



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.*

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



One Hundred Year Study: Intended Audiences

Stanford Digital Archive

Stanford University

2015

Standing committee

Study panel

Convey results
to multiple
audiences

AI researchers

General public

Industry

Policy makers

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

Process

- What should the output be?
 - Report? Webpages? Youtube videos? Multimedia?
- What topics?
- Monthly phone calls: outline for a background doc
 - The parts with **real** impact would come later
- In-person meeting in February: intensive draft writing
- Months of back and forth with a professional writer
 - Became clear the document was going to be our sole output
- Final in-person meeting (callouts)
- Lots of feedback and polishing

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: Short term

Public Safety and Security: Medium term

Meeting needs of low-resource communities: Longer term



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

➔ Public Safety and Security: Medium term

Meeting needs of low-resource communities: Longer term



Public Safety and Security: Prediction & Intervention

Fraud Detection

Problem: How do we tell which transactions are fraudulent?
Credit cards, compromised accounts, etc.

AI Technology: Classification via machine learning

How:

Data: Records of fraudulent and legitimate transactions

Predictive Model: Identify features of fraudulent transactions

Decisions: Autonomous termination of transaction when fraud

Risk: encoding, even magnifying, human biases

Public Safety and Security: Prediction

Predictive policing


Problem: Crime analysis or predict crimes for effective police allocation
Crime locations, times of higher risk of crime

AI technology: Predictive modeling via machine learning.

How:

Data: Crime data

Predictive model: Probability of high-crime areas



Areas of Focus in the Study Panel Report

Transportation: Short term

Public Safety and Security: Medium term

 Meeting needs of low-resource communities: Longer term



Low Resource Communities: Prediction Machine Learning for Preventing Lead Poisoning in Children

Problem:

Lead: Severe motor skill/brain impairment in children

Lead found in many homes in paint


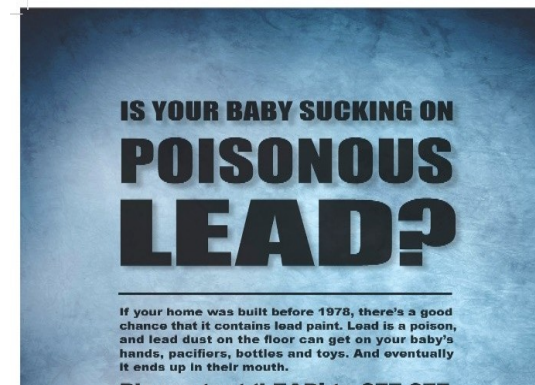
- ▶ Which houses to inspect for lead

AI technology: Machine learning predictive models

How:

Data: 20 year data on children's blood tests & house inspections

Predictive model: Pre-birth prediction if house needs lead inspections




Low Resource Communities: Prediction & Intervention

Social Networks for Raising HIV awareness in homeless youth

Problem:

- Homeless youth HIV: 10x infection rate of general population
- Distrust authority; peer interventions to spread HIV information
 - ▶ Social network: Uncertain and dynamic



AI technology: Influence maximizing peer leaders

How:

- Data:* Network information
- Decision making:* Most influential peer leaders in network
 - Gather more information about network, reduce uncertainty



Decision aid piloted in Los Angeles showing increase in HIV testing

Areas of Focus in the Study Panel Report

Transportation

Service Robots

Healthcare

Public Safety and Security

Education

Meeting needs of low-resource communities

Entertainment

Employment and Workplace

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.