Hot Science Cool Talks

UT Environmental Science Institute

68

Autonomous Robots Playing Soccer and Traversing Intersections

Dr. Peter Stone

October 15, 2010

A Goal of AI

Robust, **fully autonomous**
agents in the real world

How?

- Build complete solutions to relevant challenge tasks
  Complete agents: sense, decide, and act — closed loop
  Challenge tasks: specific, concrete objectives

- Drives research on component algorithms, theory
  - Improve from experience (Machine learning)
  - Interact with other agents (Multiagent systems)

- A top-down, empirical approach

  "Good problems . . . produce good science" [Cohen, '04]

Good Problems Produce Good Science

Manned flight 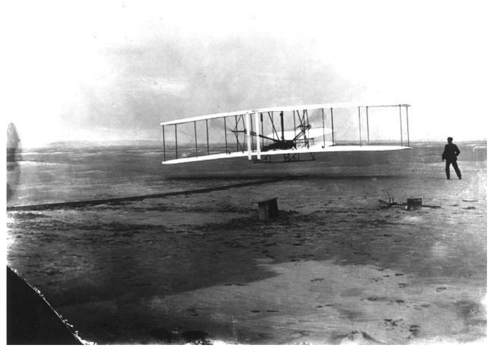

Apollo mission 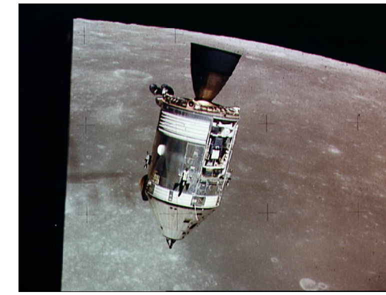

Manhattan project 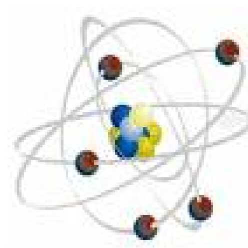

RoboCup soccer 

Goal: By the year 2050, a team of humanoid robots that can beat the human World Cup champion team.
[Kitano, '97]

RoboCup Soccer

- Still in the early stages
- Many virtues:
  - Incremental challenges, closed loop at each stage
  - Relatively easy entry
  - Multiple robots possible
  - Inspiring to many
- Visible progress

Small-sized League
Middle-sized League
Legged Robot League
Simulation League
Humanoid League

The Early Years

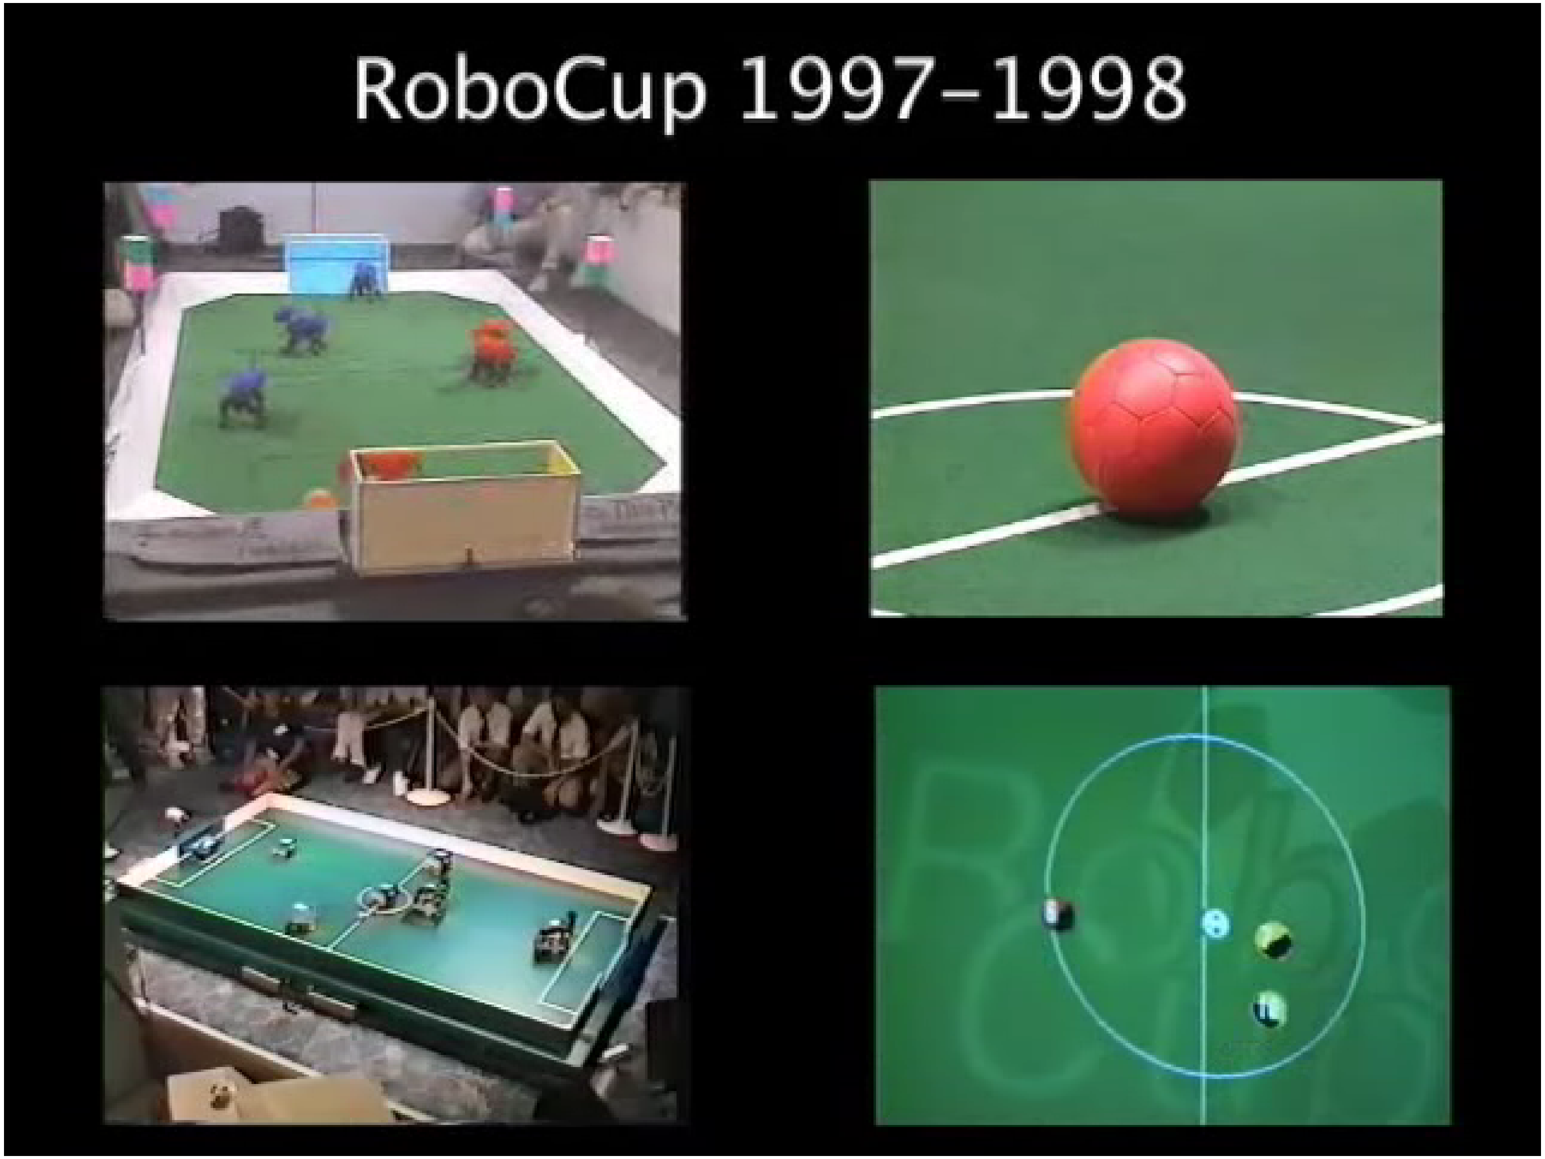

A Decade Later

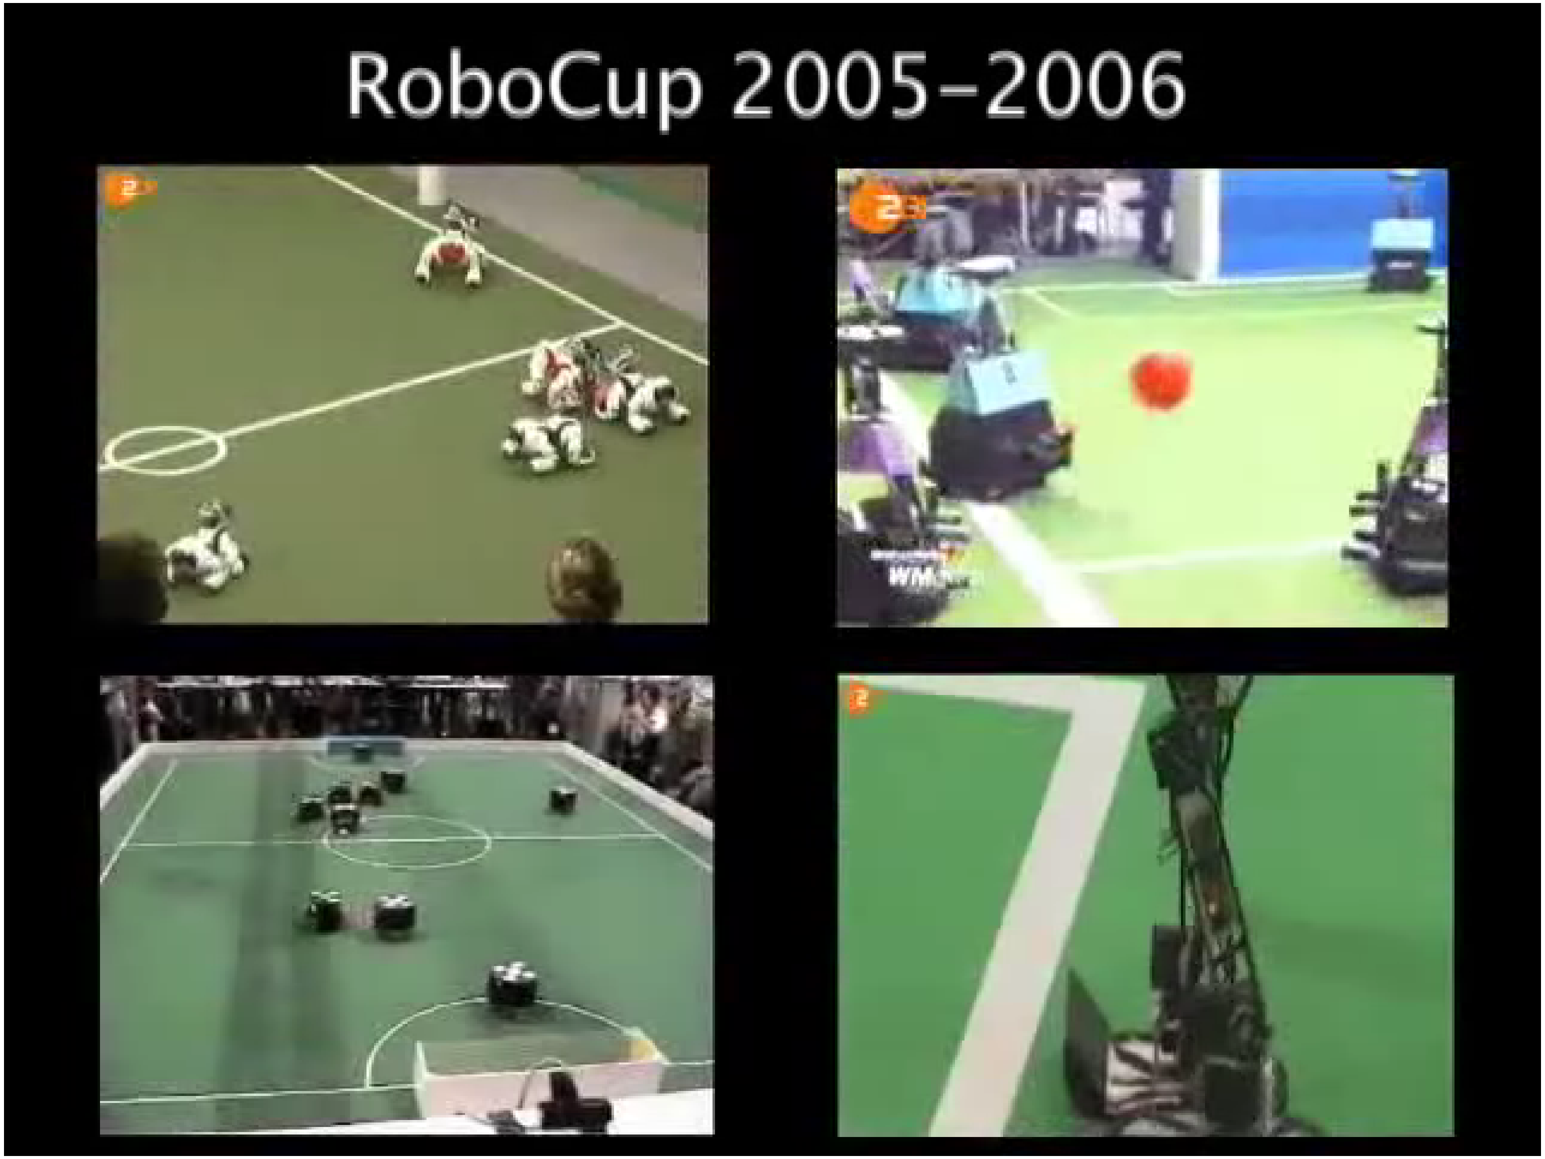

Advances due to RoboCup

- Drives **research** in many areas:
  - Control algorithms; computer vision, sensing; localization;
  - Distributed computing; real-time systems;
  - Knowledge representation; mechanical design;
  - Multiagent systems; machine learning; robotics
- **200+ publications** from simulation league alone
- **200+** from 4-legged league
- **15+** Ph.D. theses

Layered Learning

- For domains too complex for tractably mapping state features $S \longmapsto$ outputs O
- Hierarchical subtask decomposition given: $\{L_1, L_2, \ldots, L_n\}$
- Machine learning: exploit data to train, adapt
- Learning in one layer feeds into next layer

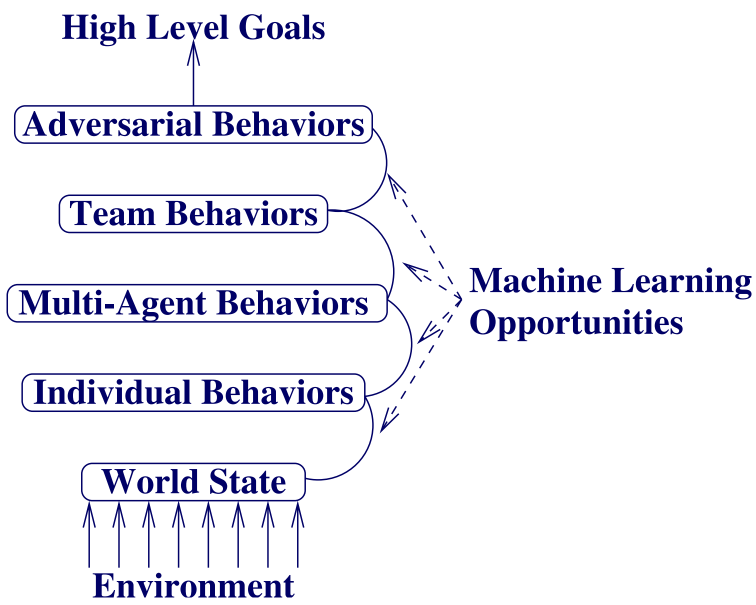

Layered Learning in Practice

First applied in simulated robot soccer [Stone & Veloso, '97]

	Strategic Level	Example
L_1	individual	ball interception
L_2	multiagent	pass evaluation
L_3	team	pass selection

Recently applied on real robots [Stone, Kohl, & Fidelman, '06]

	Strategic Level	Example
L_1	individual	fast walking
L_2	individual	ball control

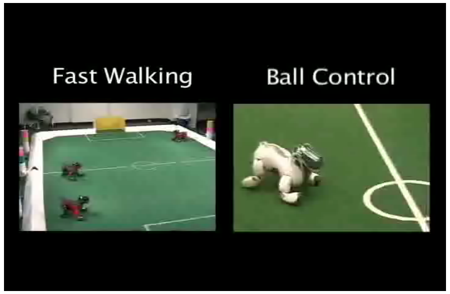

Robot Vision

- Great progress in computer vision
  - Shape modeling, object recognition, face detection. . .
- Robot vision offers new challenges 
  - Mobile camera, limited computation, color features
- **Autonomous color learning** [Sridharan & Stone, '05]
  - Learns color map based on known object locations
  - Recognizes and reacts to illumination changes
  - Object detection in real-time, on-board a robot

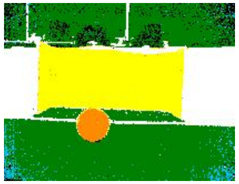

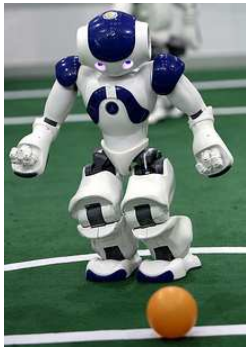

Other Good AI Challenges

Trading agents 

Autonomous vehicles 

Autonomic computing 

Socially assistive robots 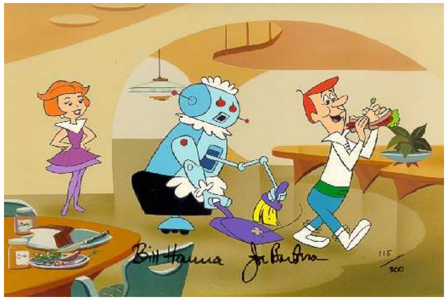

Challenge Problems Drive Research

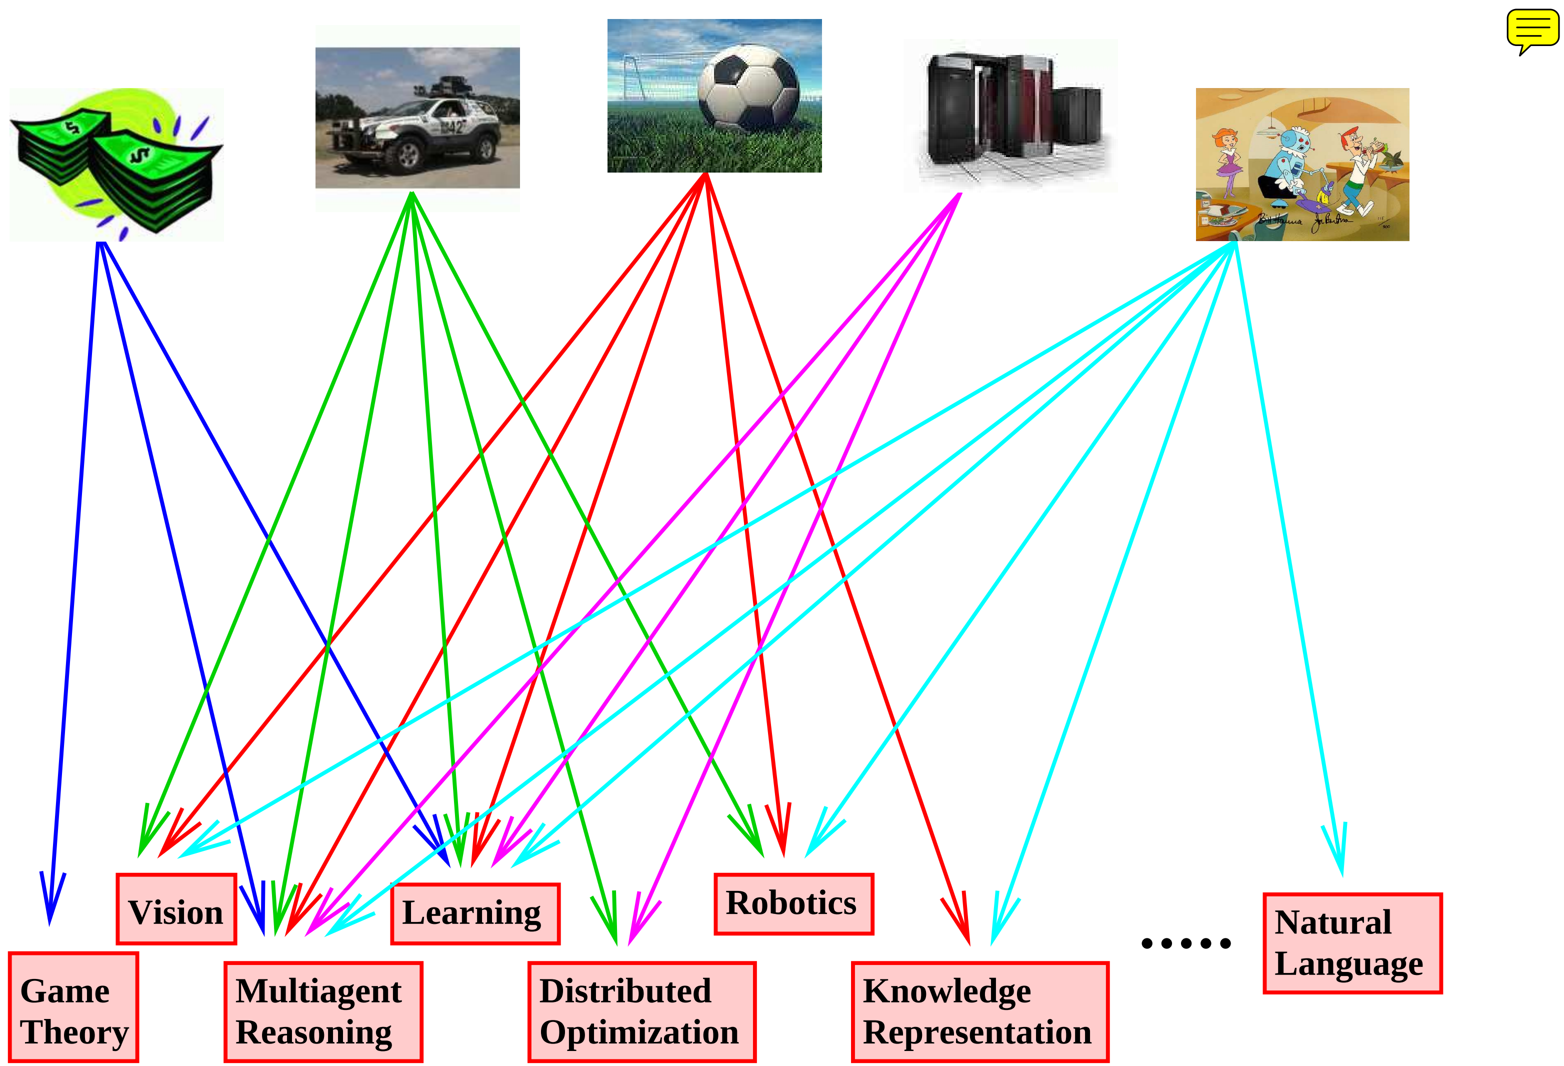

Learning and Multiagent Reasoning

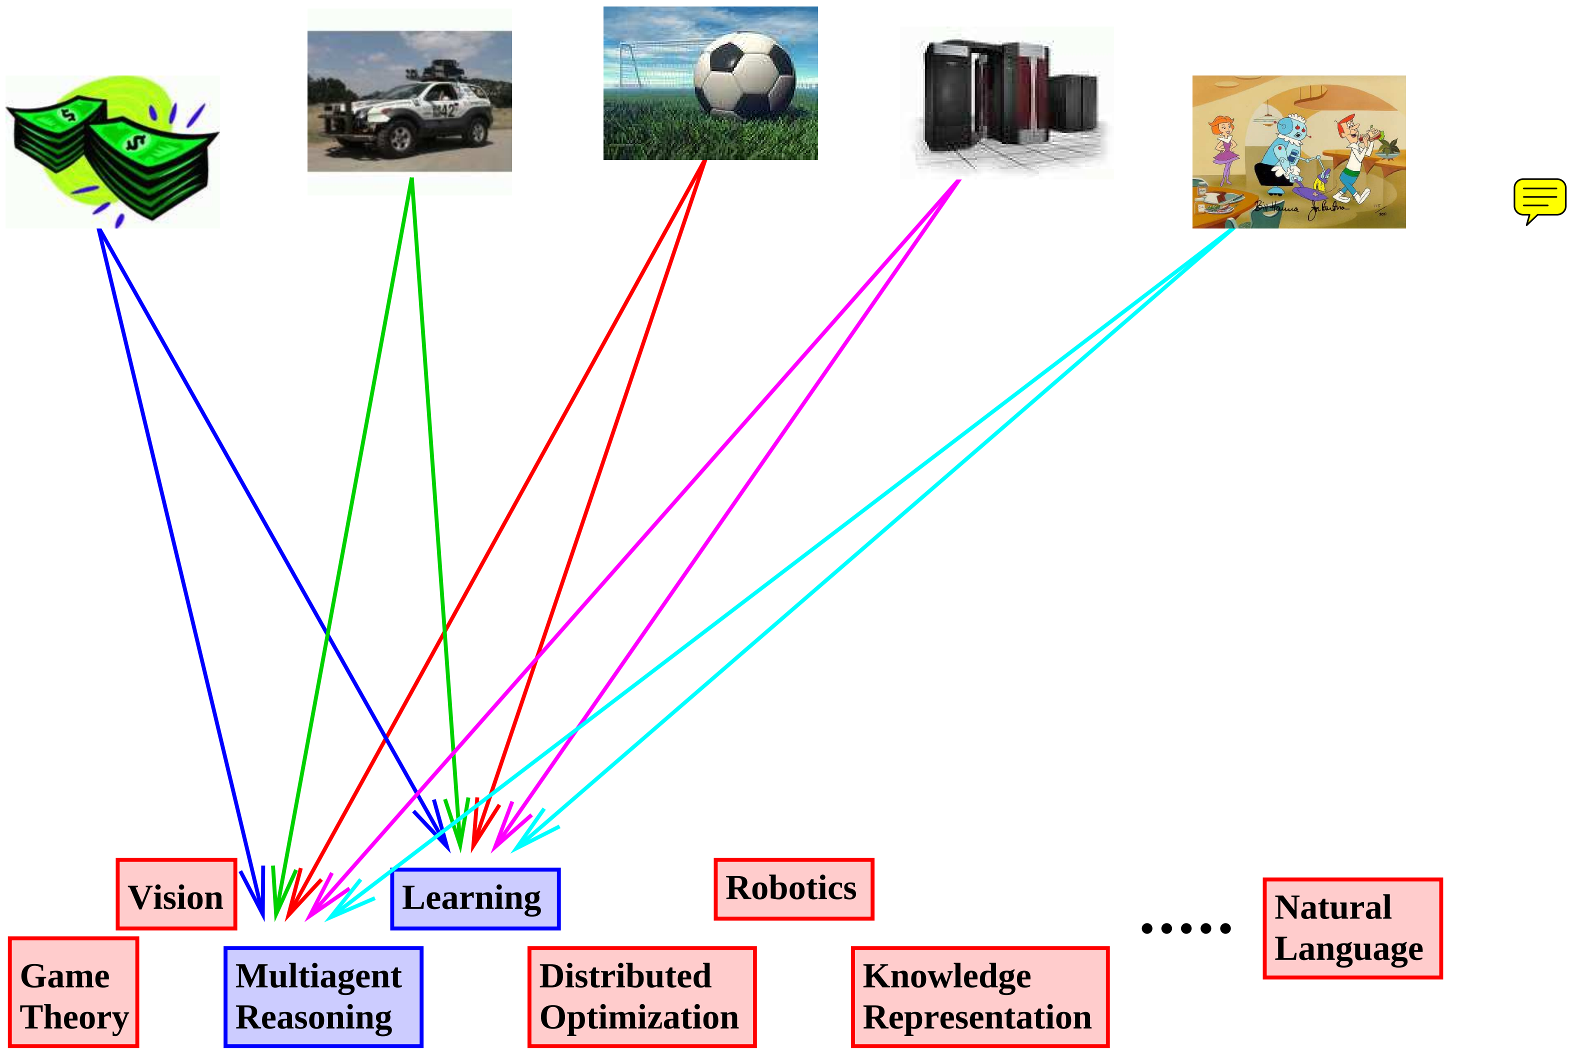

Outline

- Build complete solutions to relevant challenge tasks
  - Drives research on component algorithms, theory
- **Learning Agents**
- Multiagent reasoning

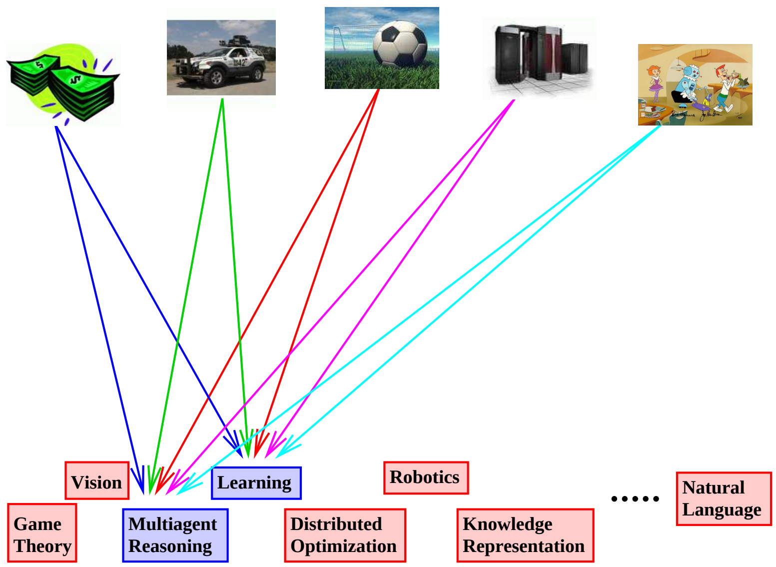

Machine Learning

"...resurgence of interest in machine learning" [Mitchell, '83]

Supervised learning mature [Kaelbling, '97]

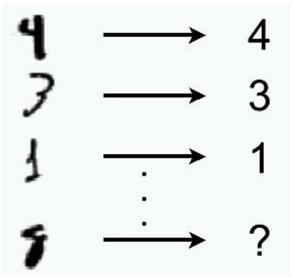

For agents, reinforcement learning most appropriate

– Foundational theoretical results
– Challenge problems require innovations to scale up

RL Theory

Success story: Q-learning converges to π^* [Watkins, 89]

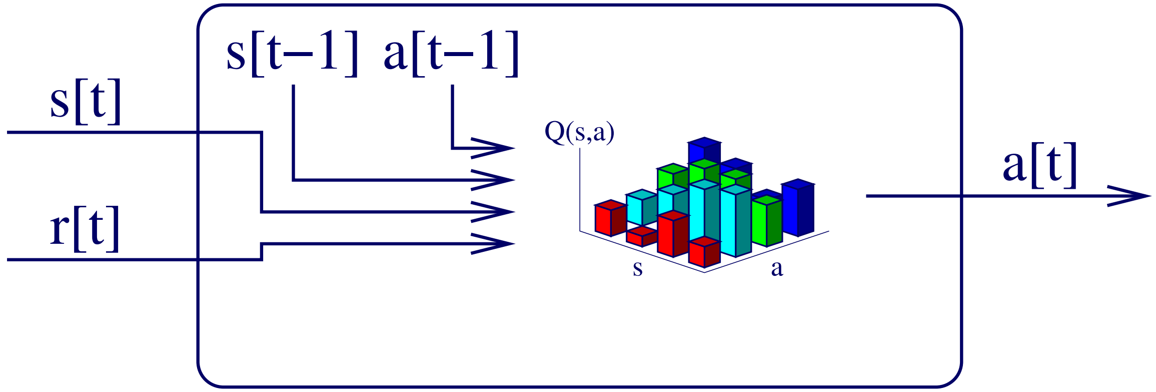

– Table-based representation
– Visit every state infinitely often

Scaling Up

- Backgammon [Tesauro, '94]
- Helicopter control [Ng et al., '03]

- **RoboCup Soccer Keepaway** [Stone & Sutton, '01]
  - Play in a small area (20m $\times$ 20m)
  - Keepers try to keep the ball
  - Takers try to get the ball
  - Performance measure: average possession duration

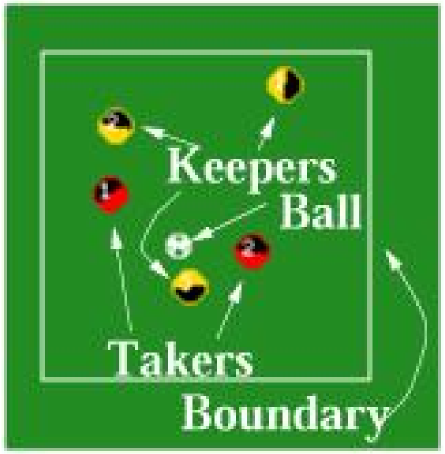

Function Approximation

In practice, visiting every state impossible

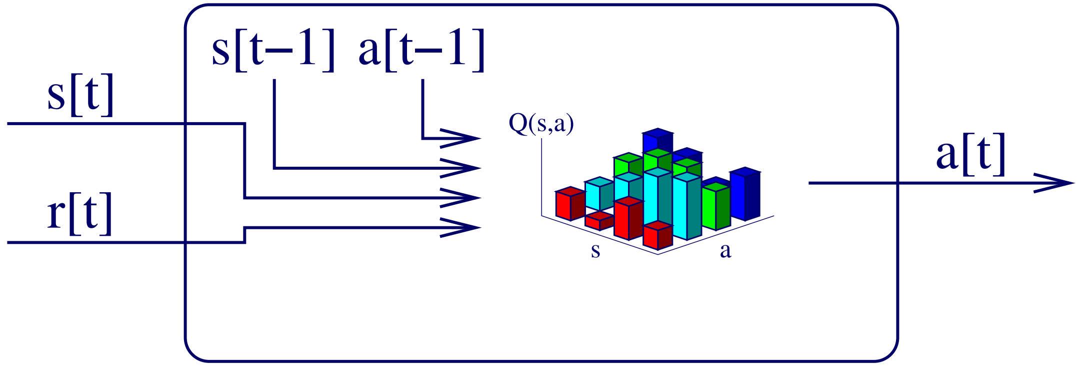

Function approximation of value function

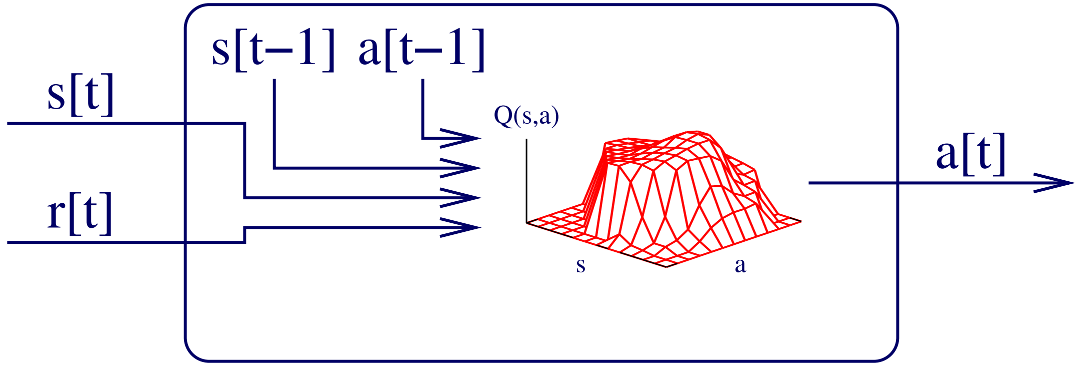

Theoretical guarantees harder to come by

Main Result

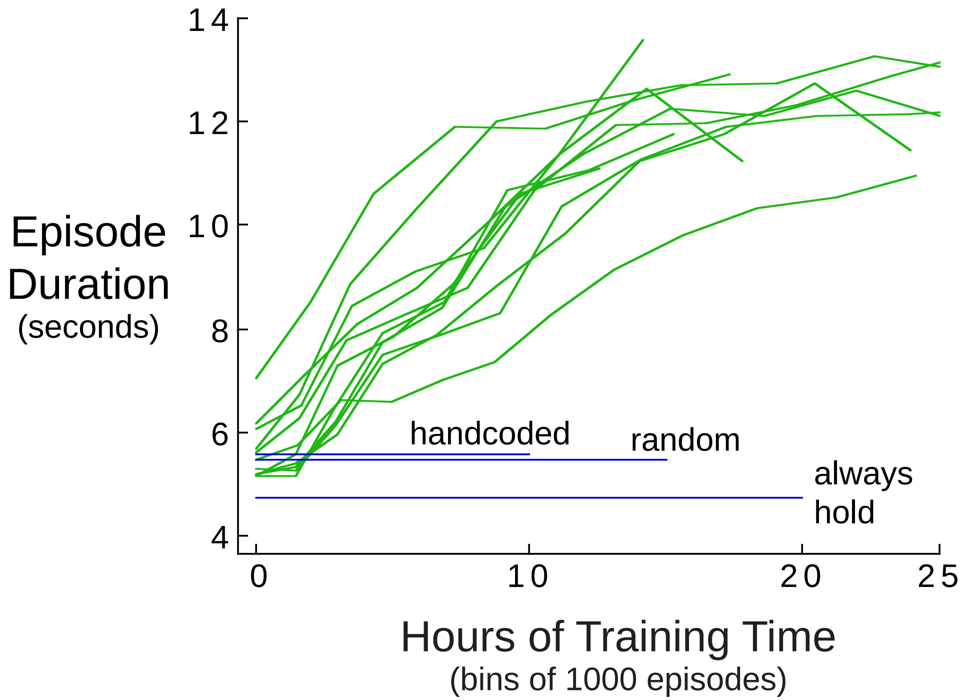

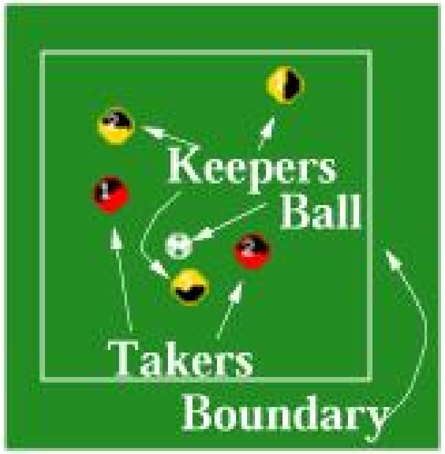

Learning: Distributed SMDP SARSA(λ) with CMACs

– Algorithm modified to enable distributed updates

1 hour = 720 5-second episodes

Batch Methods

In practice, often experience is scarce

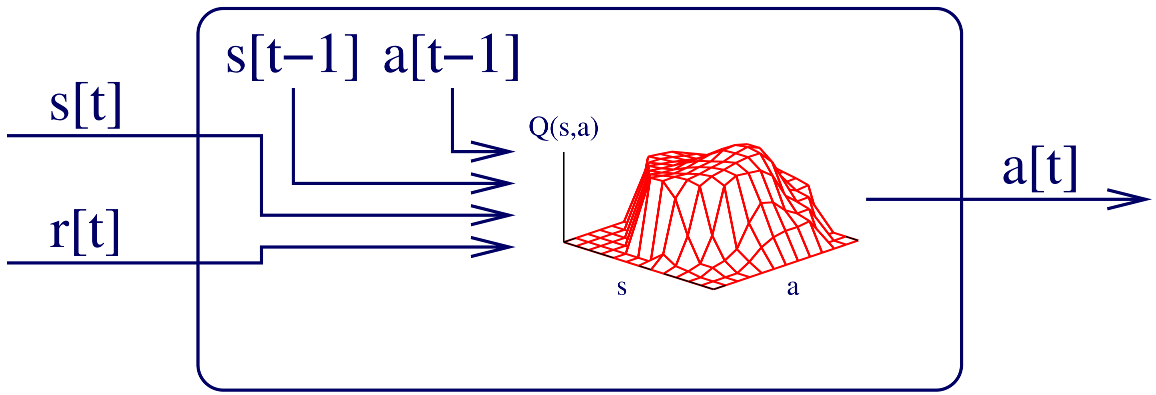

Save transitions:

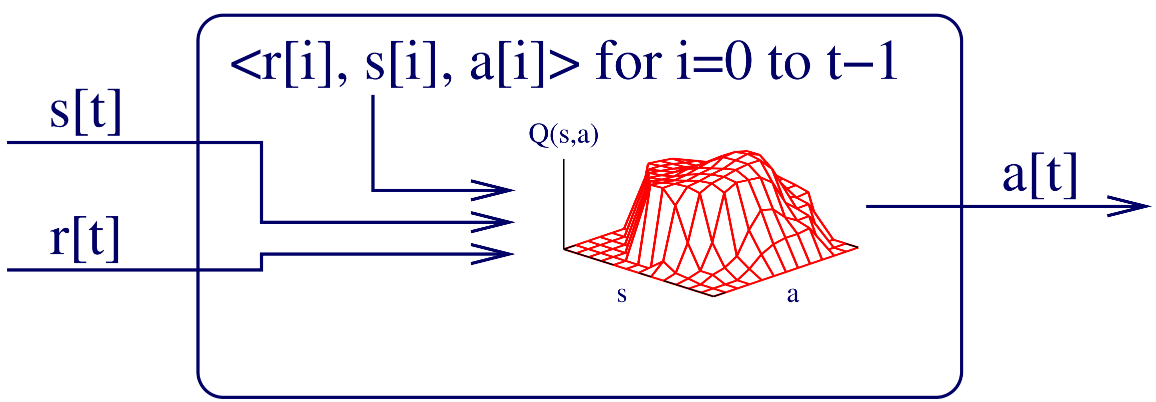

"Few Zeroes" [Kalyanakrishnan & Stone, '07]

Experience replay [Lin, '92], Fitted Q Iteration [Ernst et al., '05]

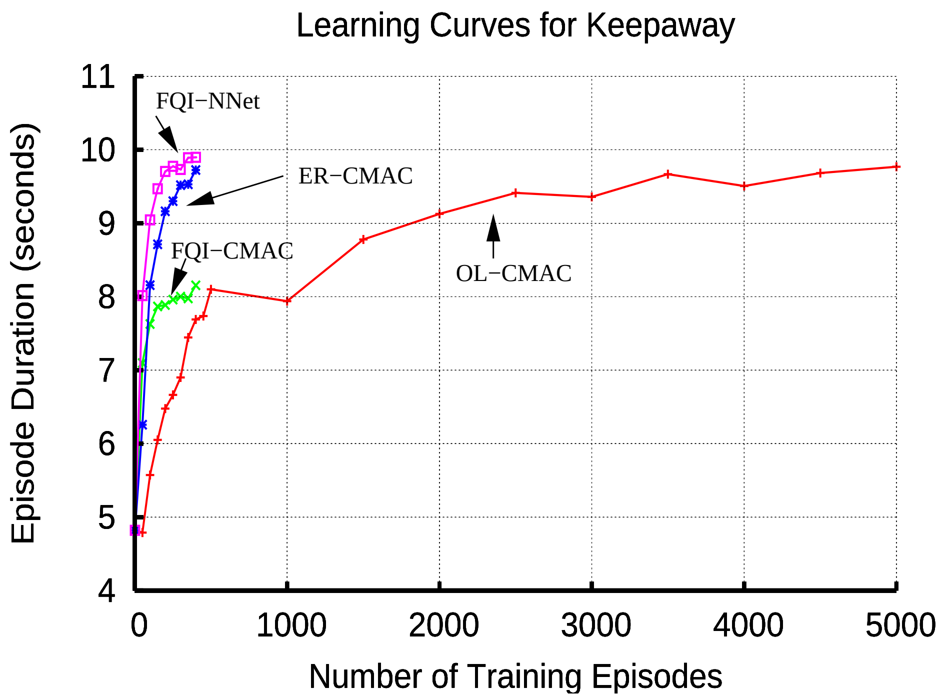

Other ways to scale up

- Advice/demonstration/TAMER, state/temporal abstraction
- Transfer learning, adaptive/hierarchical representations

Outline

- Build complete solutions to relevant challenge tasks
  - Drives research on component algorithms, theory

- Learning Agents

- **Multiagent reasoning**

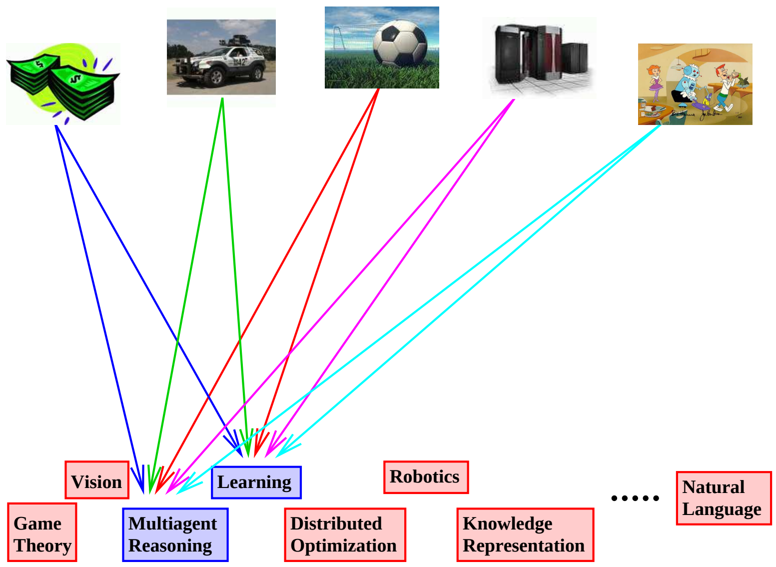

Multiagent Reasoning

Robust, fully autonomous agents in the real world

- Once there is one, there will soon be many
- To coexist, agents need to interact
- Example: **autonomous vehicles**
  - DARPA "Grand Challenge" was a great first step
  - "Urban Challenge" continues in the right direction
  - **Traffic lights and stop signs still best?** [Dresner & Stone, '04]

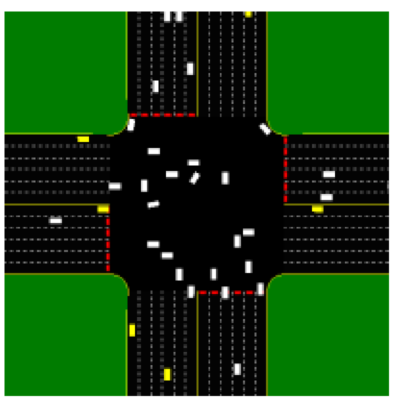

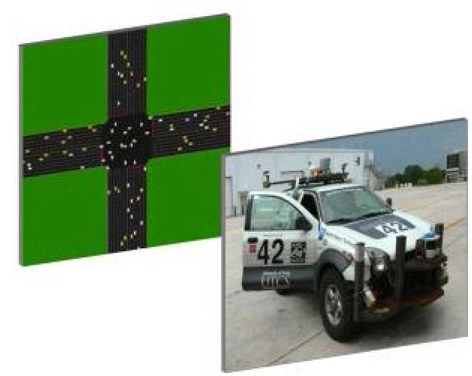

Autonomous Bidding Agents

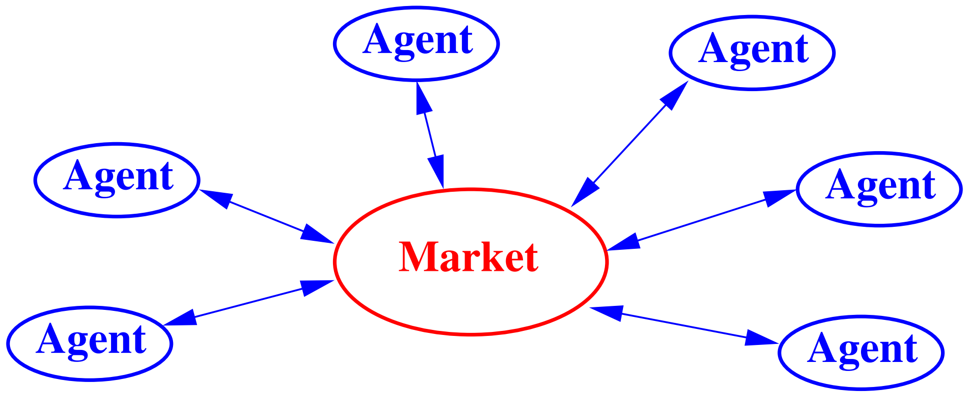

- Usual assumption: rational agents
- In practice, must prepare for the unexpected
  - Other agents created by others
  - Teammate/opponent modeling
  - Especially in competition scenarios

Trading Agent Competitions

ATTac: champion travel agent [Stone et al., '02]

– Learns model of auction closing prices from past data
– Novel algorithm for conditional density estimation

TacTex: champion SCM agent [Pardoe & Stone, '06]

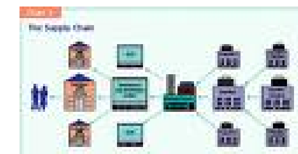

– Adapts procurement strategy based on recent data
– Predictive planning and scheduling algorithms

Common multiagent tradeoff:
learn detailed static model vs.
adapt minimally on-line

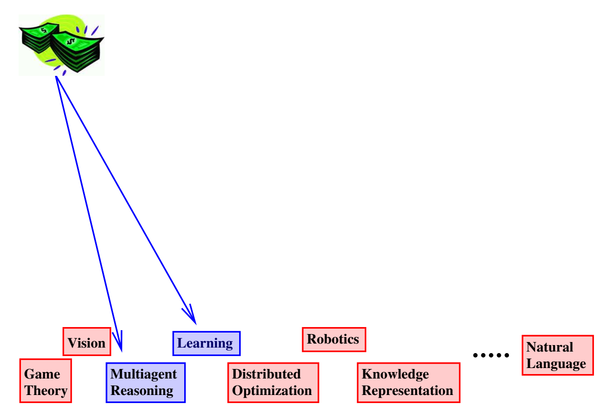

Outline

- Build complete solutions to relevant challenge tasks
  - Drives research on component algorithms, theory
- Learning Agents
- Multiagent reasoning
- **Implications**

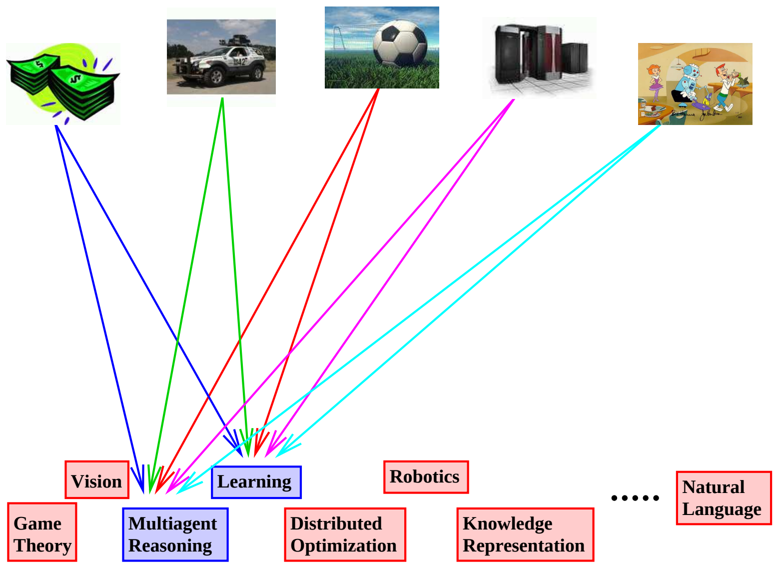

A Goal of AI

Robust, fully autonomous agents in the real world

What happens **when** we achieve this goal

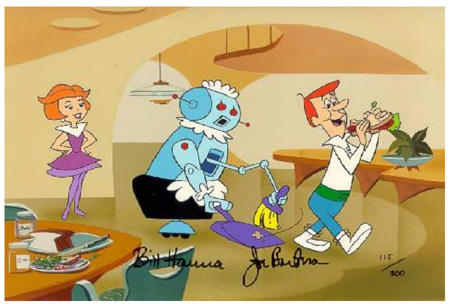 ?

?

- Question: Would you rather live
  - 100 years ago? Or 100 years in the future?
- Not clear — world changing in many ways for the worse

AI can be a part of the solution

Summary

Robust, fully autonomous agents in the real world

- Build complete solutions to relevant challenge tasks
  - Good problems drive research

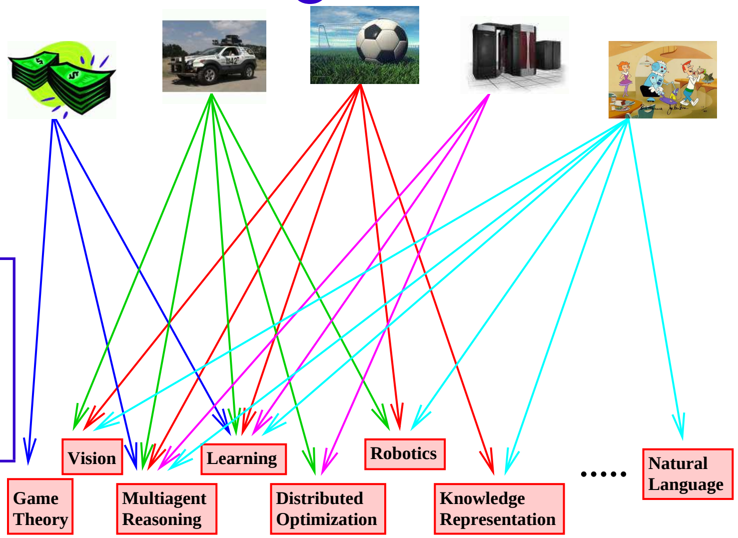

Combine algorithmic research, problem-oriented approaches

Summary

Robust, fully autonomous agents in the real world

- Build complete solutions to relevant challenge tasks
  - Good problems drive research

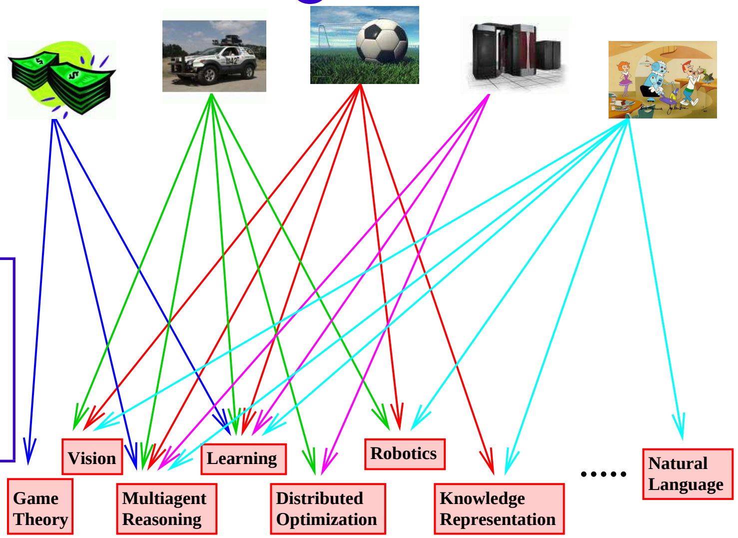

Combine algorithmic research, problem-oriented approaches

- Current challenges need learning, multiagent reasoning

ESI Environmental Science Institute UNIVERSITY OF TEXAS AT AUSTIN

JACKSON SCHOOL OF GEOSCIENCES

COLLEGE OF Natural Sciences THE UNIVERSITY OF TEXAS AT AUSTIN

This presentation is provided as part of the Hot Science - Cool Talks Outreach Series, and is supported by the Environmental Science Institute, the Jackson School of Geosciences, and the College of Natural Sciences at The University of Texas at Austin.

Additional support provided in part by a grant from the Motorola Foundation.

motorola foundation

Dr. Peter Stone

Peter Stone is the founder and director of the Learning Agents Research Group within the Artificial Intelligence Laboratory in the Department of Computer Science at The University of Texas at Austin. His main research interest in AI is understanding how we can best create complete intelligent agents. He considers adaptation, interaction, and embodiment to be essential capabilities of such agents.