Humans vs. AI:
Robot Soccer and Gran Turismo

Peter Stone

Learning Agents Research Group (LARG)
Department of Computer Science
The University of Texas at Austin

(also, Executive Director of Sony AI America)

UT Austin: Exciting Times — Year of AI!

- NSF Institute for Foundations of Machine Learning (IFML)
  - **Machine Learning Laboratory** (MLL)
- Bridging Barriers: **Good Systems**
  - Inspired by **One Hundred Year Study on AI (AI100)**
- Director of **Texas Robotics**
  - Ribbon cutting on new building in 2021

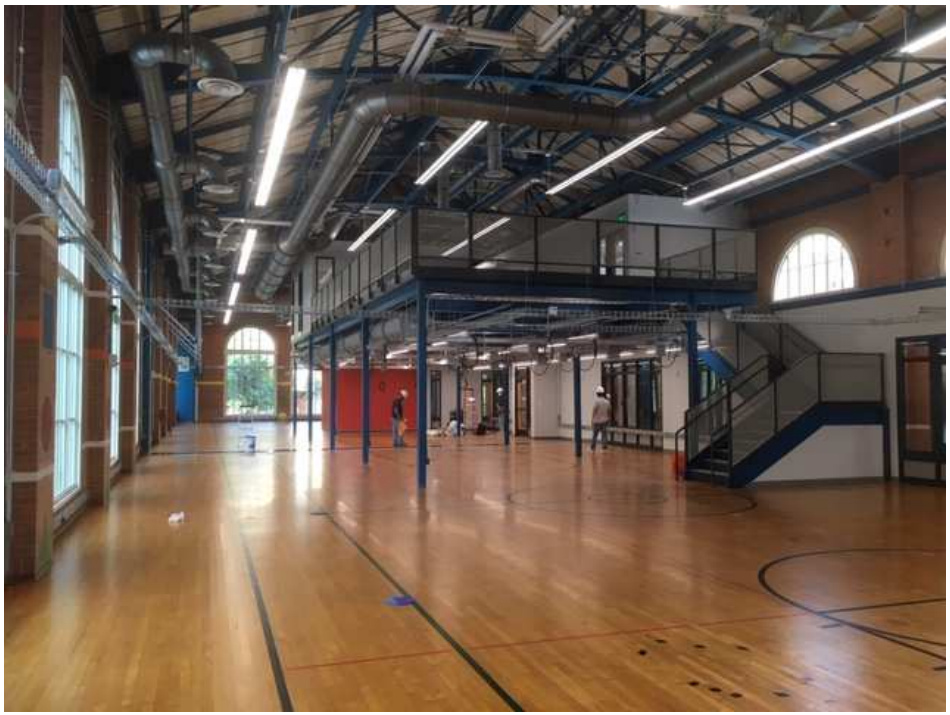

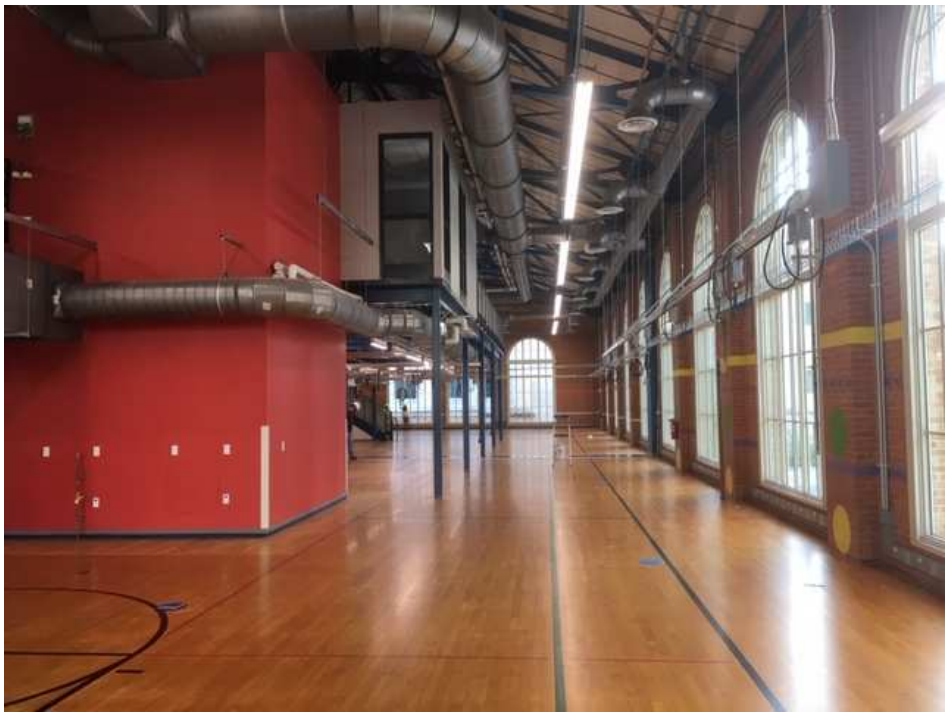

The Big Scientific Questions of our Time

- How did the **universe** originate?
- How did **life** on Earth originate?
- What is the nature of **intelligence**?

The Nature of Intelligence

How Can we Study it?

- Think about it — Philosophy
- Study human (or other animal) behavior — Psychology
- Study human (or other animal) brains — Neuroscience
- Build and analyze intelligent artifacts — Computer Science

A Goal of AI and Robotics

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

How?

- Build complete agents to perform increasingly complex tasks
  Complete agents: sense, decide, and act — closed loop

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

"Good problems produce good science"

My Research Problem

To what degree can autonomous intelligent agents learn in the presence of teammates and/or adversaries in real-time, dynamic domains?

Research Areas

- Autonomous agents
- Multiagent systems
- Robotics
- Machine learning
  - **Reinforcement learning**

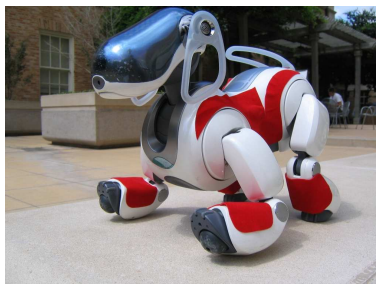
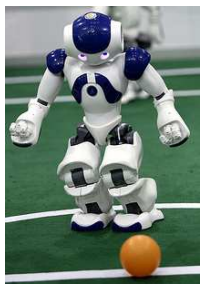

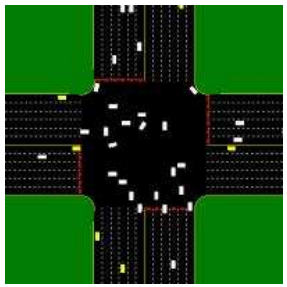

RoboCup Soccer

- Grand challenge: beat World Cup champions by 2050
- Still in relatively early stages
- Many virtues as a challenge problem:
  - Incremental challenges, closed loop at each stage
  - Robot design to multi-robot systems
  - Relatively easy entry
  - Inspiring to many
- Visible progress


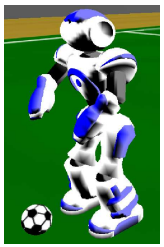



Small-sized League


Middle-sized League


Legged Robot League


Simulation League


Humanoid League

RoboCup@Home

Reinforcement Learning

Supervised learning mature [PyTorch]

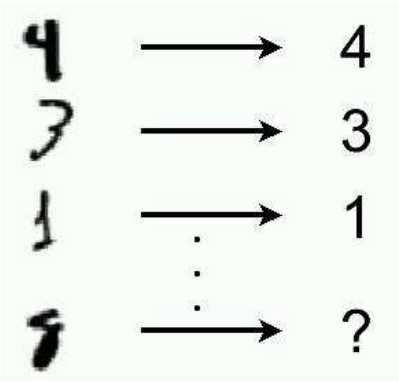

For agents, reinforcement learning most appropriate

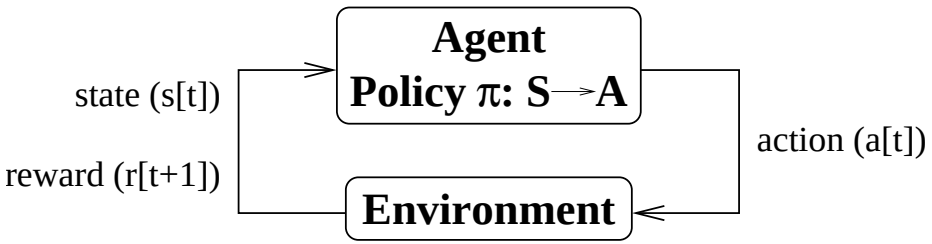

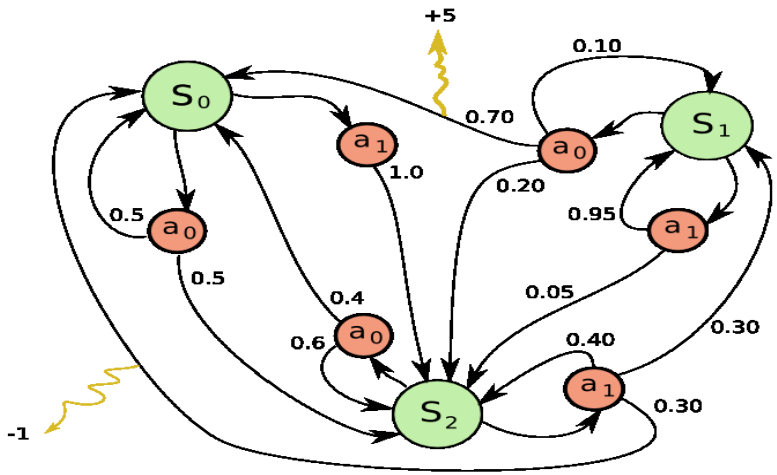

- Foundational theoretical results
- Applications require innovations to scale up

BE a reinforcement learner

- You, as a group, act as a learning agent
- **Actions**: Wave, Stand, Clap
- **Observations**: colors, reward
- **Goal**: Find an optimal *policy*
  - Way of selecting actions that gets you the most reward

How did you do it?

- What is your policy?
- What does the world look like?

Formalizing What Just Happened

Knowns:

- $\mathcal{O} = \{\text{Blue}, \text{Red}, \text{Green}, \text{Black}, \ldots\}$
- Rewards in $\mathbb{R}$
- $\mathcal{A} = \{\textit{Wave}, \textit{Clap}, \textit{Stand}\}$

$$o_0, a_0, r_0, o_1, a_1, r_1, o_2, \ldots$$

Unknowns:

- $\mathcal{S} =$ 4x3 grid
- $\mathcal{R} : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$
- $\mathcal{T} = \mathcal{S} \mapsto \mathcal{O}$
- $\mathcal{P} : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$

$$s_0, o_0, a_0, r_0, s_1, o_1, a_1, r_1, s_2, o_2, \ldots$$

$$o_i = \mathcal{T}(s_i) \qquad r_i = \mathcal{R}(s_i, a_i) \qquad s_{i+1} = \mathcal{P}(s_i, a_i)$$

My Research Problem

To what degree can autonomous intelligent agents learn in the presence of teammates and/or adversaries in real-time, dynamic domains?

Research Areas

- Autonomous agents
- Multiagent systems
- Robotics
- Machine learning
  - **Reinforcement learning**

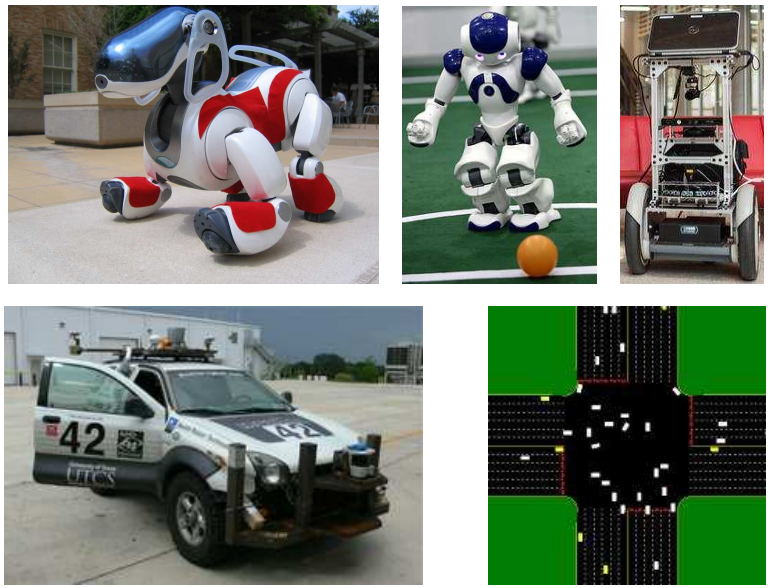

GT Sophy

Where are we now?

Me: Lots of open research challenges

- Professor of Computer Science at UT Austin
- Machine Learning Lab, Good Systems,
- Director of Texas Robotics
- Executive Director of Sony AI America

You: Lots of choices and opportunities

- Where will you go to college?
- What will you study?
- What will your lifelong challenge be?
- Opportunities: UTCS, Turing Scholars, FRI, TR

AI: Thriving, but with concerns

A Goal of AI

Robust, fully autonomous agents in the real world

What happens **when** we achieve this goal?

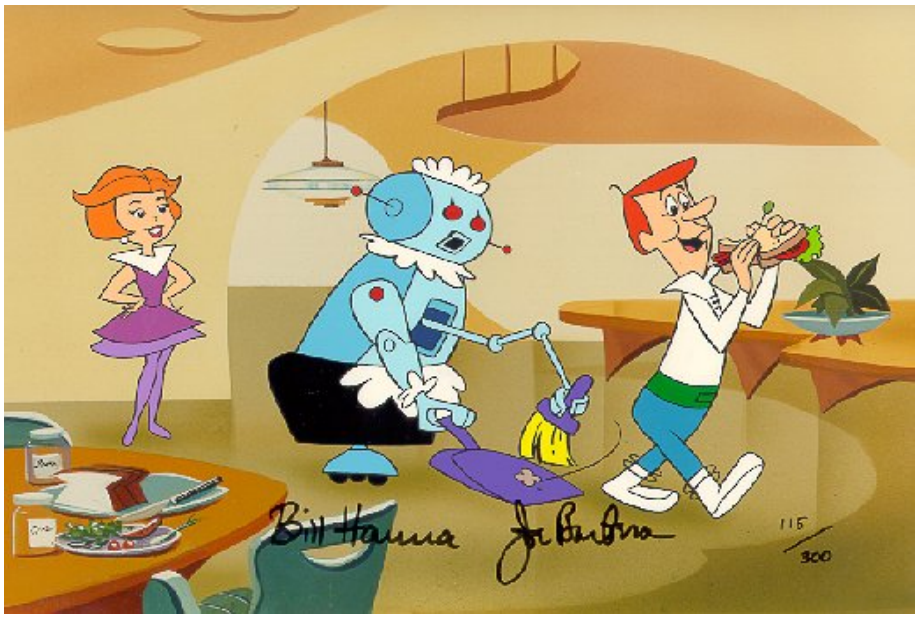

?

?

- Question: Would you rather have been born
  - 50 years earlier? Or 50 years later?
- Not clear — world changing in many ways for the worse

AI can be a part of the solution

Reality

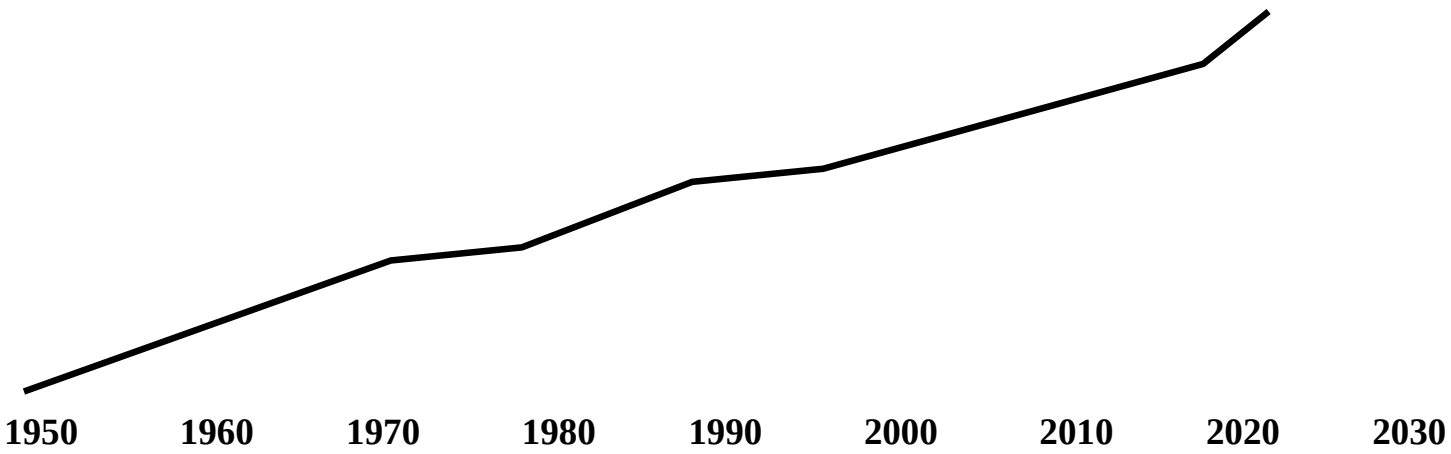

Perceptions

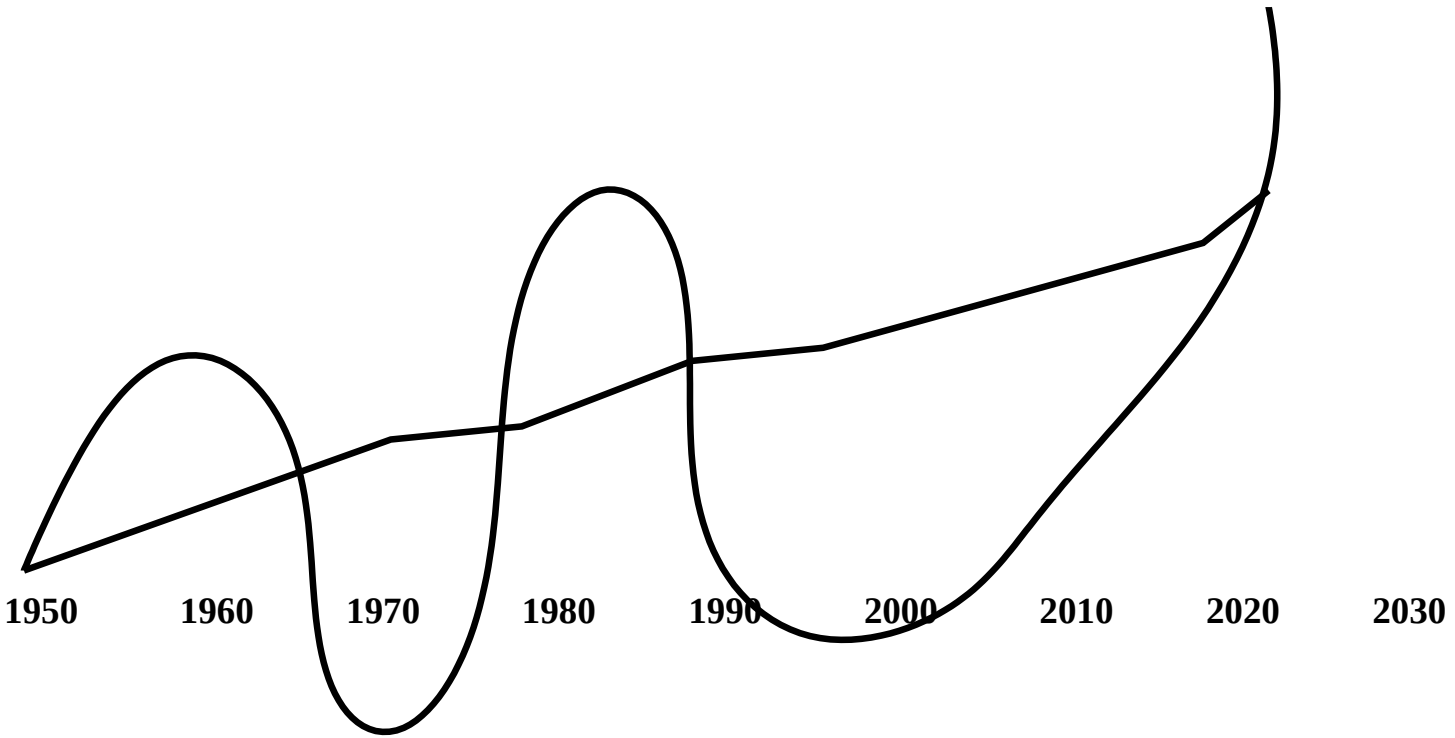

Uncertainty