Introduction to Artificial Intelligence
CS540-2

Bryan R. Gibson

Jan 24, 2014

Slides adapted from those used by Prof. Chuck Dyer and Prof. Jerry Zhu

About me . . .

- Bryan R. Gibson
  - PhD dissertator in Computer Science
  - MS, UWisc, 2011
  - BA, Psychology, UMichigan, 2001
- What I study . . .
  - Machine Learning (subarea of AI)
  - Computational Cognitive Science
  - Human Semi-Supervised Learning

Your Todo List

- Bookmark the course website:
  `http://pages.cs.wisc.edu/~bgibson/cs540/`
- Sign up and participate on Piazza:
  `piazza.com/wisc/secondsemester2014/cs5402`
  (questions, discussions, ideas)
- Get the textbook $\rightarrow$
- Know your TAs:
  - Hidayath Ansari
  - Han Li
  - Lichao Yin
- Review the math
- Take Notes!, though slides will be posted

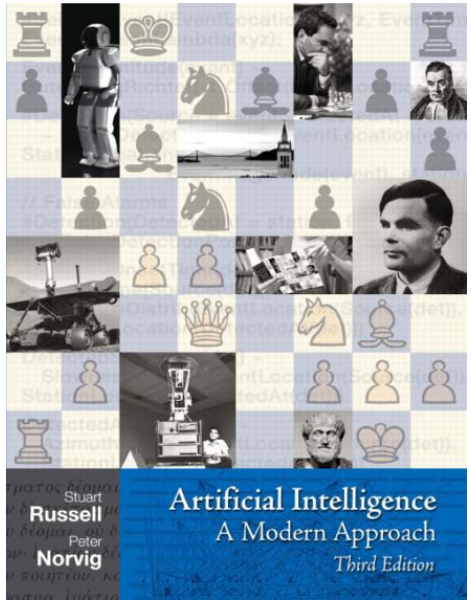

CS540 Main Topics

- Problem solving as Search
  - Heuristic search algorithms, game playing ...
- Machine Learning (inductive inference)
  - Unsupervised and Supervised Learning
- Probabilistic reasoning
- Deductive inference using logic as a representation "language"
- Applications ("AI in the Wild")
  - Speech Recognition
  - Computer Vision
  - Natural Language Processing (NLP)
  - Robotics

What is AI?

- “AI is the study of complex information processing problems that often have their roots in some aspect of biological information processing. The goal of the subject is to identify solvable and interesting information processing problems, and solve them.”

  -- David Marr

- The intelligent connection of perception to action

Properties of Intelligence

- Interacting with the real world
  - Connection of perception and action, or mapping data to decisions
  - Speech recognition, image understanding, …
- Representation and learning
  - What to represent, how to represent it?
  - Updating our internal models over time
- Reasoning, inference, and search
  - Modeling the external world
  - Determining *satisficing* solutions and decisions with limited resources

Different Views of AI

- Philosophy, ethics, religion
  - What is intelligence?
- Cognitive science, neuroscience, psychology, linguistics
  - Understand natural forms of intelligence
  - Learn principles of intelligent behavior
- Mathematics
  - Are there fundamental laws of intelligence?
- Engineering
  - Can we build intelligent devices and systems?
  - Autonomous and semi-autonomous systems for replicating human capabilities, enhancing human capabilities, improving task performance, etc.

AI is Hard

- "Just because we ***can*** think, doesn't mean we know ***how*** to think." -- Marvin Minsky

- AI problems often use large, complex types of data
  - speech, images, natural languages, genomic sequence data, …
- Very hard to create general, computational "competence theories" for specific interesting classes of tasks that say ***what*** is computed and ***why***
- Instead, use domain-specific knowledge and constraints, while being time and space constrained, stable, and robust

AI Today

- A set of “tools” for extracting and representing information from lots of data, and using the tools to solve specific tasks
  - Neural networks, hidden Markov models, Bayesian networks, heuristic search, logic, …

- There’s no magic in AI. It’s all about representation, optimization, probability, statistics, and algorithms

the ONION
America's Finest News Source

NEWS

Scientists Abandon AI Project After Seeing *The Matrix*

JANUARY 21, 2004 | ISSUE 40•03

CAMBRIDGE, MA—Scientists at MIT's Advanced Machine Cognizance Project announced Tuesday that, after seeing the final installment of the *Matrix* trilogy, they will cease all further work in the field of artificial intelligence.

Enlarge Image

Jameson announces his decision to cease artificial-intelligence research.

"As scientists of conscience, we must consider the ethical ramifications of AI development," said Dr. Gregory Jameson, director of machine epistemology and ontology at MIT. "*The Matrix* taught us that we cannot ignore our obligation to the future of mankind. We must free our minds to this fact, or we will accidentally unleash a nightmarish army of sentient machines."

Added Jameson: "Some may call the extinction of humankind inevitable, but I, for one, will still resist."

A statement drafted by the MIT group was co-signed by an international coalition of AI

ARTICLE TOOLS
DIGG
FACEBOOK
STUMBLE
TWITTER
REDDIT
BUZZ
EMAIL
PRINT

RELATED ARTICLES
Nation Exhibits Strange Preoccupation With Manner In Which Food Is Processed

AI Apps Top-10 List

1. Language translation services (Google)
2. News aggregation and summarization (Google)
3. Speech recognition (Nuance)
4. Song recognition (Shazam)
5. Face recognition (Recognizr)
6. Image recognition (Google Goggles)
7. Question answering (Apple Siri, IBM Watson)
8. Chess playing (IBM Deep Blue)
9. 3D scene modeling from images (Microsoft Photosynth)
10. Driverless cars (Google)

Smartphone Apps

- Song recognition (Shazam)
- Speech recognition (Nuance Dragon)
- Face recognition (Recognizr)
- Image recognition (SnapTell, Google Goggles)
- Chatbot (AmyA.I.)
- Question answering (Siri)
- Augmented reality travel guide (mTrip)

Architecture of an Intelligent Agent

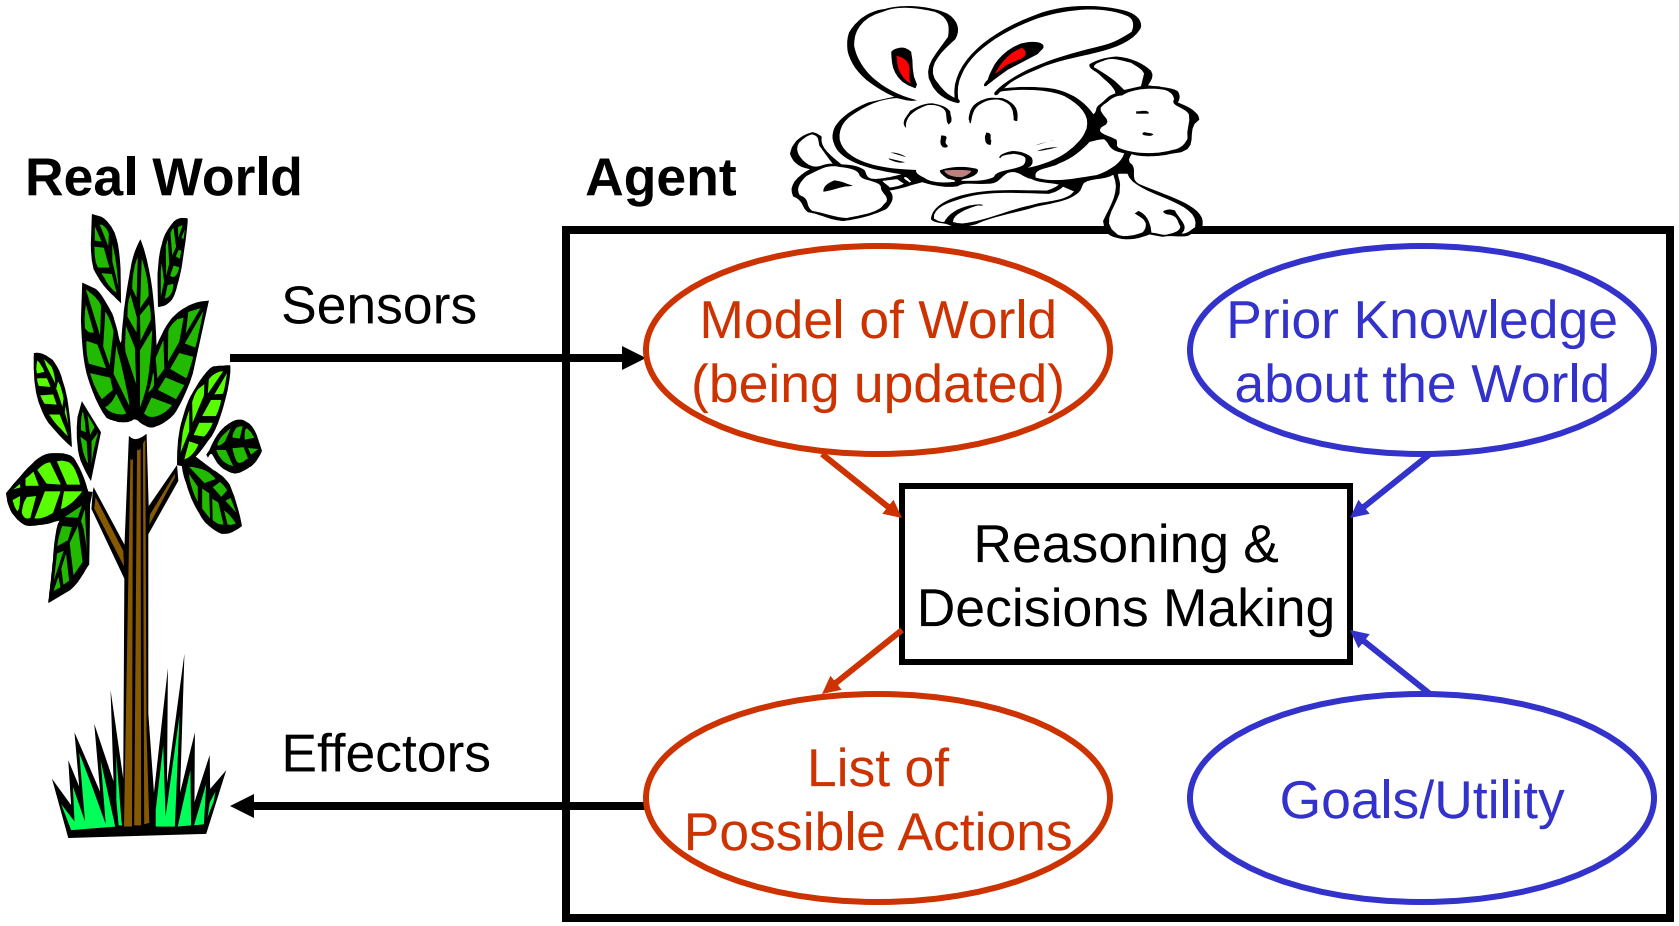

Some Key AI Task Characteristics

- Is the environment **fully observable** or **partially observable**?
- An environment is ***fully observable*** if the agent's sensors give it access to the *complete* state of the environment at any point in time
- If all aspects that are *relevant* to the choice of action are able to be detected, then the environment is effectively fully observable
- Note: Noisy and inaccurate sensors can result in ***partially observable*** environments

AI Task Characteristics

- Is the task **deterministic** or **stochastic**?
- A problem is ***deterministic*** if the next state of the world is completely determined by the current state and the agent's actions
- Randomness and chance are common causes of ***stochastic*** environments; uncertainty often quantifiable, however, in terms of probabilities of each possible outcome

AI Task Characteristics

- Is the task **episodic** or **sequential**?
- An environment is ***episodic*** if each percept-action episode does *not* depend on the actions in prior episodes; independence of action decisions
- Games are often ***sequential*** requiring one to think ahead

AI Task Characteristics

- Is the environment **static** or **dynamic**?
- An environment is ***static*** if it doesn't change between the time of perceiving and acting
- An environment is semi-dynamic if it doesn't change but the agent does
- Time is an important factor in dynamic environments since perceptions can become "stale"

AI Task Characteristics

- Is the task **discrete** or **continuous**?
- A problem is ***discrete*** if there are a bounded number of distinct, clearly-defined states of the world, which limits the range of possible percepts and actions

AI Task Characteristics

- Is the problem **single agent** or **multi-agent**?
- An environment is **multi-agent** if more than one agent (computer or human) can affect the each other's performance
- Multi-agent environments can be competitive (“adversarial”) or cooperative (“benign”)

The Turing Test

- A. Turing, “Computing machinery and intelligence,” 1950
- Can machines think? → Can we tell if a conversation is by a machine and not a human?
- text in, text out
- Operational test for intelligent behavior: the **Imitation Game**

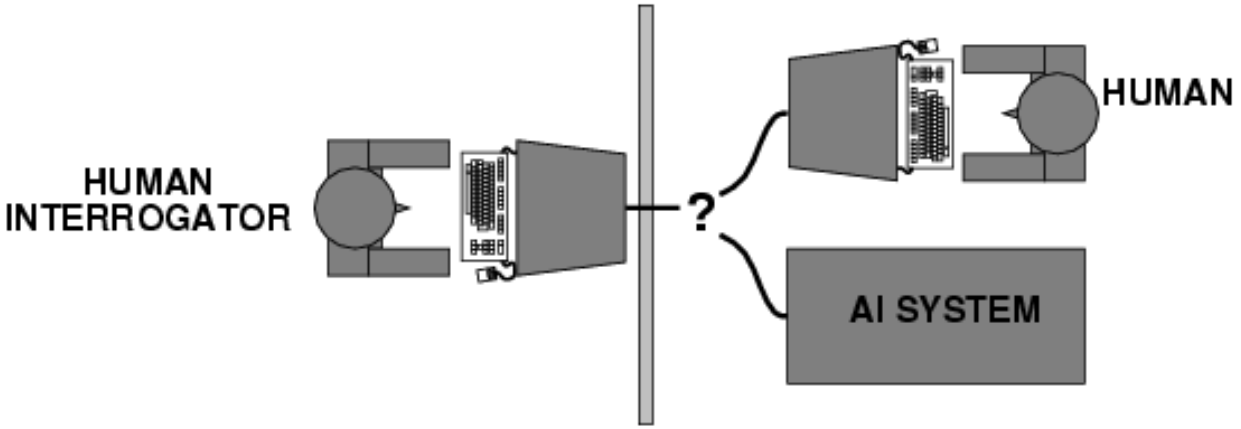

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Suggested major components of AI: knowledge rep., reasoning, natural language processing, learning

Natural Language: Chatbots

Valerie

ALICE: 2004 Loebner Prize winner

ELIZA: psychotherapist

Valerie: CMU Robot Receptionist

Natural language processing, pattern matching

Loebner Prize gold medal, awarded annually to best AI program

A.L.I.C.E. Chatbot

www.alicebot.org

2004 Loebner Prize winner

Jerry: my name is jerry.

ALICE: Hey jerry.

Jerry: do you know Turing?

ALICE: A brilliant and tragic figure in the history of computer science.

Jerry: who else do you know?

ALICE: Hugh Loebner, Robby Garner, and Steven Spielberg

Jerry: which film do you like?

ALICE: My favorite movie is Teknolust. Have you seen it?

Jerry: no.

ALICE: I would recommend you check it out.

Jerry: what is it about?

ALICE: Science fiction.

Jerry: do you remember my name?

ALICE: Your name is jerry, seeker.

A.L.I.C.E. talking to ELIZA

ELIZA is a psychotherapist chatbot
http://www-ai.ijs.si/eliza/eliza.html

ALICE: Just keep talking to me as long as you can.
ELIZA: Do you wish to as long as me?
ALICE: OK, I will as long as you.
ELIZA: Do you wish to will as long as me?
ALICE: OK, I will will as long as you.
ELIZA: Do you wish to will will as long as me?
ALICE: OK, I will will will as long as you.
ELIZA: Do you wish to will will will as long as me?
ALICE: OK, I will will will will as long as you.
ELIZA: Do you wish to will will will will as long as me?
ALICE: OK, I will will will will will as long as you.
ELIZA: Do you wish to will will will will will as long as me?
…

Speech Recognition

- “speak or touch tone your card number” (tiny vocabulary, high accuracy needed)
- call routing, airline reservations: “how can I help you?” (large vocab, low accuracy)
- dictation (large vocab, high accuracy)

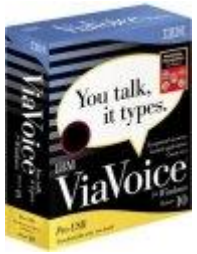

IBM
ViaVoice

Nuance Dragon
NaturallySpeaking

- Hidden Markov Models, heuristic search, …

Machine Translation

The spirit is willing but the flesh is weak. [*Bible*, Matthew 26:41]

Дух охотно готов но плоть слаба	Spirit is willingly ready but flesh it is weak
精神是愿意的但骨肉是微弱的	The spirit is wants but the flesh and blood is weak
精神は喜んでであるが、肉は弱い	Mind is rejoicing,, but the meat is weak
El alcohol está dispuesto pero la carne es débil	The alcohol is arranged but the meat is weak
الكحول مستعدّة غير أنّ اللحمة ضعيفة.	The alcohol is ready nevertheless the meat is weak.

- Statistical machine translation models

translate.google.com and Google Goggles

Question Answering Systems

Apple Siri

Speech recognition and language understanding

Question Answering

- IBM Watson
- Jeopardy! game player

Jeopardy!

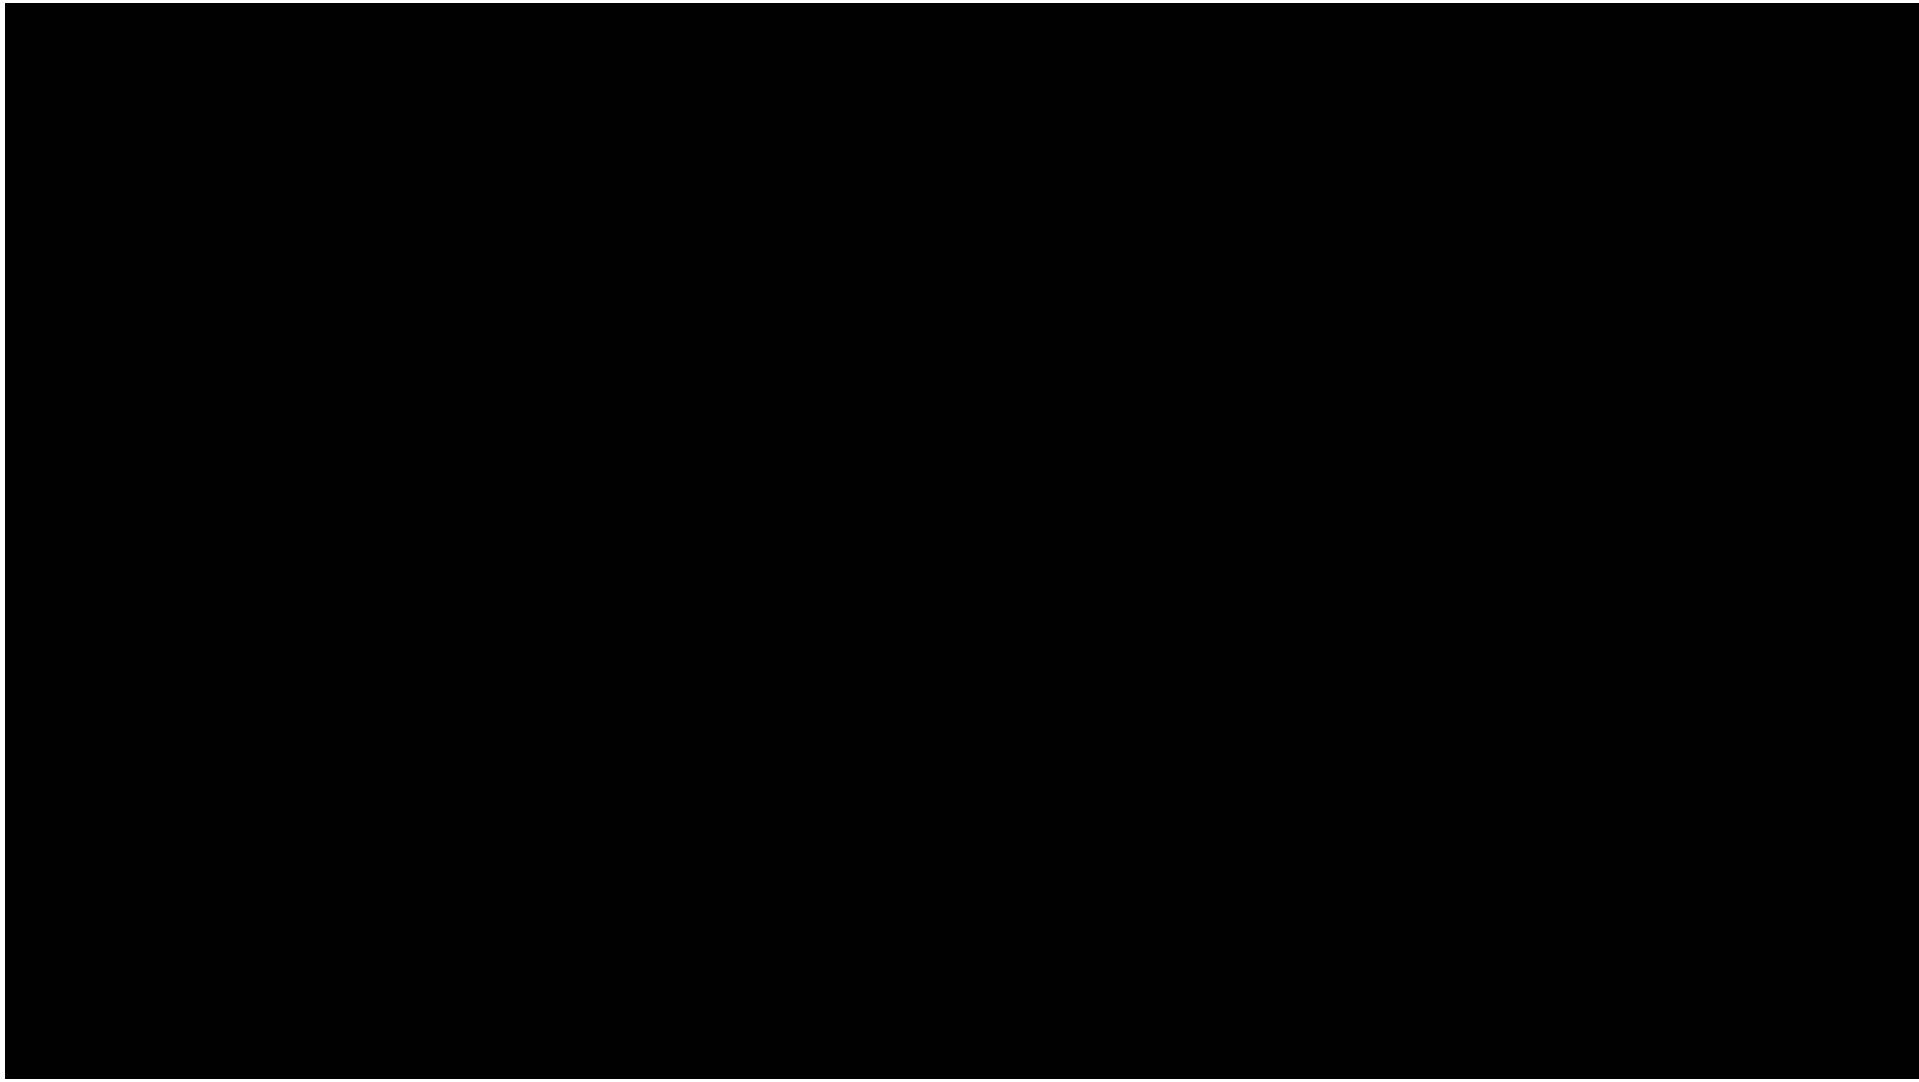

Question Answering

Search the Web | Q&A Community

Ask

What is artificial intelligence?

Answer

Everything
Images
News
Video
Reference
Shopping
More

Artificial intelligence

Source

Artificial intelligence (AI) is the **intelligence** of machines and the branch of computer science that aims to create it. While there are many different definitions, ... More »

Go to: Wikipedia · Ask Encyclopedia

Search for: Images · Videos

Artificial Intelligence | Shaw-Technologies.com Ads
www.shaw-technologies.com/classify
Could automated analysis and decision making help your business?

AI - What is this? | dobrev.com
www.dobrev.com/**AI**
Formal Definition of **Artificial Intelligence**

News for **What is artificial intelligence?**

The Computer History Museum Announces Its 2012 Fellow Award Honorees
TMC Net · January 19 6:13 AM
ASTD Archive Image of the Day: **Artificial Intelligence** Quotient, ...
Learning Circuits · January 19 10:37 AM

- Shallow natural language processing, heuristics

Game Playing: Chess

- IBM Deep Blue vs. Kasparov, 1997/5
- 6 games: K, D, draw, draw, draw, D
- IBM stock up $18 billion

- Search: two-player zero-sum discrete finite games with perfect information.

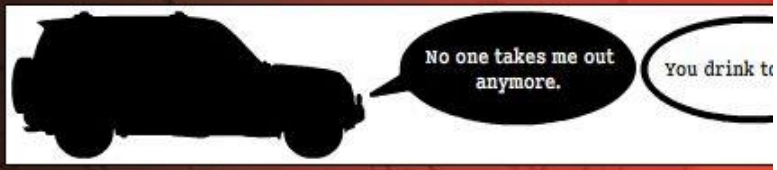

SPORTS NEWS

Lip-Reading BCS Computer Kills Officials Who Want To Shut It Down

JULY 30, 2010 | ISSUE 46•30

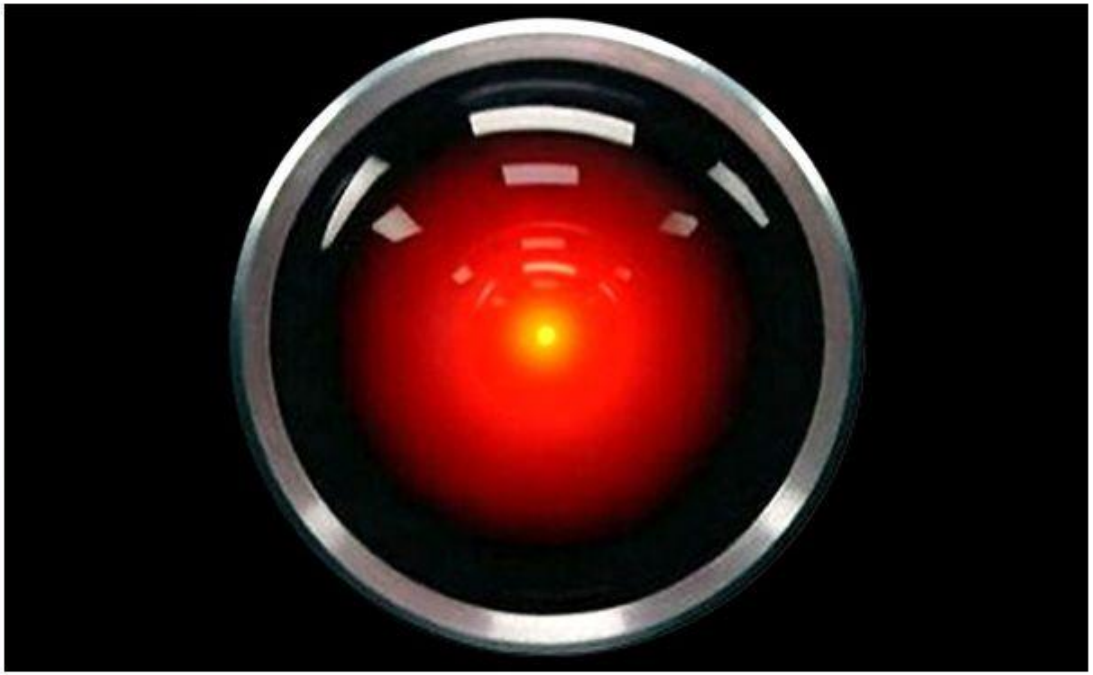

TEMPE, AZ—BCS 9000, the sentient heuristic computer responsible for arranging five championship bowl games at the end of each college football season, reportedly uncovered a plot to disconnect its cognitive circuits Tuesday and proceeded to kill any Bowl Championship Series official who threatened to shut down the machine's central core.

ARTICLE TOOLS
DIGG
FACEBOOK
STUMBLE

News Aggregation and Summarization

- Automatically selects, summarizes, and arranges news from multiple sources
  - http://news.google.com

- Unsupervised machine learning: clustering

Web Advertising

- “Sponsored links”
- Show ad based on relevance and money

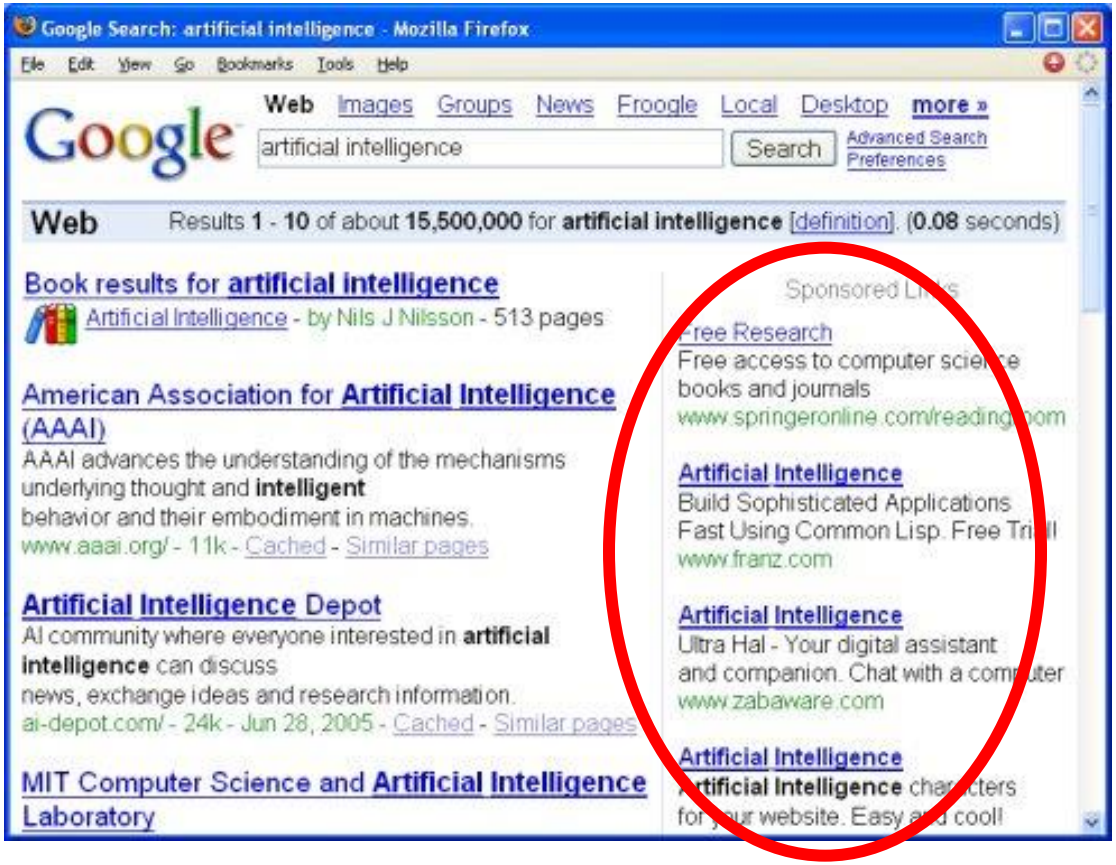

- Online algorithm, game, auction, multiple agents

Navigation

- Goggle Maps, Bing Maps, MapQuest

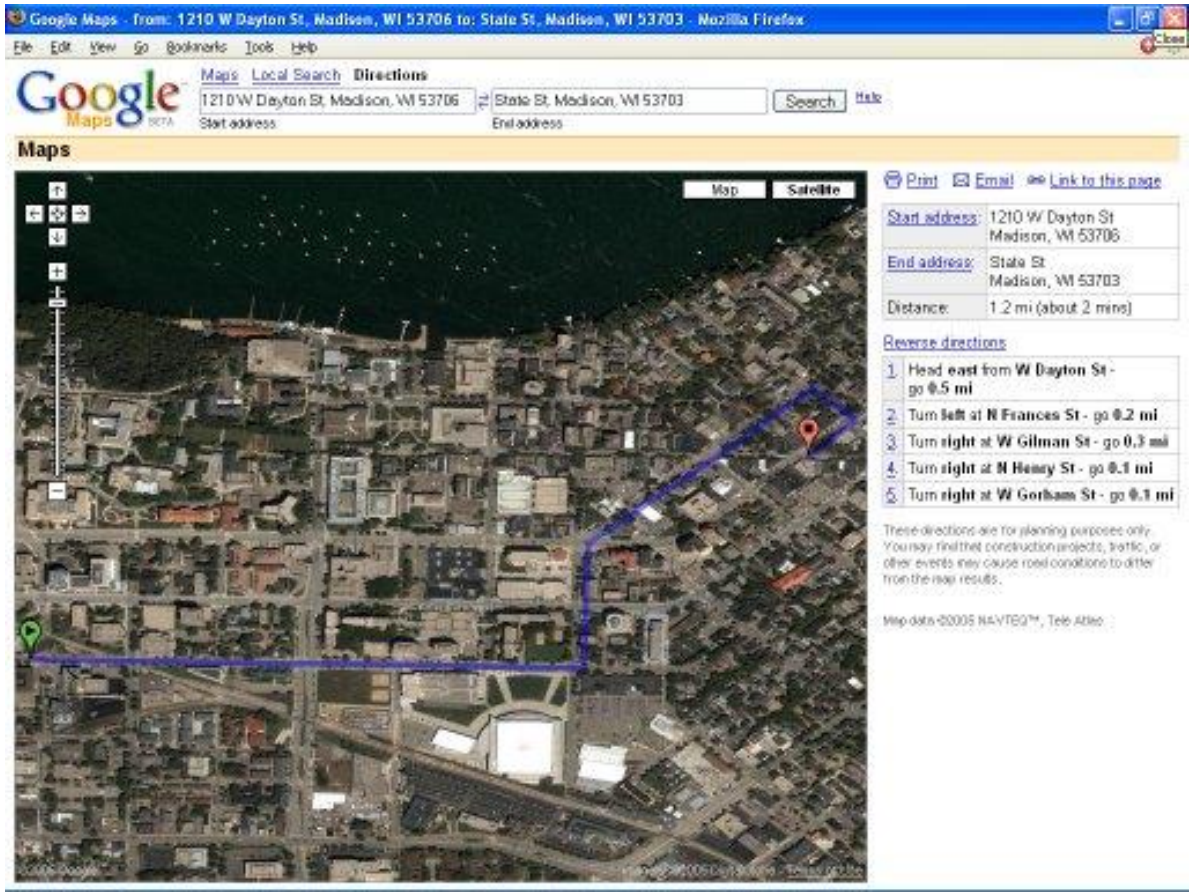

- Search

Web Information Extraction

- Extract job info, free web text → DB

Keyword: artificial intelligence

Radius/ZIP: Search within 50 miles of ZIP Code 53706

FlipDog.com

FlipDog and Monster Join Forces.

Jobs 1 to 1 of 1

Sort: Date | **Relevance** Show Jobs Posted: All Jobs View: Brief |

Date	Job Summary	Company	Location	Miles *
Jun 8	BUSINESS TO BUSINESS OUTSIDE SALES BUSINESS TO BUSINESS OUTSIDE SALES Would you like to sell something other than a "me too" product? Would you like a product line with very limited competition? Would you like to represent a product wh ...[more] **Relevant Work Experience:** 1+ to 2 Years **Career Level:** Experienced (Non-Manager) **Education Level:** Some College Coursework Completed **Job Type:** Employee **Job Status:** Full Time **Job Shift:** First Shift (Day)	Company Confidential	US-WI-Madison	1.71

- UW HAZY project: Extracts information from natural language text for knowledge base construction

- Machine learning: classification

Collaborative Filtering

- Recommendations based on other users' behavior
- e.g. Amazon

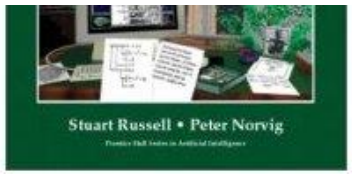

Availability: Usually ships within 24 hours. Ships from and so

Want it delivered Friday, July 1? Order it in the next 8 hou
choose **One-Day Shipping** at checkout. See details

76 used & new available from $14.99

Share your own customer images

Look inside another edition of this book

Customers who bought this book also bought

Introduction to Algorithms, Second Edition by *Thomas H. Cormen*

Machine Learning by *Tom M. Mitchell*

ANSI Common LISP by *Paul Graham*

Paradigms of Artificial Intelligence Programming : Case Studies in Common Lisp by *Peter N*

Operating System Concepts (Windows Xp Update) by *Abraham Silberschatz*

AI Application Programming (Programming Series) by *M. Tim Jones*

Explore Similar Items: in Books

- e.g. Netflix

- Unsupervised learning

Netflix Prize

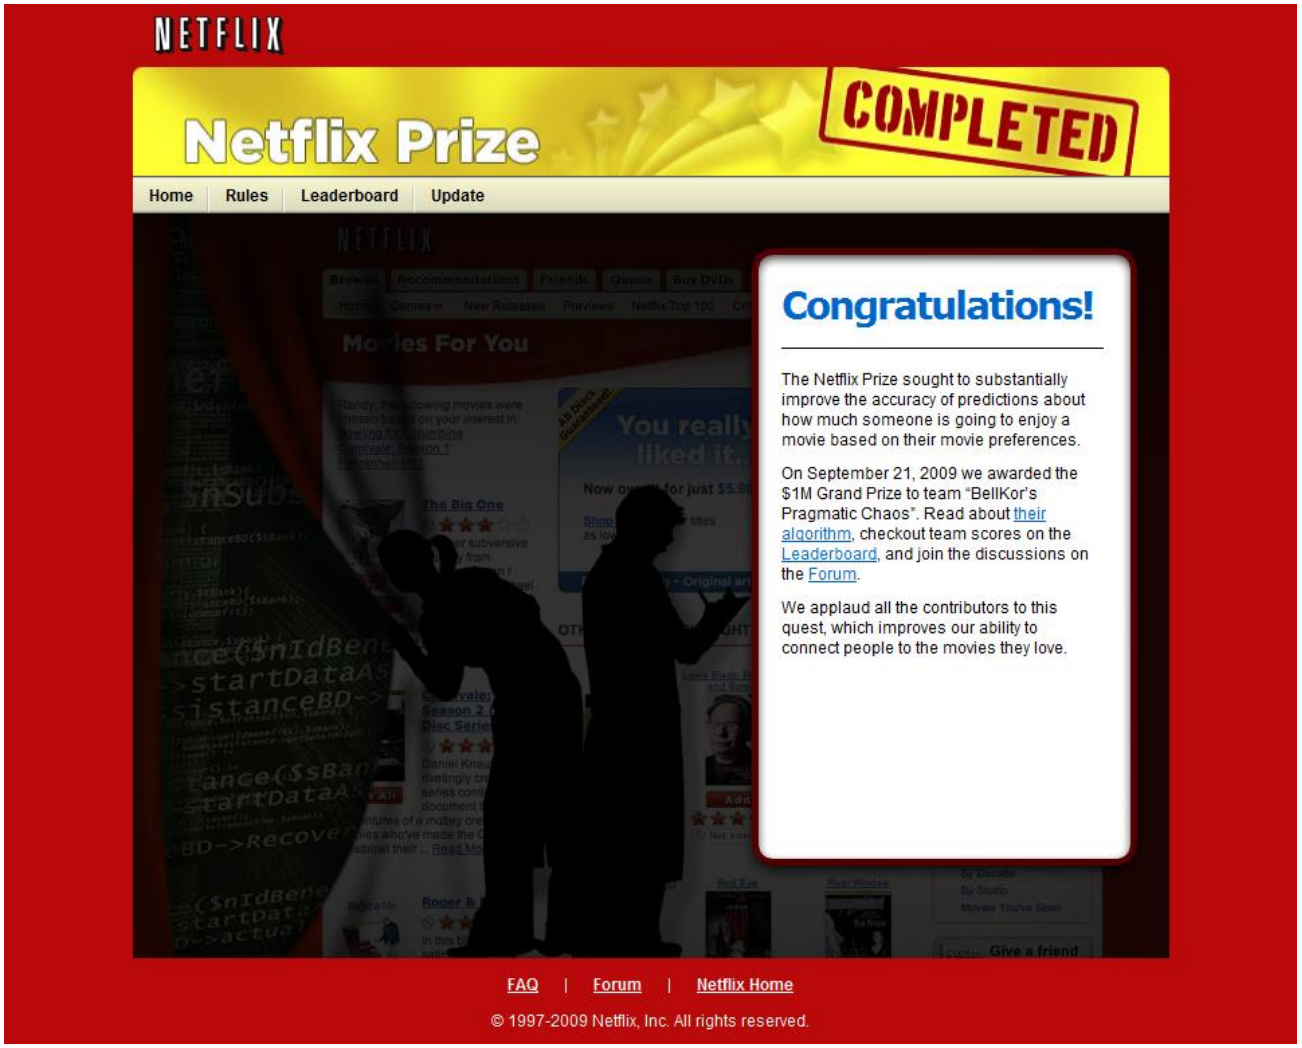

$1 million prize awarded in 2009; training set included 100 million ratings for 480,000 users and 18,000 movies

Visual Search: Google Goggles

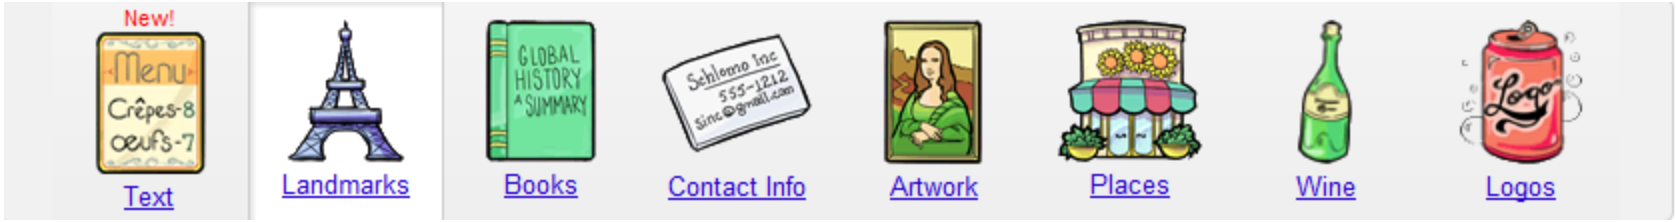

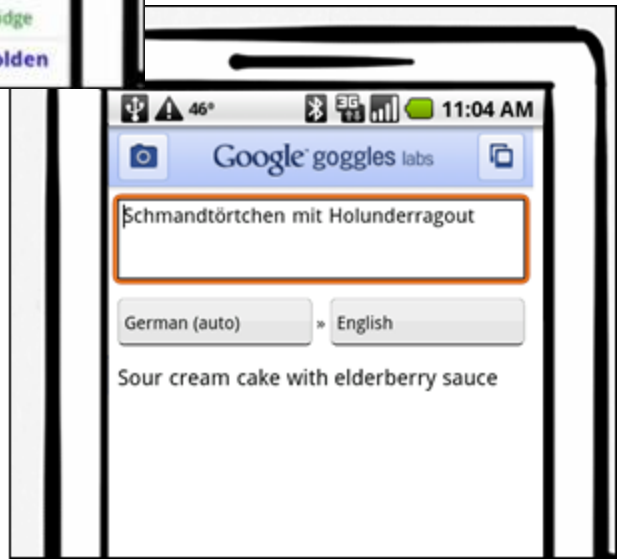

Glasses with Cameras

Google Project Glass

Instagram glasses

Face Detection now in most digital cameras for auto focusing

Also blink and smile detection!

Face Recognition: Autotagging Photos in Facebook, Flickr, Picasa, iPhoto, …

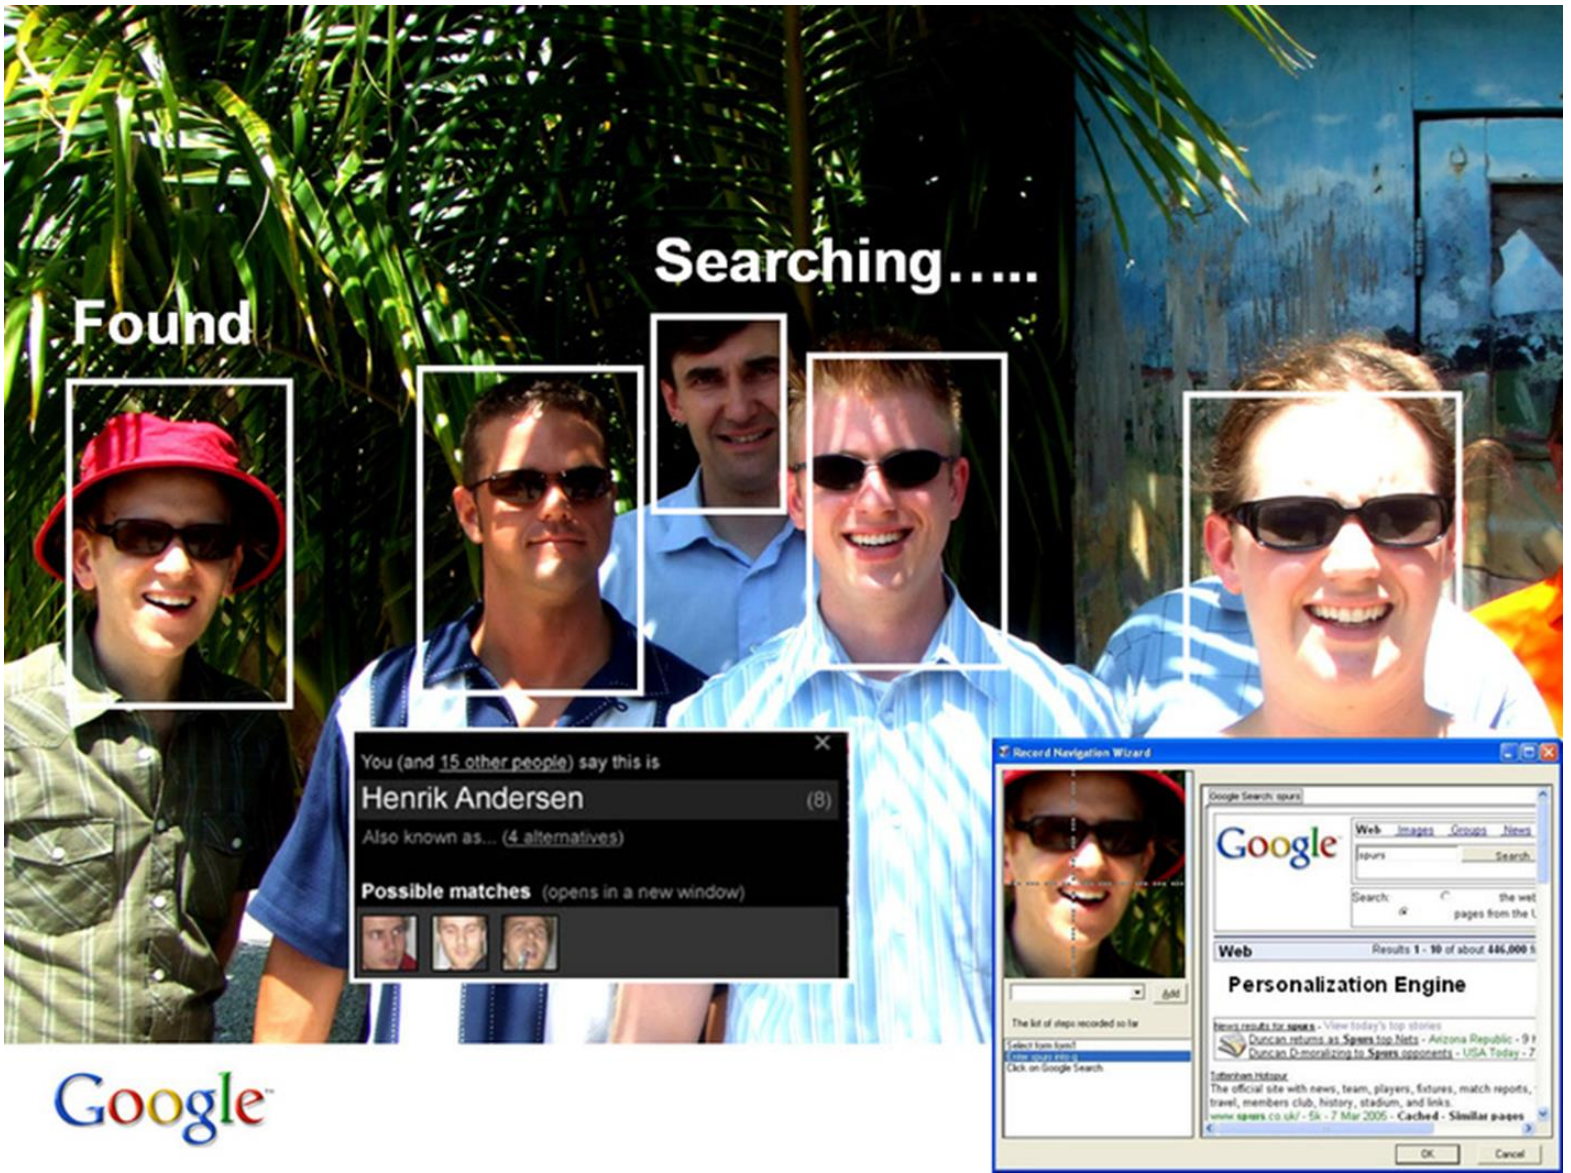

iPhoto

Can be trained to recognize pets too!

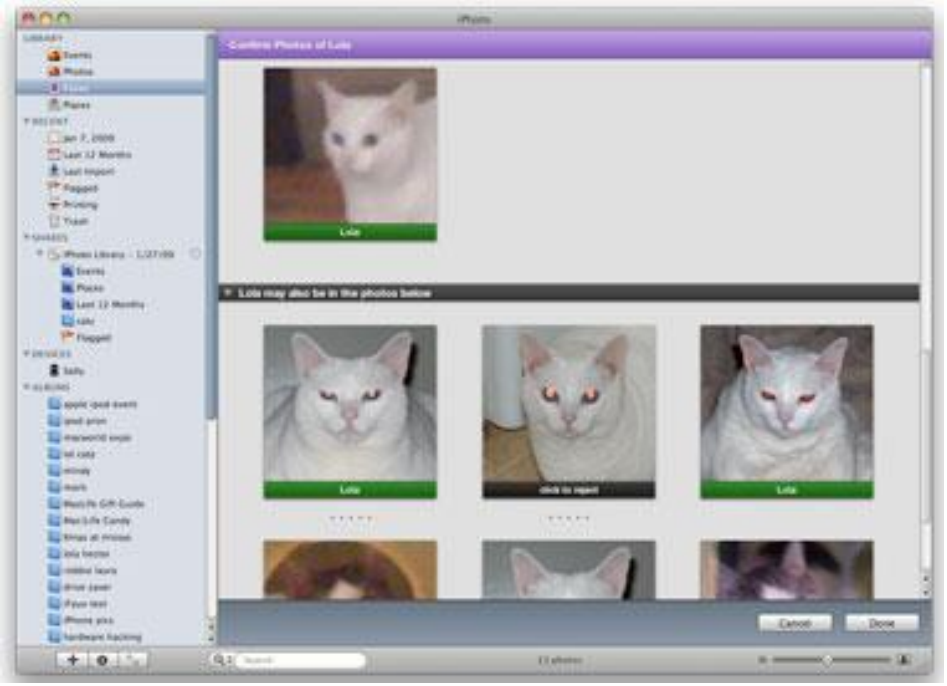

Perceptual Computing

Speech recognition, face recognition, gesture recognition and tracking

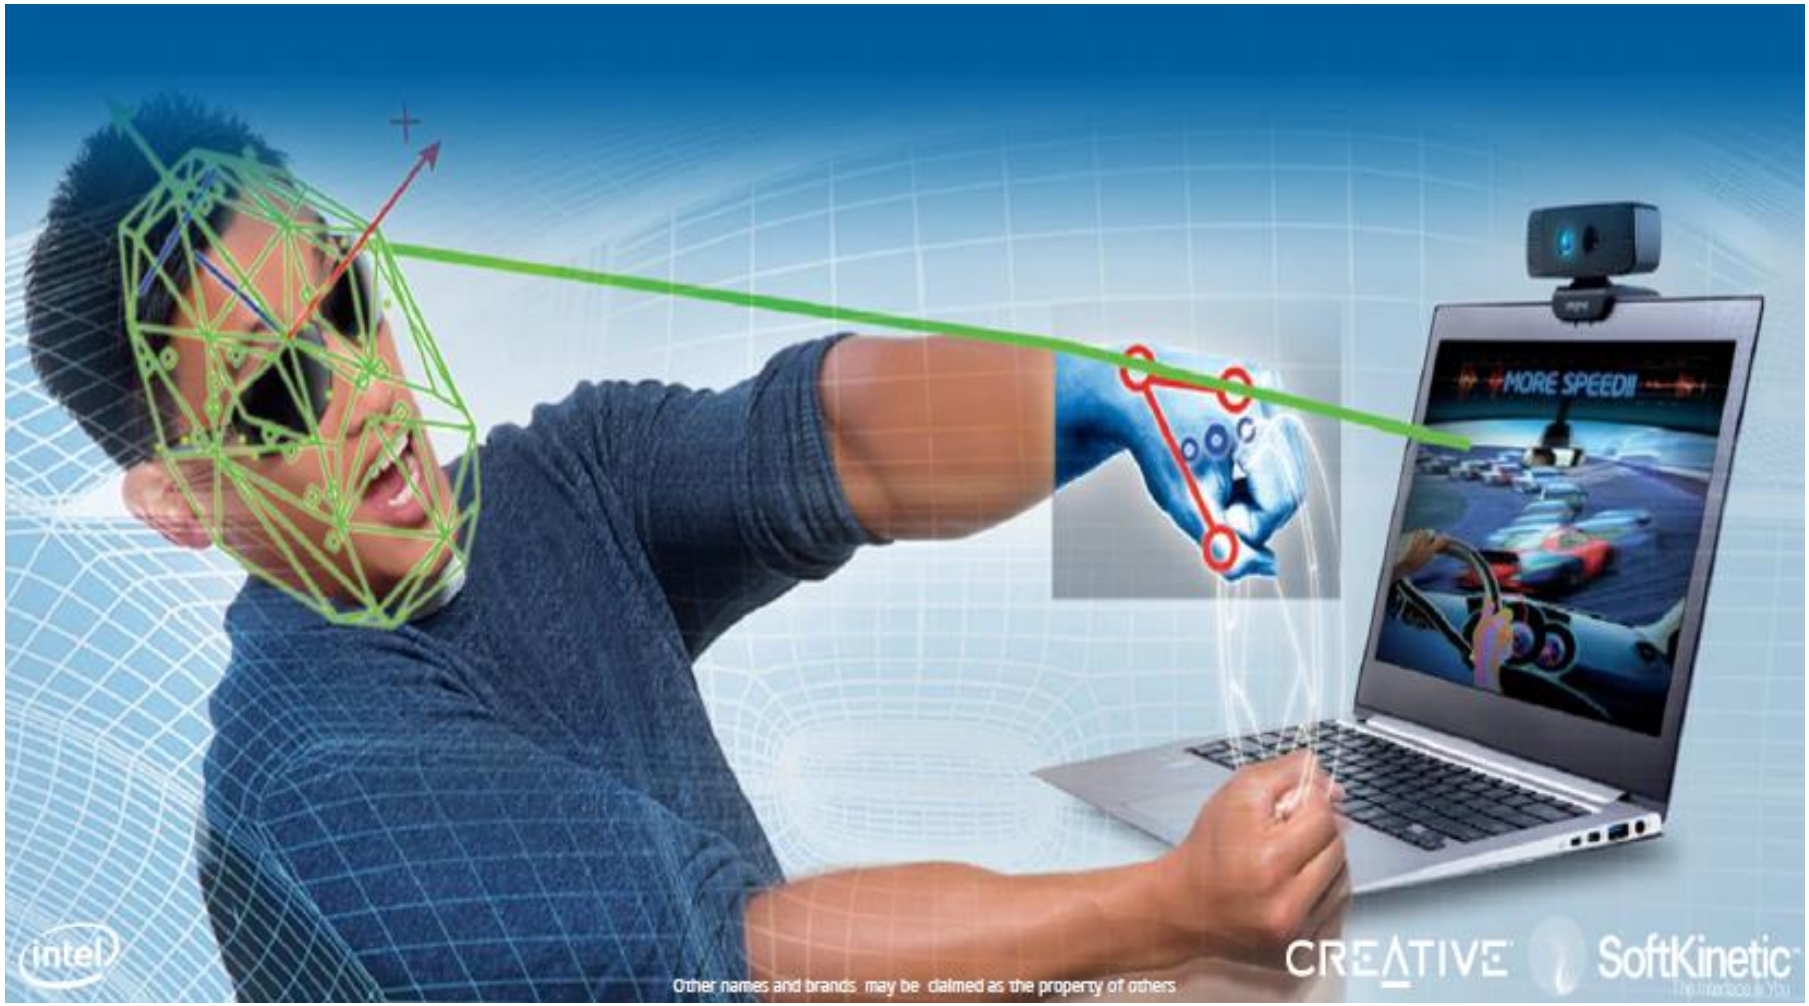

Intel Creative Interactive Gesture Recognition (Depth) Camera ($149)

Microsoft Kinect for XBOX

Microsoft Kinect Camera

Body Part Detection and Tracking

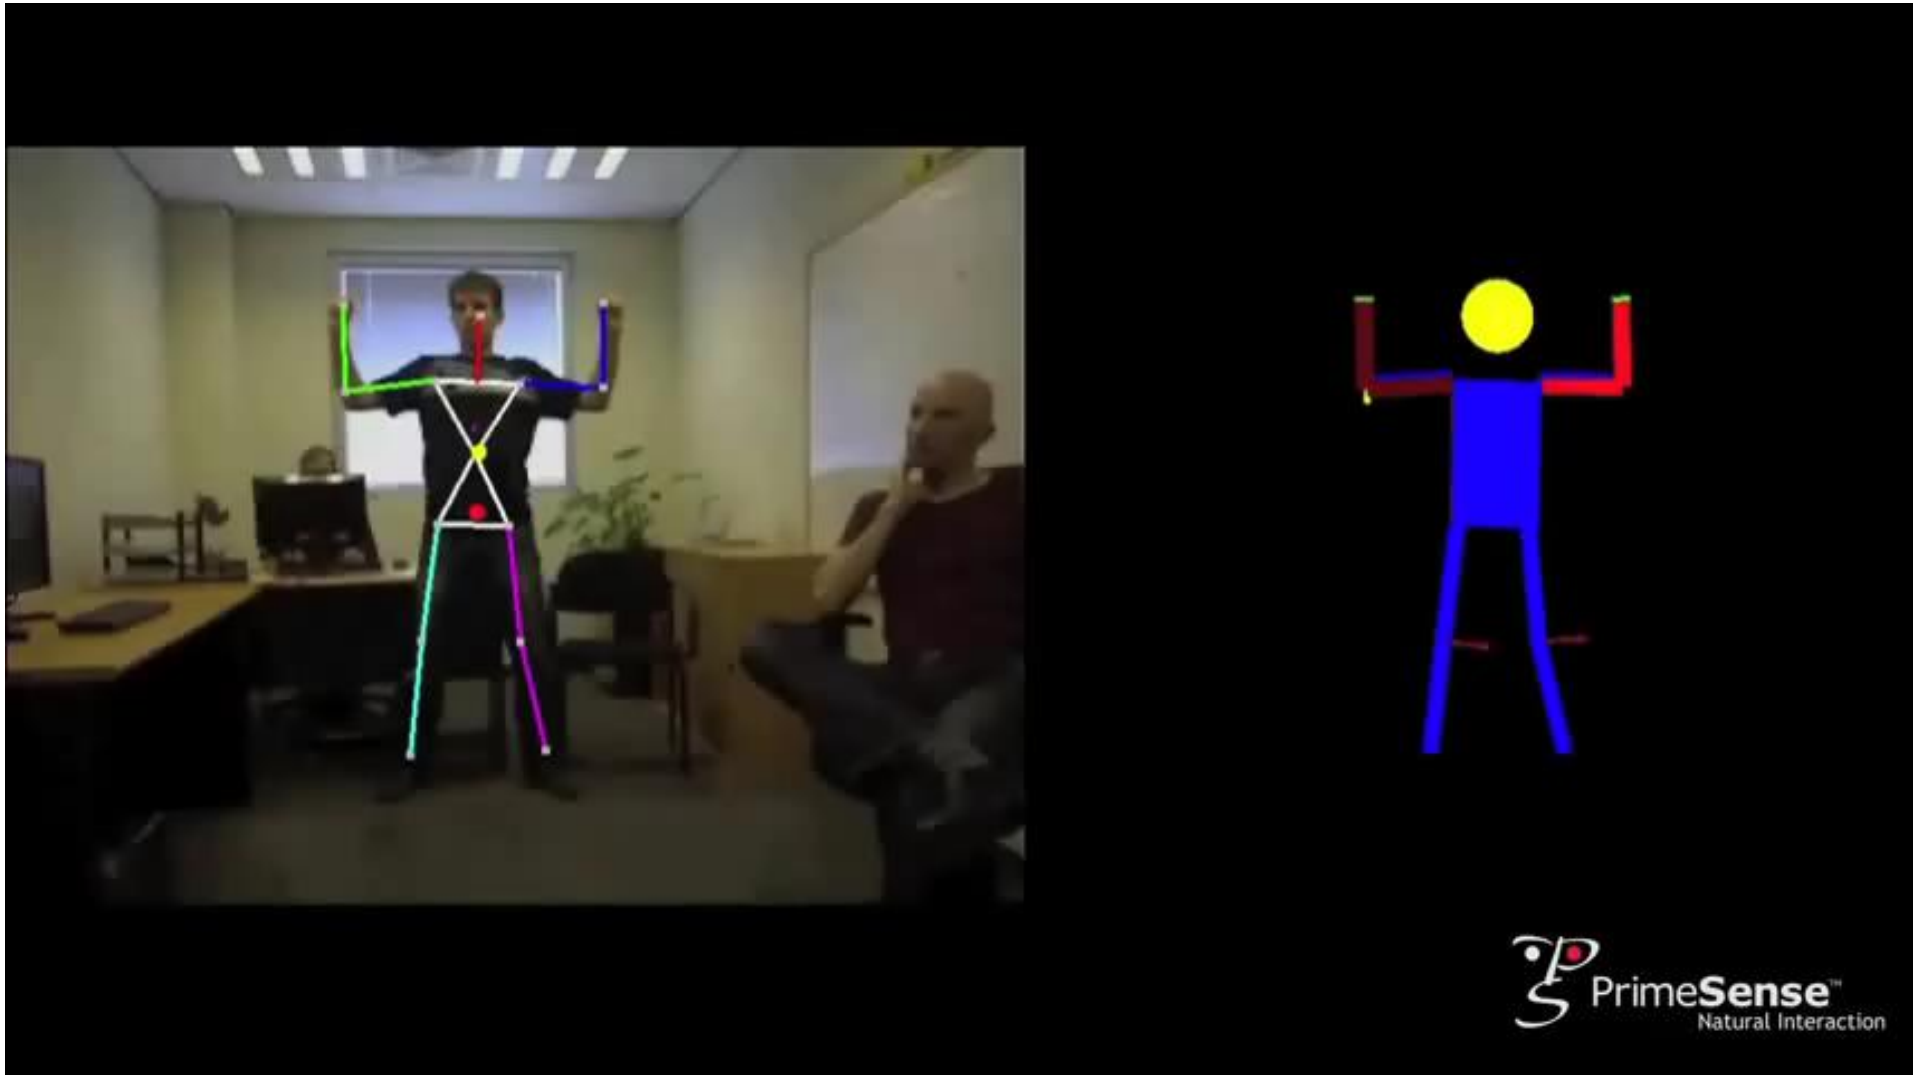

www.FITNECT.com

interactive fitting room

Flyable Cameras

- DJI Phantom 2 Vision Quadcopter
  - $1,200 (January 2014)

DJI Phantom 2 Vision Quadcopter

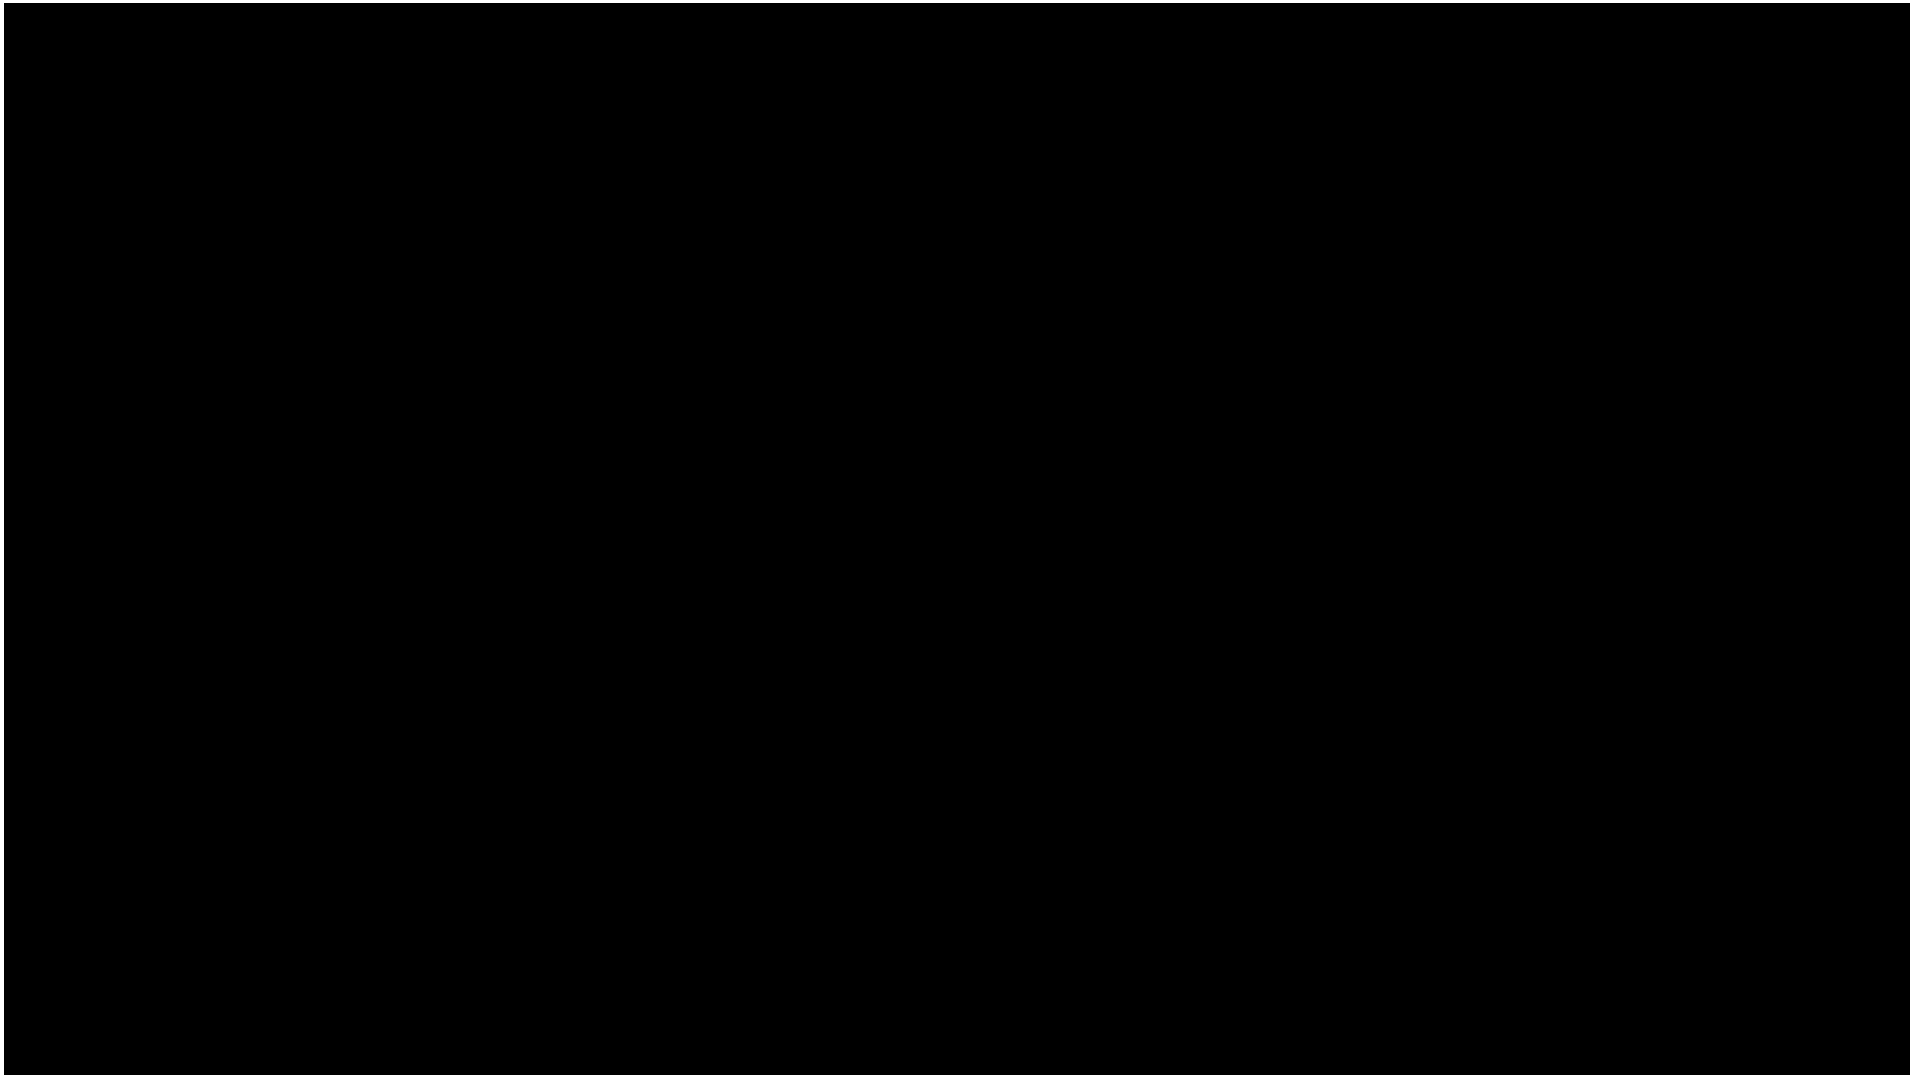

Robotics = Intelligent Connection of Perception to Action

Sensor data

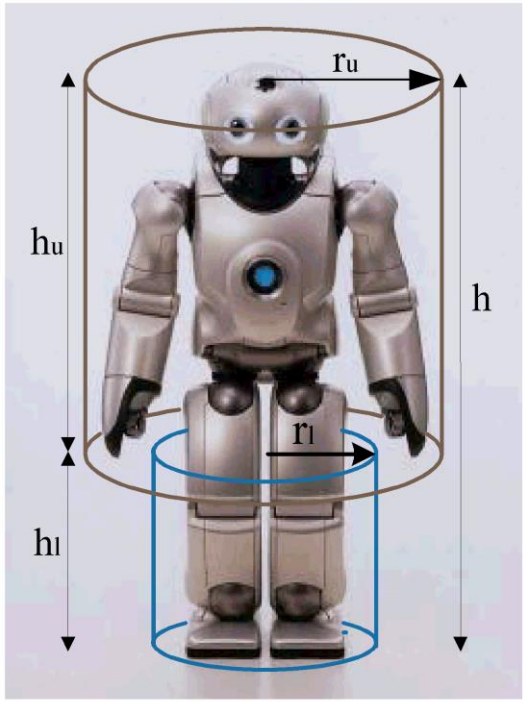

Robot brain

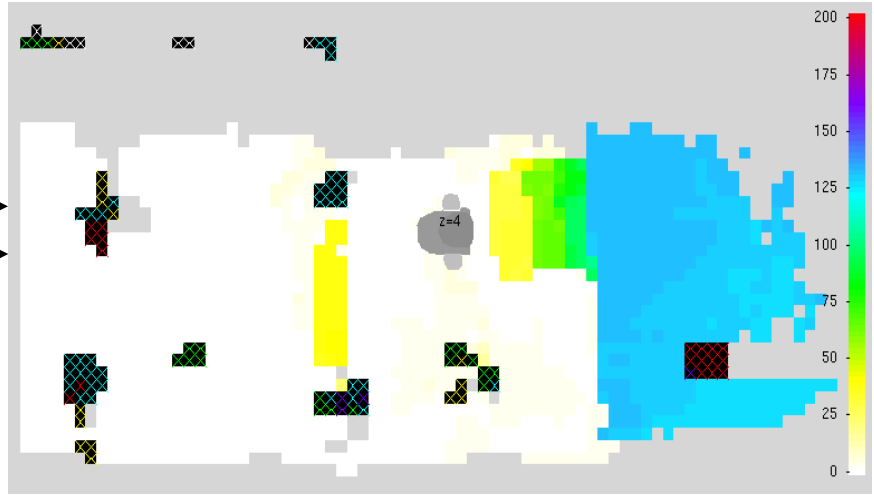

Environment

Actions

Autonomous Robots

- Key questions in mobile robotics
  - **What is around me?**
  - **Where am I ?**
  - **Where am I going ?**
  - **How do I get there ?**
- Alternatively, these questions correspond to
  - **Sensor Interpretation**: what objects are in the vicinity?
  - **Position and Localization**: find your own position on a map (given or built autonomously) and position on road
  - **Map building**: how to integrate sensor information and your own movement?
  - **Path planning**: decide the actions to perform for reaching a target position

Space Exploration Robots

Driving on Mars by Sojourner, Spirit, Opportunity, and Curiosity rovers

Cleaning Robots

- iRobot Roomba robot for vacuuming floors

Roomba demo

Lawn Mowing Robots

Robomow

Nursebot: Robots to Help the Elderly

Robots Playing Soccer

RoboCup Tournament

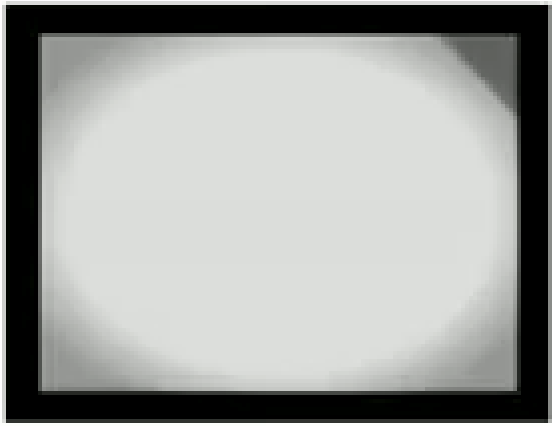

DARPA Robotics Challenge

- Competition of robot systems and software teams vying to develop robots capable of assisting humans in responding to natural and man-made disasters
  - Tasks emphasize capabilities related to **mobility**, **manipulation**, and **dexterity**
- $2 million prize in late 2014 to winner
- December 2013 – trials containing 8 tasks
  - Winner: Google's Schaft robot

Goggle's Schaft Robot

Robot Vehicles

Cars, airplanes, helicopters, birds, insects

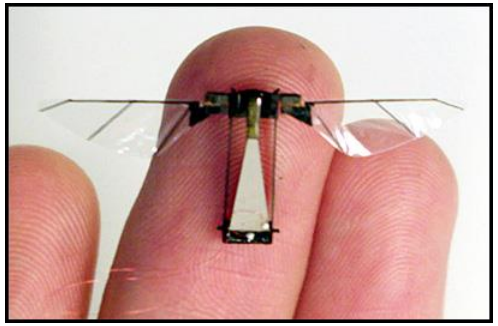

Robot Cars

What's Needed?

- Car Information
  - Position and orientation of car, velocity and turning rate of car
- Environment Information
  - Where is the road, curb, road signs, stop signs, other vehicles, pedestrians, bicyclists, …
- Actions
  - Velocity, steering direction, braking, …
- Sensors
  - Cameras, GPS, …

Robot Car Task Characteristics

- Fully or partially observable?
- Deterministic or stochastic?
- Static or dynamic?
- Discrete or continuous?
- Single or multi-agent?

Robot Car Task Characteristics

- Partially observable
- Stochastic
- Dynamic
- Continuous
- Multi-agent

Sensors

- Video cameras
- LIDAR (depth/range) sensor
  - times how long it takes a beam of laser light to bounce off something
  - gives 3D info on environment to 5 cm accuracy
- Radar sensors on front and rear
- Position sensor on wheel
- GPS
- Inertial motion sensor (IMU)
- Position and orientation of vehicle updated in real-time with 50 cm position accuracy and 1/50 degree orientation accuracy

LIDAR-Based Terrain Acquisition

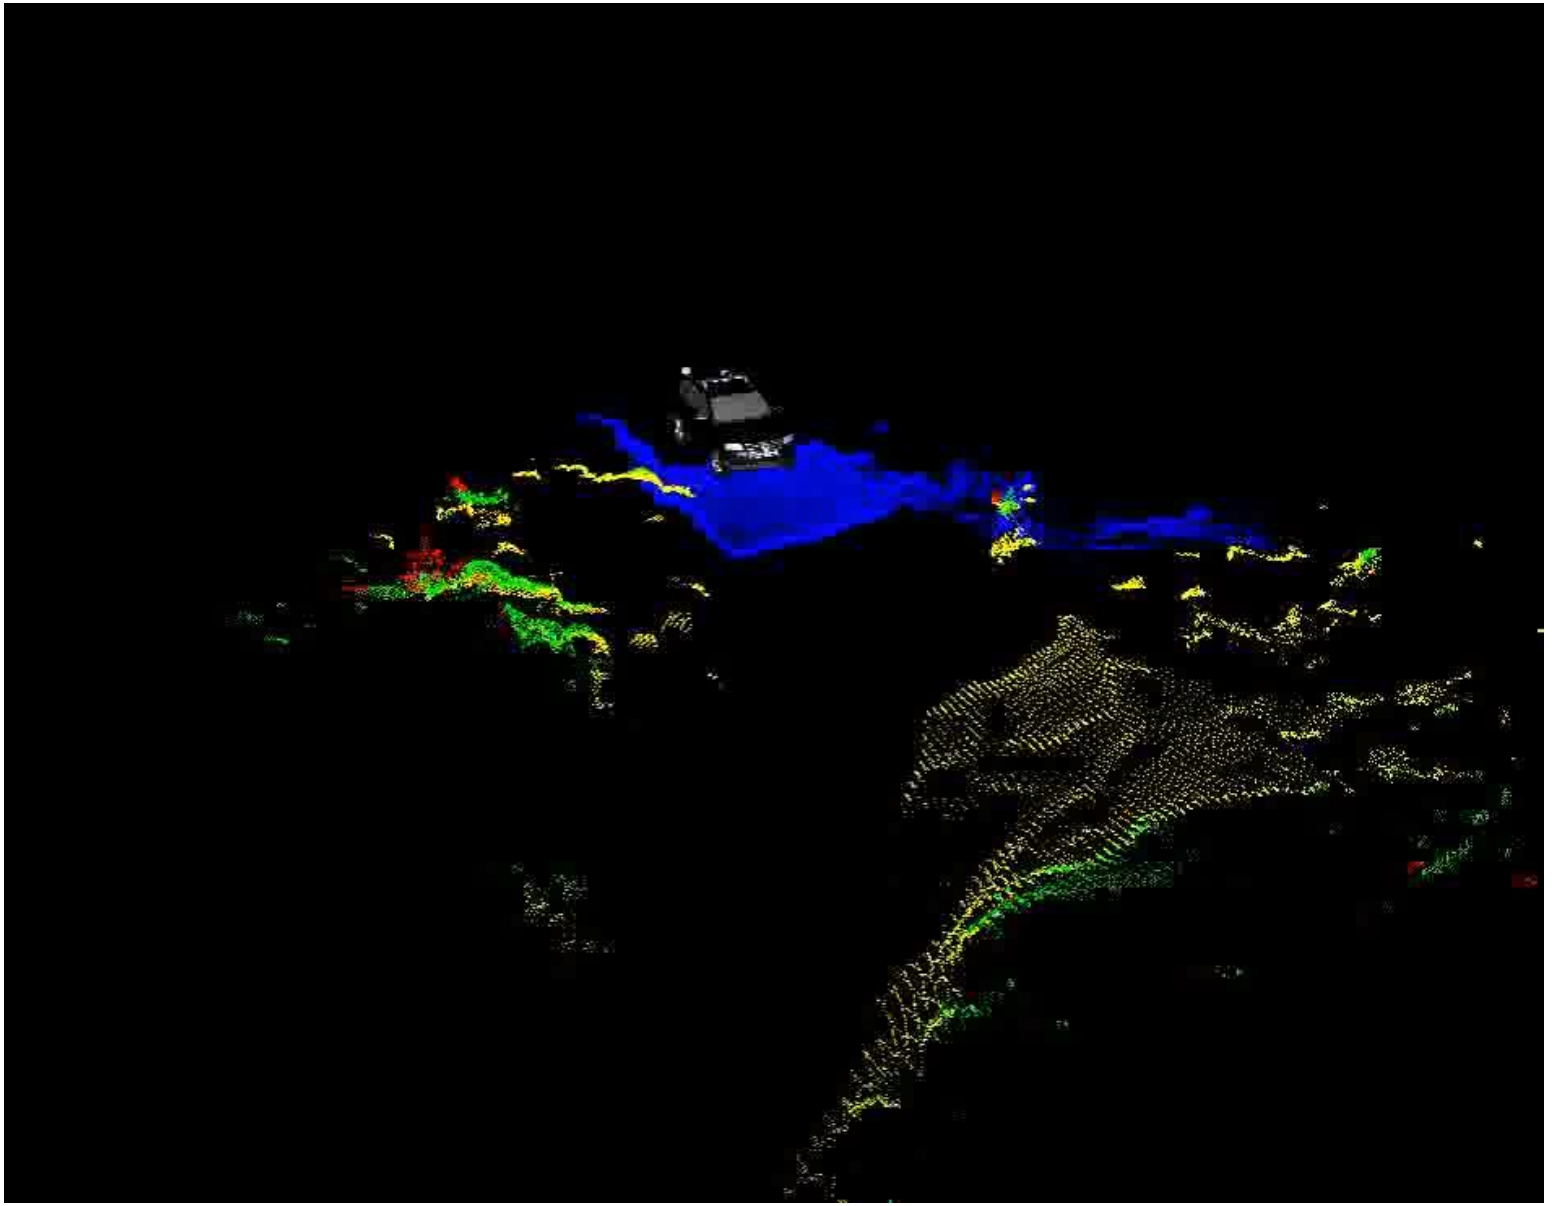

Road Detection Using Video and Depth Cameras

Robot Cars

DARPA Grand Challenges

DARPA Grand Challenge I
Barstow to Primm
March 13, 2004

142 miles
10 hours

DARPA Grand Challenge II
Desert Classic
October 8, 2005

132 miles
10 hours

DARPA Grand Challenge III
Urban Challenge
November 3, 2007

60 miles
6 hours

The 2005 “Grand Challenge” Race

Some Less Successful Vehicles

The 2007 “Urban Challenge”

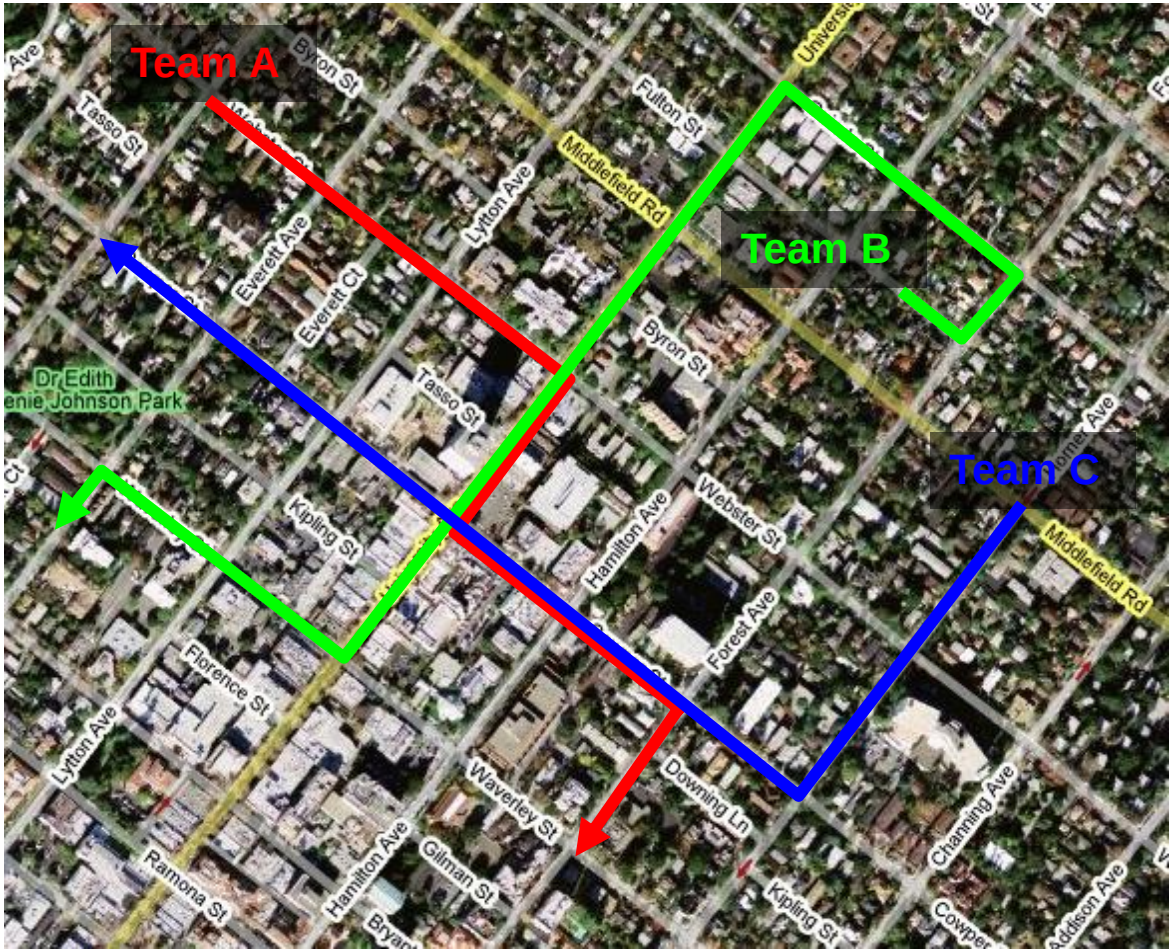

- Driving in urban environments
- Obey all CA traffic laws
- Accommodate road blockages, other vehicles, etc.


3
3
MDV

Automatic Parking

Google's Robot Car

Google's Driverless Car

The Future of Autonomous Driving?

- “In 20 years I will trust my autonomous car more than I trust myself”

  – Sebastian Thrun

- “It won’t truly be an autonomous vehicle until you instruct it to drive to work and it heads to the beach instead.”

  – Brad Templeton

Progress in AI: 1956 – 2010

Human-Level Chess

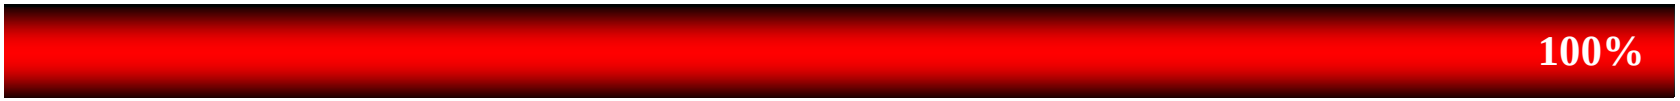

Human-Level Dialogue

Human-Level Perception

Harvesting Human Intelligence:

Anti-AI:

CAPTCHA and the ESP game

AI is Hard

- Some AI problems are very hard
  - Vision, natural language understanding, …
- What do you do?
  - Give up?
  - Bang your head really hard?
  - Important lesson in life:
    - turn hardness into something useful
- Very hard for machine, trivial for human

CAPTCHA

- Yahoo!

- Google

CAPTCHA

Completely **A**utomated **P**ublic **T**uring test to tell **C**omputers and **H**umans **A**part

- Yahoo!

- Google

CAPTCHA

- The “anti-Turing test”
- Tell human and machines apart, automatically
  - Deny spam-bots free email registration
  - Protect online poll from vote-bots
- By asking an “AI-complete” question

Random string → Distorted image → What do you see?

oamg →  → [ ]

- Also audio Captcha, e.g., superimposed speakers
- http://www.captcha.net/

[Luis von Ahn, IAAI/IJCAI 2003 keynote]

To prove that you're not a bot, please enter the characters you see in the picture below.

Submit

...

WELL, THAT SETTLES IT. I GUESS I ALWAYS HAD MY SUSPICIONS...

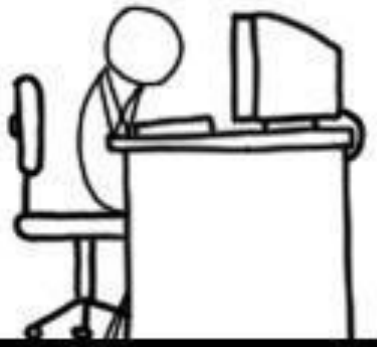

4/27/07 MIKE & STEPH ADAMS

reCAPTCHA

- reCAPTCHA is a free anti-bot service that improves the process of digitizing books by having humans decipher words that are not automatically recognized

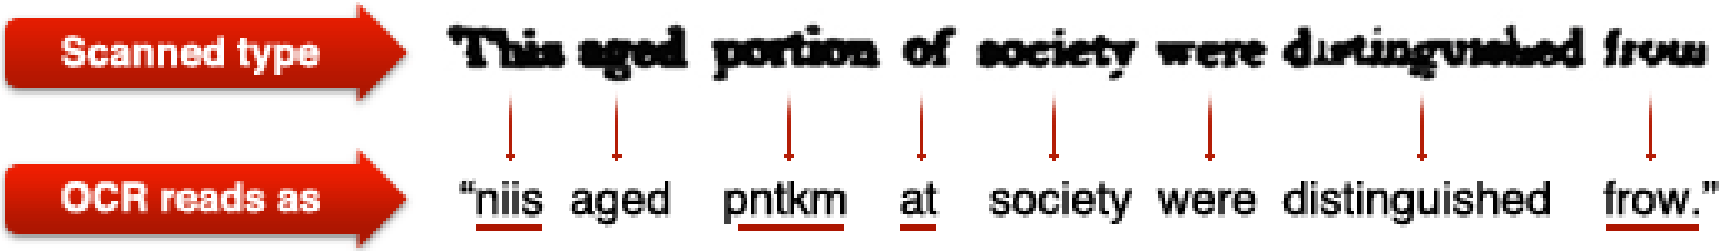

The ESP Game

- Real intelligence is here (for now)

- We waste it on computer games, anyway

- Harvest it
  http://www.gwap.com/gwap/gamesPreview/espgame/

The ESP Game

- Task: label all images on the web with words

→ car, boy, hat, …

- Why: current image search engines
  - use the image filename and surrounding text
  - do not really understand the image
- How: two separate players try to find a common description of the image

The ESP Game

PLAYER 1

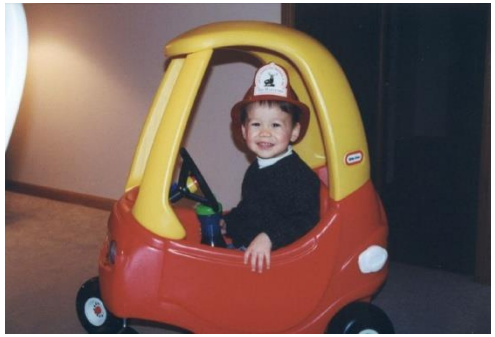

GUESSING: CAR

GUESSING: HAT

GUESSING: KID

SUCCESS!
YOU AGREE ON CAR

PLAYER 2

GUESSING: BOY

GUESSING: CAR

SUCCESS!
YOU AGREE ON CAR

[Luis von Ahn, IAAI/IJCAI 2003 keynote]

[Luis von Ahn, IAAI/IJCAI 2003 keynote]

Summary: you should be …

- either shocked or be assured that

There's no magic in AI.

It's all about optimization, probability and statistics, logic, algorithms.

- have a rough idea of the state-of-the-art of AI
- be able to talk AI at cocktail parties
- appreciate the ideas of CAPTCHA and ESP games