CS295-Z: Robot Learning and Autonomy

Chad Jenkins - Spring 2010
Brown Computer Science

http://www.cs.brown.edu/courses/csci2950-z.html

Big Question

- What does society want to do with robots?
  - Informally, what is the “killer app” of robotics?
  - Note: not the “killer robot app”

Big Question

- What does society want to do with robots?
- Problems:
  - Society has little idea what robots can do
  - Programming robots requires significant technical expertise
  - Chicken-egg problem -> scifi notions disparate from reality

One Possible Answer

- Program robots from human demonstration
- Research Problems:
  - <u>Algorithms</u>: learn policy from data (exper., exprl., guidance, etc)
  - <u>Data collection</u>: “lifelong” human supervision and robot performance
  - <u>Usability</u> by humans; interruptions

Course Structure

- Group project for entire class
- Cover research papers in robot learning and object manipulation
  - cover 2-3 papers per class
  - student paper presentations (20 mins max, minus questions)
  - everyone must summarize each paper

Group Project

- Massive-scale learning from demonstration
- Implement in ROS; do Create tutorial:
  - http://code.google.com/p/brown-ros-pkg/
- Learn three tasks from demonstration
  - Create robot soccer
  - Nao magneto assembly
  - PR2 intern challenge
- One learning alg, one infrastructure box
- Human subjects study

WHY ROBOT LEARNING?

“Any controller that has been learned could have been programmed in less time and performed better”

- anonymous big name in robotics

A GOAL FOR ROBOTICS

Collaborators for human endeavors

- Robot → tool for user productivity
- path of least resistance for doing physical tasks
- user-developed applications through learning
- critical path tasks?
- societal utility?

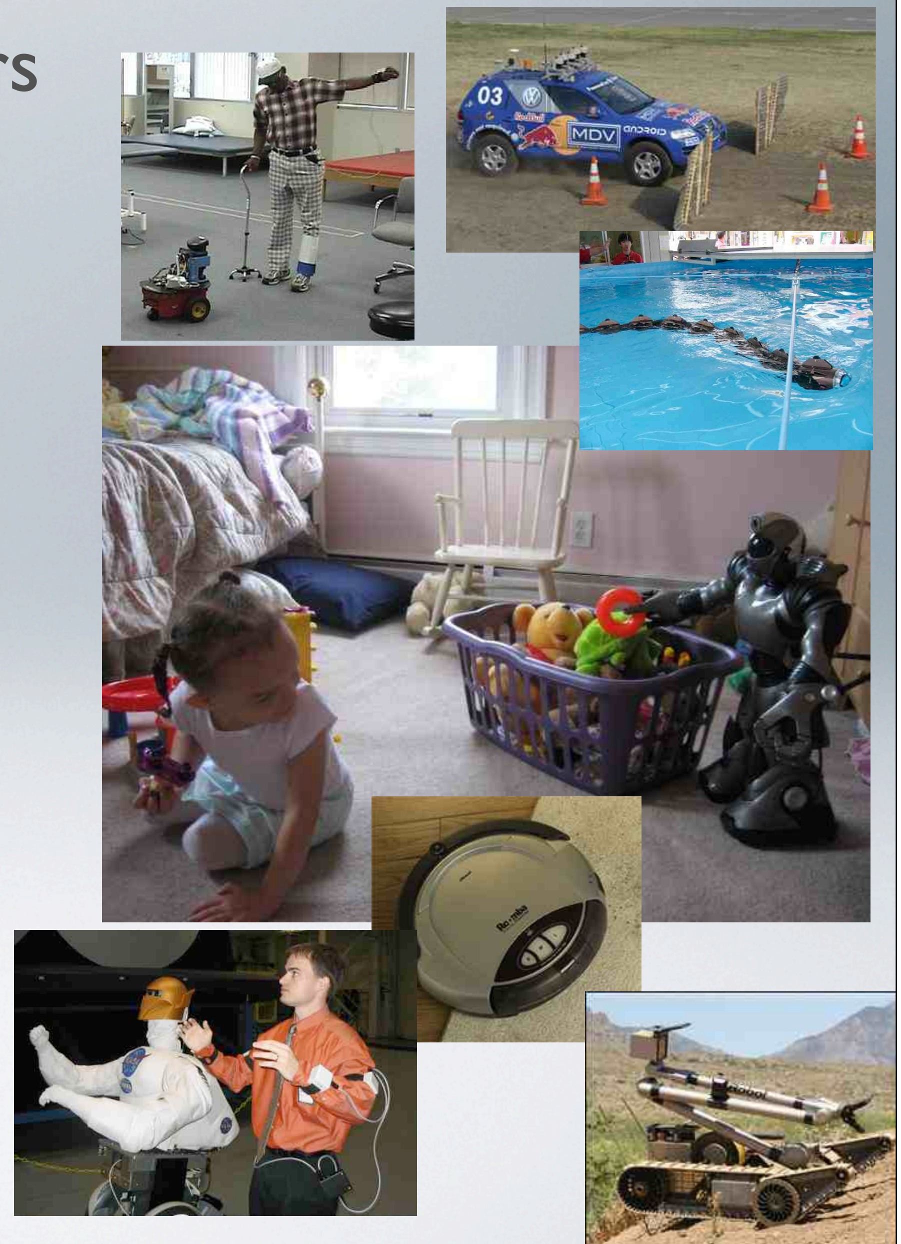

Personal Computing

ENIAC

Apple II

Laptop

OLPC

“technology exponentials”, e.g., Moore’s Law; mentioned by Brooks and others

Research

Novelty tech

Pervasive tools

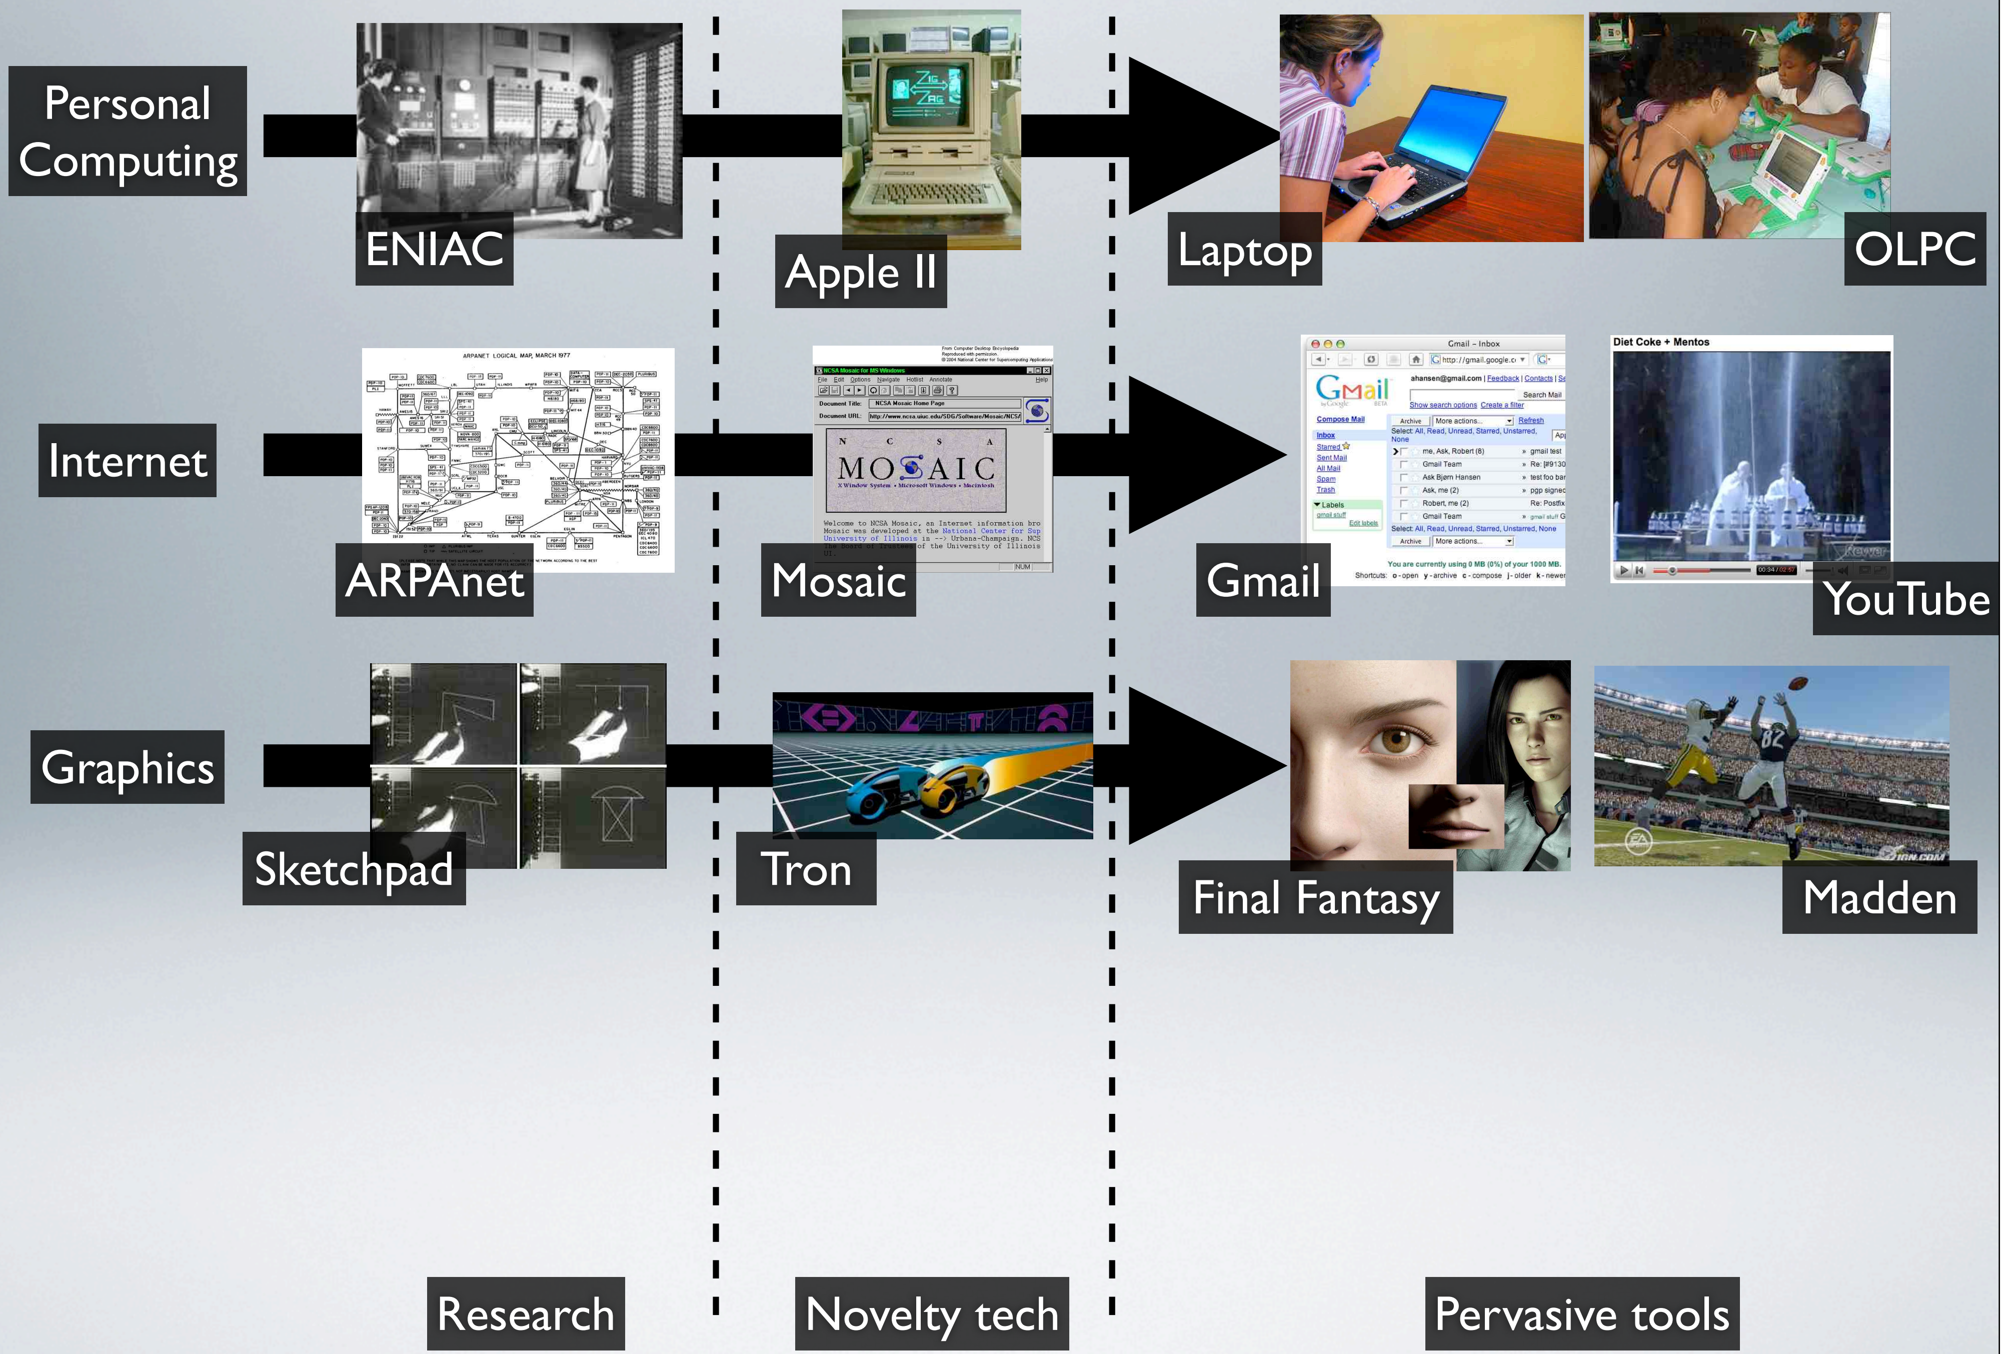
Personal Computing
ENIAC
Apple II
Laptop
OLPC
Internet
ARPAnet
Mosaic
Gmail
YouTube
Graphics
Sketchpad
Tron
Final Fantasy
Madden
Research
Novelty tech
Pervasive tools

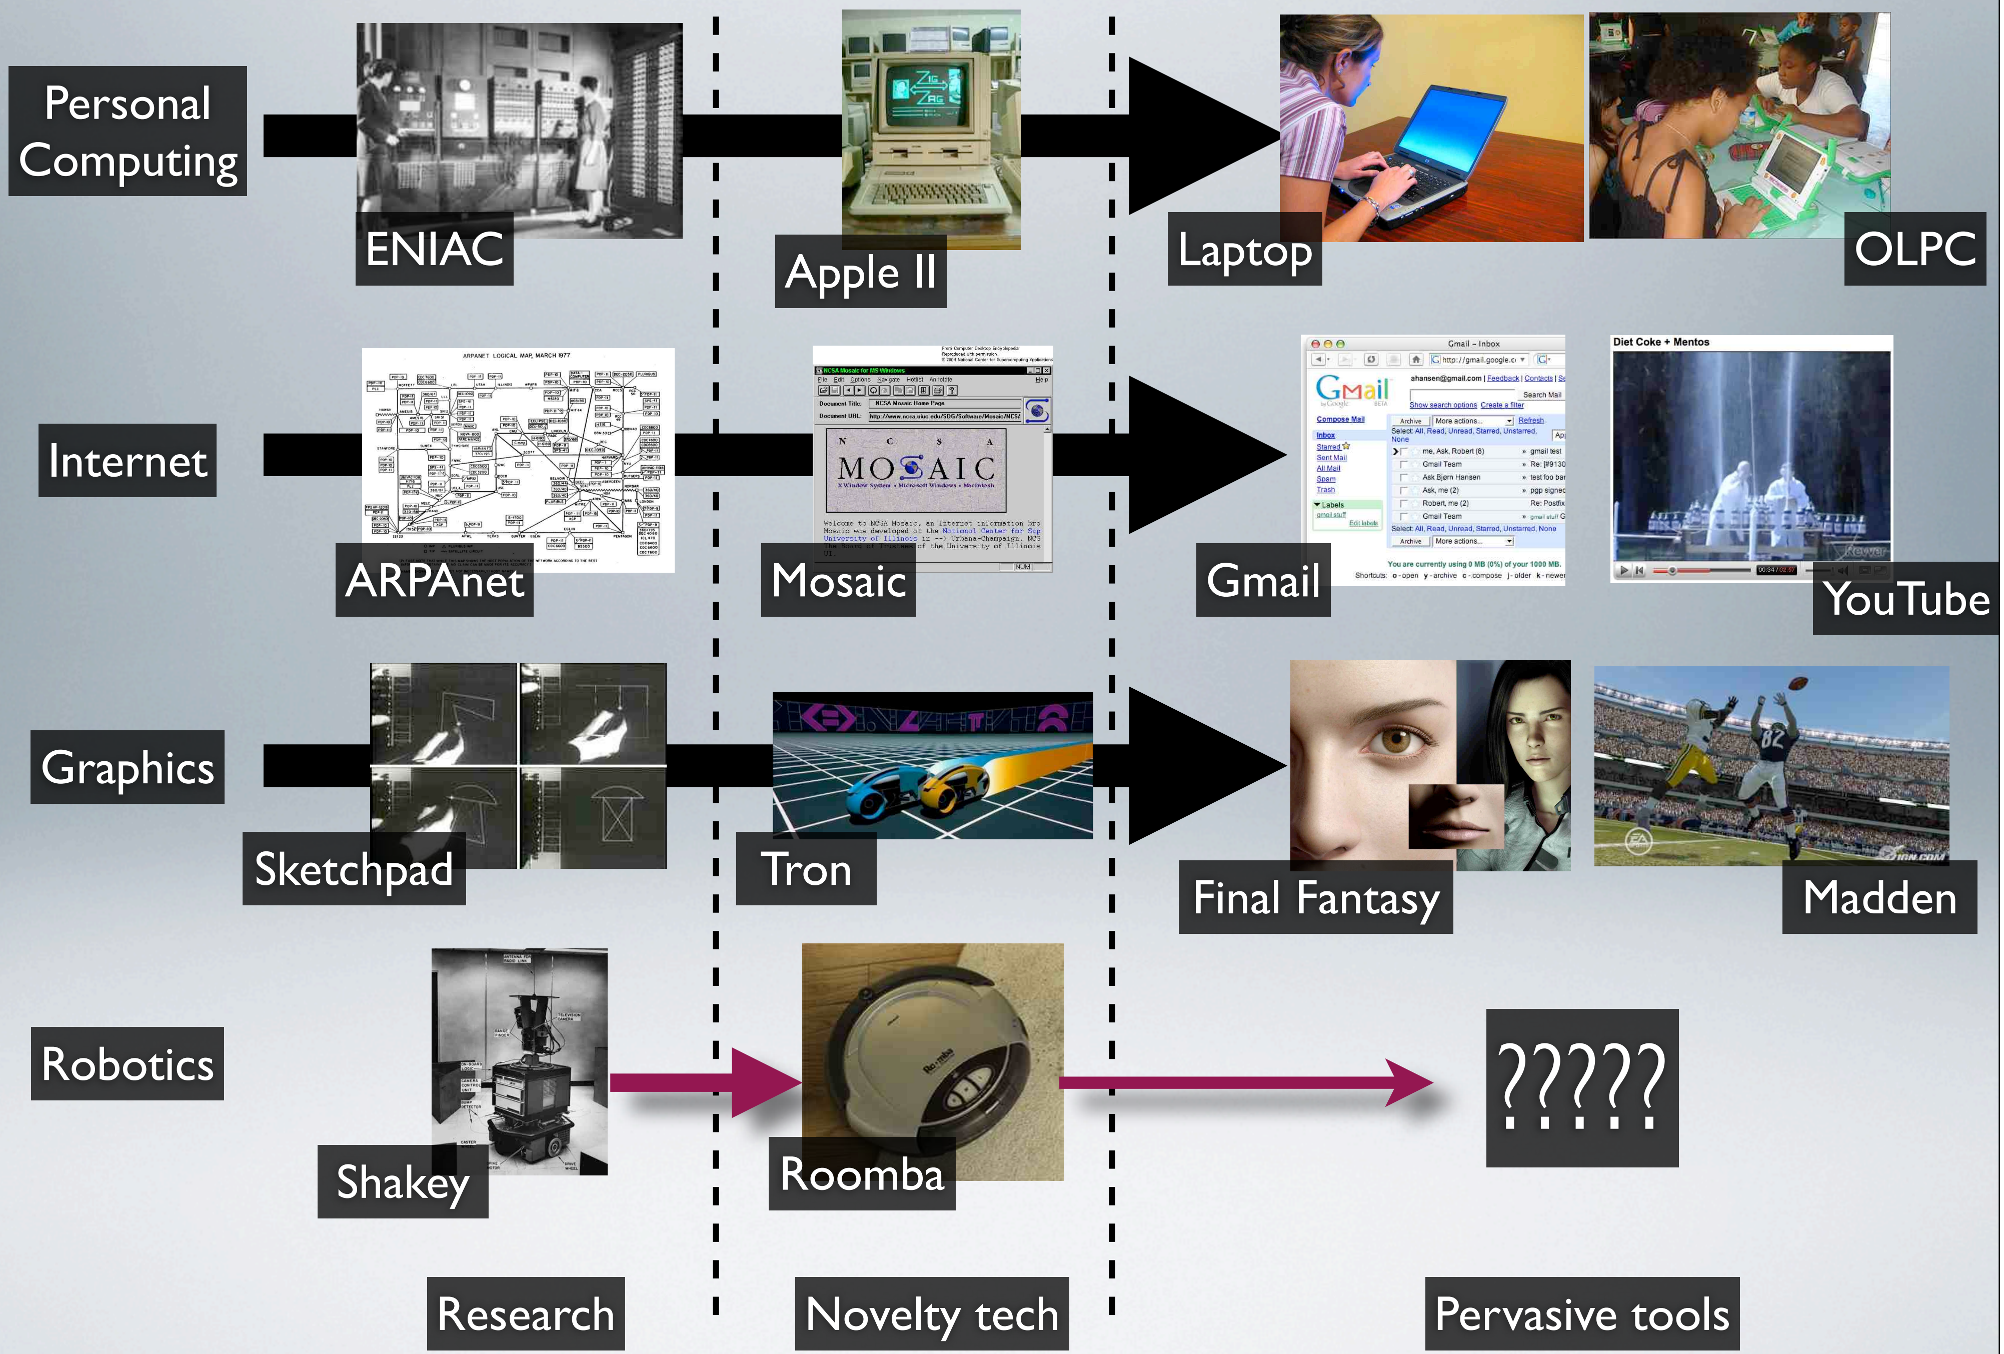
Personal Computing
ENIAC
Apple II
Laptop
OLPC
Internet
ARPAnet
Mosaic
Gmail
YouTube
Graphics
Sketchpad
Tron
Final Fantasy
Madden
Robotics
Shakey
Roomba
?????
Research
Novelty tech
Pervasive tools

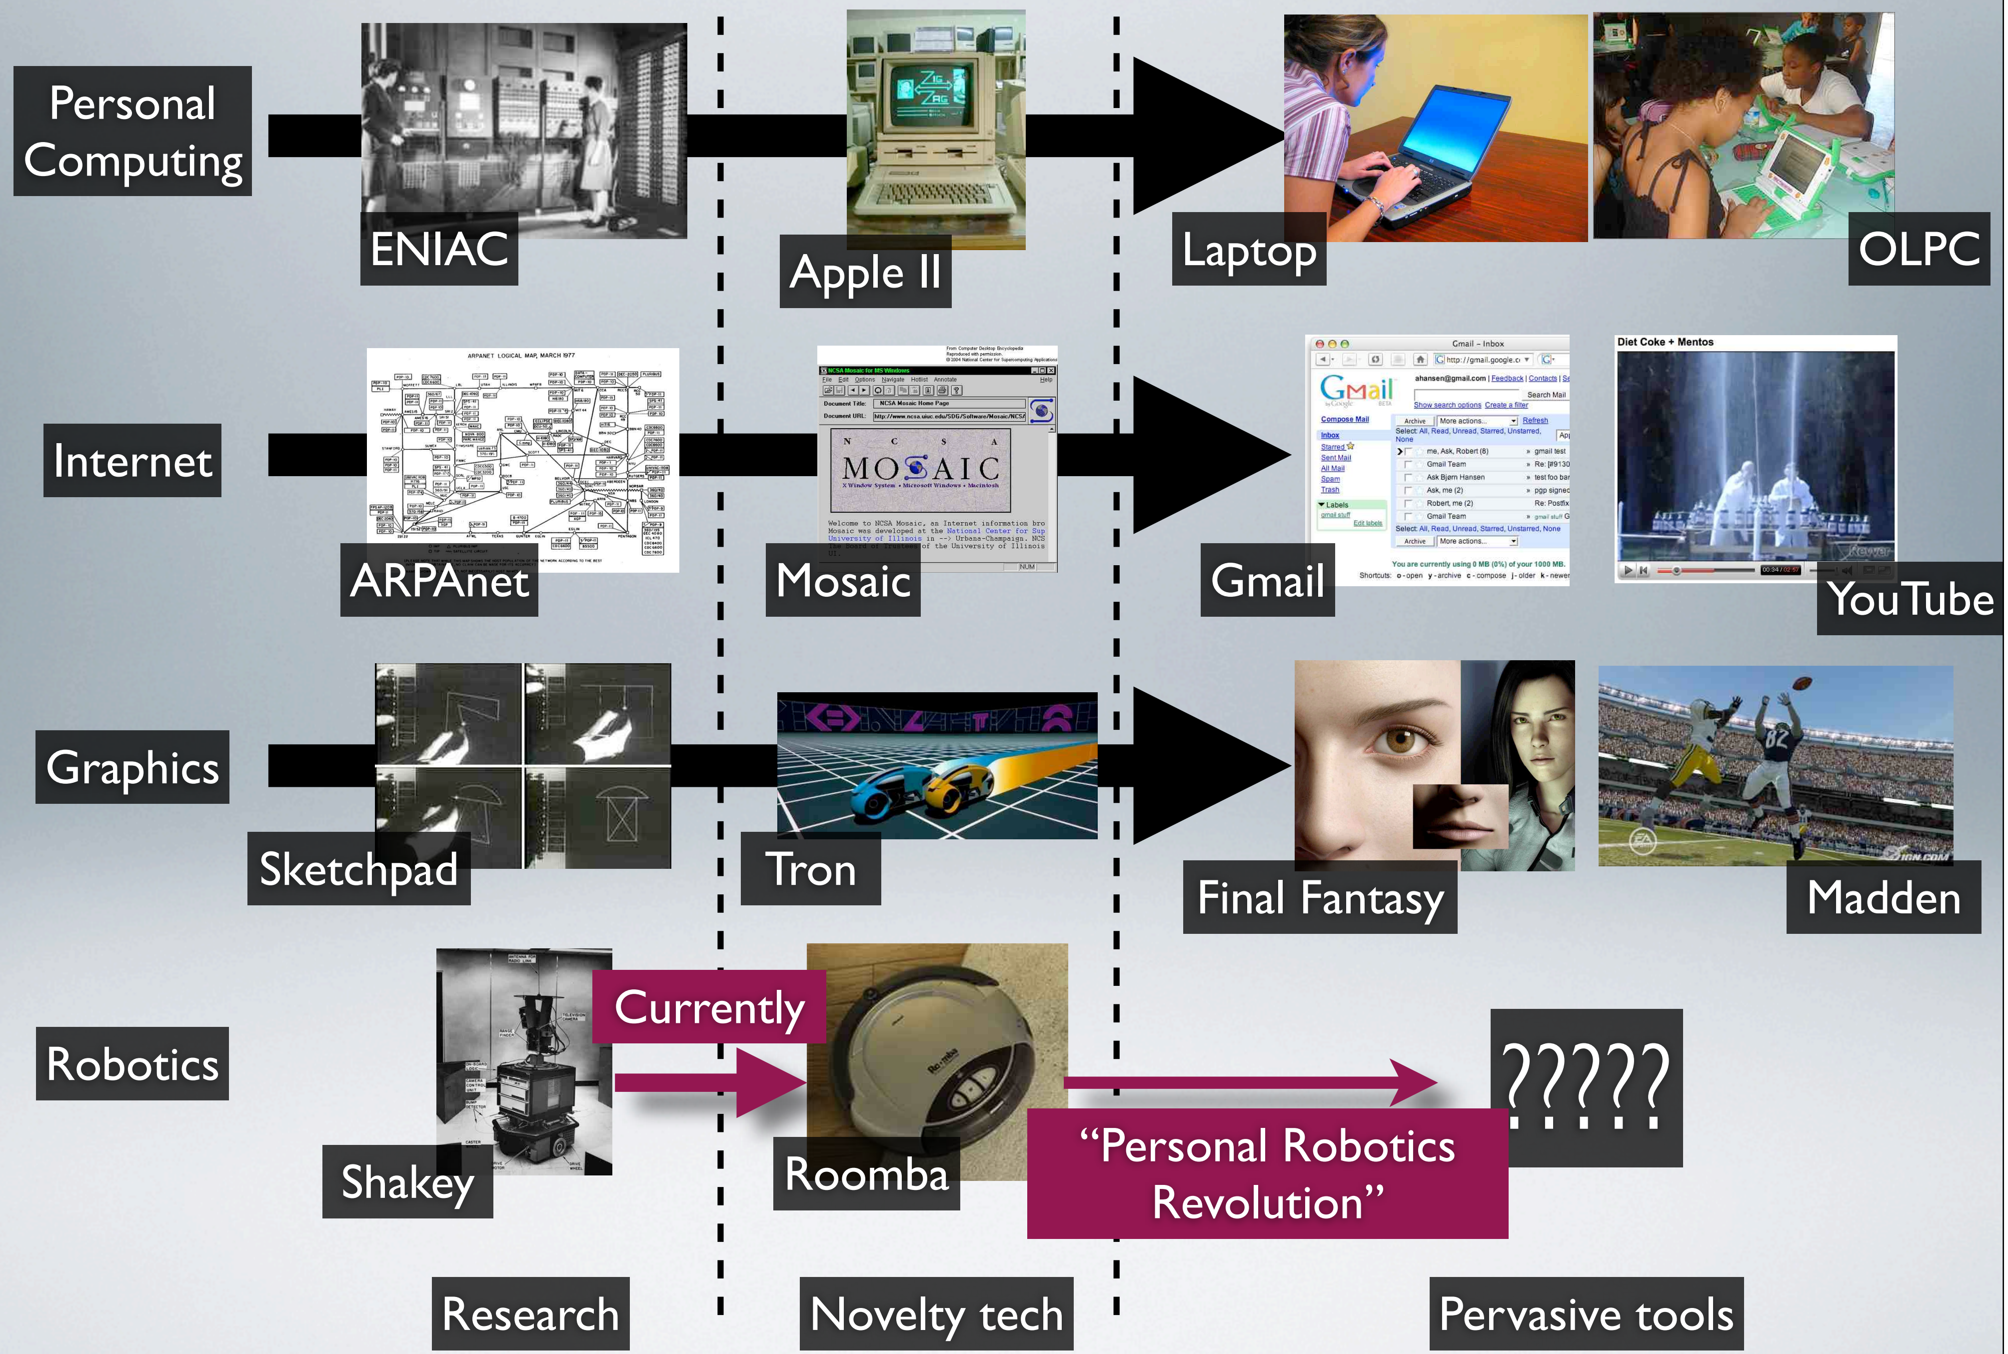
Personal Computing
ENIAC
Apple II
Laptop
OLPC
Internet
ARPAnet
Mosaic
Gmail
YouTube
Graphics
Sketchpad
Tron
Final Fantasy
Madden
Robotics
Shakey
Currently
Roomba
"Personal Robotics Revolution"
?????
Research
Novelty tech
Pervasive tools

DISTINCT CHALLENGES

Other exponentials predicated on deterministic manipulation of state

Enables "write local, run global" development

Variance and uncertainty in tasks, users, and environments limits this model for robotics

WHY ROBOT LEARNING?

When does learning make sense compared to teleop or manual programming?

- Discovery of controllers difficult to phrase analytically
- Enabling non-technical users to express robot controllers

WHY ROBOT LEARNING?

Either way, expression of computing required: FSMs, MDPs, objective functions, likelihoods etc.

- Discovery of controllers difficult to phrase analytically

Trained users fluent in expressing models of computing

- Enabling non-technical users to express robot controllers

Non-technical users might not gain such programming fluency

BROADER VIEW

Casted in FSMs, learn as a whole:

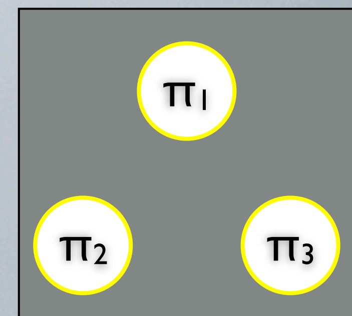

1) Policies for states/primitives

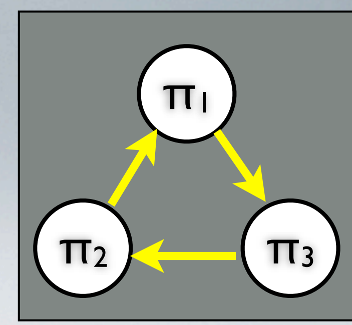

2) Transitions between states

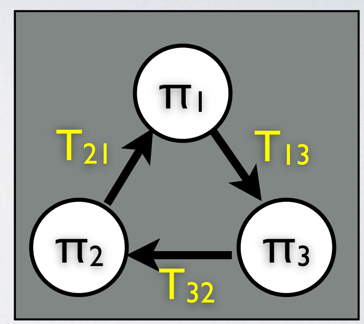

3) State pre/postconditions

INFLAMMATORY STATEMENT:

Computational models learned for robots are significantly more limited than handcoded models b/c learning focuses on individual issues above

BROADER VIEW

Casted in FSMs, learn as a whole:

1) Policies for states/primitives

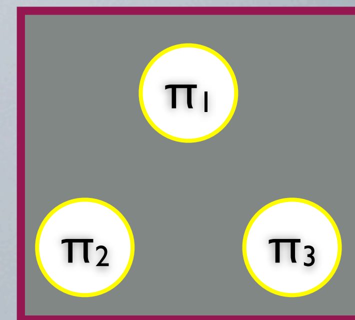

Our use of pairwise kernels to learning primitives from human demonstration

2) Transitions between states

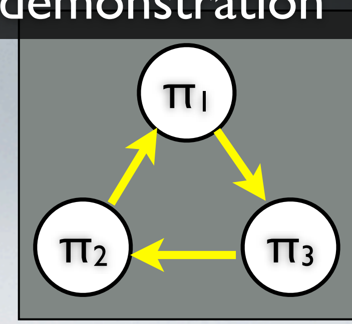

3) State pre/postconditions

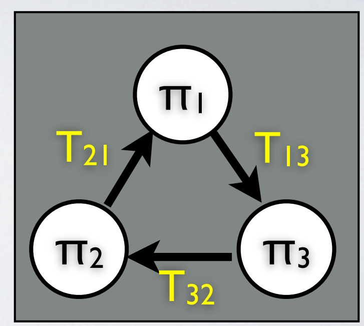

BEGINNINGS: ROBOT IMITATION

[Fod, Mataric, Jenkins 2002]
[Jenkins, Mataric 2004]

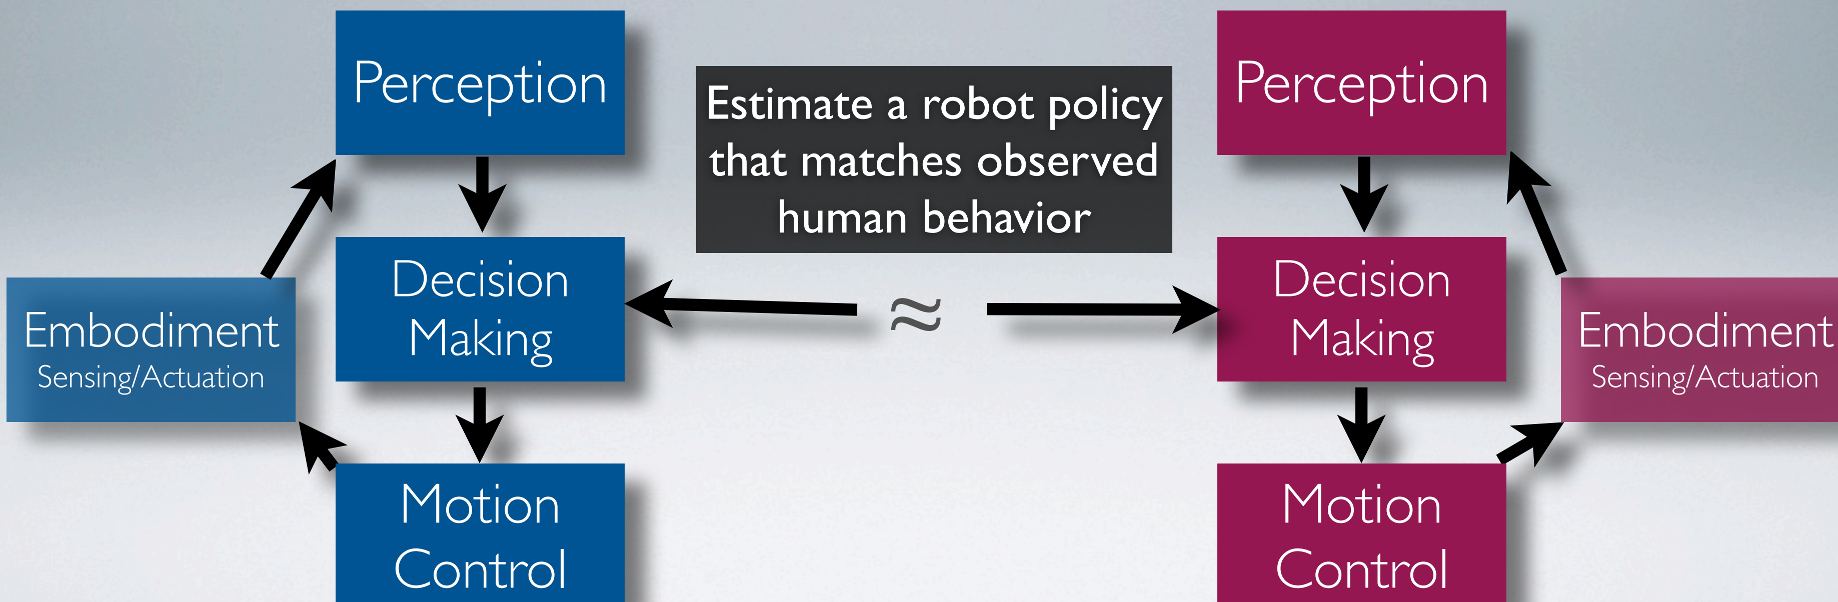

[Fod, Mataric, Jenkins 2002]
[Jenkins, Mataric 2004]

BEGINNINGS: ROBOT IMITATION

Estimate a robot policy that matches observed human behavior

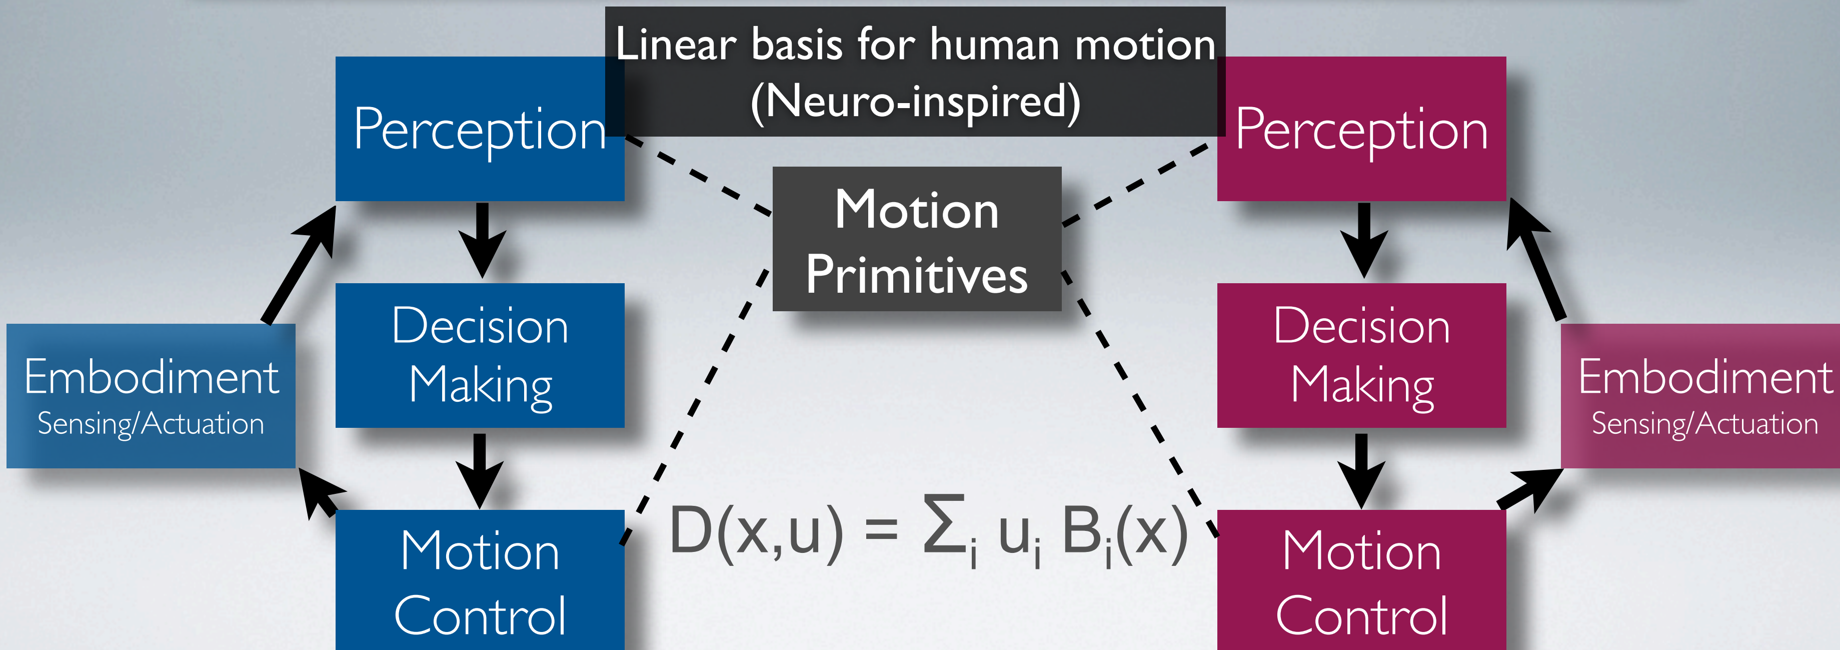

BEGINNINGS: ROBOT IMITATION

[Fod, Mataric, Jenkins 2002]
[Jenkins, Mataric 2004]

Estimate a robot policy that matches observed human behavior

Linear basis for human motion (Neuro-inspired)

Perception

Decision Making

Motion Control

Embodiment
Sensing/Actuation

Motion Primitives

Perception

Decision Making

Motion Control

Embodiment
Sensing/Actuation

Motion Dynamics

$$\underline{D(x,u)} = \Sigma_i \underline{u_i} \underline{B_i(x)}$$

Mirror Neurons

Spinal Fields

[Rizzolatti et al]

[Giszter et al]

[Fod, Mataric, Jenkins 2002]
[Jenkins, Mataric 2004]

BEGINNINGS: ROBOT IMITATION

Estimate a robot policy that matches observed human behavior

Linear basis for human motion
(Neuro-inspired)

Perception

Decision Making

Motion Control

Embodiment
Sensing/Actuation

Motion Primitives

Motion Dynamics

$$D(x,u) = \Sigma_i u_i B_i(x)$$

Perception

Decision Making

Motion Control

Embodiment
Sensing/Actuation

Classify human motion

Mirror Neurons

[Rizzolatti et al]

Spinal Fields

[Giszter et al]

Predict human motion

[Grollman, Jenkins 09]

LEARNING FSMs FROM DEMONSTRATION?

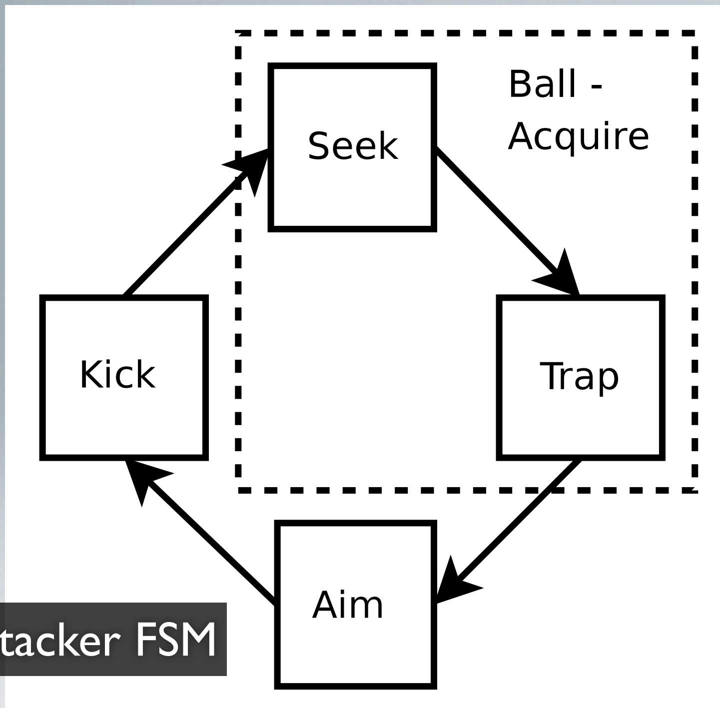

Attacker FSM

Basic robot soccer attack move

PERCEPTUAL ALIASING

Standard attack is 2 overlapping policies

- distinguished by latent context variable

Unimodal attacker is much less efficient

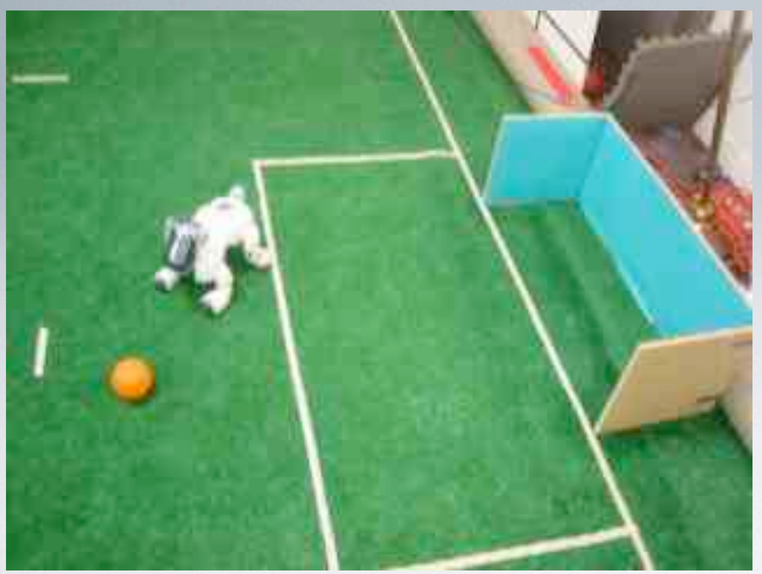

Standard offensive move:
acquire ball, find goal, shoot

Unimodal attacker:
line up ball and goal, then shoot

assume only camera in nose and priopriocpetion

SQUARE ROOT EXAMPLE

- **Consider** y **= sqrt(**x**)**
  - **averaging outputs will be incorrect**
  - **2 regressors needed for pos. and neg.**

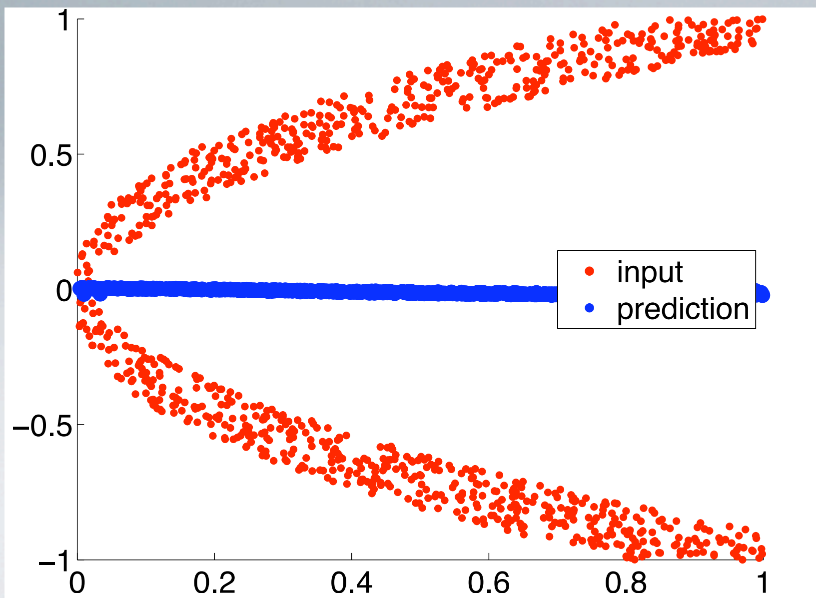

Locally Weighted Projection Regression
or Gaussian Process Regression

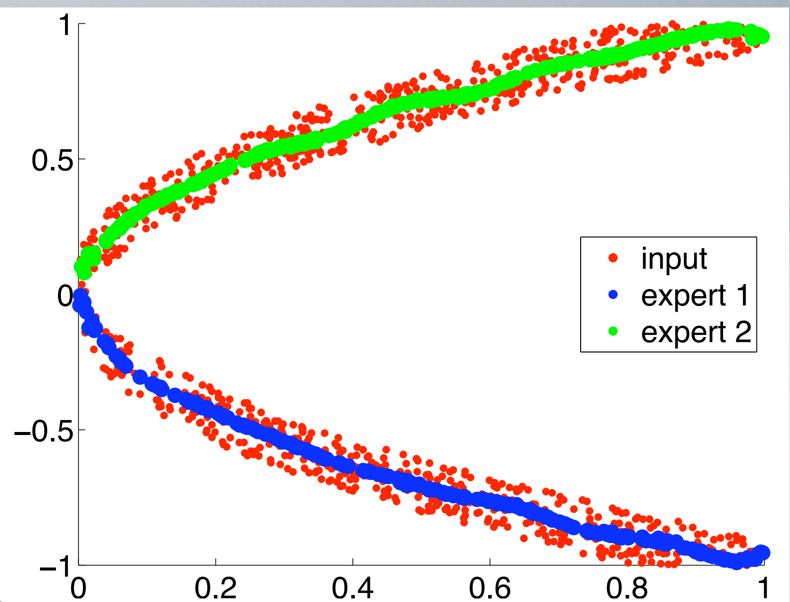

Multimap Regression

[Grollman, Jenkins 09]

INFINITE MIXTURES OF EXPERTS

Infer $\pi : X \rightarrow Y$

Given $(x_i, y_i)_{i=1..t}$

cluster inputs into models

predict output given input

$$p(X, Y, Z) \propto p(Z)p(X|Z)p(Y|X, Z)$$

Z: space of mixture models

prior over models

mixture model

regressor for each model

INFINITE MIXTURES OF EXPERTS

Infer $\pi : X \rightarrow Y$

Given $(x_i, y_i)_{i=1..t}$

User demonstration

$$p(X, Y, Z) \propto \underbrace{p(Z)}_{\text{prior over models}} \underbrace{p(X|Z)}_{\text{mixture model}} \underbrace{p(Y|X, Z)}_{\text{regressor for each model}}$$

- cluster inputs into models
- predict output given input
- Z: space of mixture models
- prior over models
- mixture model
- regressor for each model

[Grollman, Jenkins 09]

INFINITE MIXTURES OF EXPERTS

Infer $\pi : X \rightarrow Y$

Given $(x_i, y_i)_{i=1..t}$

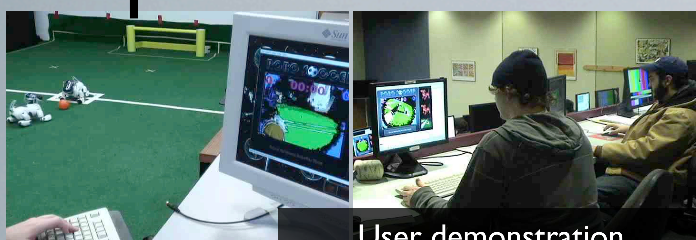
User demonstration

cluster inputs into models

predict output given input

$$p(X,Y,Z) \propto p(Z)p(X|Z)p(Y|X,Z)$$

Z: space of mixture models

prior over models

mixture model

regressor for each model

Sparse (Pairwise) Gaussian Process Regression

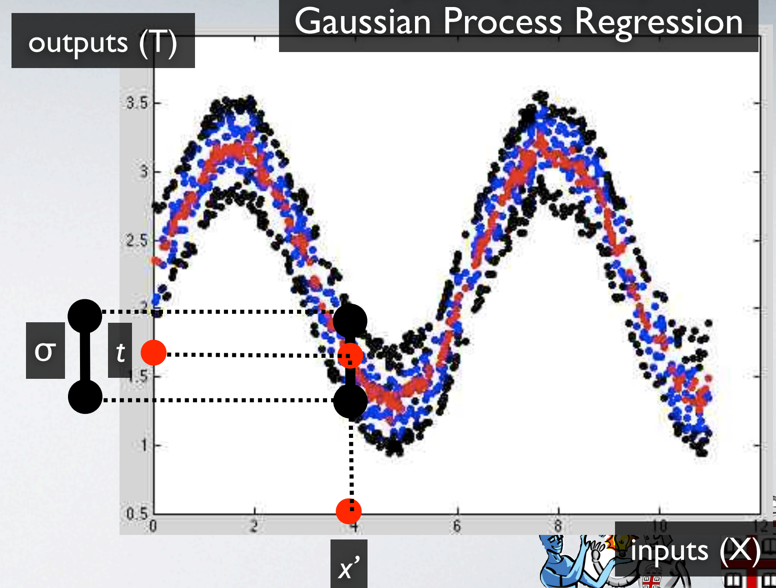

[Grollman, Jenkins 09]

INFINITE MIXTURES OF EXPERTS

Infer $\pi : X \rightarrow Y$

Given $(x_i, y_i)_{i=1..t}$

cluster inputs into models

predict output given input

$$p(X, Y, Z) \propto p(Z)p(X|Z)p(Y|X, Z)$$

Z: space of mixture models

prior over models

mixture model

regressor for each model

Gaussian Mixture Model

$$p(X|Z) = \prod_{k=1}^{K_+} p(x_k|z = k) = \prod_{k=1}^{K_+} p(x_k|\mu_k, \Sigma_k)$$

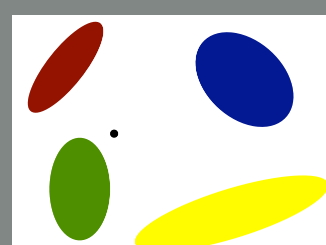

Sparse (Pairwise) Gaussian Process Regression

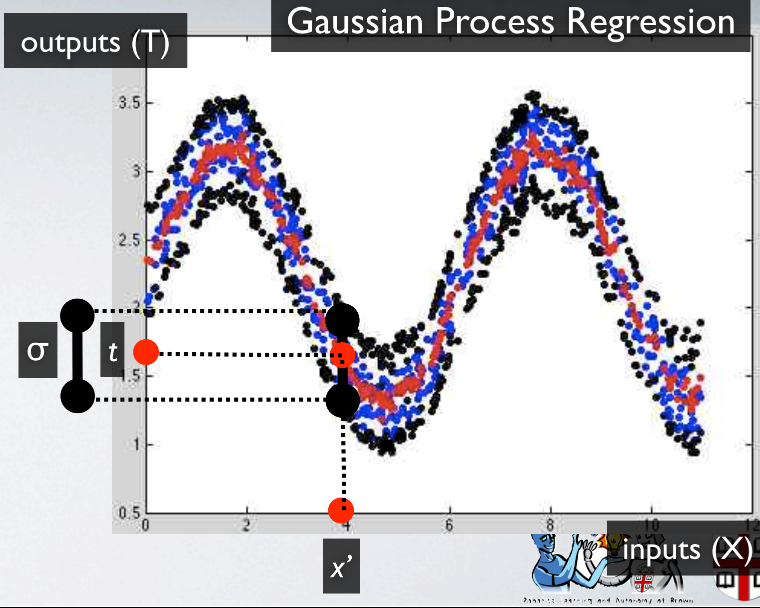

[Grollman, Jenkins 09]

INFINITE MIXTURES OF EXPERTS

Infer $\pi : X \rightarrow Y$

Given $(x_i, y_i)_{i=1..t}$

User demonstration

cluster inputs into models

predict output given input

$$p(X,Y,Z) \propto p(Z)p(X|Z)p(Y|X,Z)$$

Z: space of mixture models

prior over models

mixture model

regressor for each model

Chinese Restaurant Process

$$P(z_i = k|\mathbf{z}_{-i}) = \begin{cases} \frac{m_k}{N+\alpha-1}, & k \leq K_+ \\ \frac{\alpha}{N+\alpha-1}, & k = K_+ + 1 \end{cases}$$

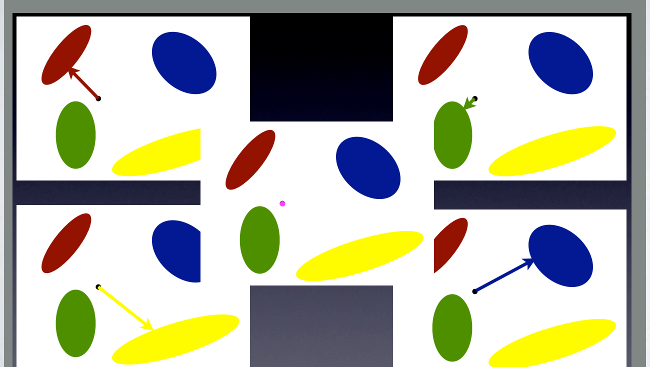

Gaussian Mixture Model

$$p(X|Z) = \prod_{k=1}^{K_+} p(x_k|z=k)$$
$$= \prod_{k=1}^{K_+} p(x_k|\mu_k, \Sigma_k)$$

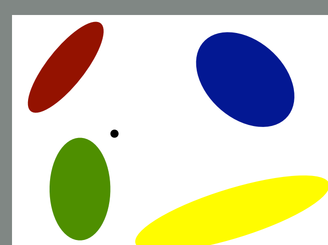

Sparse (Pairwise) Gaussian Process Regression

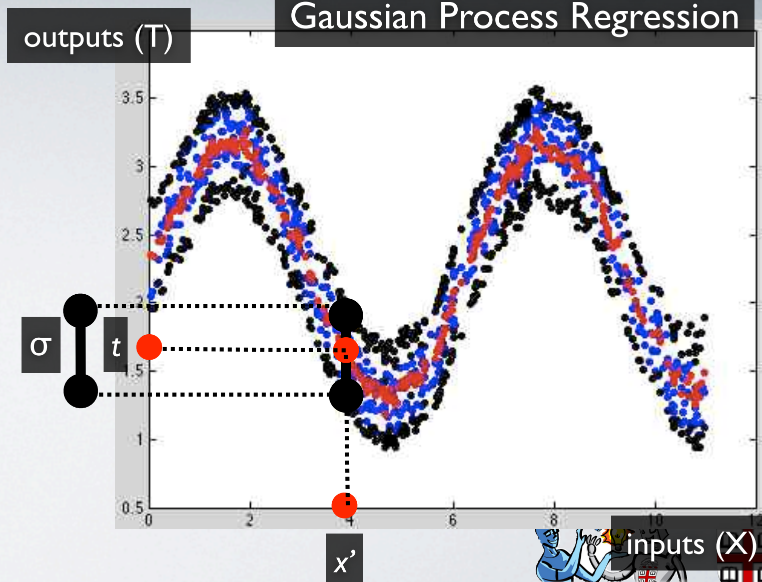

[Grollman, Jenkins 09]

INFINITE MIXTURES OF EXPERTS

Infer $\pi : X \to Y$

Given $(x_i, y_i)_{i=1..t}$

cluster inputs into models

predict output given input

$$p(X,Y,Z) \propto p(Z)p(X|Z)p(Y|X,Z)$$

Z: space of mixture models

prior over models

mixture model

regressor for each model

Sparse incremental inference with particle filter

Chinese Restaurant Process

$$P(z_i = k|\mathbf{z}_{-i}) = \begin{cases} \frac{m_k}{N+\alpha-1}, & k \leq K_+ \\ \frac{\alpha}{N+\alpha-1}, & k = K_+ + 1 \end{cases}$$

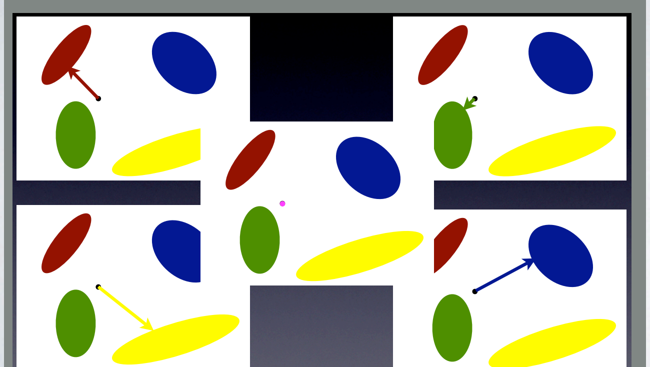

Gaussian Mixture Model

$$p(X|Z) = \prod_{k=1}^{K_+} p(x_k|z=k)$$

$$= \prod_{k=1}^{K_+} p(x_k|\mu_k, \Sigma_k)$$

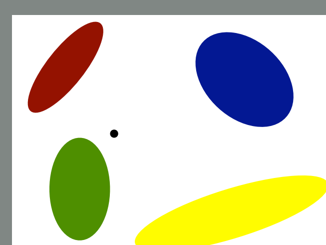

Sparse (Pairwise) Gaussian Process Regression

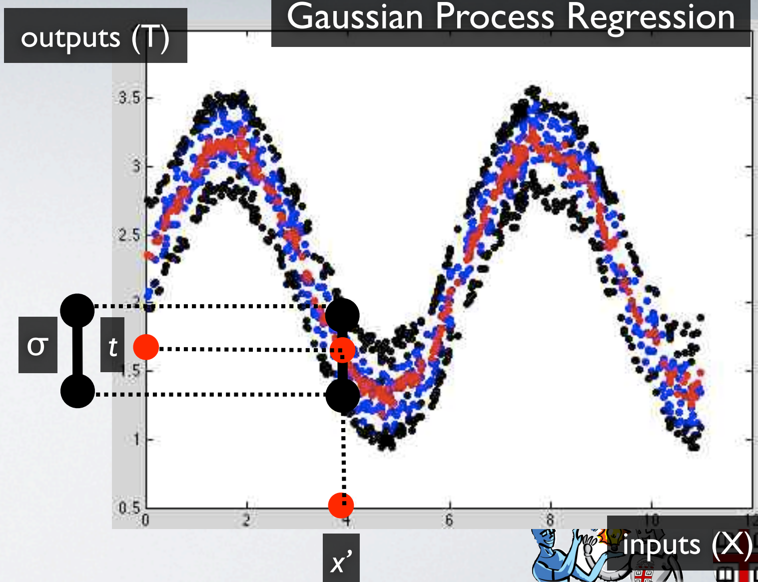

LEARNED GOAL SCORER

OVERVIEW

Why robot learning?

Learning from demonstration

- Human motion primitives through dimension reduction
- Decision making primitives through infinite mixtures of experts

Learning → the path of least resistance?

Group Project

- Massive-scale learning from demonstration
- Implement in ROS; do Create tutorial:
    - http://code.google.com/p/brown-ros-pkg/
- Learn three tasks from demonstration
    - Create robot soccer
    - Nao magneto assembly
    - PR2 intern challenge
- One learning alg, one infrastructure box
- Human subjects study

Create CS148 Task

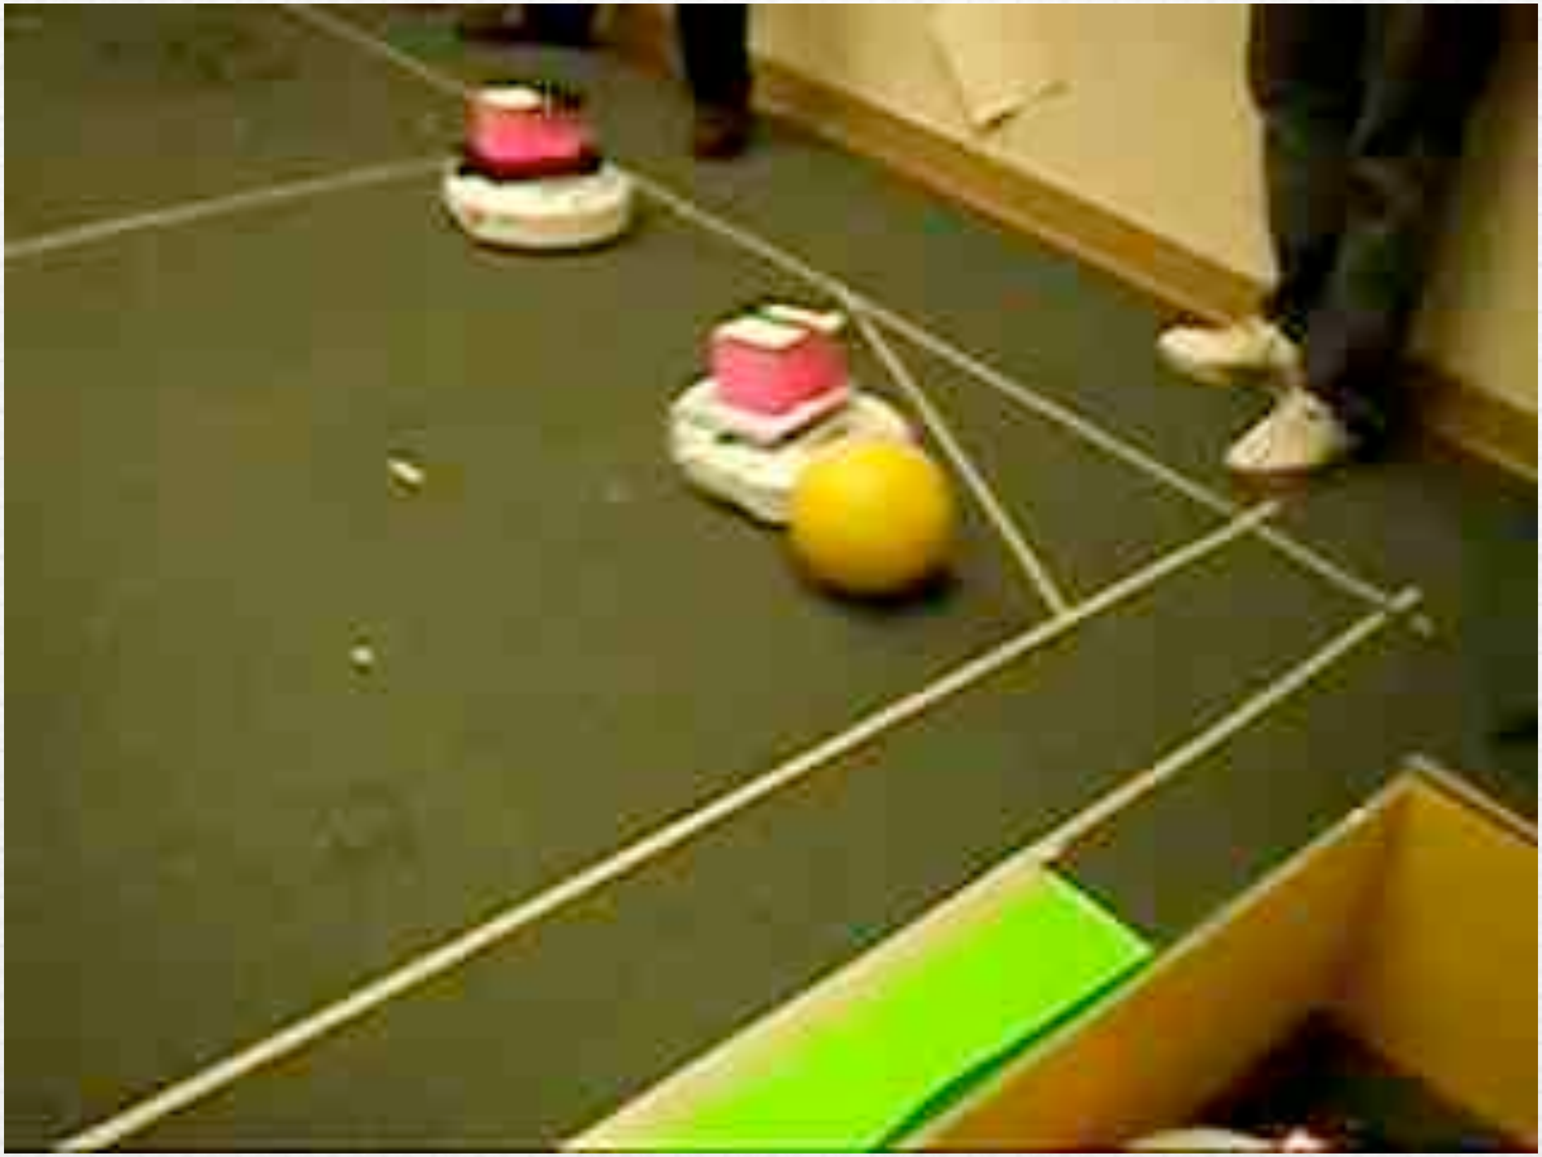

Robot Soccer

Create: Current Status

single objective possible, but

multiple objectives remains problem

Goal scoring with regression

PR2 Intern Challenge

Serving Drinks

PR2: Current Status

PR2 simulator (Gazebo)
running on maria/rlab

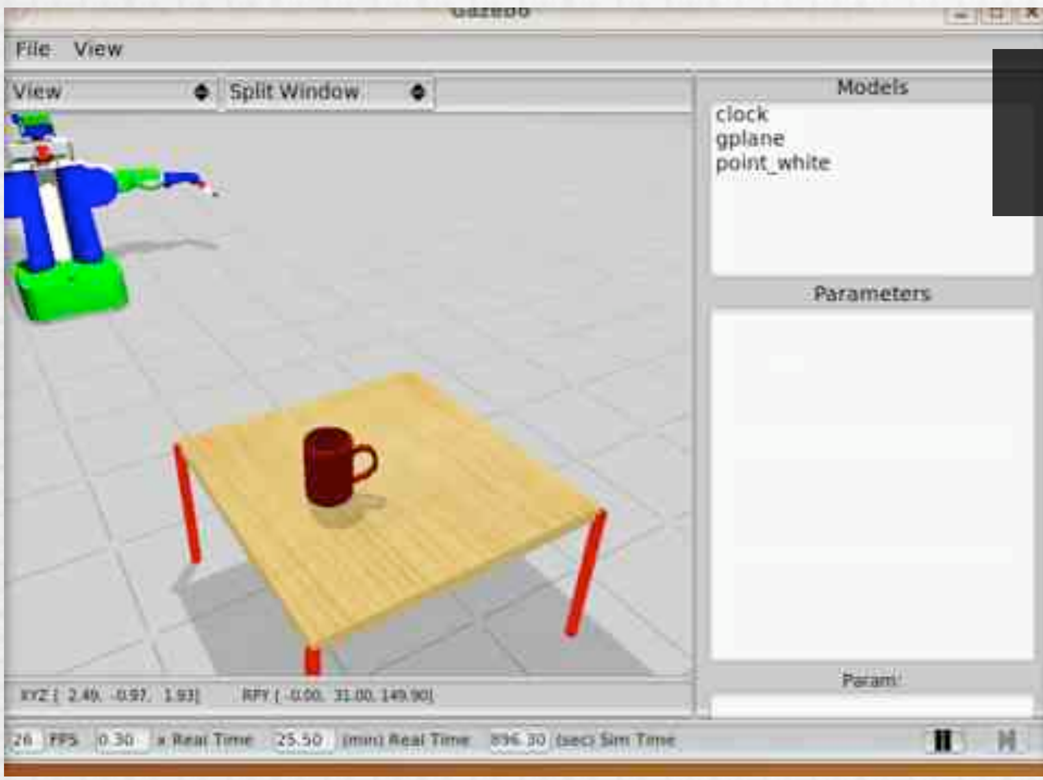

issues: experimental setup,
code interface

Getting a PR2 at Brown!
Applying for PR2 Beta Program

Nao: Current Status

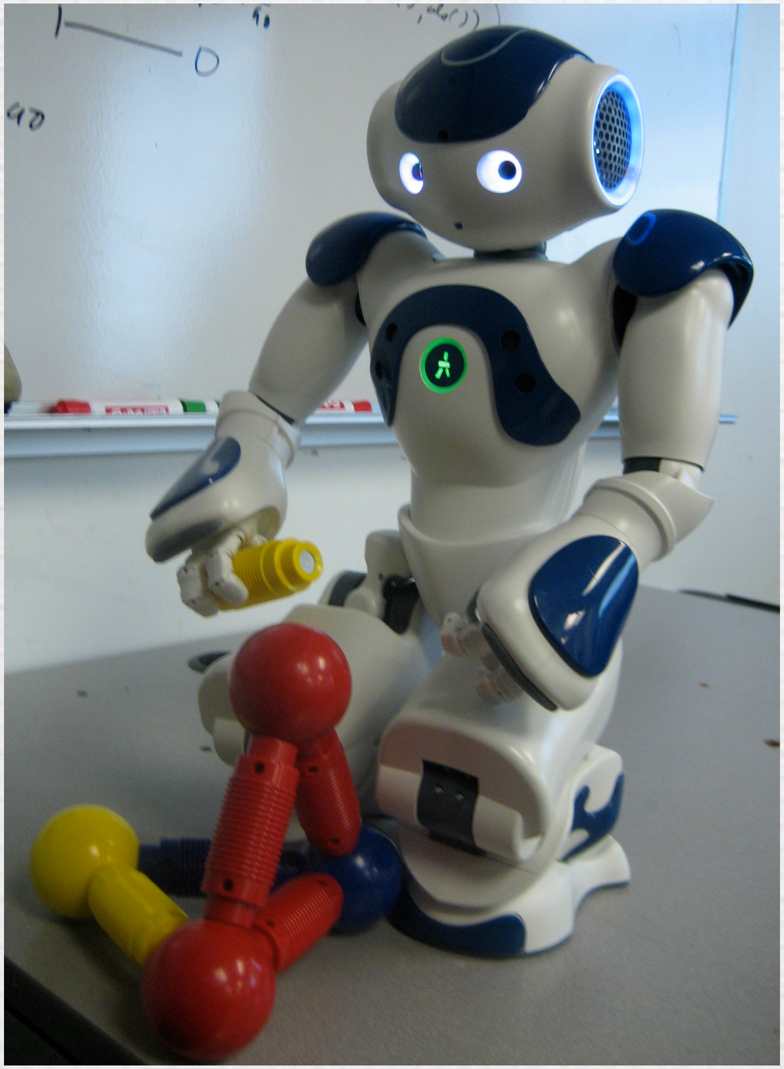

issues: teleoperation interface, object recognition

Create: Current Status

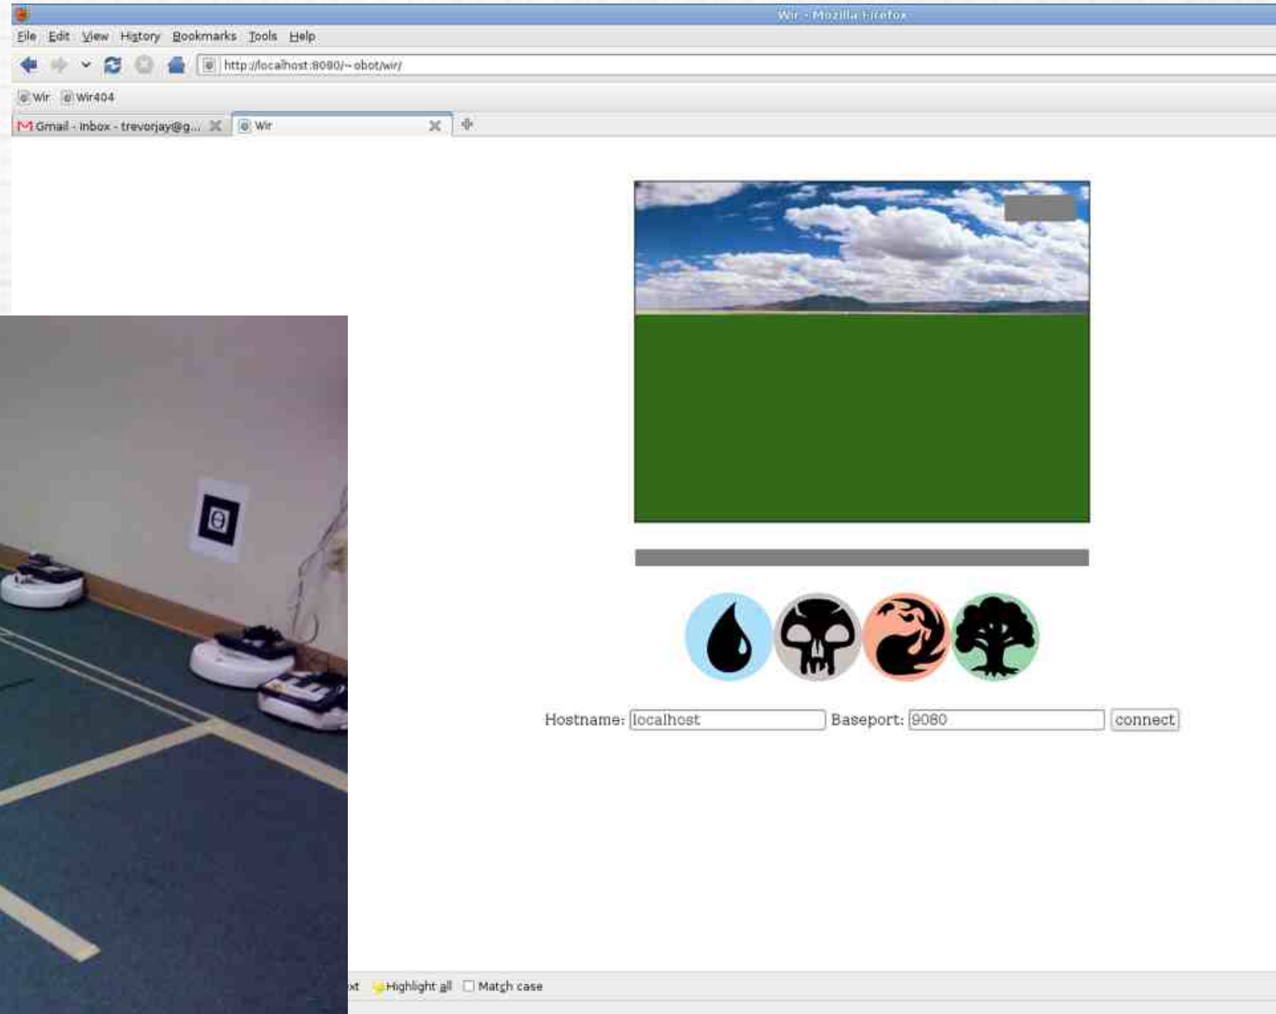

issues: arintegration, localization

Teleop interface with ARtags