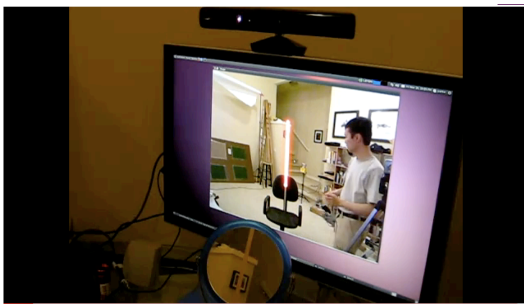

http://www.youtube.com/watch?v=3EeJCln5KYg

Robotics

CS311, Spring 2013
David Kauchak

Some material adapted from slides from Zach Dodds

Admin

- Assignment 5 graded
- Exam #2 available later today
  - To be done by Sunday at midnight

What is a robot?

"I can't define a robot, but I know one when I see one."
--Joseph Engelberger (1966)

Justice Potter Stewart wrote in Jacobellis v. Ohio (1964), "I can't define pornography, but I know it when I see it."

Robot Defined

Word robot was coined by a Czech novelist Karel Capek in a 1920 play titled Rossum's Universal Robots (RUR)

Robota in Czech is a word for worker or servant

Karel Capek

Definition of robot:

Any machine made by one our members: Robot Institute of America ☺

A robot is a reprogrammable, multifunctional manipulator designed to move material, parts, tools or specialized devices through variable programmed motions for the performance of a variety of tasks: Robot Institute of America, 1979

What is a Robot

Manipulator

What is a Robot

Legged Robot

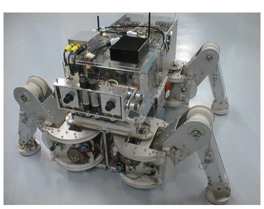

Wheeled Robot

What is a Robot

Autonomous Underwater Vehicle

Unmanned Aerial Vehicle

Robot Plot

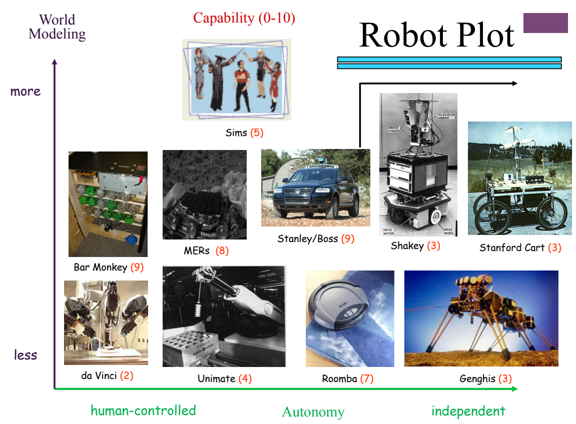

Robot timeline?

1921 — 1950 — 2020 — 2150 — 2421 …

Fictional Robot timeline

Put these robots in chronological order?

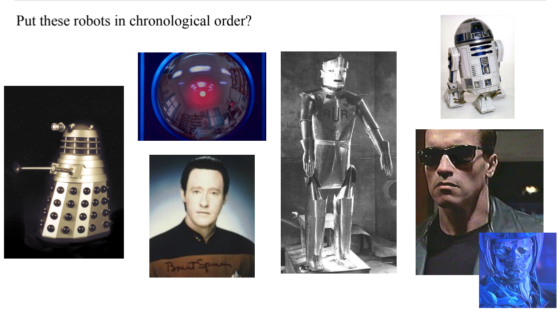

1921 — 2020 — 2150 — 2421 …

Fictional robot timeline

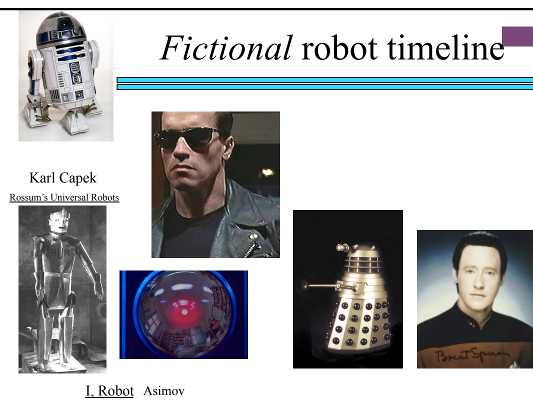

1921 — 1950 — 2020 — 2150 — 2421 …

Real robot timeline

1951 — 1968 — 1976 — 1985 …

Real robot timeline

Tortoise "Elsie"

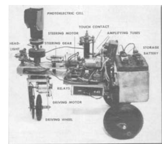

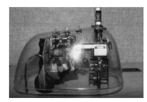

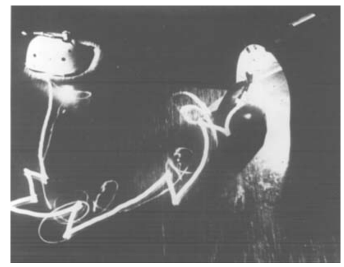

by Neurophysiologist Grey Walter

1951 …

http://www.frc.ri.cmu.edu/~hpm/talks/revo.slides/1950.html

Shakey

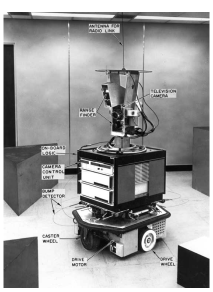

Nils Nilsson @ Stanford Research Inst.

first "general-purpose" mobile platform

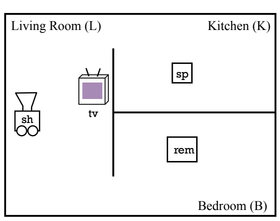

… 1968 …

Robotics's *Shakey* start

START

At(sh,L) ∧ At(sp,K) ∧ At(rem,B) ∧ At(tv,L)

Go(L,B)

Go(L,K)

Push(tv,L,K)

Push(tv,L,B)

At(sh,K) ∧ At(sp,K) ∧ At(rem,B) ∧ At(tv,K)

At(sh,L) ∧ At(sp,L) ∧ At(rem,L) ∧ At(tv,L)

GOAL

<u>ACTIONS</u>

Go(from,to)
Preconditions: At(sh,from)
Postconditions: At(sh,to)

Push(obj,fr,to)
Preconditions: At(sh,fr) ∧ At(obj,fr)
Postconditions: At(sh,to) ∧ At(obj,to)

+ Shakey in video

http://www.youtube.com/watch?v=qXdn6ynwpiI

Stanford Cart: *SPA*

Hans Moravec @ SAIL

"functional" task decomposition →
"horizontal" subtasks

SENSING | perception | world modeling | Planning | task execution | motor control | ACTING

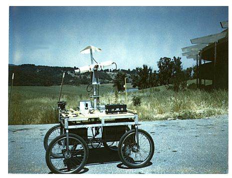

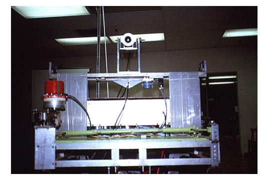

… 1976 …

Cartland (outdoors)

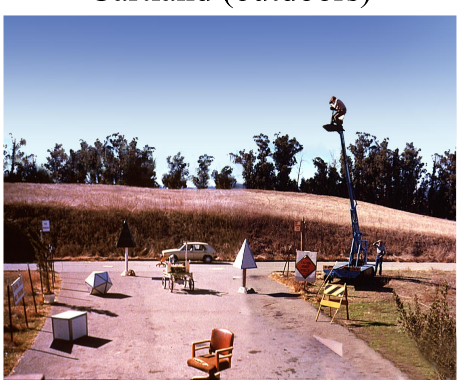

Cartland (indoors)

"Robot Insects"

Rodney Brooks @ MIT

"behavioral" task decomposition →
"vertical" subtasks

SENSING

planning and reasoning
identify objects
build maps
explore
wander
avoid objects

ACTING

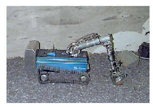

… 1985 …

Robotics

What are the challenges?
How do these relate to AI?

AI

- Search
  - planning
- Game playing
- CSPs
- Bayesian
- HMMs
- Machine learning
  - neural nets
- Knowledge representation
- Natural Language processing
- Computer vision

Autonomy/behavior

how much of the world do we need to represent internally ?

Robot Architecture

how should we internalize the world ?

what outputs can we effect ?

what inputs do we have ?

what algorithms connect the two ?

how do we use this "internal world" effectively ?

Robot Architecture

how much / how do we represent the world internally ?

As much as possible!

SPA paradigm

sense → plan → act

Not at all

Reactive paradigm

Task-specific

Behavior-based architecture

As much as *possible*.

Hybrid approaches

history…

Sense - Plan - Act

Shakey

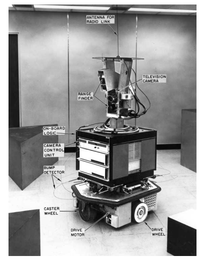

SENSING | perception | world modeling | planning | task execution | motor control | ACTING

sense → plan → act

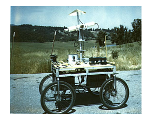

Stanford Cart

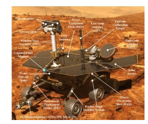

MERs

… 1968 — 1976 — … - 2009 …

Mars Exploration Rovers

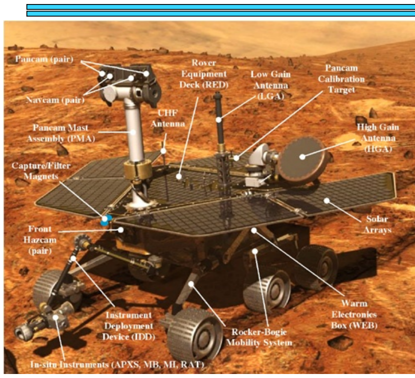

Sense – Plan – Act
"deliberative"
architecture

Mars Science Lab

2011 - lasers, lifebio, and maybe nuclear-powered

Robot Architecture

how much / how do we represent the world internally ?

As much as possible!

SPA paradigm

sense → plan → act

Not at all

Reactive paradigm

sense ↔ act

stimulus - response

Task-specific

Behavior-based architecture

As much as *possible*.

Hybrid approaches

Biological Inspiration

Ethology: describing animal behavior

Getting to the ocean?

Digger wasps' nest-building sequence

AI reasoning systems abstract too much away: *frame problem*

"The world is its own best model"

sense ⟷ act

Decision-making is based only on current sensor inputs.

Analog reactive robots

"Tortoise" — Gray Walter — Valentino Braitenberg — Mark Tilden (commercial products…) — "BEAM"

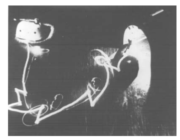

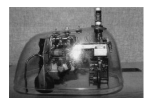

"light-headed" behavior

http://people.cs.uchicago.edu/~wiseman/vehicles/

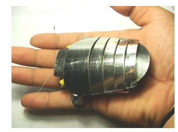

http://haroldsbeambugs.solarbotics.net/mercury.htm

robot made from Playstation pieces…!

1951 — 1984 — 1989- …

stateless...

Robot Architecture

how much / how do we represent the world internally ?

As much as possible!

SPA paradigm: sense → plan → act

Not at all

Reactive paradigm: sense ⟷ act

stimulus – response == "behavior"

Task-specific

Behavior-based architecture → Subsumption paradigm / Potential Fields } different ways of composing behaviors

As much as possible.

Hybrid approaches

Behavior-based control

Behavior — a direct mapping of sensory inputs to a pattern of task-specific motor actions

sense ⟷ act

extinguish
approach
wander

little explicit deliberation except through system state

"Vertical" task decomposition

SENSING

planning and reasoning
identify objects
build maps
explore
wander
avoid objects

ACTING

Genghis

… 1985 …

Subsumption

Subsumption builds intelligence incrementally in layers

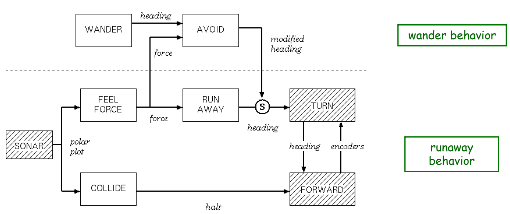

Subsumption

Where would a light-seeking behavior/layer connect?

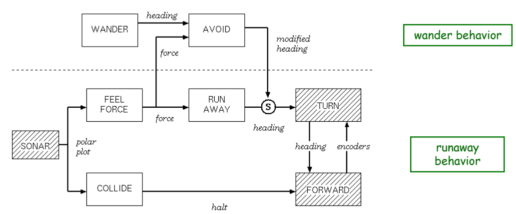

Subsumption

Where would a light-seeking behavior/layer connect?

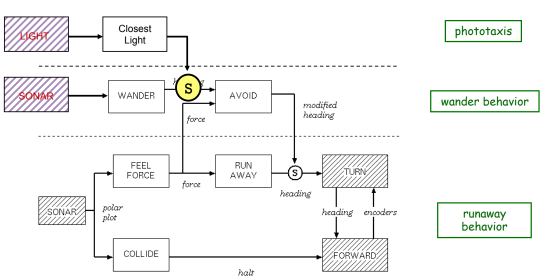

Subsumption - Limits

Reaching the end of the subsumption architecture and purely reactive approaches.

Herbert, a soda-can-collecting robot

http://www.youtube.com/watch?v=YtNKuwiVYm0

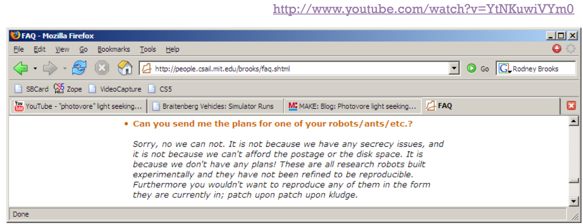
FAQ - Mozilla Firefox

• Can you send me the plans for one of your robots/ants/etc.?

Sorry, no we can not. It is not because we have any secrecy issues, and it is not because we can't afford the postage or the disk space. It is because we don't have any plans! These are all research robots built experimentally and they have not been refined to be reproducible. Furthermore you wouldn't want to reproduce any of them in the form they are currently in; patch upon patch upon kludge.

Success of behavior-based systems depends on how well-tuned they are to their environment. This is a huge strength, but it's also a weakness …

Subsumption *limits*: Genghis

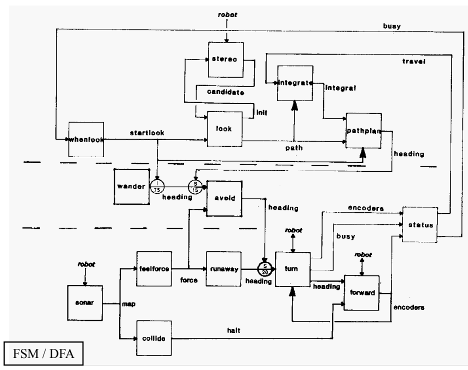

navigate behavior

wander behavior

runaway behavior

FSM / DFA

Unwieldy!

Larger example -- Genghis

1) *Standing* by tuning the parameters of two behaviors: the leg "swing" and the leg "lift"
2) *Simple walking*: one leg at a time
3) *Force Balancing*: via incorporated force sensors on the legs
4) *Obstacle traversal*: the legs should lift much higher if need be
5) *Anticipation*: uses touch sensors (whiskers) to detect obstacles
6) *Pitch stabilization*: uses an inclinometer to stabilize fore/aft pitch
7) *Prowling*: uses infrared sensors to start walking when a human approaches
8) *Steering*: uses the difference in two IR/range sensors to follow

57 modules wired together !

Robot Architecture

how much / how do we represent the world internally ?

As much as possible!

SPA paradigm

sense → plan → act

Not at all

Reactive paradigm

sense ↔ act

Task-specific

Behavior-based architecture → Subsumption paradigm / Potential Fields } different ways of composing behaviors

As much as possible.

Hybrid approaches

Potential Fields

Potential fields compose simple behaviors by *adding* the outputs that each sensor/input sends the robot

Individual potential fields (motor schemas) contain state

A sequencing process (FSM/DFA) updates the potential fields and/or decides which ones to run next…

Ron Arkin @ Georgia Tech

Motor Schemas / Potential Fields

Direct mapping from the environment to a control signal

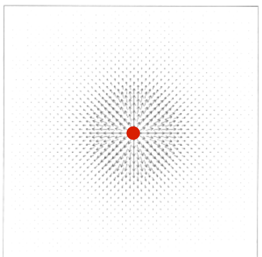

obstacle-avoiding schema

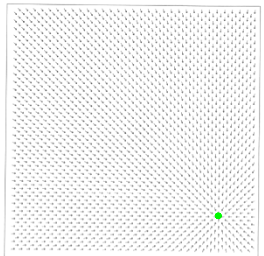

goal-seeking schema

note that the complete environmental vector fields are only for visualization!

combine?

Behavior Summer

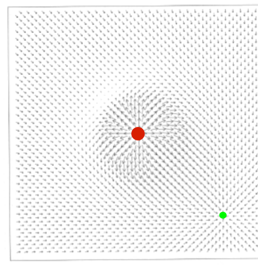

vector sum of the avoid and
goal motor schemas

path taken by a robot
controlled by the resulting field

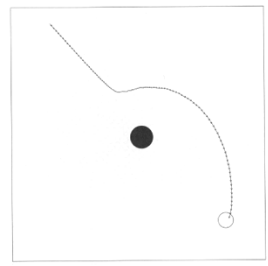

Implementation details

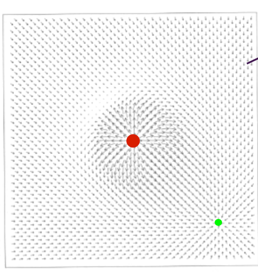

the extent to which potential field
force drops off with distance…

what crucial assumption is being
made here?

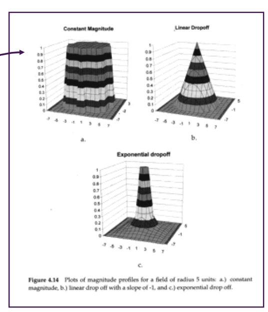

corridor-following schema(s)?

Additional behavior primitives

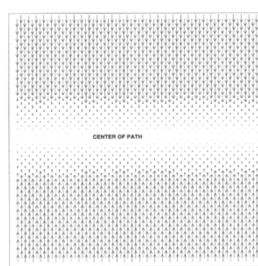

corridor-centering schema

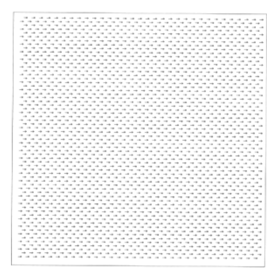

go! schema

A more complex task

Direct mapping from the environment to a control signal

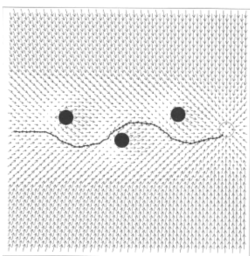

larger composite task

How many individual fields are summed in this task?

Not necessarily all at one time!

Local minima

A potential-field-based system can get stuck!

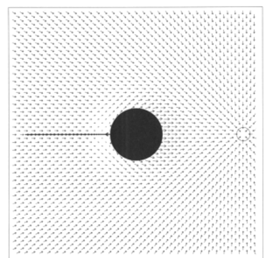

the problem

What would happen if a robot came in in the middle on the left?

a solution?

Local minima

A potential-field-based system can get stuck!

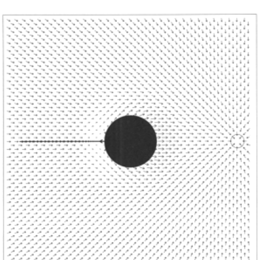

the problem

Why is the “local minimum” problem, as illustrated to the left, *not* likely to actually cause a robot to get stuck in practice?

robots controlled by summing goal/obstacle potential fields *can* get stuck in practice -- draw an example of an environment with both obstacle(s) and goals(s) in which getting stuck might actually occur.

Suggest how a robot might overcome the problem of getting stuck in such cases…

Local minima

A potential-field-based system can get stuck!

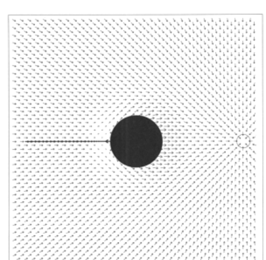

the problem

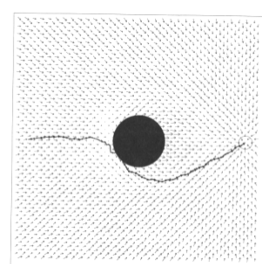

a solution

Bigger deadends...

How to get out of larger wells ?

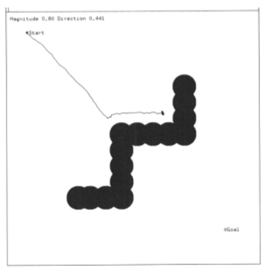

Bigger deadends...

uses memory of where the robot has been

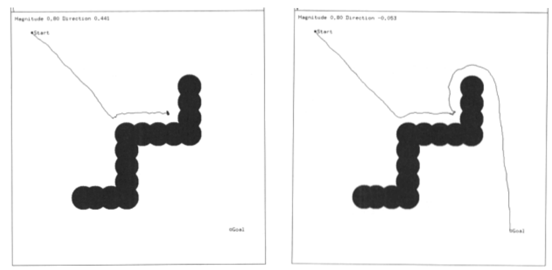

past-avoiding motor schema

Another example

Keeping away from past locations...

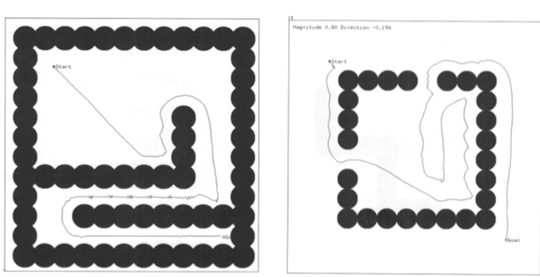

Pfields in Practice

Steathy navigation @ USC (Ashley Tews, Gaurav S. Sukhatme, and Maja J. Mataric)

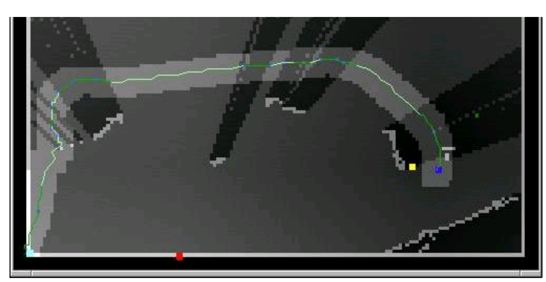

http://robotics.usc.edu/interaction/?l=Research:Projects:stealth:index#experiments

part of the potential field… What's going on here?

Docking with potential fields

Why might a simple attractive force not be sufficient for docking (plugging-in, etc.)?

example goals

How does the idea of docking, e.g., with an electrical outlet change the requirements for a potential field?

Docking with potential fields

The key insight is the need to establish an approach *direction*

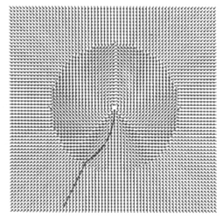

Figure 4.25 Docking potential field showing path of robot entering from slightly off course.

example goals

Docking with potential fields

The key insight is the need to establish an approach *direction*

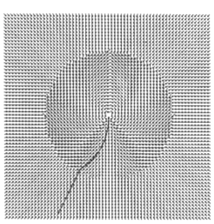

Figure 4.25 Docking potential field showing path of robot entering from slightly off course.

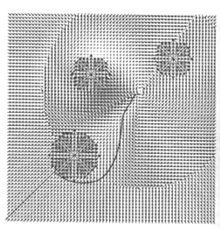

Figure 4.26 Visualization of the docking behavior with obstacles.

+ Review

- Machine learning
  - general learning concepts
    - supervised vs. unsupervised
    - features/feature-based problems/feature space
    - bias/variance
    - overfitting
    - hyperplanes/linear seperability
  - Supervised learning
    - applications
    - approaches
      - k-NN
      - decision trees
      - NB
      - SVM (large margin classifiers)
    - Ensemble approaches (boosting)

+ Review

- Machine learning (continued)
  - unsupervised learning
    - application
    - issues
      - number of clusters
      - flat vs. hierarchical
      - soft vs. hard clustering
    - approaches
      - k-means
      - EM
        - word alignment
        - clustering (mixture of gaussians)
      - spectral clustering (min-cut)

+ Review

- Neural networks (Machine learning?)
  - perceptrons/neurons
    - activation functions (threshold vs. sigmoid)
    - perceptron learning
  - multi-layer networks
- Knowledge representation
  - basic logic
  - ontology
  - NELL

+ Review

- CSPs
  - problem formulation
    - variables
    - domain
    - constraints
  - why CSPs? applications?
  - constraint graph
  - CSP as search
    - backtracking algorithm
    - forward checking
    - arc consistency
  - heuristics
    - most constrained variable
    - least constrained value
    - …

+ Review

- Natural language processing
  - Applications
  - Problem areas
  - Why it's hard?
  - Machine translation setup

Guest speaker

- Rodney Brooks
  - Professor at MIT (was previous director of CSAIL)
  - Founder of iRobot

- http://www.youtube.com/watch?v=B79D9nW2AFA