MERL, ...
- Lots of companies!
  - Google, Microsoft, Honeywell, Teknowledge, SAIC, MITRE, Fujitsu, Global InfoTek, BodyMedia, ...



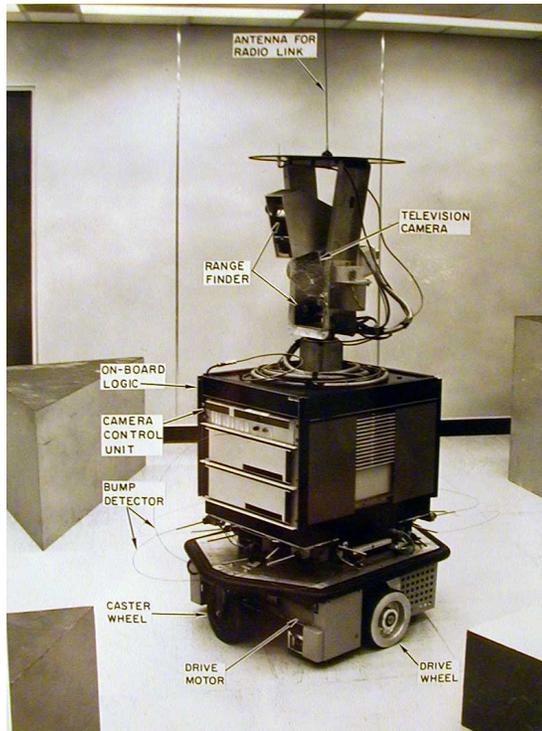
# **Evolutionary optimization, virtual life**

**Evolved Virtual  
Creatures**

**Examples from  
work in progress**

# Robotics

Shakey (1966-1972)



Cog (90s)



Kismet (late 90s, 2000s)



Robocup Soccer (2000s)



Boss (2007)

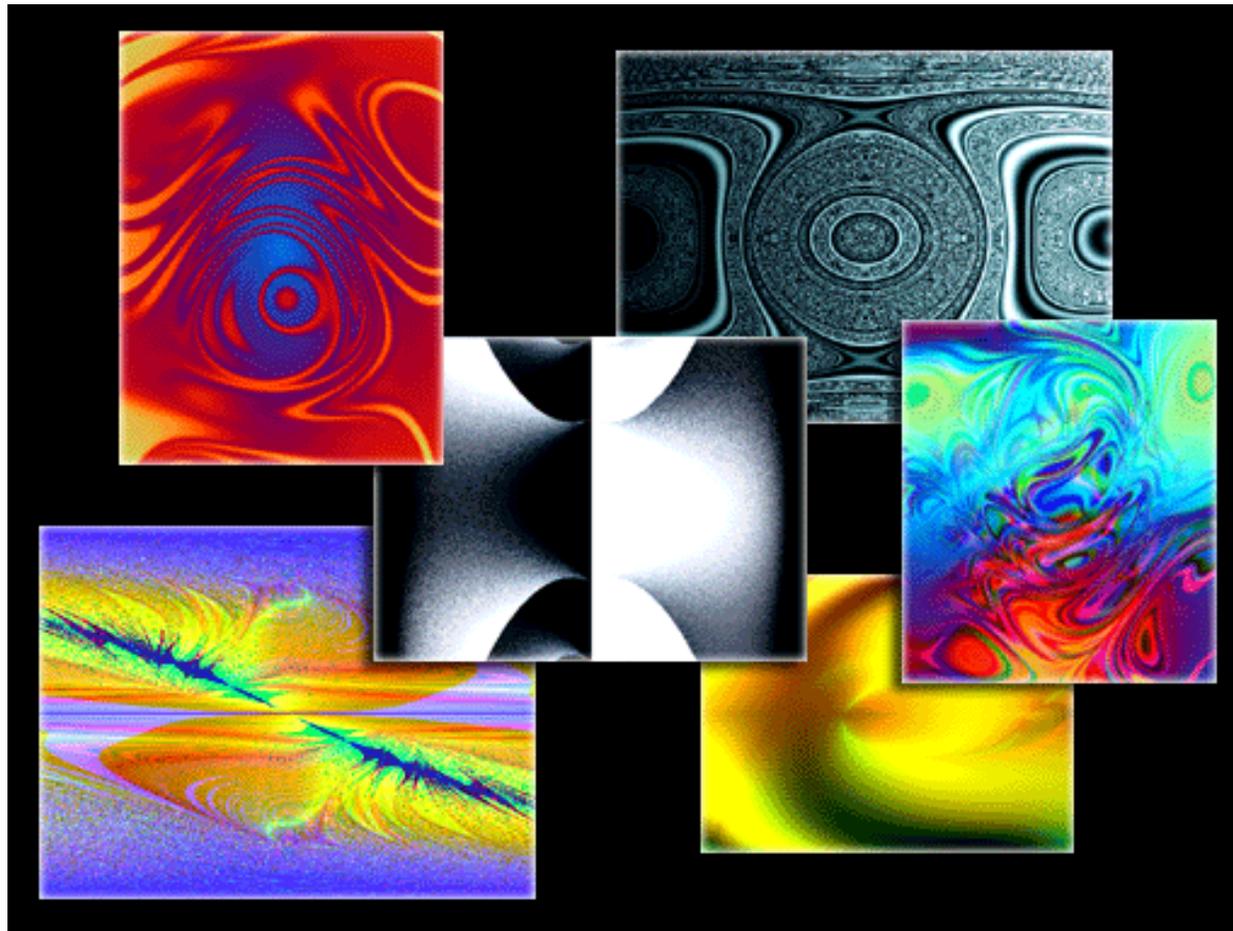


# Applications

- A sample from recent IAAI conferences:
  - Real-Time Identification of Operating Room State from Video
  - Developing the next-generation prosthetic arm
  - Automatically mapping planetary surfaces
  - Automated processing of immigration applications
  - Crops selection for optimal planting
  - Heart wall motion abnormality detection
  - Classifying handwriting deficiencies
  - Personal assistants
  - Emergency landing planner for damaged aircraft
  - Airspace deconfliction
  - Art print authentication
  - Price prediction for Ebay online trading

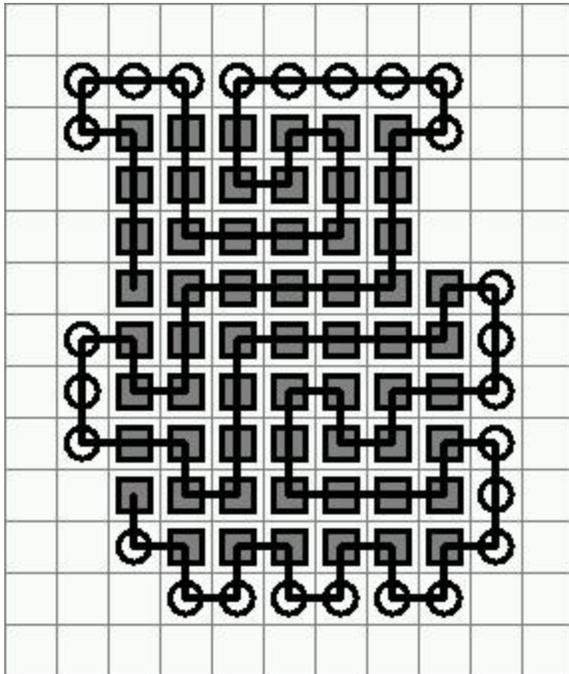
# AI & art: NEvAr

- Neuro-evolutionary Art
  - See <http://eden.dei.uc.pt/~machado/NEvAr>

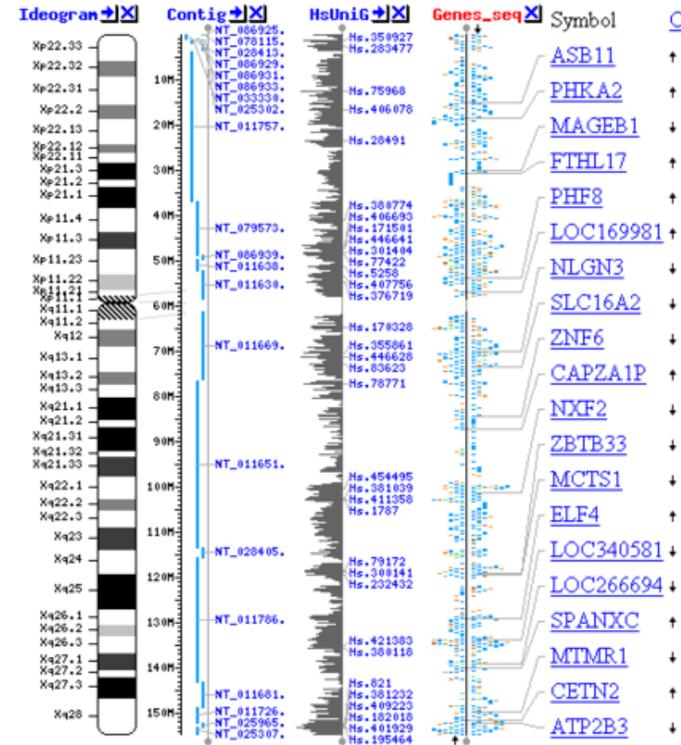


# Bioinformatics

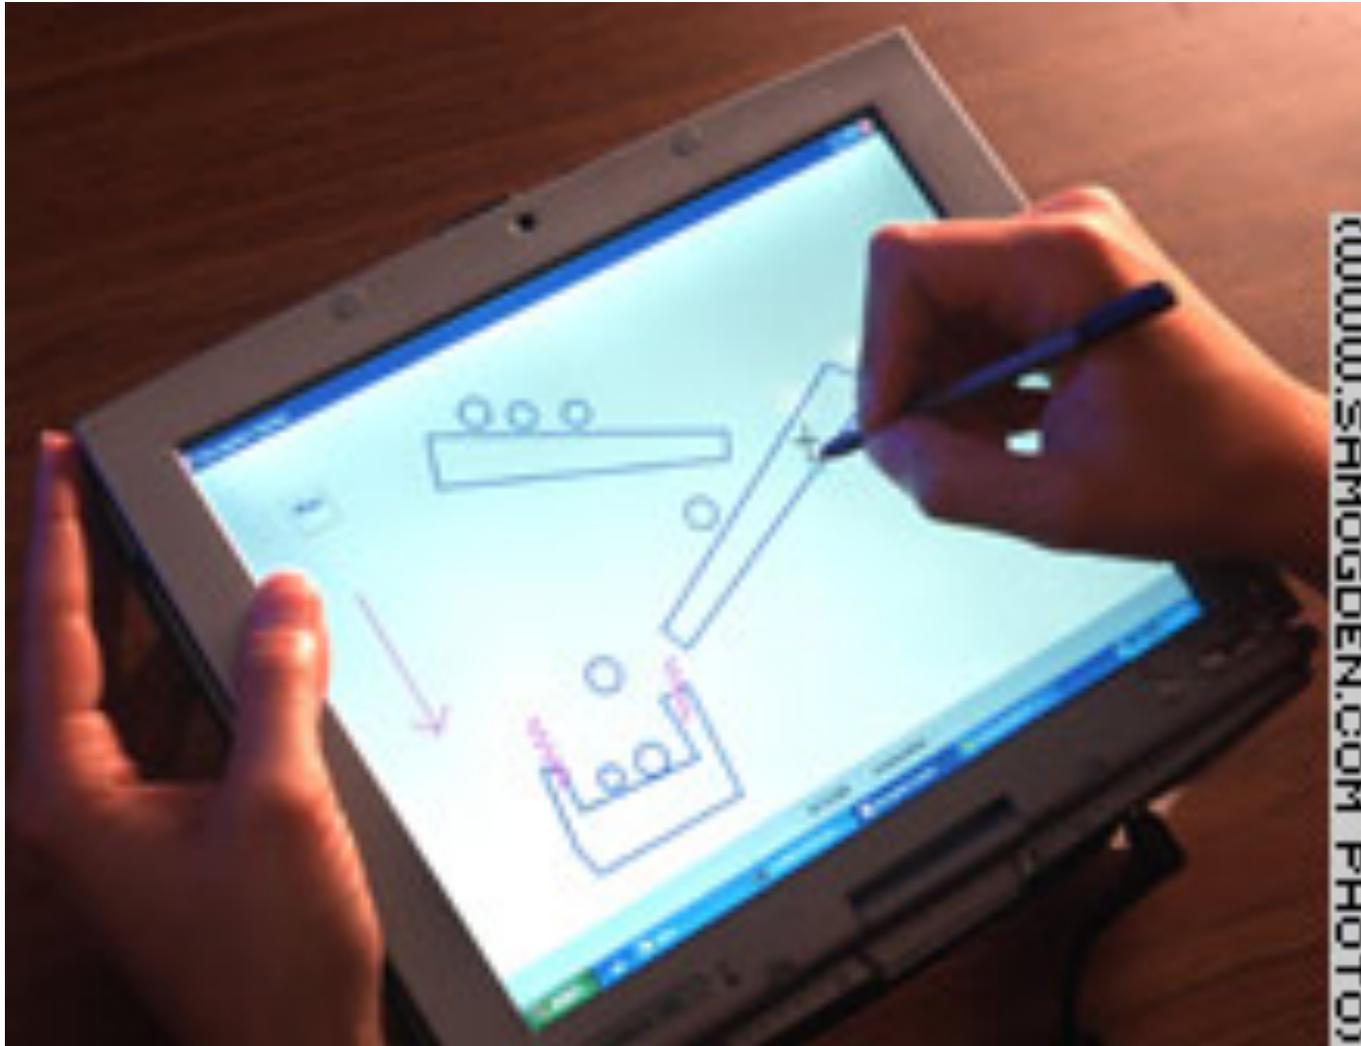
- MERL: constraint-based approach to protein folding



- Genetic Motif Discovery and Mapping



# Interaction: MIT Sketch Tablet

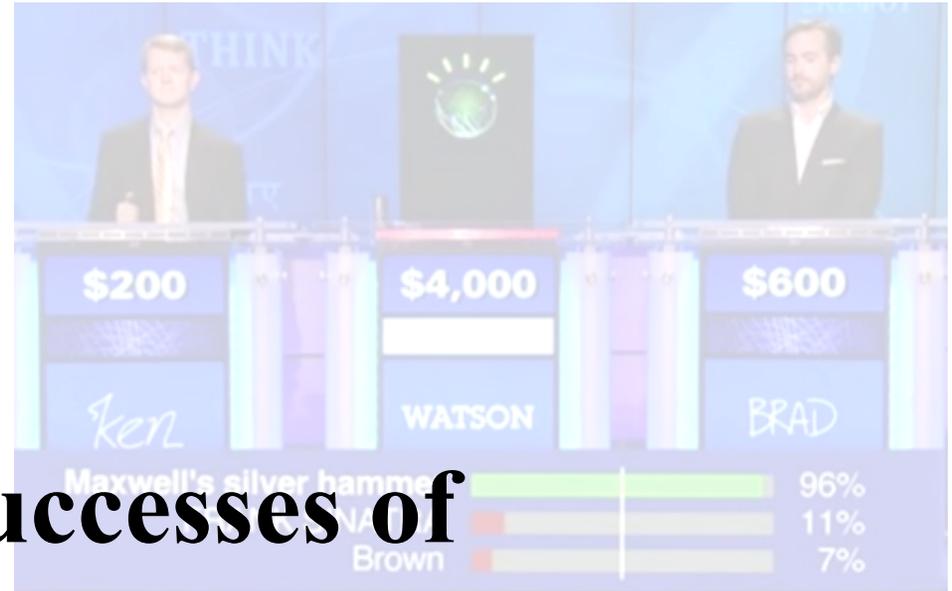


# Swarm Systems and Emergence

- Boids demo!

# Other topics/paradigms

- Intelligent tutoring systems
- Agent architectures
- Mixed-initiative systems
- Embedded systems / mobile autonomous agents
- Machine translation
- Statistical natural language processing
- Object-oriented software engineering / software reuse



# Recent Successes of Artificial Intelligence



# AI's Recent Successes

- The IBM Deep Blue chess system beats the world chess champion Kasparov (1996)
- Checkers is solved (July 2007)
- Autonomous cars complete the DARPA Grand and Urban Challenges (2005/2007)
- Computer beats humans at Jeopardy! (2011)



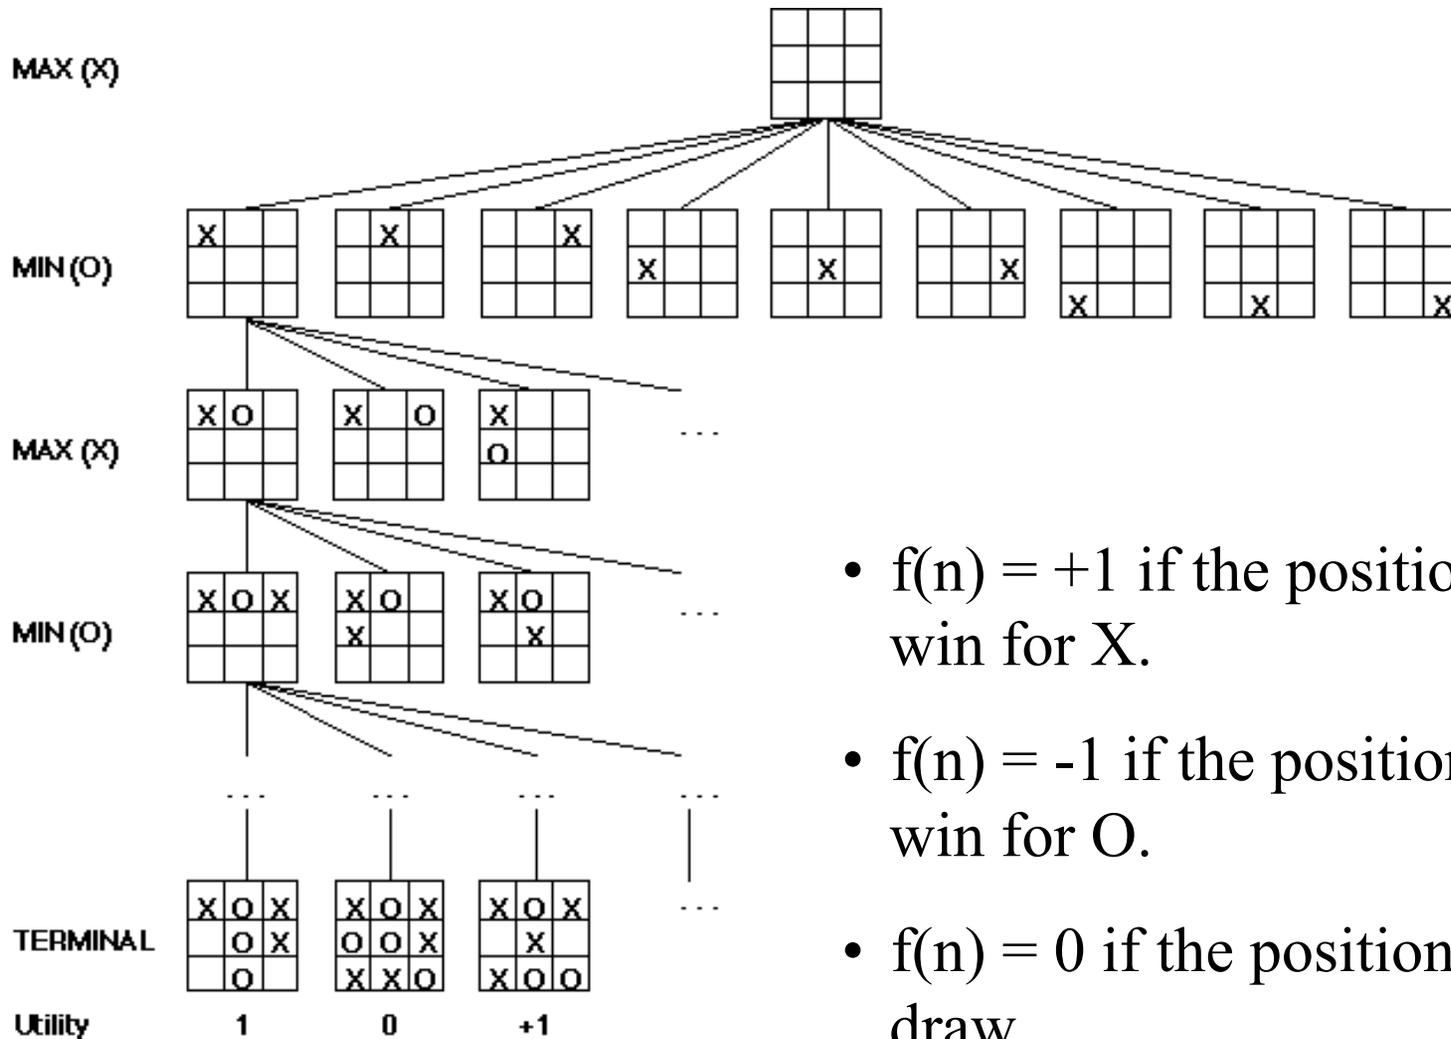
# IBM's Deep Blue versus Kasparov

- On May 11, 1997, Deep Blue was the first computer program to beat reigning chess champion Kasparov in a 6 game match (2 : 1 wins, with 3 draws)
- Massively parallel computation (259<sup>th</sup> most powerful supercomputer in 1997)
- Evaluation function criteria learned by analyzing thousands of master games



- Searched the game tree from 6-12 ply usually, up to 40 ply in some situations.
  - One ply corresponds to one turn of play.

# Partial Game Tree for Tic-Tac-Toe



- $f(n) = +1$  if the position is a win for X.
- $f(n) = -1$  if the position is a win for O.
- $f(n) = 0$  if the position is a draw.

# Checkers is Solved – It's a Draw!

## (July 2007)

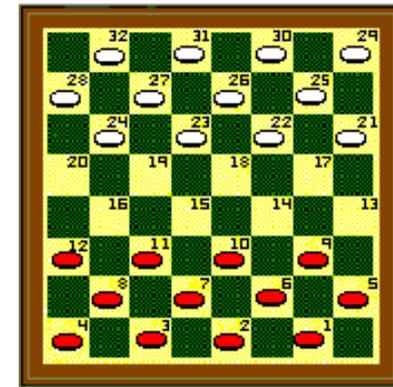
- Researchers at the University of Alberta proved that perfect play on both sides in checkers results in a draw.
- Dozens of computers have been working in parallel since 1989 to get this result.
- Checkers has approximately 500 billion billion possible positions ( $5 \times 10^{20}$ ).
- Deep Blue used heuristics to win.
- This research solves the game of checkers, yielding a perfect player that no longer needs heuristics.



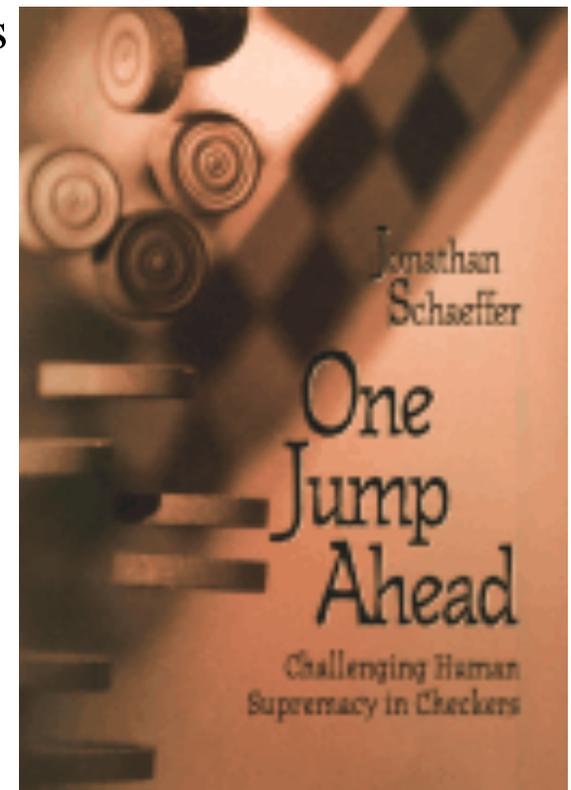
# Chinook

- Chinook is the World Man-Machine Checkers Champion, developed by researchers at the University of Alberta.
- It earned this title by competing in human tournaments, winning the right to play for the (human) world championship, and eventually defeating the best players in the world.
- Visit <http://www.cs.ualberta.ca/~chinook/> to play a version of Chinook over the Internet.
- The developers have fully analyzed the game of checkers and have the complete game tree for it.
  - Perfect play on both sides results in a tie.
- “One Jump Ahead: Challenging Human Supremacy in Checkers” Jonathan Schaeffer, University of Alberta (496 pages, Springer. \$34.95, 1998).

The board set for play



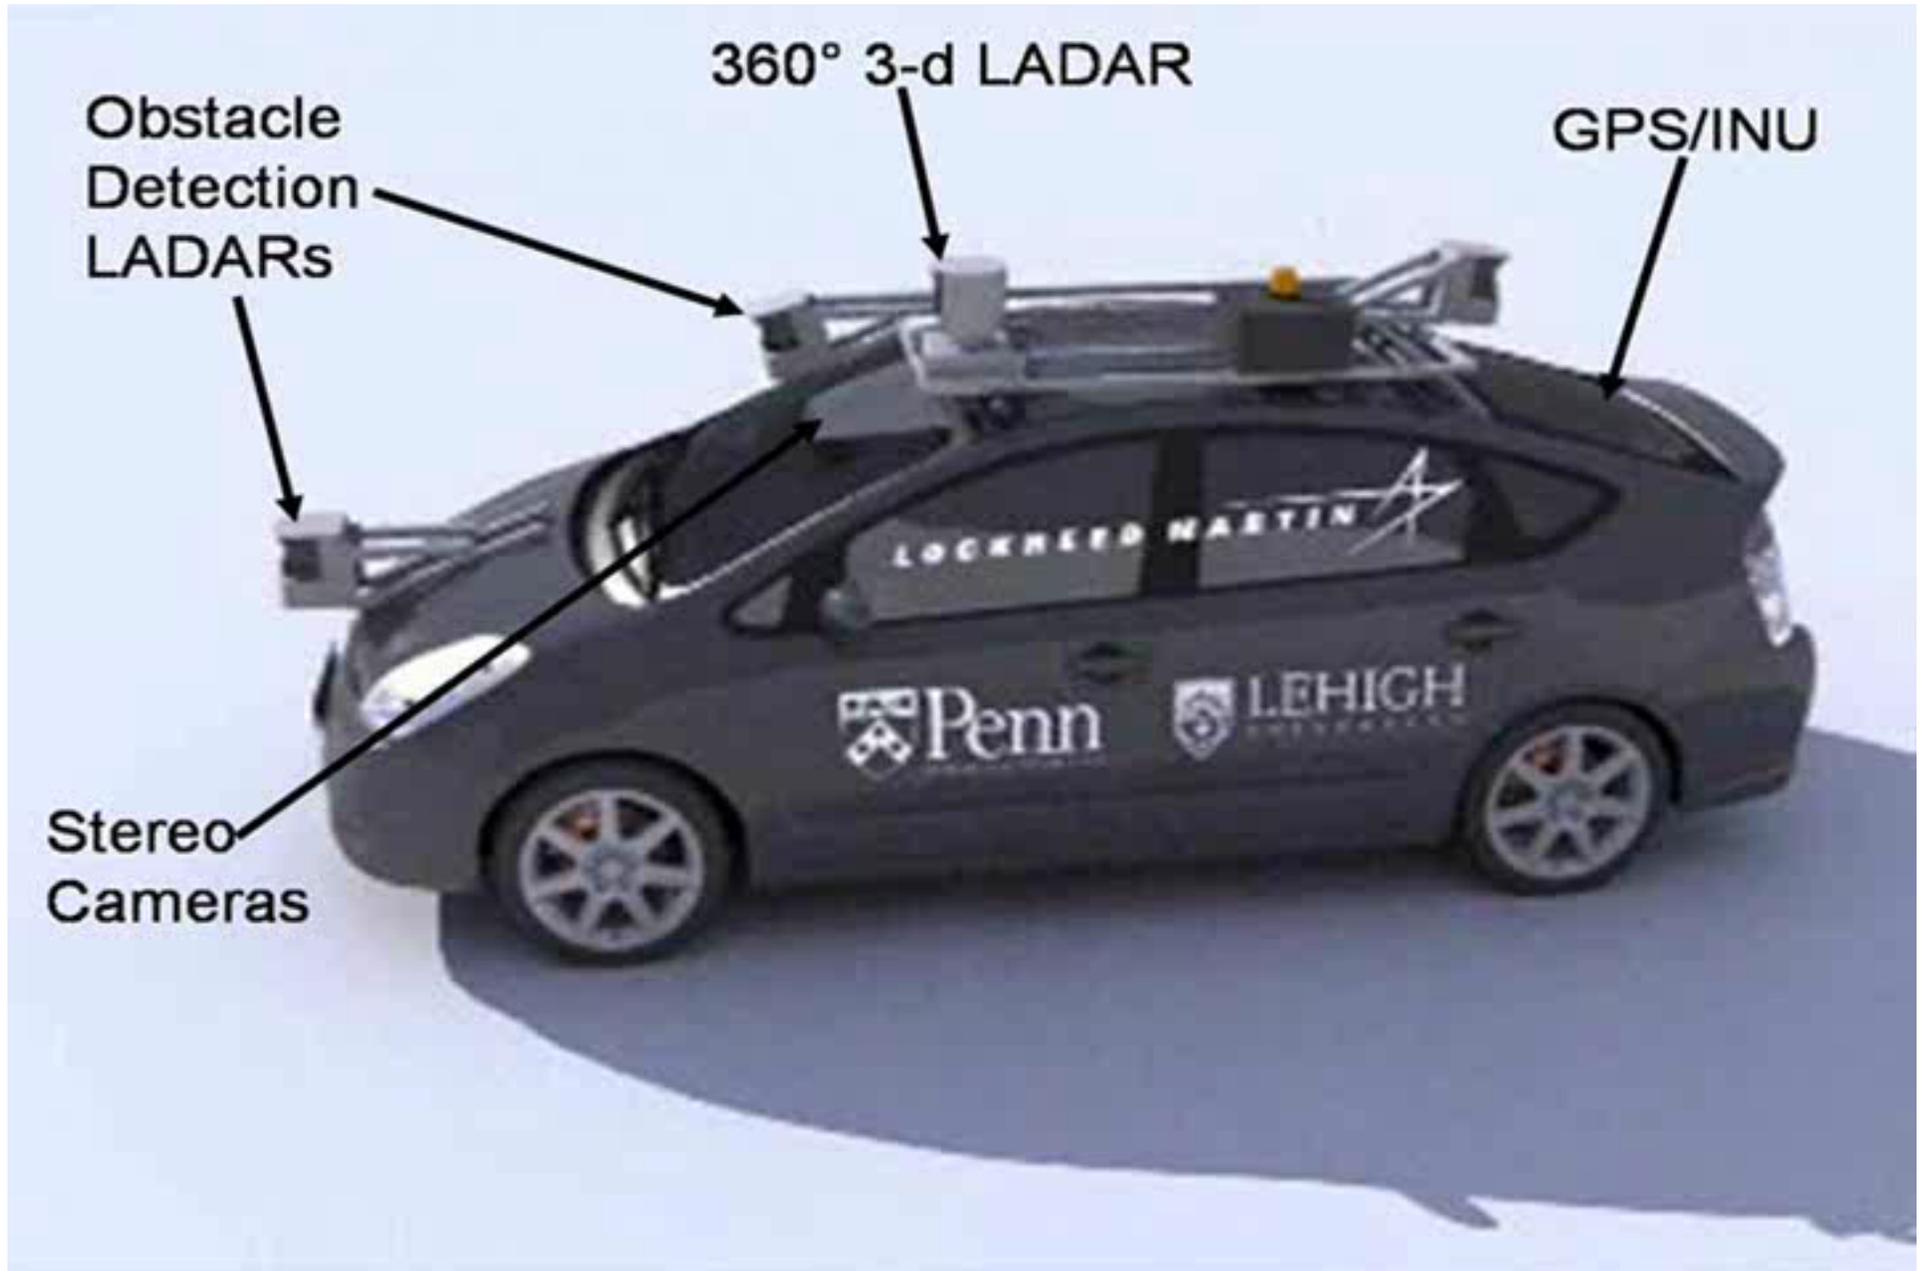
Red to play



# 2005 DARPA Grand Challenge

- A race of autonomous vehicles through the Mojave desert, including 3 narrow tunnels and winding paths with steep drop-offs.
- The route was provided 2 hrs before the start in the form of GPS waypoints every 72 meters.
- The Stanford Racing Team won with a time of 6:54 hrs, closely followed by two teams from CMU (7:05hrs, 7:14 hrs) and the Gray Insurance Company (7:30 hrs). Next closest was 12:51 hrs.





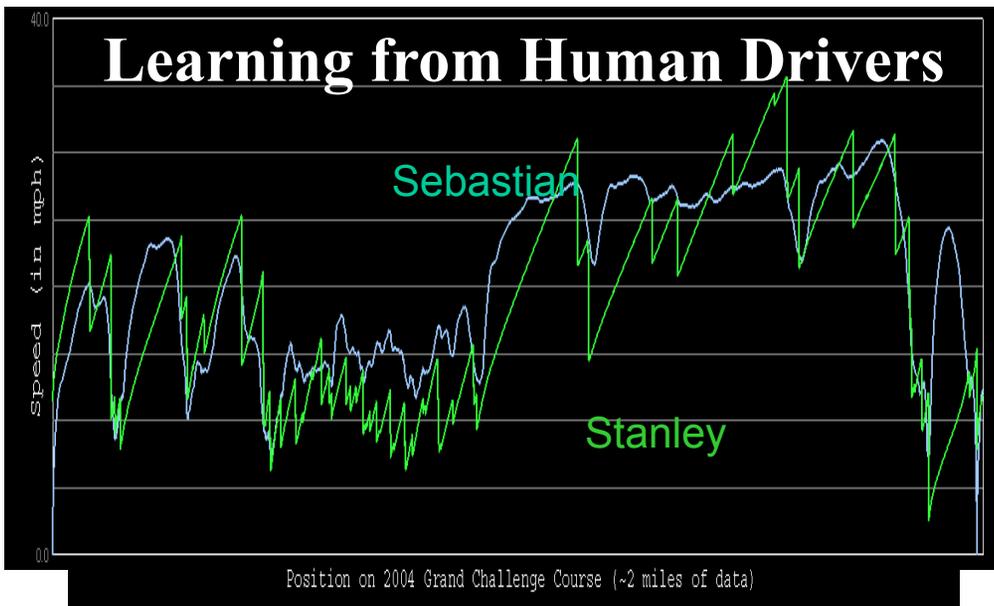
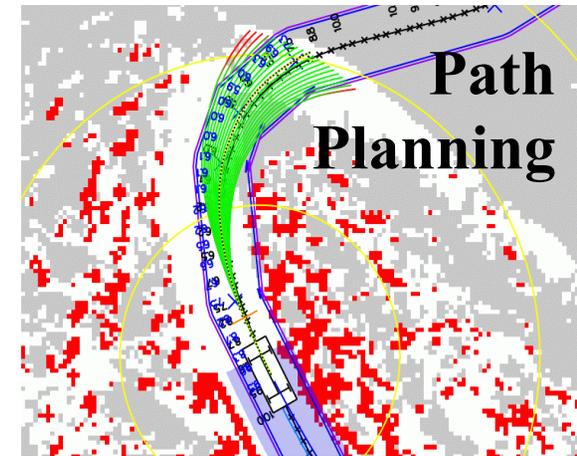
360° 3-d LADAR

GPS/INU

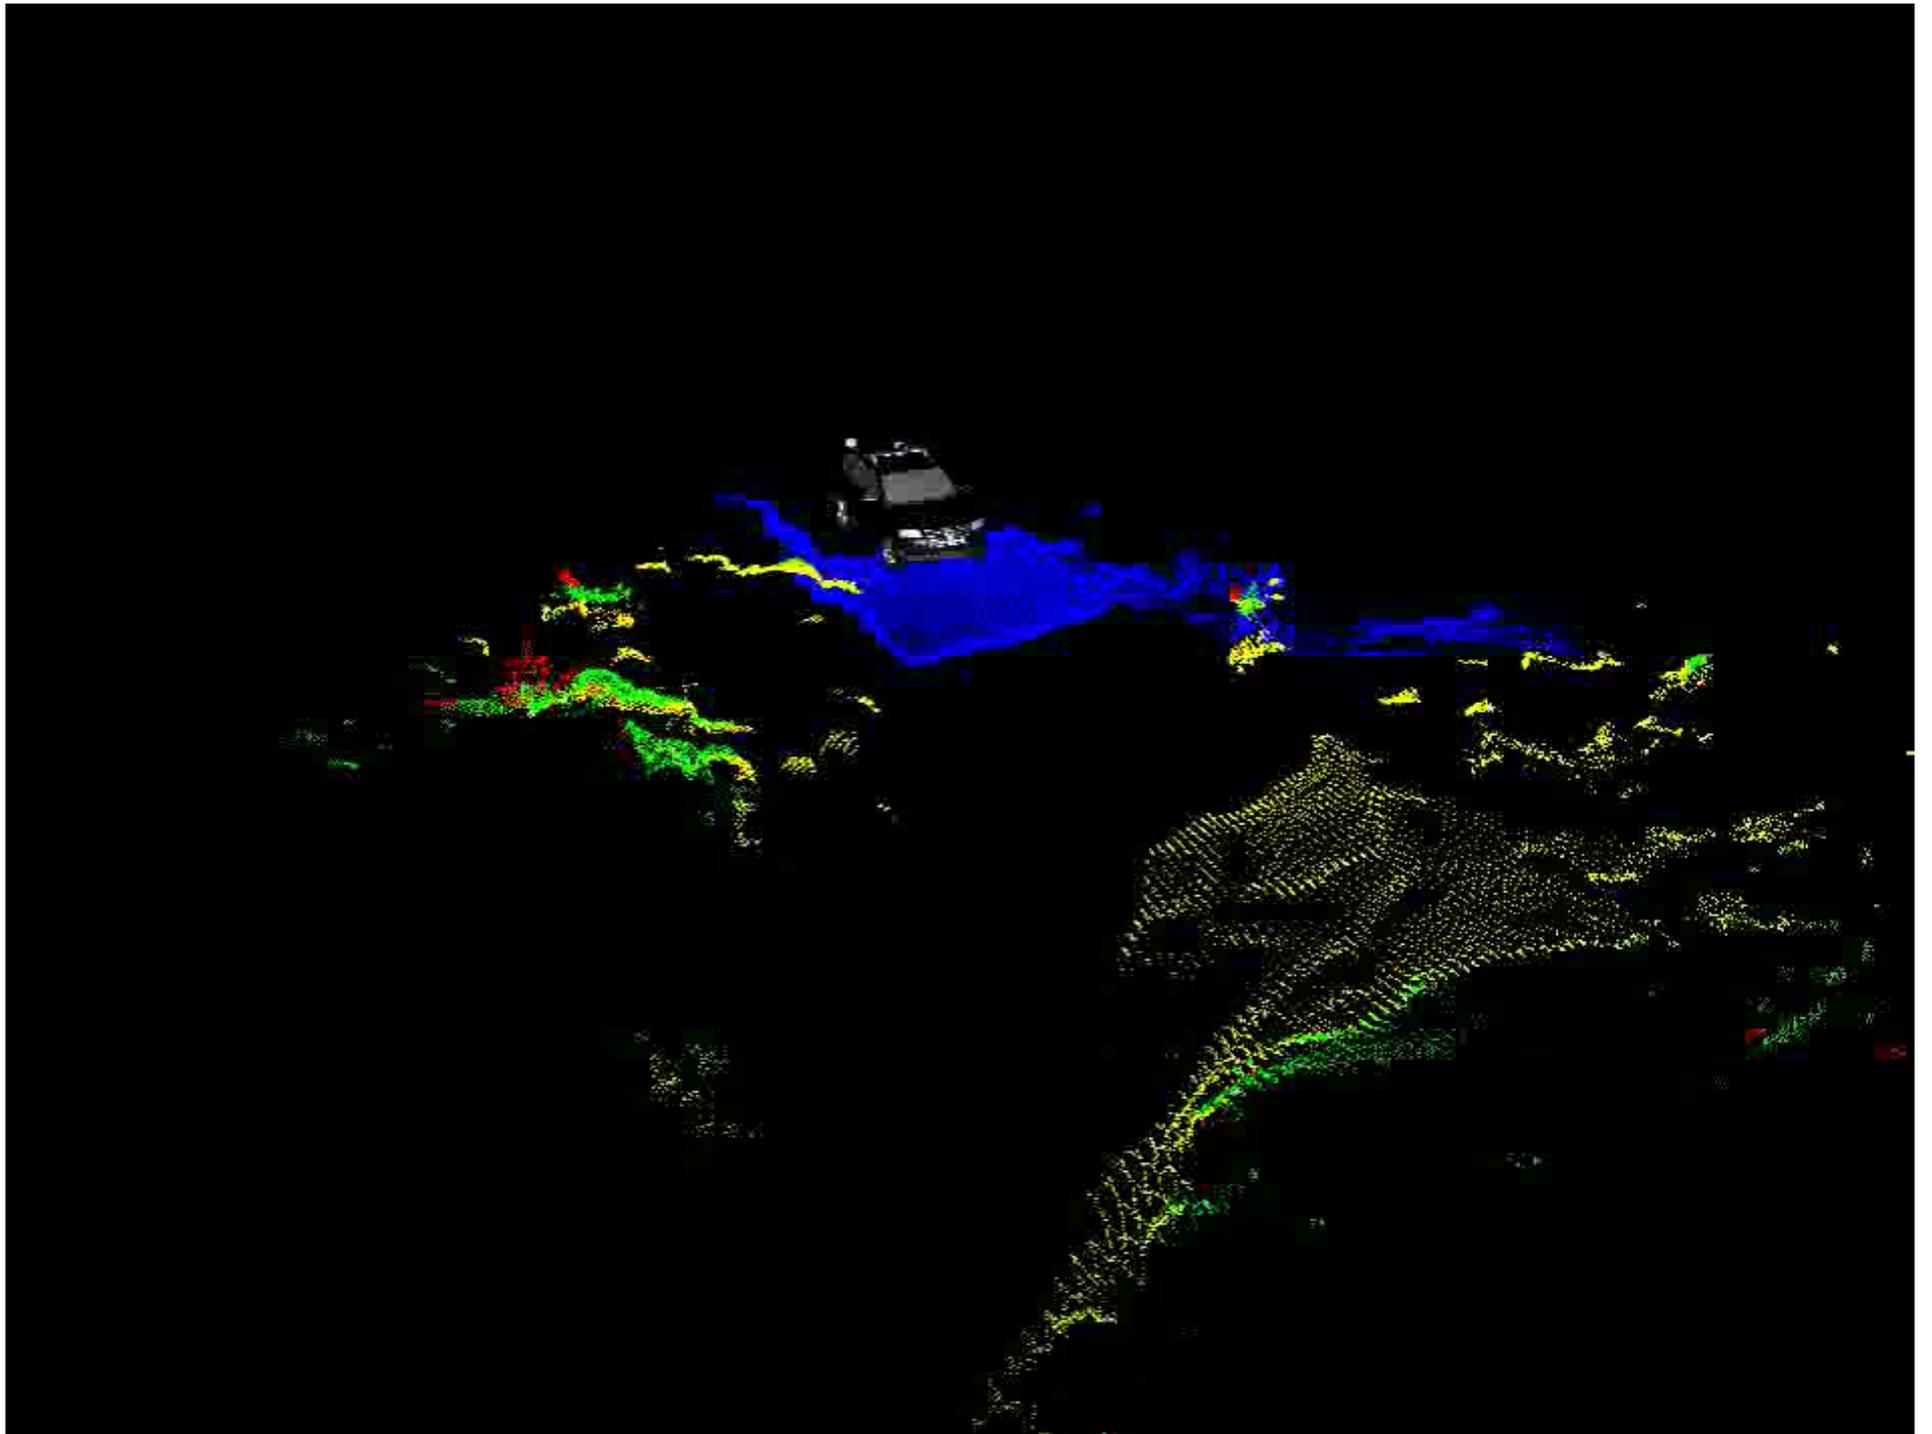
Obstacle  
Detection  
LADARs

Stereo  
Cameras

# Stanley's Technology



Images and movies taken from Sebastian Thrun's multimedia website.



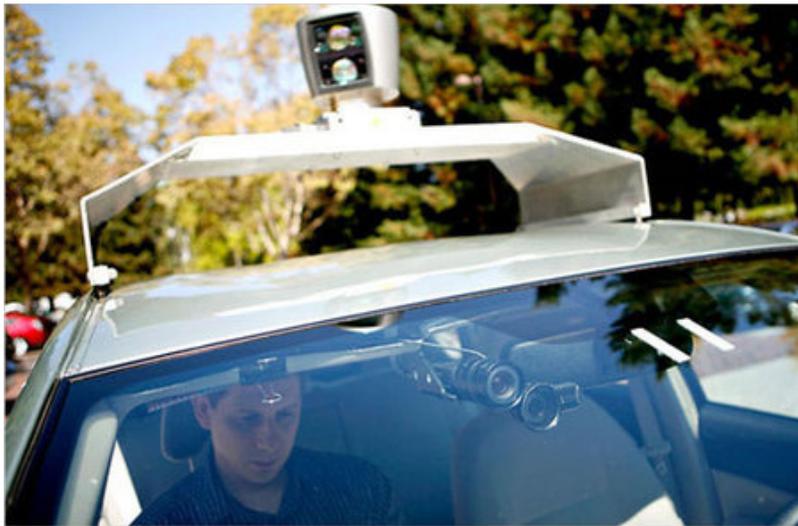
# 2007 DARPA Urban Challenge

- Took place on a closed Air Force base in California
- 11 teams made the finals, 6 vehicles finished:
  - #1: Carnegie Mellon University
  - #2: Stanford University
  - #3: Virginia Tech



- Vehicles navigated using predetermined waypoints while obeying all California driving laws, including:
  - Speed limits
  - Right-of-way
  - U-turns
  - Stop / Yield signs
  - Passing
  - Parking
- Other vehicles (autonomous cars and stunt cars) were present on the road, along with other obstacles (cones, barrels, utility poles)

# Google's Autonomous Car



- Nevada made it legal for autonomous cars to drive on roads in June 2011





Boston Dynamics

# 2011 Jeopardy!



- In February 2011, IBM Watson bested Brad Rutter (biggest all-time money winner) and Ken Jennings (longest winning streak)
- IBM is currently applying Watson's technology to medical diagnosis and legal research

**What are the next  
challenges for AI?**

# What's Next for AI?

- DARPA Machine Reading  
Build a system that read natural language texts and acquires that knowledge in a form suitable for answering questions.
- DARPA Deep Learning  
Learn layered structures that represent important aspects of the real world. Pushes unsupervised learning to be competitive with supervised learning.
- Poker: Many research universities are working on agents for poker.
  - AAI-07 in Vancouver held the first ever man vs. machine poker competition. The humans won 3:1 matches with 1 draw.
  - In the 2008 “rematch,” the humans won 2 matches, lost 3, and tied 2.