Machine Learning and Artificial Intelligence for Autonomous Robots

Peter Stone

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

(Also, Cogitai Inc.)

A Goal of AI and Robotics

A Goal of AI and Robotics

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

A Goal of AI and Robotics

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

How?

A Goal of AI and Robotics

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

How?

- Build complete agents to perform increasingly complex tasks

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

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

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)

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)

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"

Research Question

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

Research Question

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**

Research Question

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**

Research Question

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**

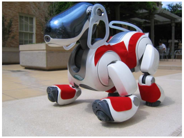
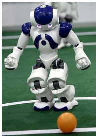

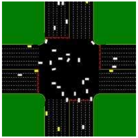

Research Question

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**
  - **Cogitai**

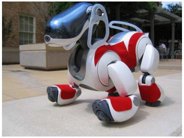
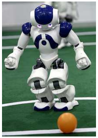

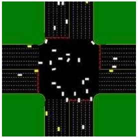

RoboCup Soccer

RoboCup Soccer

- Grand challenge: beat World Cup champions by 2050

RoboCup Soccer

- Grand challenge: beat World Cup champions by 2050
- Still in relatively early stages

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

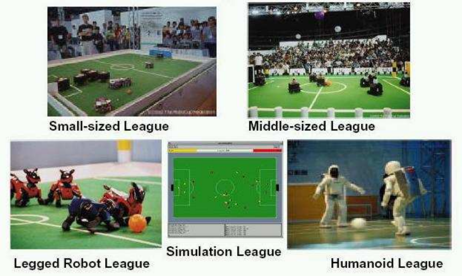

RoboCup Soccer

- 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  

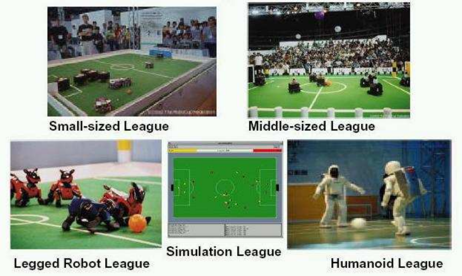

RoboCup Soccer

- 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  

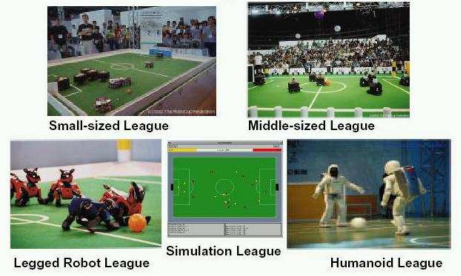

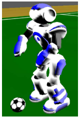

Research Advances due to RoboCup

Research 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

Research 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
- **400+ publications** from simulation league alone
- **200+** from 4-legged league
- **Dozens** (at least) of Ph.D. theses

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 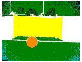

RoboCup@Home

RoboCup@Home

Reinforcement Learning

Supervised learning mature [TensorFlow]

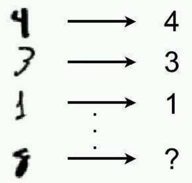

Reinforcement Learning

Supervised learning mature [TensorFlow]

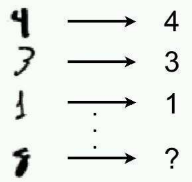

For agents, reinforcement learning most appropriate

Reinforcement Learning

Supervised learning mature [TensorFlow]

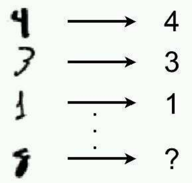

For agents, reinforcement learning most appropriate

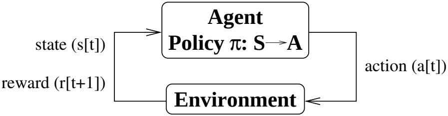

Reinforcement Learning

Supervised learning mature [TensorFlow]

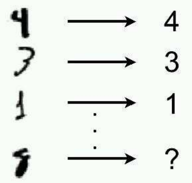

For agents, reinforcement learning most appropriate

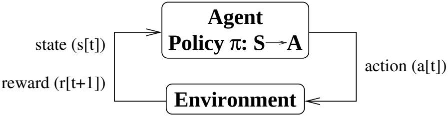

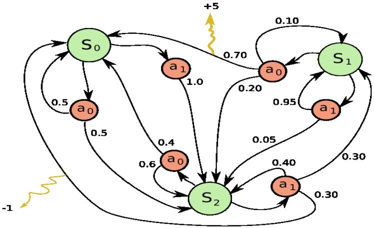

Reinforcement Learning

Supervised learning mature [TensorFlow]

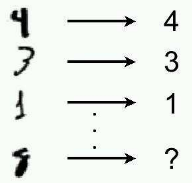

For agents, reinforcement learning most appropriate

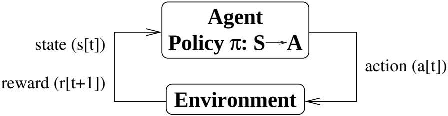

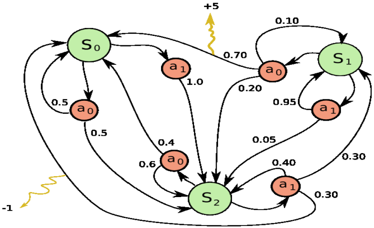

– Foundational theoretical results
– Applications require innovations to scale up

RL Theory

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

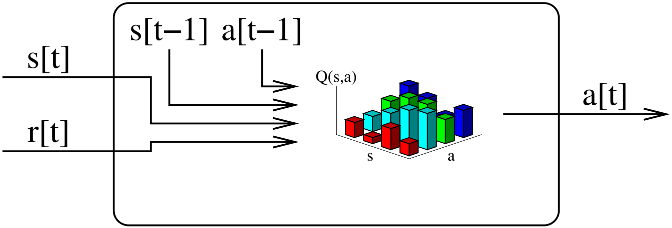

RL Theory

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

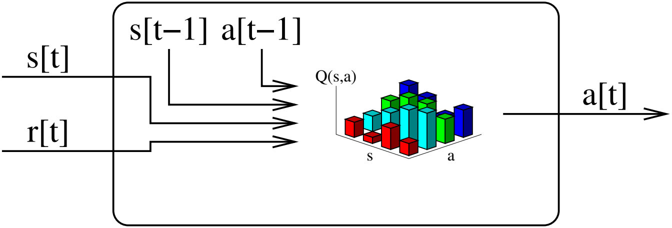

– Table-based representation
– Visit every state infinitely often

Function Approximation

In practice, visiting every state impossible

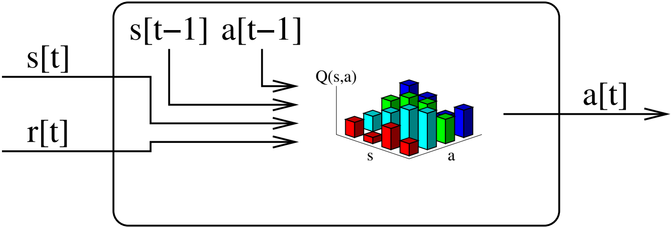

Function Approximation

In practice, visiting every state impossible

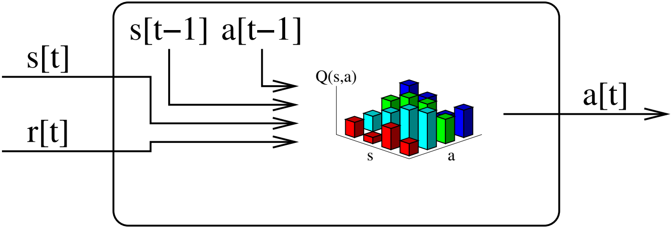

Function approximation of value function

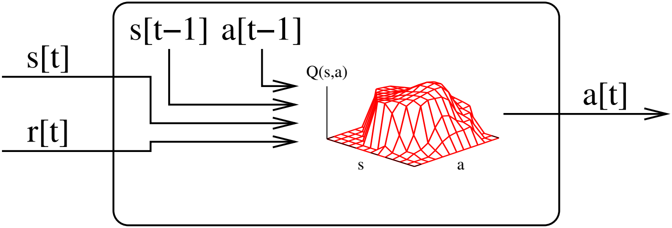

Theoretical guarantees harder to come by

Applications: Towards a Useful Tool

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

Applications: Towards a Useful Tool

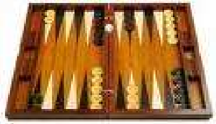

- Backgammon [Tesauro, '94]
- Helicopter control [Ng et al., '03]
- Adaptive treatment of epilepsy [Pineau et al., '08]

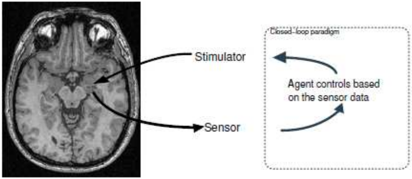

- Invasive species management, wildfire suppression [Dietterich et al., '13]

Applications: Towards a Useful Tool

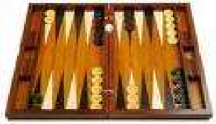

- Backgammon [Tesauro, '94]
- Helicopter control [Ng et al., '03]
- Adaptive treatment of epilepsy [Pineau et al., '08]

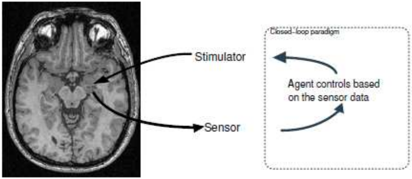

- Invasive species management, wildfire suppression [Dietterich et al., '13]
- Google DeepMind beats human go champion, [Silver et al., '16]

Selected RL Contributions

- Human interaction

Selected RL Contributions

- Human interaction
  - Advice, Demonstration 

Selected RL Contributions

- Human interaction
  - Advice, Demonstration
  - Positive/Negative Feedback

[Knox & Stone, '09]

Selected RL Contributions

- Human interaction
  - Advice, Demonstration
  - Positive/Negative Feedback

[Knox & Stone, '09]

- Transfer learning for RL [Taylor & Stone, '07]
- Curriculum Learning [Narvekar et al., '16]

Selected RL Contributions

- Human interaction
  - Advice, Demonstration
  - Positive/Negative Feedback [Knox & Stone, '09]
- Transfer learning for RL [Taylor & Stone, '07]
- Curriculum Learning [Narvekar et al., '16]
- RL for musical playlist recommendation [Liebman et al., '15]

Selected RL Contributions

- Human interaction
  - Advice, Demonstration
  - Positive/Negative Feedback

[Knox & Stone, '09]

- Transfer learning for RL [Taylor & Stone, '07]
- Curriculum Learning [Narvekar et al., '16]
- RL for musical playlist recommendation [Liebman et al., '15]
- TEXPLORE for Robot RL [Hester & Stone, '13]

Selected RL Contributions

- Human interaction
  - Advice, Demonstration
  - Positive/Negative Feedback

[Knox & Stone, ’09]

- Transfer learning for RL [Taylor & Stone, ’07]
- Curriculum Learning [Narvekar et al., ’16]
- RL for musical playlist recommendation [Liebman et al., ’15]
- TEXPLORE for Robot RL [Hester & Stone, ’13]
  - Sample efficient; real-time
  - Continuous state; delayed effects

Selected RL Contributions

- Human interaction
  - Advice, Demonstration
  - Positive/Negative Feedback

[Knox & Stone, '09]

- Transfer learning for RL [Taylor & Stone, '07]
- Curriculum Learning [Narvekar et al., '16]
- RL for musical playlist recommendation [Liebman et al., '15]
- TEXPLORE for Robot RL [Hester & Stone, '13]
  - Sample efficient; real-time
  - Continuous state; delayed effects
- Deep RL in continuous action spaces [Hausknecht & Stone, '16]

Artificial Intelligence and Life in 2030

100 Year Study on AI: 1st Study Panel Report

Prof. Peter Stone*
Study Panel Chair
Department of Computer Science
The University of Texas at Austin
*Also Cogitai, Inc.

https://ai100.stanford.edu

September 2016

One Hundred Year Study
Goals of the Endowment

"To support a longitudinal study of influences of AI advances on people and society,

centering on periodic studies of developments, trends, futures, and potential disruptions associated with the developments in machine intelligence, and

on formulating assessments, recommendations and guidance on proactive efforts."   (July 2014)

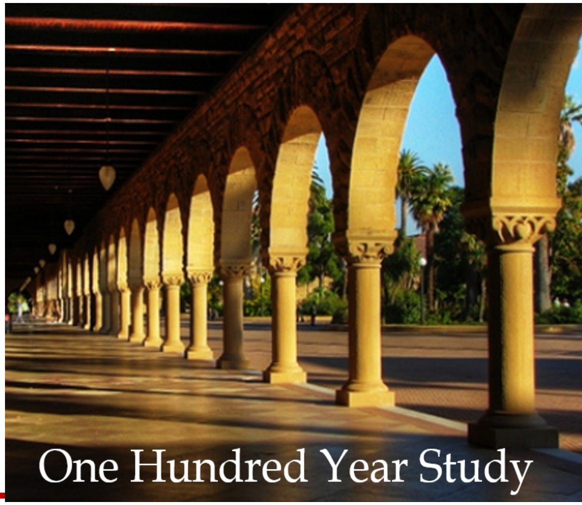

Standing Committee

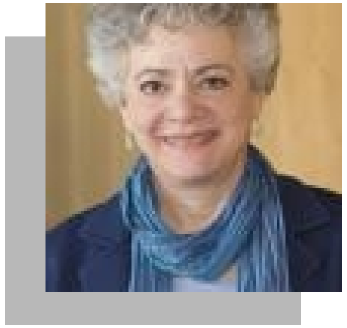

Barbara Grosz, Chair

Russ Altman

Eric Horvitz

Alan Mackworth

Tom Mitchell

Deirdre Mulligan

Yoav Shoham

One Hundred Year Study: Timeline of Studies

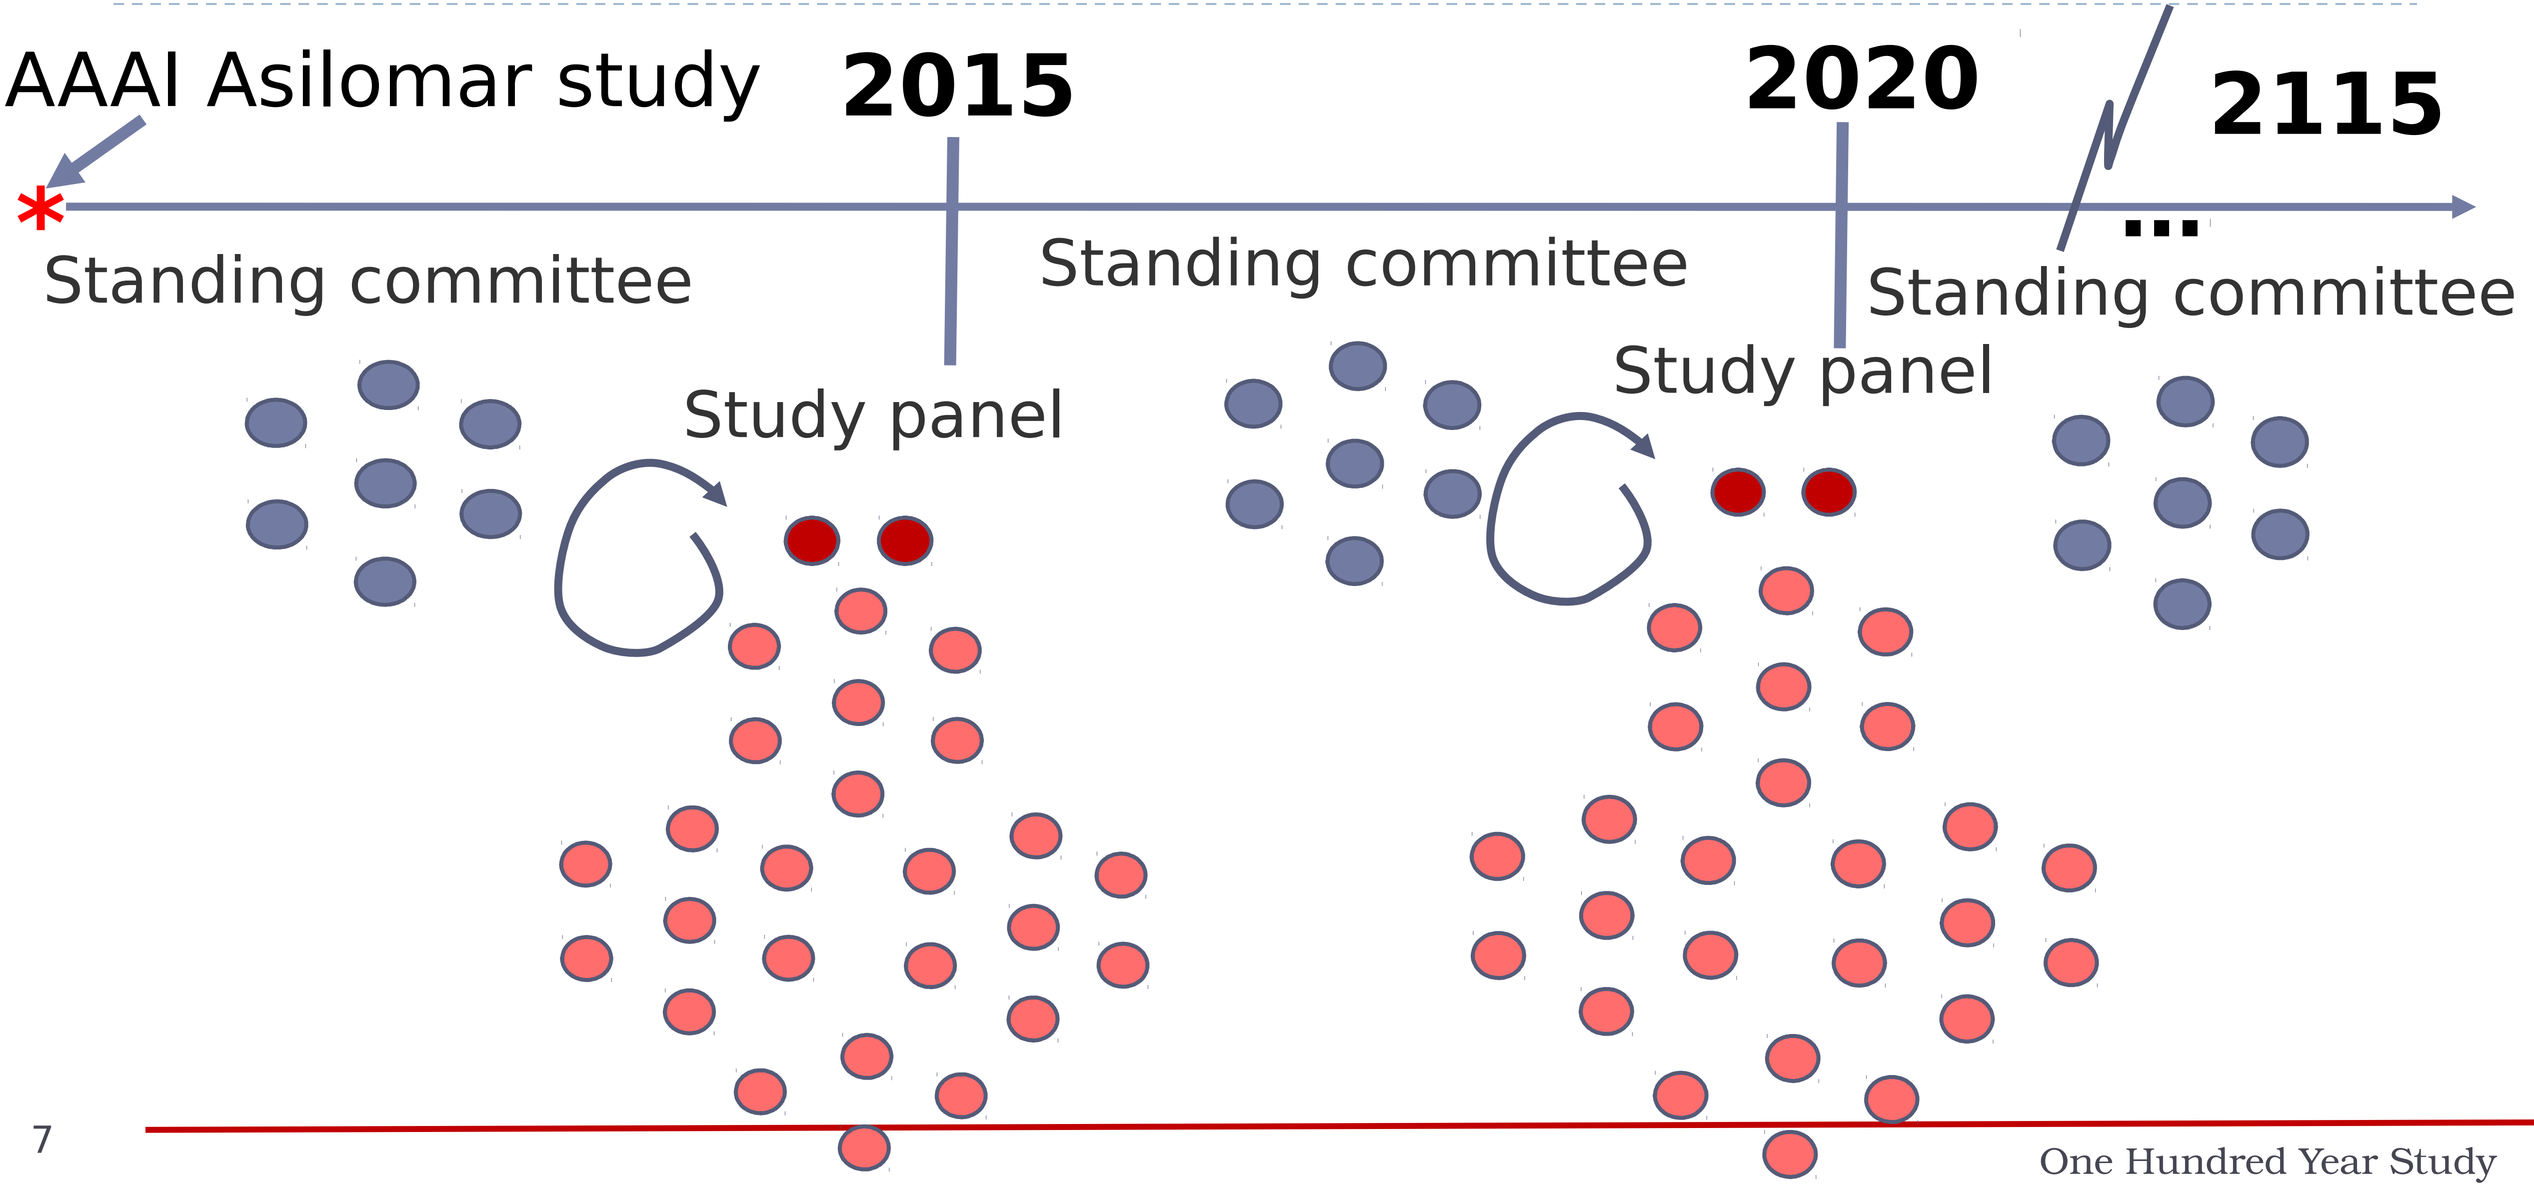

One Hundred Year Study: Intended Audiences

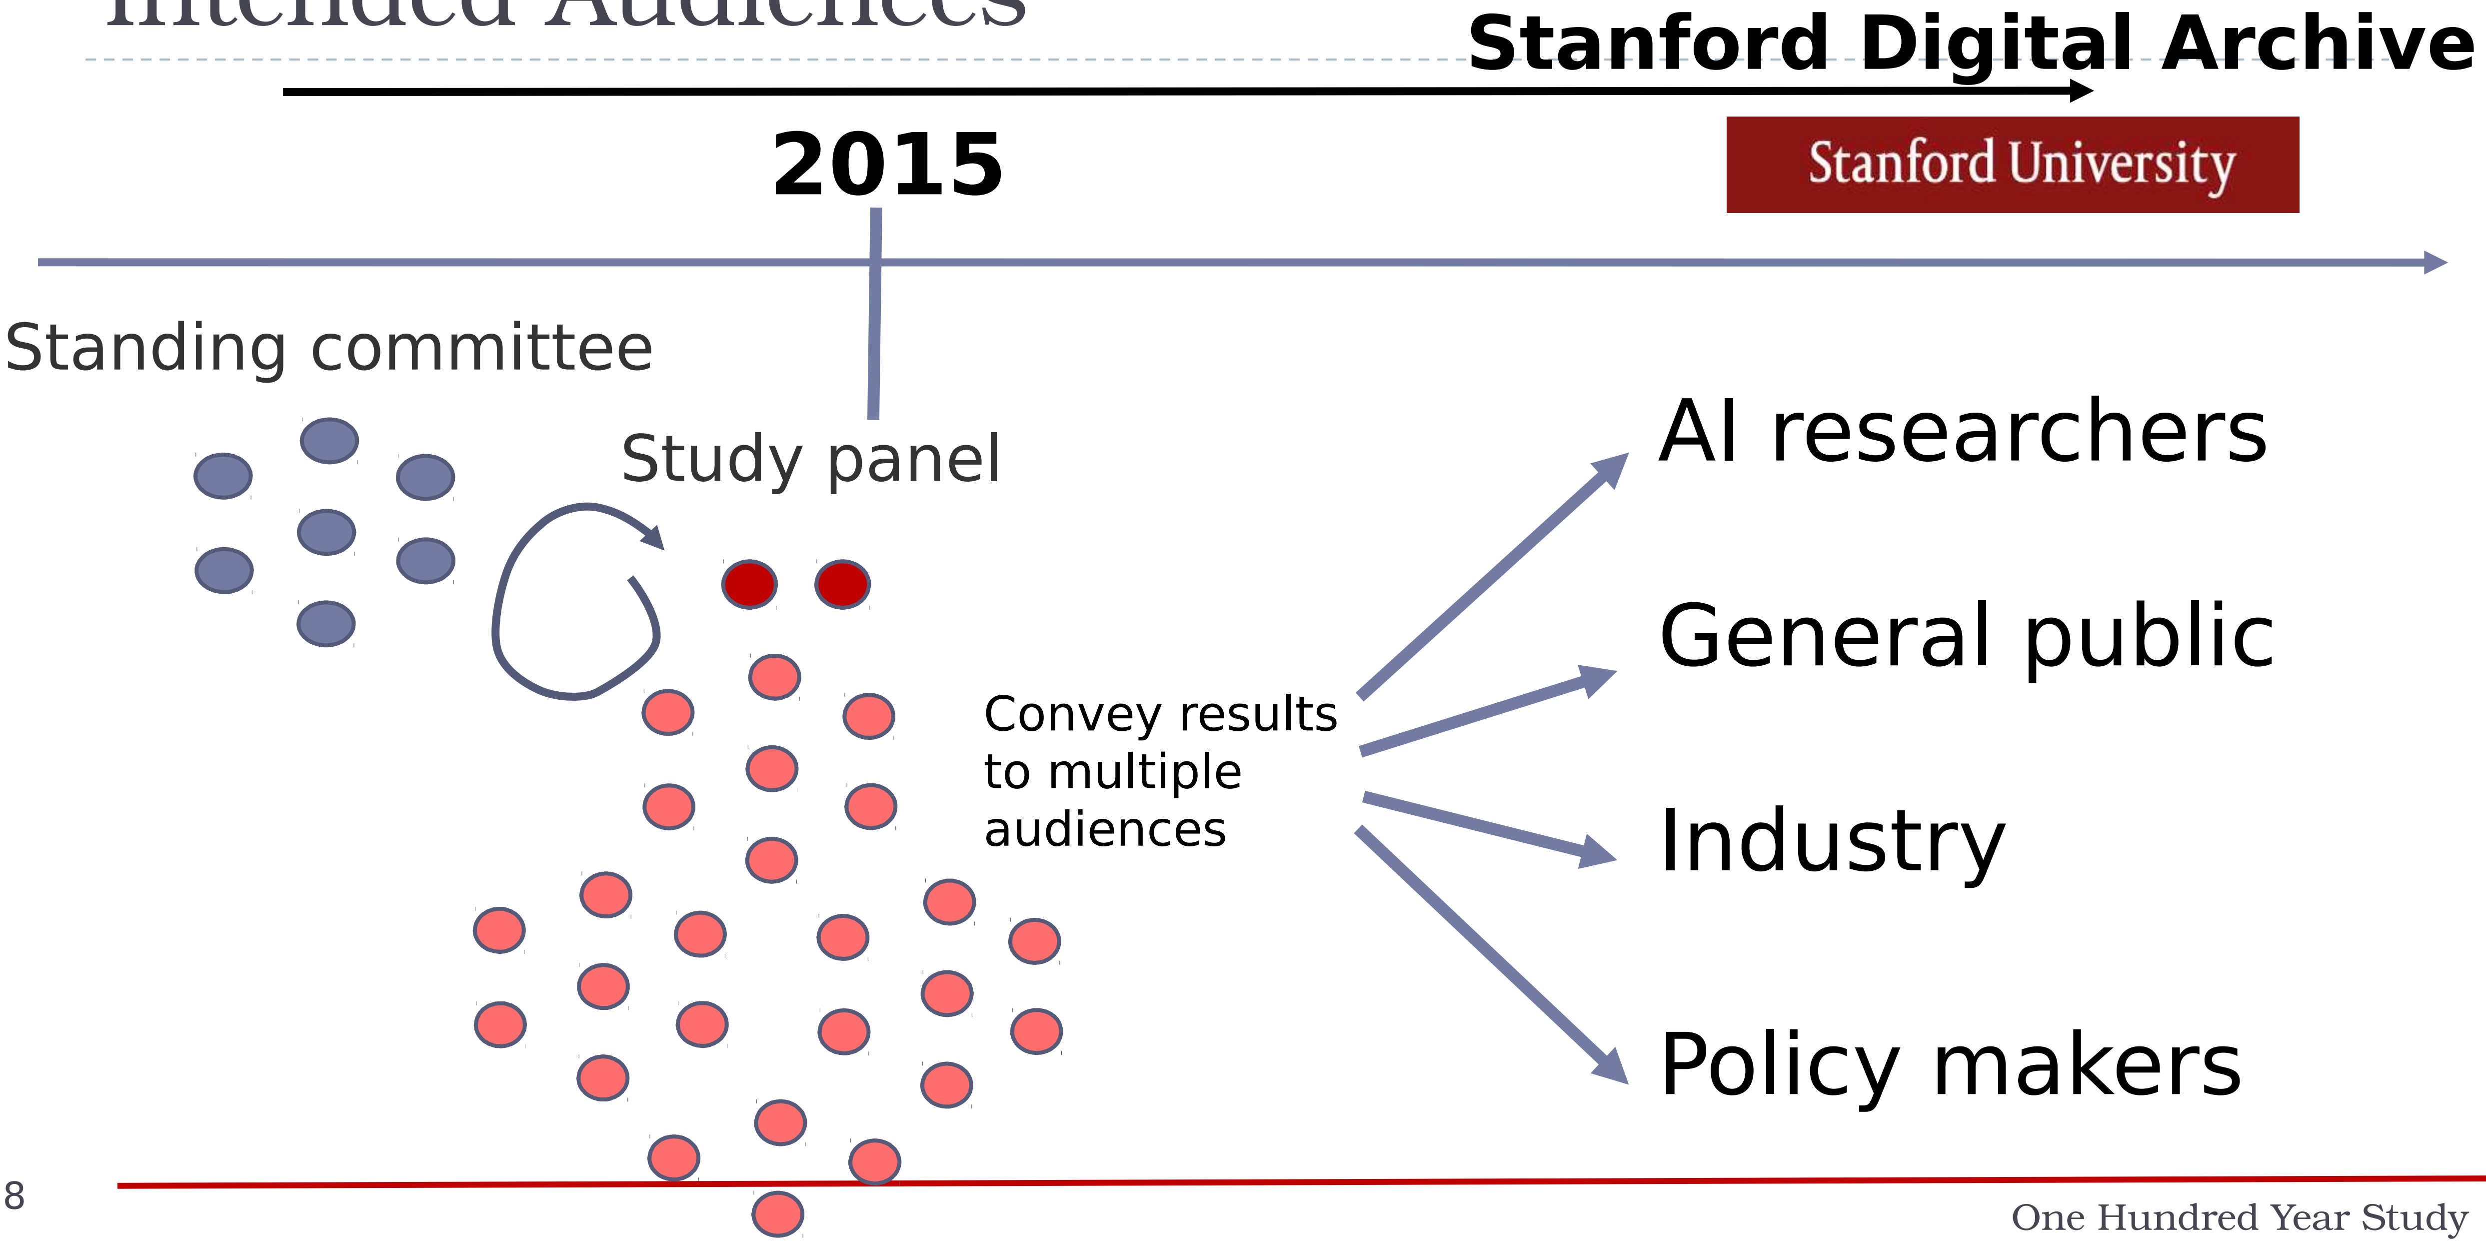

Charge to the Inaugural Study Panel: Artificial Intelligence and Life in 2030

Identify possible *advances* in AI over next 15 years and *their potential influences on daily life*.

Specify *scientific, engineering, and legal efforts* needed to realize these developments.

Consider actions needed to shape outcomes for *societal good*, deliberating *design, ethical and policy challenges.*

Focus: *large urban regions* (typical North American city), grounding the examination of AI technologies in a context that highlights

- *potential influences on a wide variety of activities*
- *interdependencies and interactions among AI technologies*

Members of the Inaugural Study Panel
Artificial Intelligence and Life in 2030

Chair: Peter Stone, UT Austin

- Rodney Brooks, Rethink Robotics
- Erik Brynjolfsson, MIT
- Ryan Calo, University of Washington
- Oren Etzioni, Allen Institute for AI
- Greg Hager, Johns Hopkins
- Julia Hirschberg, Columbia
- Shivaram Kalyanakrishnan, IIT Bombay
- Ece Kamar, Microsoft
- Sarit Kraus, Bar Ilan
- Kevin Leyton-Brown, UBC
- David Parkes, Harvard
- William Press, UT Austin
- Julie Shah, MIT
- Astro Teller, X
- Milind Tambe, USC
- AnnaLee Saxenian, Berkeley

Structure

- Preface for context
- Executive Summary (1 page)
- Overview (5 pages)
- Introduction
  - Defining AI; Current research trends
- AI by domain
  - 8 areas with likely urban impact by 2030
  - Look backwards 15 years and forward 15 years
  - Opportunities, barriers, and realistic risks
- Policy and legal issues
  - Current status; Recommendations
- Lots of callouts in the margins

Areas of Focus in the Study Panel Report

Transportation
Home-Service Robots
Healthcare
Education
Public Safety and Security
Low-resource communities
Employment and Workplace
Entertainment

hardware

partnering with people

building trust

societal futures

interpersonal interaction

Areas of Focus in the Study Panel Report

Transportation

Home-Service Robots

Healthcare

Education

Public Safety and Security

Low-resource communities

Employment and Workplace

Entertainment

Policy and Legal Issues

Summarizing callouts in the report

Artificial Intelligence and Life in 2030

100 Year Study on AI: 1st Study Panel Report

Prof. Peter Stone*
Study Panel Chair
Department of Computer Science
The University of Texas at Austin
Also Cogitai, Inc.

https://ai100.stanford.edu

September 2016