Introduction to Artificial Intelligence

David Pinto

Fundamentals…

- As we begin the new millenium
  - science and technology are changing rapidly
  - "old" sciences such as physics are relatively well-understood
  - computers are ubiquitous

- Grand Challenges in Science and Technology
  - understanding the brain
    - reasoning, cognition, creativity
  - creating intelligent machines
    - is this possible?
    - what are the technical and philosophical challenges?
  - arguably AI poses the most interesting challenges and questions in computer science today

Outline

- What is intelligence? What is artificial intelligence?
- A very brief history of AI
  - Modern successes: Stanley the driving robot
- An AI scorecard
  - How much progress has been made in different aspects of AI
- AI in practice
  - Successful applications
- The rational agent view of AI

What is Intelligence?

- Intelligence:
  - "the capacity to learn and solve problems" (Websters dictionary)
  - in particular,
    - *the ability to solve novel problems*
    - *the ability to act rationally*
    - *the ability to act like humans*

- Artificial Intelligence
  - build and understand intelligent entities or agents
  - 2 main approaches: "engineering" versus "cognitive modeling"

What's involved in Intelligence?

- Ability to interact with the real world
  - to perceive, understand, and act
  - e.g., speech recognition and understanding and synthesis
  - e.g., image understanding
  - e.g., ability to take actions, have an effect

- Reasoning and Planning
  - modeling the external world, given input
  - solving new problems, planning, and making decisions
  - ability to deal with unexpected problems, uncertainties

- Learning and Adaptation
  - we are continuously learning and adapting
  - our internal models are always being "updated"
    - e.g., a baby learning to categorize and recognize animals

Academic Disciplines relevant to AI

• Philosophy	Logic, methods of reasoning, mind as physical system, foundations of learning, language, rationality.
• Mathematics	Formal representation and proof, algorithms, computation, (un)decidability, (in)tractability
• Probability/Statistics	modeling uncertainty, learning from data
• Economics	utility, decision theory, rational economic agents
• Neuroscience	neurons as information processing units.
• Psychology/ Cognitive Science	how do people behave, perceive, process cognitive information, represent knowledge.
• Computer engineering	building fast computers
• Control theory	design systems that maximize an objective function over time
• Linguistics	knowledge representation, grammars

History of AI

- 1943: early beginnings
  - McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing
  - Turing's "Computing Machinery and Intelligence"
- 1956: birth of AI
  - Dartmouth meeting: "Artificial Intelligence" name adopted
- 1950s: initial promise
  - Early AI programs, including
  - Samuel's checkers program
  - Newell & Simon's Logic Theorist
- 1955-65: "great enthusiasm"
  - Newell and Simon: GPS, general problem solver
  - Gelertner: Geometry Theorem Prover
  - McCarthy: invention of LISP

History of AI

- 1966—73: Reality dawns
  - Realization that many AI problems are intractable
  - Limitations of existing neural network methods identified
    - Neural network research almost disappears

- 1969—85: Adding domain knowledge
  - Development of knowledge-based systems
  - Success of rule-based expert systems,
    - E.g., DENDRAL, MYCIN
    - But were brittle and did not scale well in practice

- 1986-- Rise of machine learning
  - Neural networks return to popularity
  - Major advances in machine learning algorithms and applications

- 1990-- Role of uncertainty
  - Bayesian networks as a knowledge representation framework

- 1995-- AI as Science
  - Integration of learning, reasoning, knowledge representation
  - AI methods used in vision, language, data mining, etc

Success Stories

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- AI program proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans
- Robot driving: DARPA grand challenge 2003-2007
- 2006: face recognition software available in consumer cameras

Example: DARPA Grand Challenge

- Grand Challenge
  - Cash prizes ($1 to $2 million) offered to first robots to complete a long course completely unassisted
  - Stimulates research in vision, robotics, planning, machine learning, reasoning, etc

- 2004 Grand Challenge:
  - 150 mile route in Nevada desert
  - Furthest any robot went was about 7 miles
  - … but hardest terrain was at the beginning of the course

- 2005 Grand Challenge:
  - 132 mile race
  - Narrow tunnels, winding mountain passes, etc
  - Stanford 1st, CMU 2nd, both finished in about 6 hours

- 2007 Urban Grand Challenge
  - November in Victorville, California

Stanley Robot

Stanford Racing Team www.stanfordracing.org

Next few slides courtesy of Prof. Sebastian Thrun, Stanford University

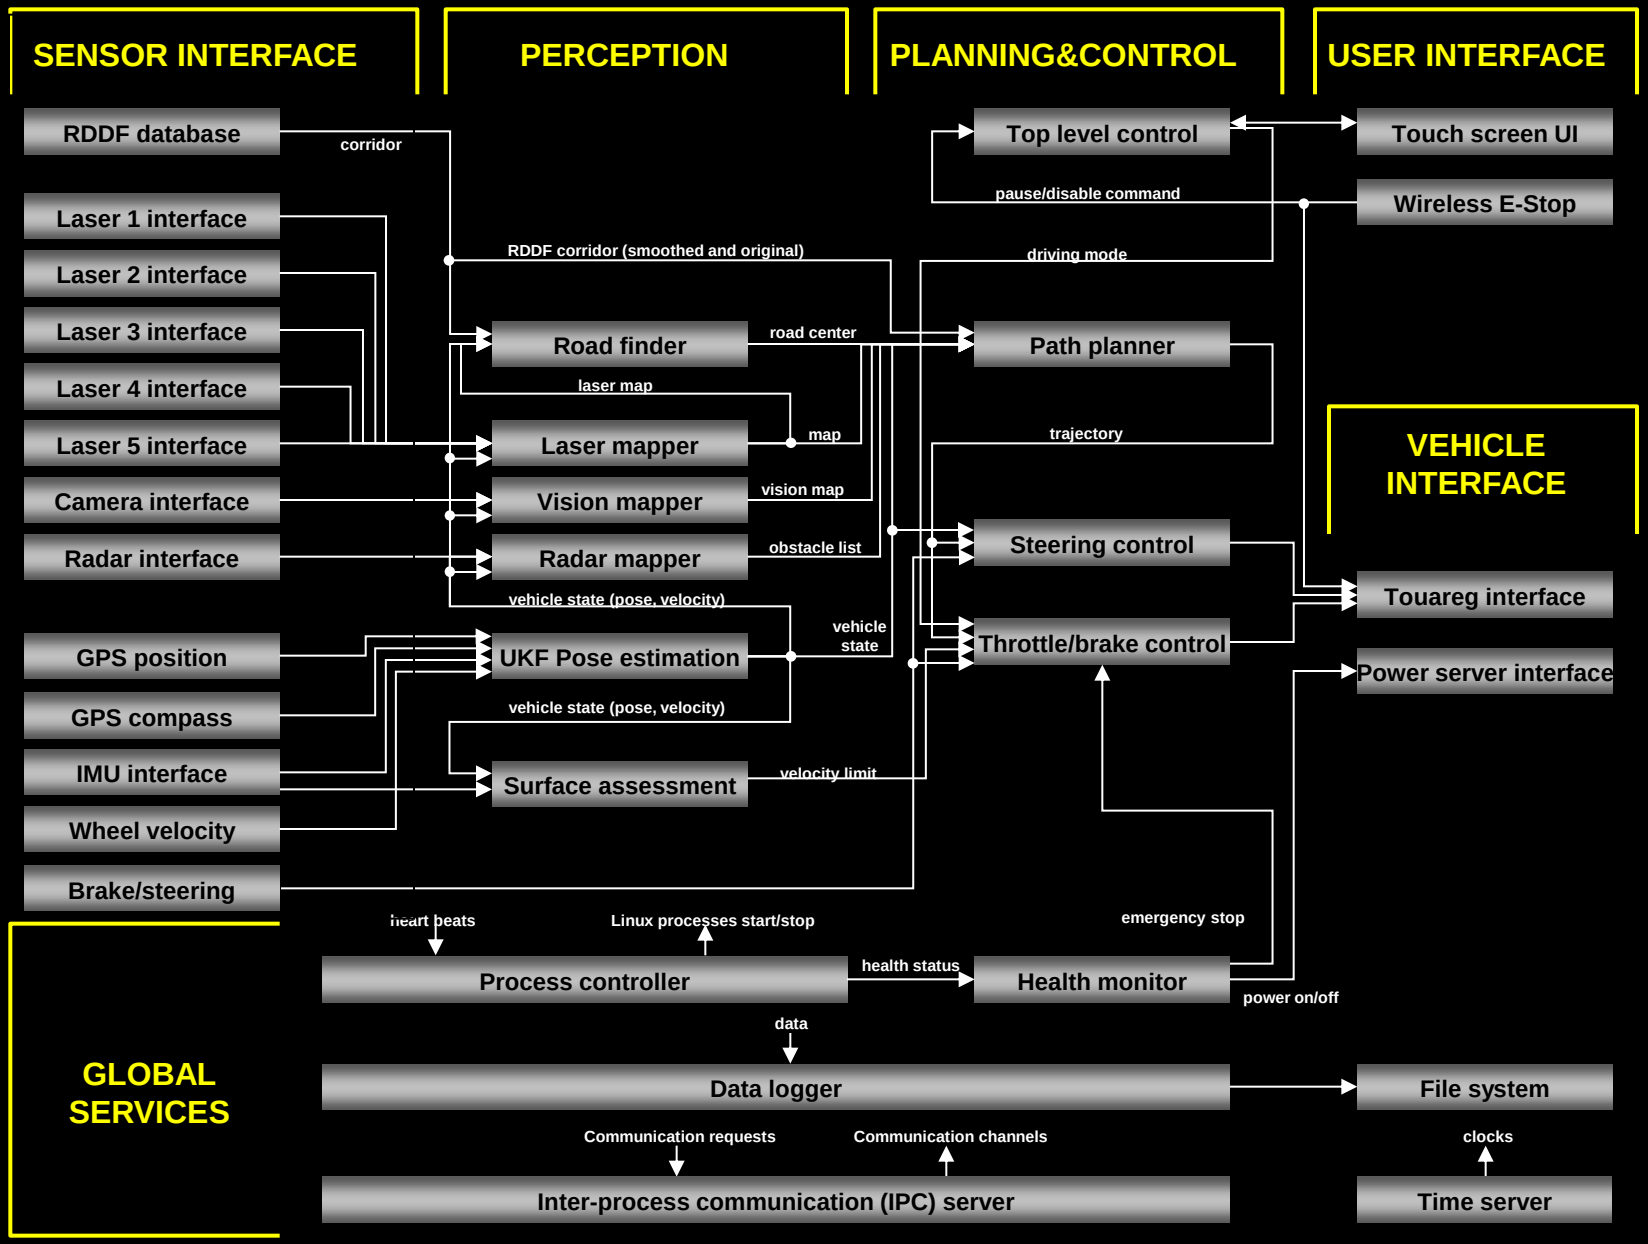
SENSOR INTERFACE
PERCEPTION
PLANNING&CONTROL
USER INTERFACE
RDDF database
corridor
Laser 1 interface
Laser 2 interface
Laser 3 interface
Laser 4 interface
Laser 5 interface
Camera interface
Radar interface
GPS position
GPS compass
IMU interface
Wheel velocity
Brake/steering
RDDF corridor (smoothed and original)
Road finder
road center
laser map
Laser mapper
map
Vision mapper
vision map
Radar mapper
obstacle list
vehicle state (pose, velocity)
UKF Pose estimation
vehicle state
vehicle state (pose, velocity)
Surface assessment
velocity limit
Top level control
pause/disable command
driving mode
Path planner
trajectory
Steering control
Throttle/brake control
Touch screen UI
Wireless E-Stop
VEHICLE INTERFACE
Touareg interface
Power server interface
heart beats
Linux processes start/stop
emergency stop
Process controller
health status
Health monitor
power on/off
data
GLOBAL SERVICES
Data logger
File system
Communication requests
Communication channels
clocks
Inter-process communication (IPC) server
Time server

Planning = Rolling out Trajectories

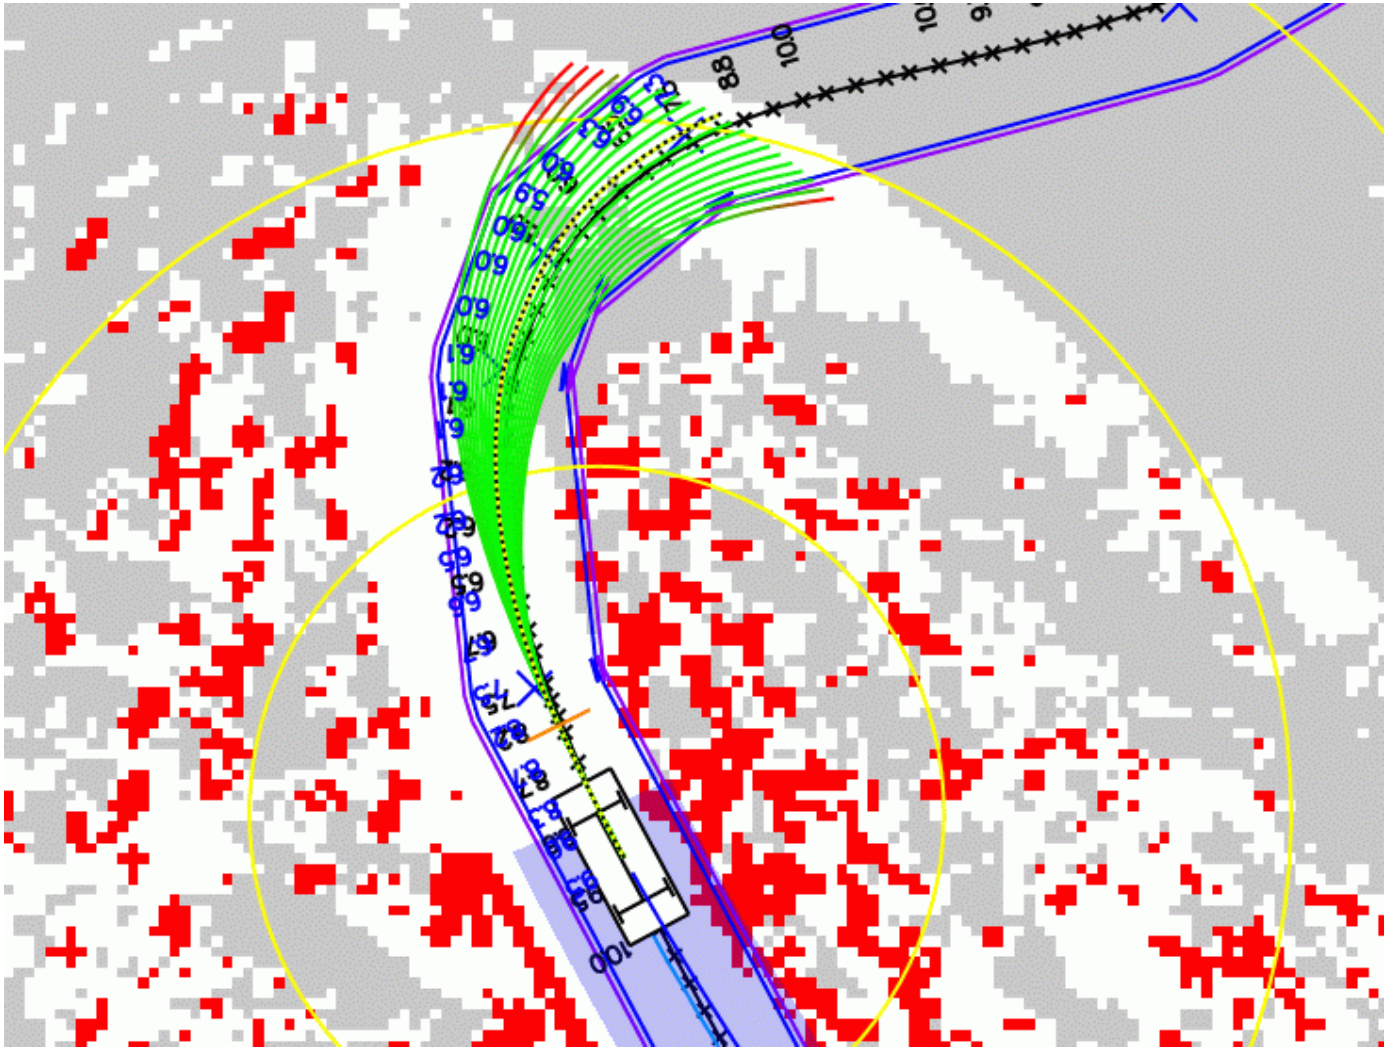

2004: Barstow, CA, to Primm, NV

- 150 mile off-road robot race across the Mojave desert
- Natural and manmade hazards
- No driver, no remote control
- No dynamic passing
- Fastest vehicle wins the race (and 2 million dollar prize)

2005 Semi-Finalists: 43 Teams

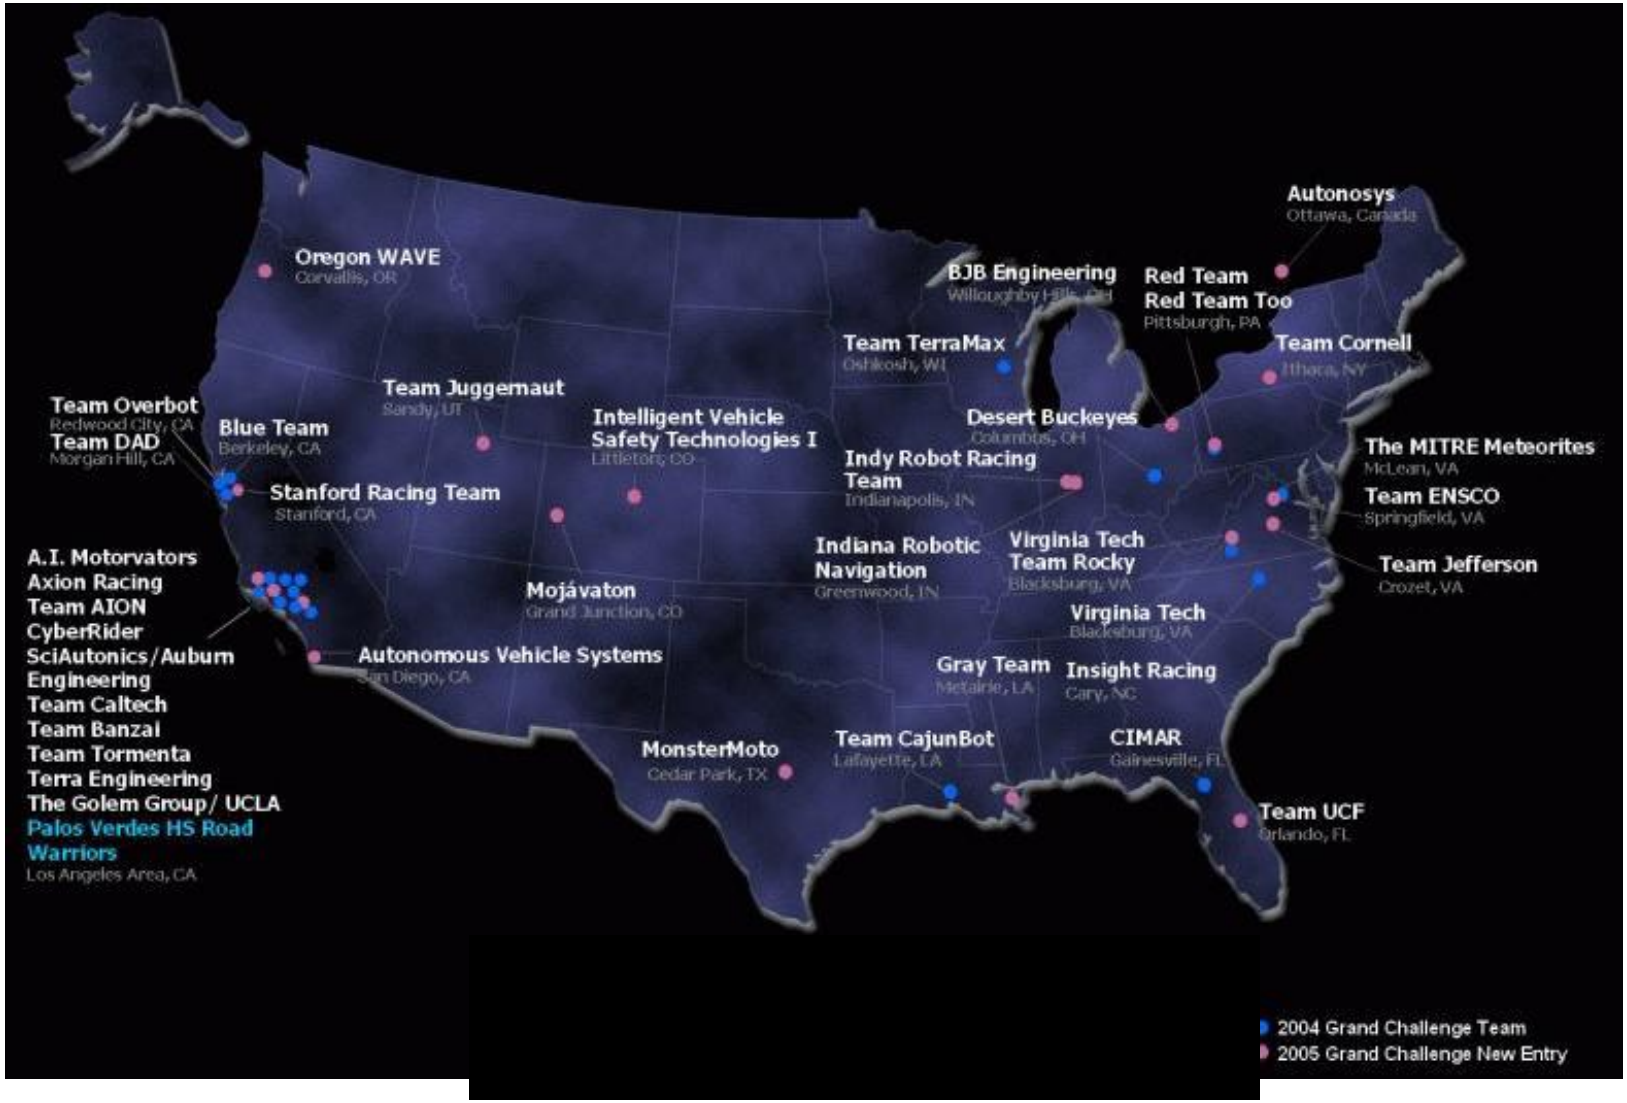

The Grand Challenge Race

HAL: from the movie 2001

- *2001: A Space Odyssey*
  - classic science fiction movie from 1969

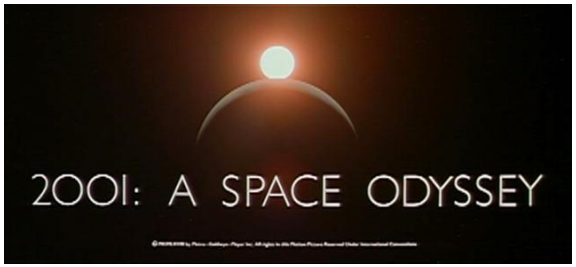

- HAL
  - part of the story centers around an intelligent computer called HAL
  - HAL is the "brains" of an intelligent spaceship
  - in the movie, HAL can
    - speak easily with the crew
    - see and understand the emotions of the crew
    - navigate the ship automatically
    - diagnose on-board problems
    - make life-and-death decisions
    - display emotions

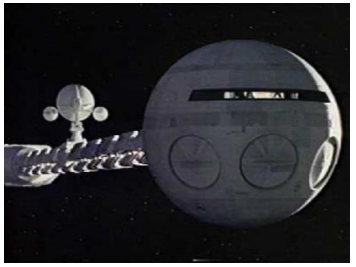

- In 1969 this was science fiction: is it still science fiction?

Hal and AI

- *HAL's Legacy: 2001's Computer as Dream and Reality*
  - MIT Press, 1997, David Stork (ed.)
  - discusses
    - HAL as an intelligent computer
    - are the predictions for HAL realizable with AI today?

- Materials online at
  - http://mitpress.mit.edu/e-books/Hal/contents.html

- The website contains
  - full text and abstracts of chapters from the book
  - links to related material and AI information
  - sound and images from the film

Consider what might be involved in building a computer like Hal….

- What are the components that might be useful?
  - Fast hardware?
  - Chess-playing at grandmaster level?
  - Speech interaction?
    - speech synthesis
    - speech recognition
    - speech understanding
  - Image recognition and understanding ?
  - Learning?
  - Planning and decision-making?

Can we build hardware as complex as the brain?

- How complicated is our brain?
  - a neuron, or nerve cell, is the basic information processing unit
  - estimated to be on the order of 10^{12} neurons in a human brain
  - many more synapses (10^{14}) connecting these neurons
  - cycle time: 10^{-3} seconds (1 millisecond)

- How complex can we make computers?
  - 10^8 or more transistors per CPU
  - supercomputer: hundreds of CPUs, 10^{12} bits of RAM
  - cycle times: order of 10^{-9} seconds

- Conclusion
  - YES: in the near future we can have computers with as many basic processing elements as our brain, but with
    - far fewer interconnections (wires or synapses) than the brain
    - much faster updates than the brain
  - but building hardware is very different from making a computer behave like a brain!

Can Computers beat Humans at Chess?

- Chess Playing is a classic AI problem
  - well-defined problem
  - very complex: difficult for humans to play well

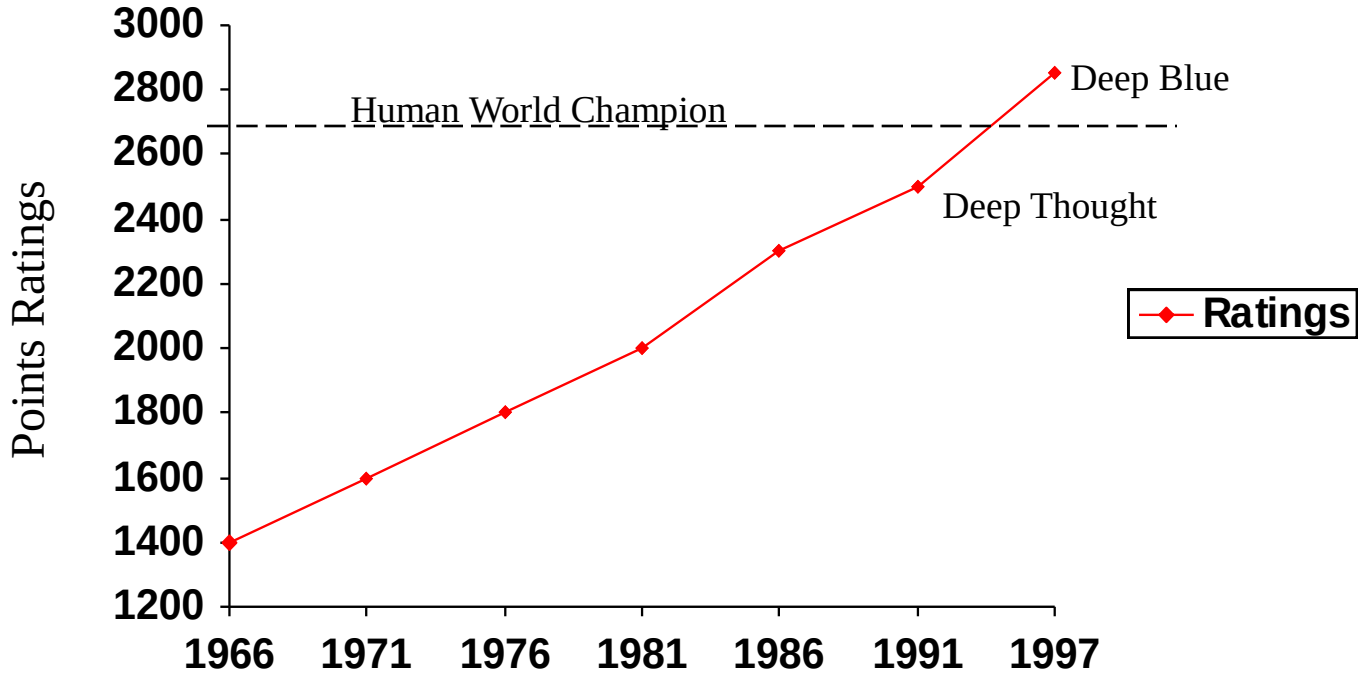

- Conclusion:
  - YES: today's computers can beat even the best human

Can Computers Talk?

- This is known as "speech synthesis"
  - translate text to phonetic form
    - e.g., "fictitious" -> fik-tish-es
  - use pronunciation rules to map phonemes to actual sound
    - e.g., "tish" -> sequence of basic audio sounds

- Difficulties
  - sounds made by this "lookup" approach sound unnatural
  - sounds are not independent
    - e.g., "act" and "action"
    - modern systems (e.g., at AT&T) can handle this pretty well
  - a harder problem is emphasis, emotion, etc
    - humans understand what they are saying
    - machines don't: so they sound unnatural

- Conclusion:
  - NO, for complete sentences
  - YES, for individual words

Can Computers Recognize Speech?

- Speech Recognition:
  - mapping sounds from a microphone into a list of words
  - classic problem in AI, very difficult
    - "Lets talk about how to wreck a nice beach"
    - (I really said "_____")
- Recognizing single words from a small vocabulary
    - systems can do this with high accuracy (order of 99%)
    - e.g., directory inquiries
      - limited vocabulary (area codes, city names)
      - computer tries to recognize you first, if unsuccessful hands you over to a human operator
      - saves millions of dollars a year for the phone companies

Recognizing human speech (ctd.)

- Recognizing normal speech is much more difficult
  - speech is continuous: where are the boundaries between words?
    - e.g., "John's car has a flat tire"
  - large vocabularies
    - can be many thousands of possible words
    - we can use **context** to help figure out what someone said
      - e.g., hypothesize and test
      - try telling a waiter in a restaurant:
        "I would like some dream and sugar in my coffee"
  - background noise, other speakers, accents, colds, etc
  - on normal speech, modern systems are only about 60-70% accurate

- Conclusion:
  - NO, normal speech is too complex to accurately recognize
  - YES, for restricted problems (small vocabulary, single speaker)

Can Computers Understand speech?

- Understanding is different to recognition:
  - "Time flies like an arrow"
    - assume the computer can recognize all the words
    - how many different interpretations are there?

Can Computers Understand speech?

- Understanding is different to recognition:
  - "Time flies like an arrow"
    - assume the computer can recognize all the words
    - how many different interpretations are there?
      - 1. time passes quickly like an arrow?
      - 2. command: time the flies the way an arrow times the flies
      - 3. command: only time those flies which are like an arrow
      - 4. "time-flies" are fond of arrows

Can Computers Understand speech?

- Understanding is different to recognition:
  - "Time flies like an arrow"
    - assume the computer can recognize all the words
    - how many different interpretations are there?
      - 1. time passes quickly like an arrow?
      - 2. command: time the flies the way an arrow times the flies
      - 3. command: only time those flies which are like an arrow
      - 4. "time-flies" are fond of arrows
    - only 1. makes any sense,
      - but how could a computer figure this out?
      - clearly humans use a lot of implicit commonsense knowledge in communication

- Conclusion: NO, much of what we say is beyond the capabilities of a computer to understand at present

Can Computers Learn and Adapt ?

- Learning and Adaptation
  - consider a computer learning to drive on the freeway
  - we could teach it lots of rules about what to do
  - or we could let it drive and steer it back on course when it heads for the embankment
    - systems like this are under development (e.g., Daimler Benz)
    - e.g., RALPH at CMU
      - in mid 90's it drove 98% of the way from Pittsburgh to San Diego without any human assistance
  - **machine learning** allows computers to learn to do things without explicit programming
  - many successful applications:
    - requires some "set-up": does not mean your PC can learn to forecast the stock market or become a brain surgeon

- Conclusion: YES, computers can learn and adapt, when presented with information in the appropriate way

Can Computers "see"?

- Recognition v. Understanding (like Speech)
  - Recognition and Understanding of Objects in a scene
    - look around this room
    - you can effortlessly recognize objects
    - human brain can map 2d visual image to 3d "map"

- Why is visual recognition a hard problem?

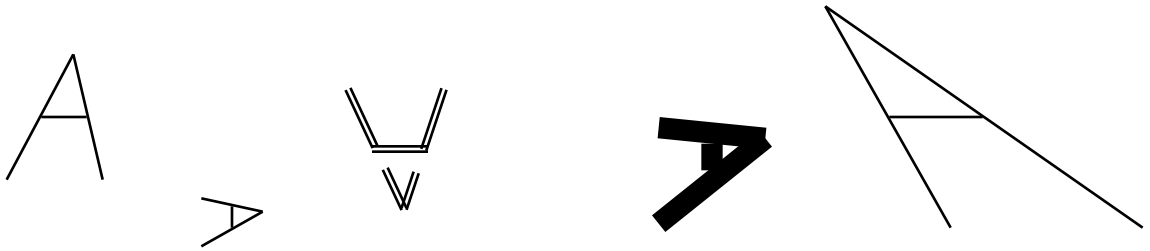

- Conclusion:
  - mostly NO: computers can only "see" certain types of objects under limited circumstances
  - YES for certain constrained problems (e.g., face recognition)

Can computers plan and make optimal decisions?

- Intelligence
  - involves solving problems and making decisions and plans
  - e.g., you want to take a holiday in Brazil
    - you need to decide on dates, flights
    - you need to get to the airport, etc
    - involves a sequence of decisions, plans, and actions

- What makes planning hard?
  - the world is not predictable:
    - your flight is canceled or there's a backup on the 405
  - there are a potentially huge number of details
    - do you consider all flights? all dates?
      - no: commonsense constrains your solutions
  - AI systems are only successful in constrained planning problems

- Conclusion: NO, real-world planning and decision-making is still beyond the capabilities of modern computers
  - exception: very well-defined, constrained problems

Summary of State of AI Systems in Practice

- Speech synthesis, recognition and understanding
  - very useful for limited vocabulary applications
  - unconstrained speech understanding is still too hard

- Computer vision
  - works for constrained problems (hand-written zip-codes)
  - understanding real-world, natural scenes is still too hard

- Learning
  - adaptive systems are used in many applications: have their limits

- Planning and Reasoning
  - only works for constrained problems: e.g., chess
  - real-world is too complex for general systems

- Overall:
  - many components of intelligent systems are “doable”
  - there are many interesting research problems remaining

Intelligent Systems in Your Everyday Life

- Post Office
  - automatic address recognition and sorting of mail
- Banks
  - automatic check readers, signature verification systems
  - automated loan application classification
- Customer Service
  - automatic voice recognition
- The Web
  - Identifying your age, gender, location, from your Web surfing
  - Automated fraud detection
- Digital Cameras
  - Automated face detection and focusing
- Computer Games
  - Intelligent characters/agents

AI Applications: Machine Translation

- Language problems in international business
  - e.g., at a meeting of Japanese, Korean, Vietnamese and Swedish investors, no common language
  - or: you are shipping your software manuals to 127 countries
  - solution; hire translators to translate
  - would be much cheaper if a machine could do this

- How hard is automated translation
  - very difficult! e.g., English to Russian
    - "The spirit is willing but the flesh is weak" (English)
    - "the vodka is good but the meat is rotten" (Russian)
  - not only must the words be translated, but their meaning also!
  - is this problem "AI-complete"?

- Nonetheless....
  - commercial systems can do a lot of the work very well (e.g.,restricted vocabularies in software documentation)
  - algorithms which combine dictionaries, grammar models, etc.
  - Recent progress using "black-box" machine learning techniques

AI and Web Search

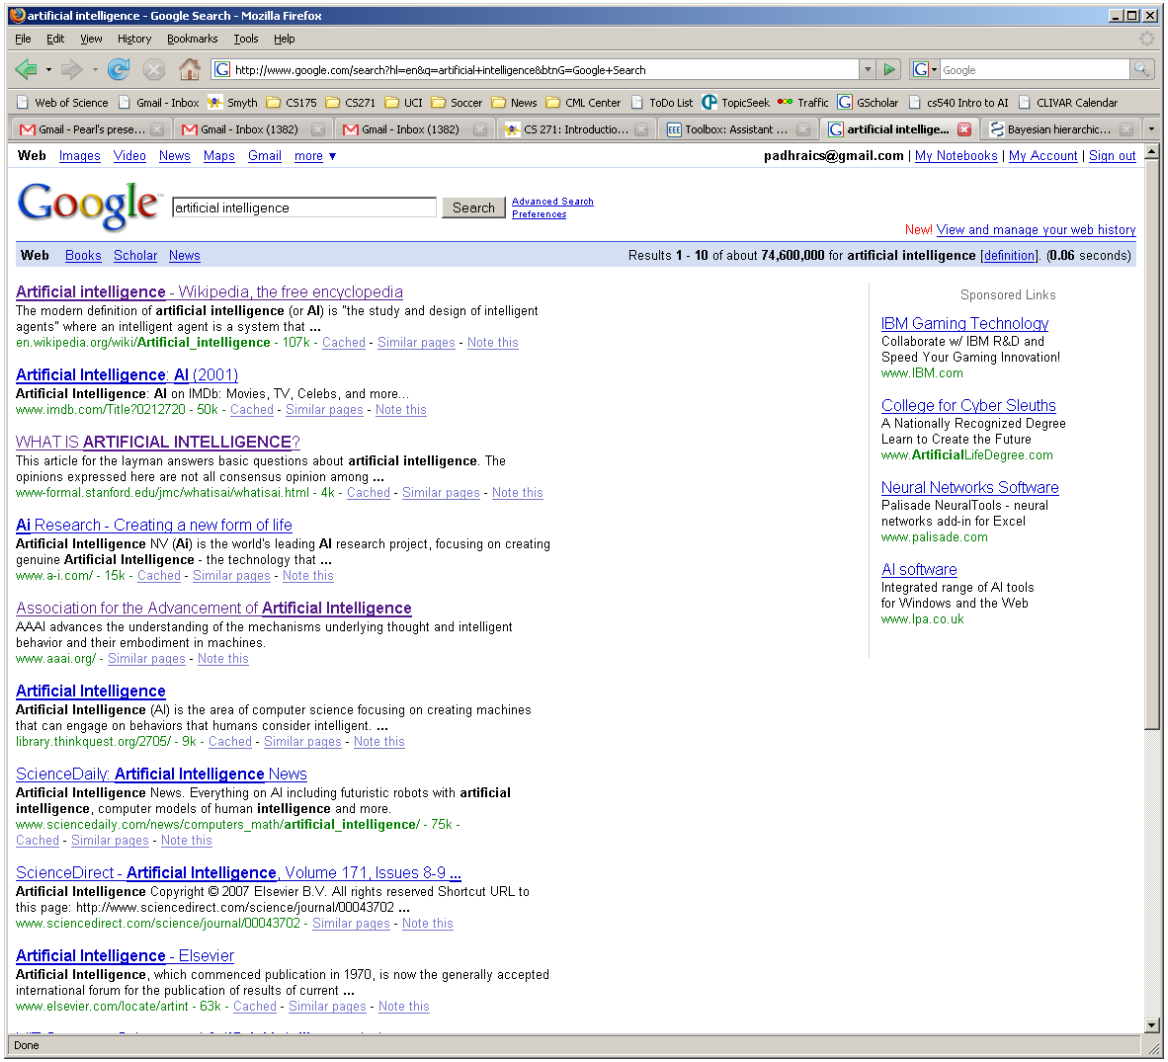

What's involved in Intelligence? (again)

- Perceiving, recognizing, understanding the real world
- Reasoning and planning about the external world
- Learning and adaptation

- So what general principles should we use to achieve these goals?

Different Types of Artificial Intelligence

1. Modeling exactly how humans actually think
2. Modeling exactly how humans actually act
3. Modeling how ideal agents "should think"
4. Modeling how ideal agents "should act"

- Modern AI focuses on the last definition
  - we will also focus on this "engineering" approach
  - success is judged by how well the agent performs

Acting humanly: Turing test

- Turing (1950) "Computing machinery and intelligence"
- "Can machines think?" → "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game

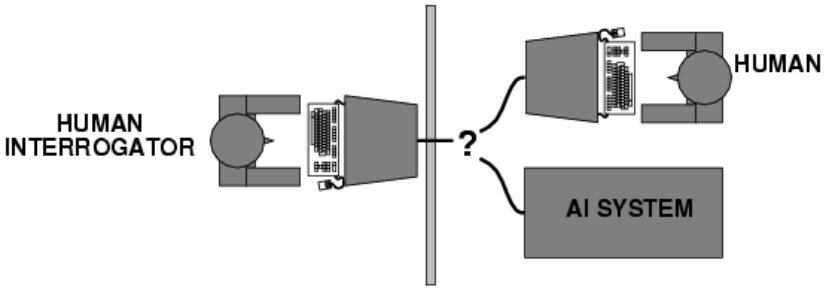

- Suggests major components required for AI:
  - knowledge representation
  - reasoning,
  - language/image understanding,
  - learning

* Question: is it important that an intelligent system act like a human?

Thinking humanly

- Cognitive Science approach
  - Try to get "inside" our minds
  - E.g., conduct experiments with people to try to "reverse-engineer" how we reason, learning, remember, predict

- Problems
  - Humans don't behave rationally
    - e.g., insurance
  - The reverse engineering is very hard to do
  - The brain's hardware is very different to a computer program

Thinking rationally

- Represent facts about the world via logic
- Use logical inference as a basis for reasoning about these facts
- Can be a very useful approach to AI
  - E.g., theorem-provers
- Limitations
  - Does not account for an agent's uncertainty about the world
    - E.g., difficult to couple to vision or speech systems
  - Has no way to represent goals, costs, etc (important aspects of real-world environments)

Acting rationally

- Decision theory/Economics
  - Set of future states of the world
  - Set of possible actions an agent can take
  - Utility = gain to an agent for each action/state pair

  - An agent acts rationally if it selects the action that maximizes its "utility"
    - Or expected utility if there is uncertainty

- Emphasis is on autonomous agents that behave rationally (make the best predictions, take the best actions)
  - on average over time
  - within computational limitations ("bounded rationality")

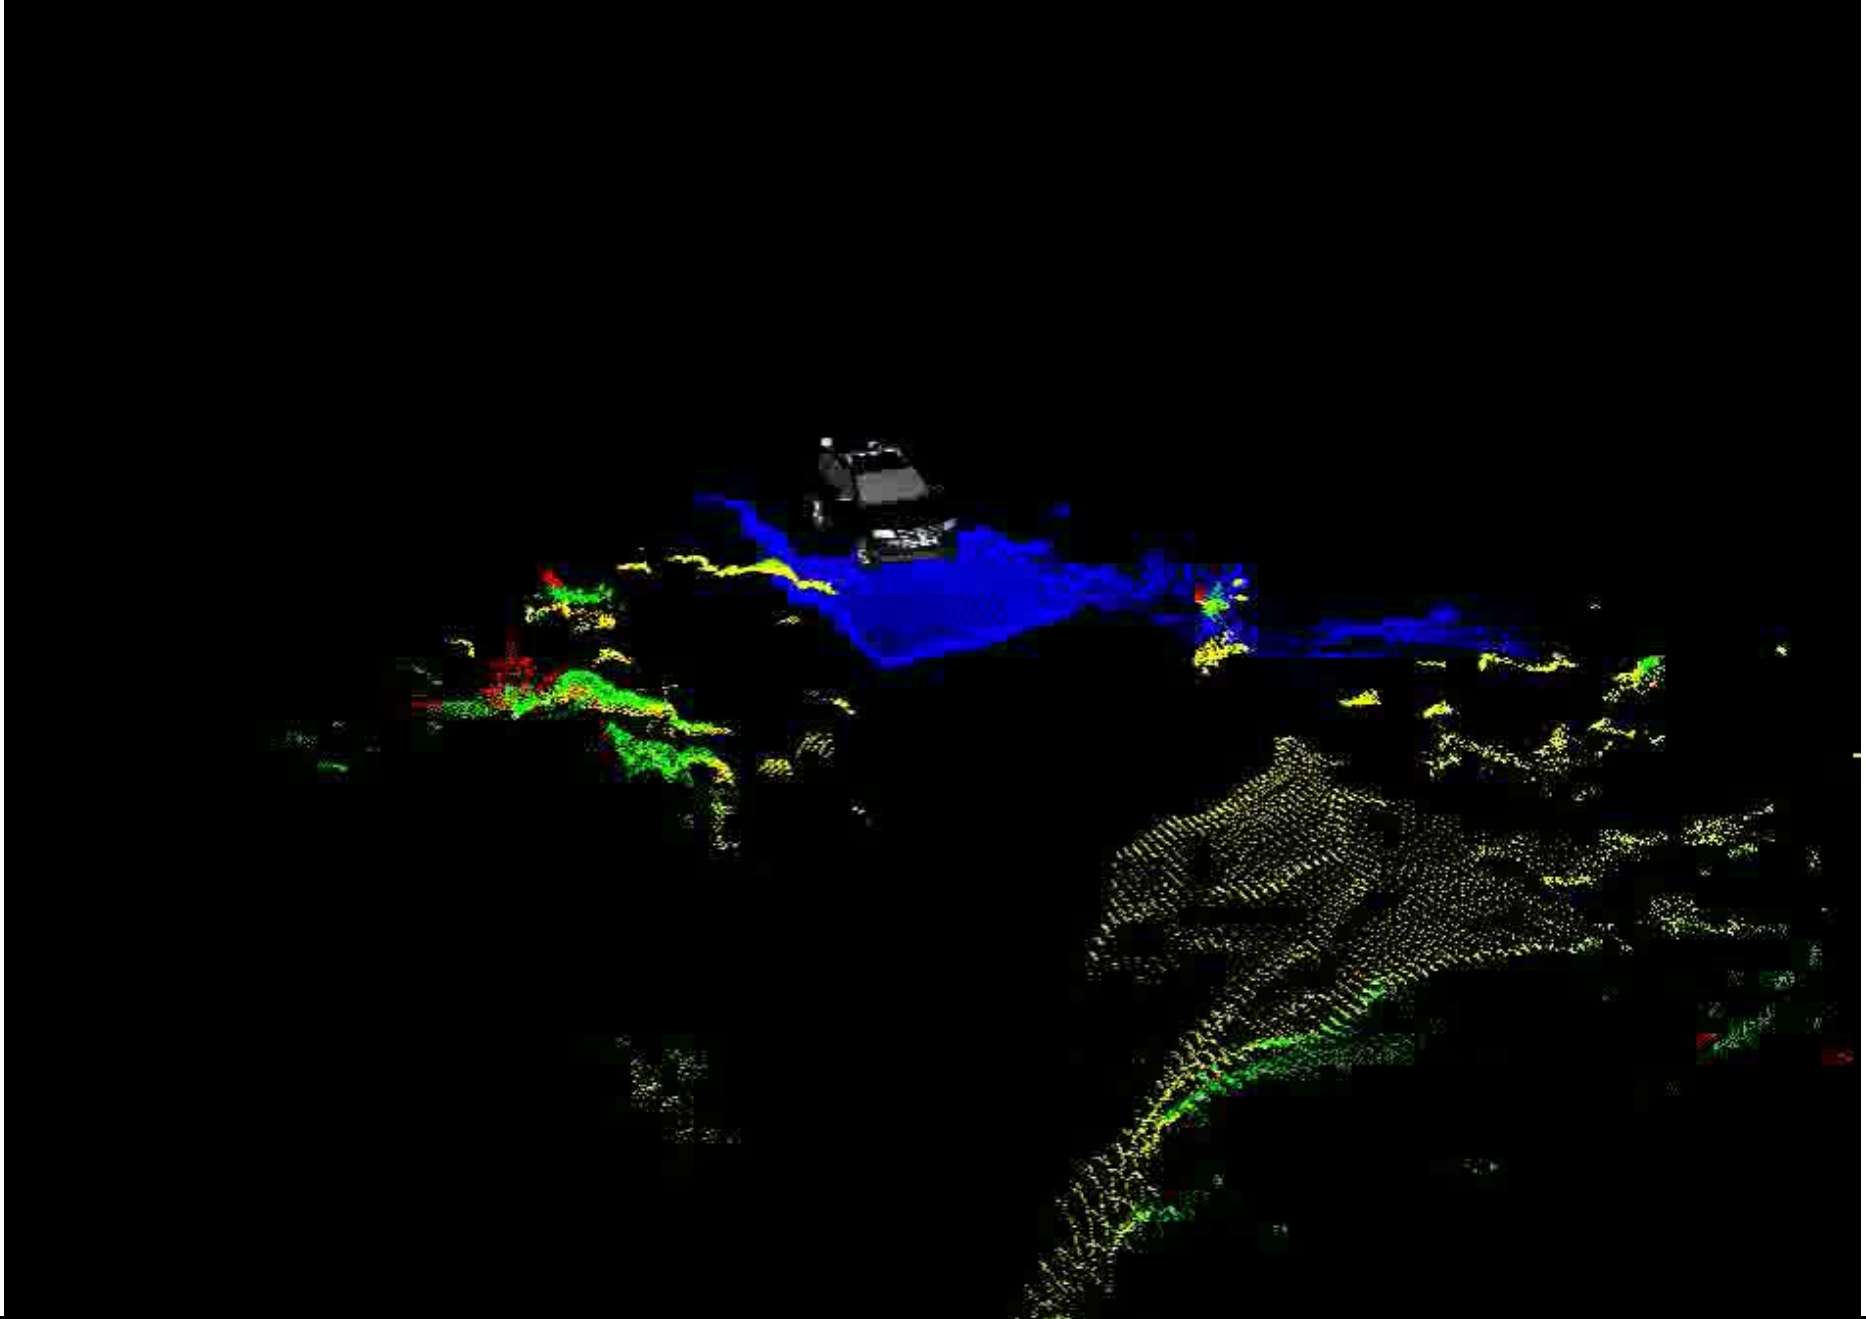

Summary

- Artificial Intelligence involves the study of:
  - automated recognition and understanding of signals
  - reasoning, planning, and decision-making
  - learning and adaptation

- AI has made substantial progress in
  - recognition and learning
  - some planning and reasoning problems
  - …but many open research problems

- AI Applications
  - improvements in hardware and algorithms => AI applications in industry, finance, medicine, and science.

- Rational agent view of AI

- Reading…