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)

One Hundred Year Study
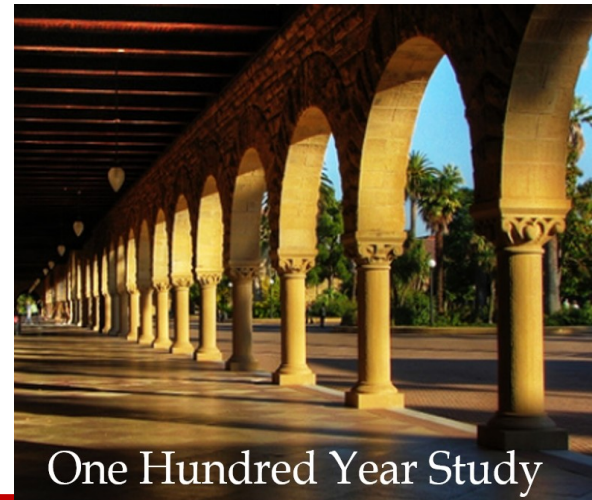
One Hundred Year Study

Roots: AAAI Presidential Panel on Long-term AI Futures (Asilomar 2008-09)

Commissioned by AAAI President, Eric Horvitz

Co-chaired by Eric Horvitz & Bart Selman

Charge: Explore potential long-term societal influences of AI advances.

Subgroups focused on

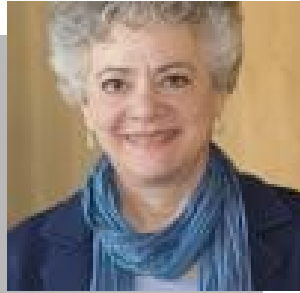
- Potential Disruptive Advances Over the Short-term

- Longer-term Pace, Concerns, Control

- Ethical and Legal Challenges

http://www.aaai.org/Organization/presidential-panel.php

Standing Committee

Barbara Grosz, Chair

Russ Altman

Eric Horvitz

Alan Mackworth

Tom Mitchell

Deirdre Mulligan

Yoav Shoham

One Hundred Year Study:
Timeline of Studies
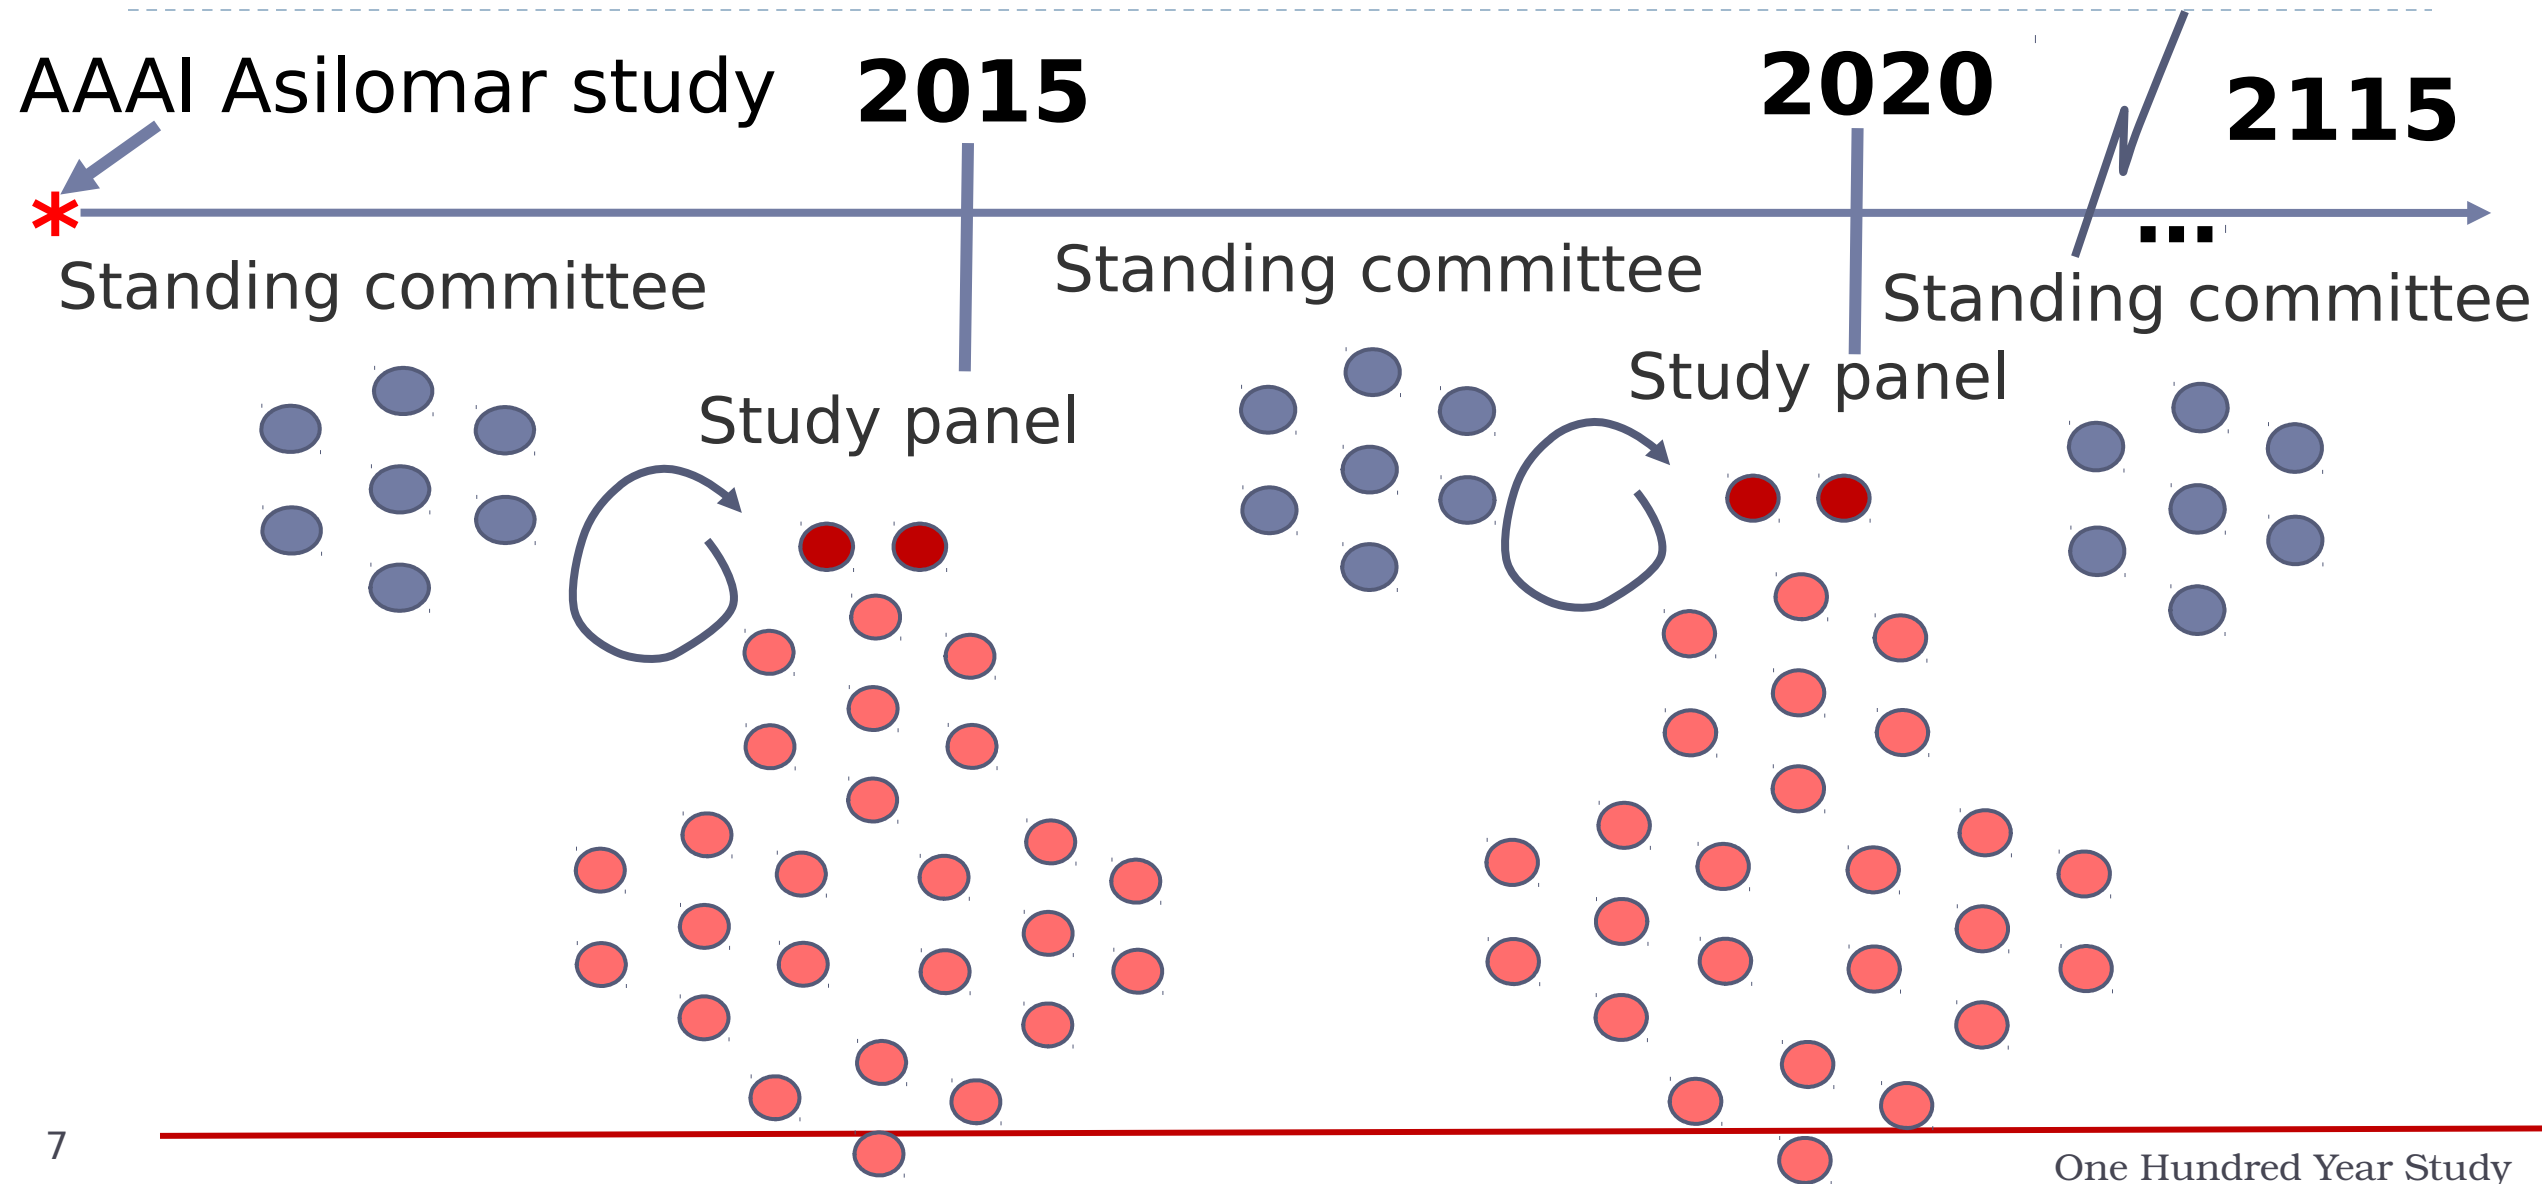
Charge to the Inaugural Study Panel: Artificial Intelligence and Life in 2030

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*

An Opportunity and a Challenge

- Carte Blanche with respect to format
  - Even flexibility on topic
- Should be topical now at a time of great interest in AI
  - Industry progress, white house studies, press, fear-mongering
  - Balanced view, taking into account possibilities, barriers, and **realistic** risks
- Also should be relevant over time
- Set a precedent for future studies

- First task: invite panelists
  - Balance of AI areas, seniority, gender, geography (to some extent)

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 ⟵ ————————— **hardware**

Healthcare

Education ⟵ ————————— **partnering with people**

Public Safety and Security ⟵ ————————— **building trust**

Low-resource communities

Employment and Workplace ⟵ ————————— **societal futures**
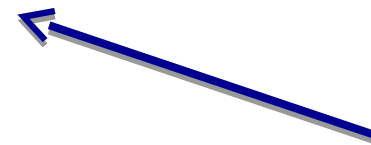
Entertainment ⟵ ————————— **interpersonal interaction**

Areas of Focus in the Study Panel Report

→ Transportation: Short term

Healthcare, Education: Medium term

Employment and Workplace: Longer term

Data → Predictive models → Decision models

Transportation: Prediction and Intervention
First domain where public asked to trust AI on a large scale

Problem: Sense surroundings, car state, lane change, parking, route plan..

AI Technologies: Advanced sensors, computer vision, machine learning, ...

Surprising progress in academia and industry:

- Autonomous vehicles (Google, Tesla, Uber, etc)
- Not just cars but drones

But now, must grapple with world full of people

In 15 years, autonomous pickup and delivery of people and packages

Areas of Focus in the Study Panel Report

Transportation: Short term

➤ Healthcare, Education: Medium term

Employment and Workplace: Longer term

Data ⟶ Predictive models ⟶ Decision models

Healthcare: prediction and intervention

Applications: clinical decision support, patient monitoring/coaching
management of healthcare systems

AI Technologies: predictive analytics
- Patients like mine (population level)
- People like me (individuals)

"AI-based applications could improve health outcomes...for millions of people - but only if they gain the trust of doctors, nurses, and patients"

Technical barrier: integrate human care, automated reasoning

Social barriers:
- Data privacy
- Trust of doctors, nurses, patients (transparency)
- Removal of regulatory and commercial obstacles

Education: Prediction & Intervention

AI promises to enhance education at all levels

- Personalization at scale: MOOCs, Intelligent Tutoring Systems
- Blur the line between formal classroom training and self-paced learning
- Opportunities for remote areas and adult retraining

Robots: Fun and engaging, but need evidence they help learning

Prediction: Engaged human teachers needed for quality education

Question: What effect will online learning, with less face-to-face interaction, have on social development?

Areas of Focus in the Study Panel Report

Transportation: Short term

Healthcare, Education: Medium term

➡ Employment and Workplace: Longer term

Data ➡ Predictive models ➡ Decision models

Employment and Workplace: crosscutting

"The current social safety net may need to evolve into better social services for everyone, such as healthcare and education, or a guaranteed basic income. Indeed, countries such as Switzerland and Finland have actively considered such measures"

Near term: AI tech will replace more tasks than jobs
- Will also create jobs, but harder to imagine what types

Ought to lower cost of goods and services
- Makes everyone effectively richer, but job loss more salient

Fear of replacing all human jobs in one generation: overblown
- But gap between rich and poor could grow

"It is not too soon for social debate on how the economic fruits of AI-technologies should be shared. As children in traditional societies support their aging parents, perhaps our artificially intelligent 'children' should support us, the 'parents' of their intelligence"

Policy and Legal Issues

Problem: AI could widen existing inequalities of opportunity if access to AI technologies…is unfairly distributed across society.

Problem: As a society, we are underinvesting resources in research on the societal implications of AI technologies.

Problem: As AI applications engage in behavior that, were it done by a human, would constitute a crime, courts and other legal actors will have to puzzle through whom to hold accountable and on what theory.

Problem: Like other technologies, AI has the potential to be used for good or nefarious purposes.

Problem: Absent sufficient technical expertise to assess safety or other metrics, national or local officials may refuse to permit a potentially promising application – or green light a sensitive application that has not been adequately vetted.

Recommedations

- Define a path toward accruing technical expertise in AI at all levels of government.

- Remove the perceived and actual impediments to research on the fairness, security, privacy, and social impacts of AI systems.

- Increase public and private funding for interdisciplinary studies of societal impacts of AI.

Whitehouse reports: Simultaneous effort

- Some overlap in participants
- Much overlap in focus and findings:
    - Potential for AI to promote public good
    - Regulation by sector, not of AI as a whole
    - More AI-literate people in positions of influence
    - Potential for AI to lead to greater wealth
- Also several differences (mainly in emphasis):
    - Concern about increased wage gap based on education (agree)
    - Ethical training for AI practitioners and students (agree)
    - Global considerations and security (out of scope)
    - More support for basic research (agree)
    - No universal basic income (but raise minimum wage)

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