



National Strategy on Artificial Intelligence



Michael Garris

Senior Scientist / National Institute of Standards and Technology
Co-Chair / National Science and Technology Council,
Subcommittee on Machine Learning and Artificial Intelligence

What is AI?

- Think and act like humans
- Think and act rationally
- Sometimes mutually exclusive. 😊

- AI Problem Space Categories (by F. Chen)
 - Logical Reasoning
 - Knowledge Representation
 - Planning and Navigation
 - Natural Language Processing
 - Perception

Artificial Intelligence (AI)

- Not new

- Coined 1956 – computer science researchers met at Dartmouth
 - (McCarthy @ Dartmouth; Minsky @ Harvard; Rochester @ IBM; Shannon @ Bell Labs)
 - “machines will use language, form abstractions/concepts, solve problems now reserved for humans, and improve themselves”
- Crystallized in Alan Turing’s 1950 paper, *“Computing Machinery and Intelligence”*
- Waves and Winters
 - Personally involved end of the 80’s with MLP “wave”
 - LeCun’s presentation of a new 4th (convolutional) layer
 - Reading address zip codes for automated mail sorting
 - Ironically just prior to AI winter of the 90’s

New Wave of AI

- Availability of Big Data
- Improved Machine Learning Algorithms
- More Powerful Computing
- Mobile Connectivity

- Pervasive and Revolutionary
 - Personal assistive agent, language translation, self-driving vehicles
...
 - Healthcare, transportation, cybersecurity, commerce, ...
- Promises much social good
- Serious challenges and risks
- Must be developed in a trustworthy manner
 - Reliability and Safety

Narrow vs. General AI

Narrow	General
○ Application specific/ task limited	○ Perform general (human) intelligent action
○ Fixed domain models provided by programmers	○ Self-learns and reasons with its operating environment
○ Learns from thousands of labeled examples	○ Learns from few examples and/or from unstructured data
○ Reflexive tasks with no understanding	○ Full range of human cognitive abilities
○ Knowledge does not transfer to other domains or tasks	○ Leverages knowledge transfer to new domains and tasks
○ Today's AI	○ Future AI?

NSTC[†] Subcommittee on Machine Learning & Artificial Intelligence (ML/AI)

- EOP[‡] -> NSTC -> Committee on Technology -> Subcommittee on ML/AI
- Chartered May 6, 2016 through July 31, 2017

- 7 Functions for AI
 - Advise COT, NSTC, and EOP on Federal efforts
 - Coordinate application in Federal Government operations
 - Support and advise NITRD's work in coordinating Federally funded research
 - Facilitate efforts across Federal agencies to address S&T policy issues, research, testing, standards, education, training, privacy, and outreach
 - Ensure ML/AI is incorporated into Federal efforts to increase diversity of the STEM workforce
 - Publish a strategic plan for unclassified ML/AI R&D (in consultation with NITRD)
 - Report to Federal officials on status of unclassified ML/AI applications across Federal Government, and development and practice in private sector

[†] National Science and Technology Council

[‡] Executive Office of the President

Subcommittee on MLAI - Membership

Co-chair – Michael Garris, Senior Scientist, NIST/DOC


Co-chair – Daniel Oates, Foreign Affairs Officer, STATE

- Departments and Agencies:
 - Department of Commerce (Co-chair)
 - Department of Defense
 - Department of Education
 - Department of Energy
 - Department of Health and Human Services
 - Department of Homeland Security
 - Department of Justice
 - Department of Labor
 - Department of State (Co-chair)
 - Department of Transportation
 - Department of Treasury
 - Department of Veterans Affairs
 - United States Agency for International Development
 - Central Intelligence Agency
 - General Services Administration
 - National Science Foundation
 - National Security Agency
 - Office of the Director of National Intelligence
 - Social Security Administration
- Components of Executive Office of the President:
 - Council of Economic Advisers
 - Domestic Policy Council
 - Office of Management and Budget
 - Office of Science and Technology Policy (Co-chair)
 - Office of the Vice President
 - National Economic Council
 - National Security Council

AI Outreach Events


- AI, Law, and Governance
 - May 24, in Seattle, co-hosted by OSTP, the National Economic Council (NEC), and the University of Washington
- AI for Social Good
 - June 7, in Washington DC, co-hosted by OSTP, the Association for the Advancement of AI (AAAI) and the Computing Community Consortium (CCC)
- Future of AI: Emerging Topics and Societal Benefit at the Global Entrepreneurship Summit
 - June 23, in Palo Alto, co-hosted by OSTP and Stanford University
- AI Technology, Safety, and Control
 - June 28, in Pittsburgh, co-hosted by OSTP and Carnegie Mellon University
- Social and Economic Impacts of AI
 - July 7, in New York, co-hosted by OSTP, NEC, and New York University
- OSTP Request for Information (RFI)
 - Seeking public comment on the topics of the workshops. The RFI received 161 responses.

Strategic Plan – Two Reports



Policy Report

- NSTC / Subcommittee on MLAI
- Drafted May–Sept; Released Oct 2016



R&D Strategic Plan

- NSTC / NITRD - Task Force on AI
- Drafted June–Sept; Released Oct 2016

Third Report: OSTP – “Artificial Intelligence, Automation, and the Economy,” Dec 2016

Preparing for the Future of AI

(23 Recommendations Across Topics)

- **Applications of AI for Public Good**
 - E.g., Health Care, Transportation, Environment, Economic Inclusion, Unemployment, School Dropouts, Homelessness, Global Poverty, ...
 - Social Service Agencies
 - Large societal impact, small R&D budgets, little S&T expertise, big challenges, and big data
- **AI and Regulation**
 - Protect public from harm; ensure fairness and economic competition
 - Regulation of AI technologies should be adaptive and incremental
 - Balance: (lower costs and barriers to innovation) without (adversely impacting safety or market fairness)
- **Monitor AI Progress**
 - Milestones that represent / foreshadow leaps in AI capability
- **New AI Workforce**
 - Training new AI Specialists & AI Users via STEAM (Science, Technology, Engineering, Art & Design, Mathematics)
 - Diversity challenge (gender and race)
- **AI, Automation and the Economy**
 - Minimize impact of certain types of jobs and sectors
 - Impact of AI on number of jobs (more or less?) is uncertain and highly debated; research needed
 - White House report, “Artificial Intelligence, Automation, and the Economy,” Dec. 2016
- **Fair, Accountable, Safe AI**
 - Ethics training and engineering
 - Transparency – decisions must be explainable
 - AI systems are probabilistic – requires testing. (How good is good enough?)
- **Global Considerations of AI**
 - AI is changing the World: health care, manufacturing, information and communication technologies, disaster preparedness and response, climate change, wildlife trafficking, the digital divide, jobs, smart cities, autonomous vehicles ...
 - Special considerations: AI in Cybersecurity, and AI in Defense

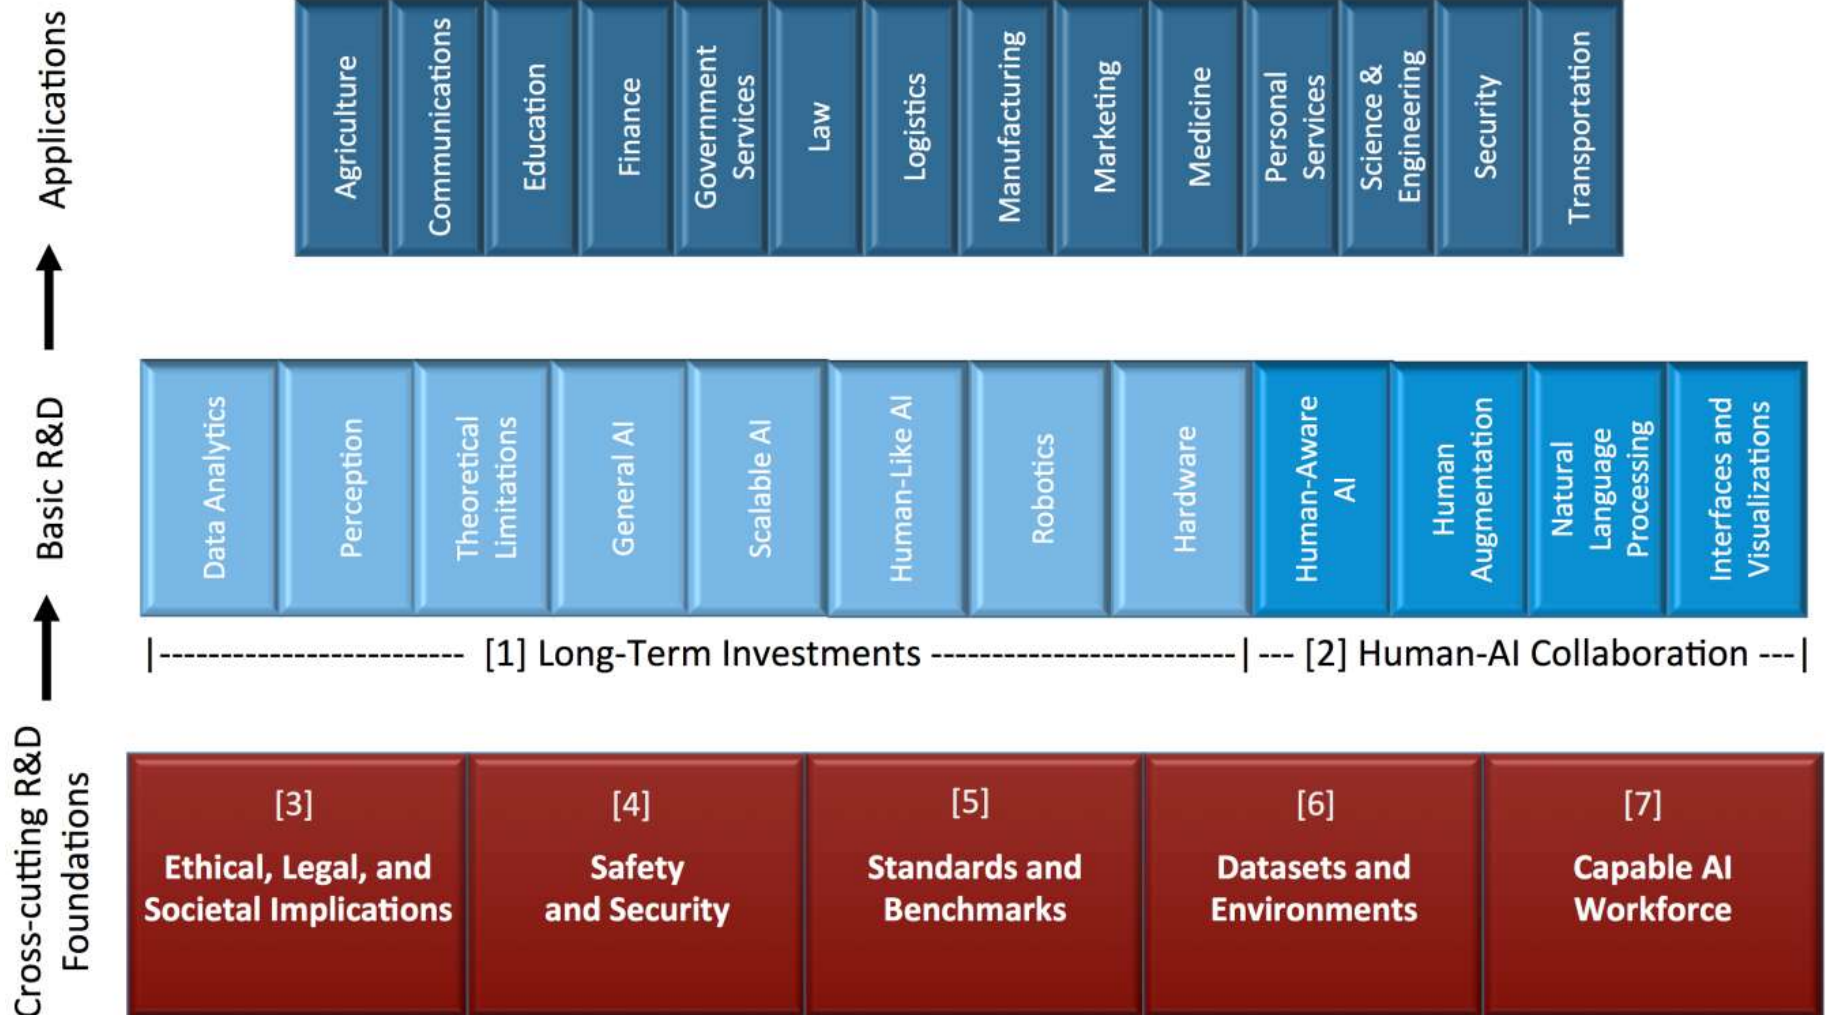
NITRD Task Force on AI

Co-Chair Lynne Parker Division Director Information and Intelligent Systems National Science Foundation	Co-Chair Jason Matheny Director Intelligence Advanced Research Projects Activity
Milton Corn National Institutes of Health	Nikunj Oza National Aeronautics and Space Administration
William Ford National Institute of Justice	Robinson Pino Department of Energy
Michael Garris National Institute of Standards and Technology	Gregory Shannon Office of Science and Technology Policy
Steven Knox National Security Agency	Scott Tousley Department of Homeland Security
John Launchbury Defense Advanced Research Projects Agency	Faisal D'Souza , <i>Technical Coordinator</i> National Coordination Office for Networking and Information Technology Research and Development
Richard Linderman Office of the Secretary of Defense	

With agency input/approval from:

NSF, NIH, DoD, DARPA, OSD, NIST, DOE, NASA, DHS, NSA, IARPA, DOT, EPA, NIJ, FBI, ODNI

Organization of the AI R&D Strategic Plan



National R&D Strategic Plan

(7 Strategies)

1. Make Long-Term Investments in AI Research

- Data methods for discovery; perceptual and theoretical capabilities; general, scalable, human-like AI; robotics; hardware ...

2. Effective Methods for Human-AI Collaboration

- Human-aware & augmentation; visualization; AI-human interfaces; improved natural language processing

3. Ethical, Legal, Societal Implications

- Fairness, transparency, accountability, ethical, ...

4. Safety and Security of AI Systems

- Explainable, trustworthy, verification/validation, secured against attacks

5. Develop Shared Public Datasets and Environments for AI Training and Testing

- Open data and promote innovation in new application domains

6. Measure and Evaluate Technologies through Standards and Benchmarks

- Competition, interoperability, reliability, ...

7. Better Understand the National AI R&D Workforce Needs

- Pipeline for future AI Specialists and AI Users (diversity & STEAM)

USG Role in AI

- Investing in basic and applied R&D
- Early customer for AI technologies and applications
- Release public datasets, create real-world testbeds, support pilot projects
- Sponsor incentive prizes and Grand Challenges
- Fund rigorous evaluations of AI applications to measure progress and impact
- Create a policy, legal, and regulatory environment allowing innovation while protecting the public

MLAI-SC AI R&D Coordination

- Implementation of the *National AI R&D Strategic Plan* recommendations
- Fundamental R&D topics (e.g.)
 - Machine Learning
 - AI planning, reasoning, and decision making
 - Human-AI interaction
 - AI evaluation through metrics, standards, benchmarks, testbeds, and datasets
- Focus areas (e.g.)
 - Reliability, safety, scalability, privacy, ethics, harmful bias avoidance, and interpretability/explainability
- Coordination with other Federal R&D coordinating bodies (e.g.)
 - Cybersecurity, Privacy, Robotics, Software Design, Cyber Physical Systems, High Performance Computing, ...



Advancing Artificial Intelligence @ NIST



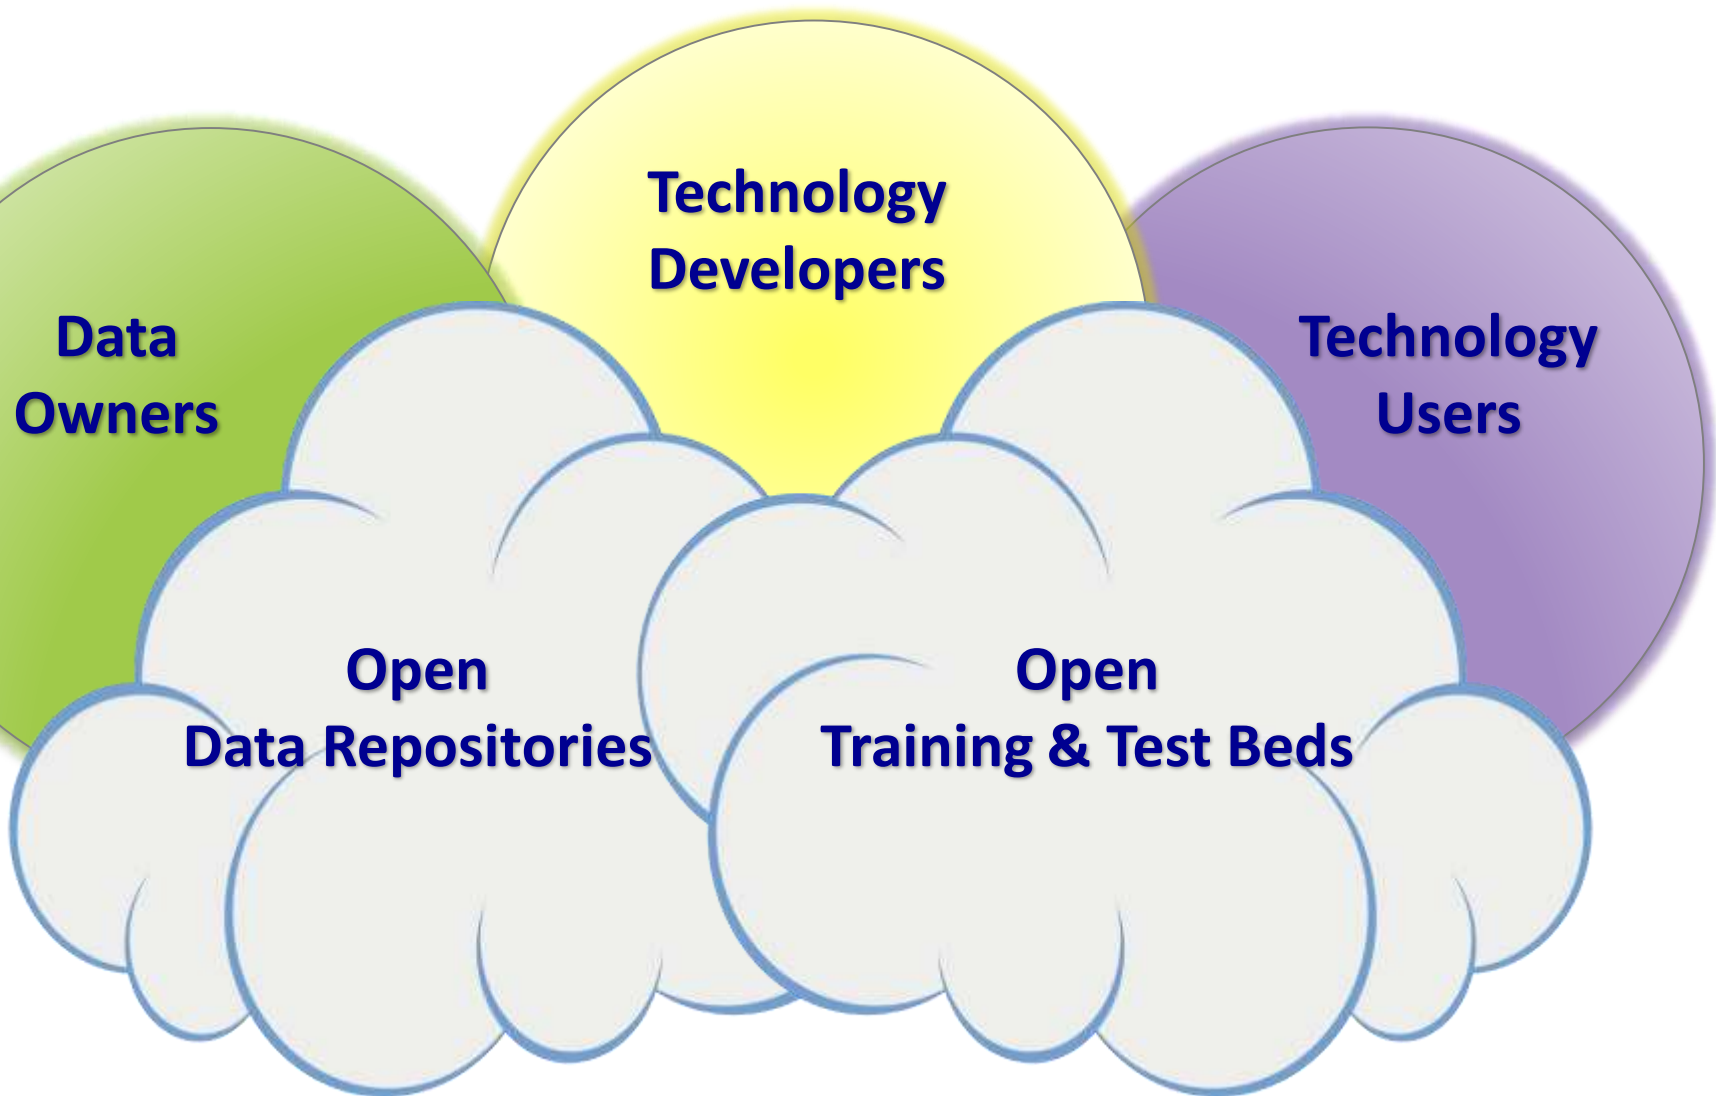
NIST cultivates trust in technology by developing and deploying standards, tests and metrics that make technology more secure, usable, interoperable and reliable, and by strengthening measurement science.

Trustworthiness

(a challenge to capturing the benefits of AI)

- AI must be assured to be:
 - Safe
 - Secure
 - Reliable
 - Private
 - Scalable
 - Explainable
 - Resilient
 - Not Harmfully Biased
- How?
 - Provable by design
 - Methods of empirical testing

Benchmarks / Challenge Problems



NIST

Drive AI Competencies

Reliability



Scalability



Harmful Bias



Safe



Ethical



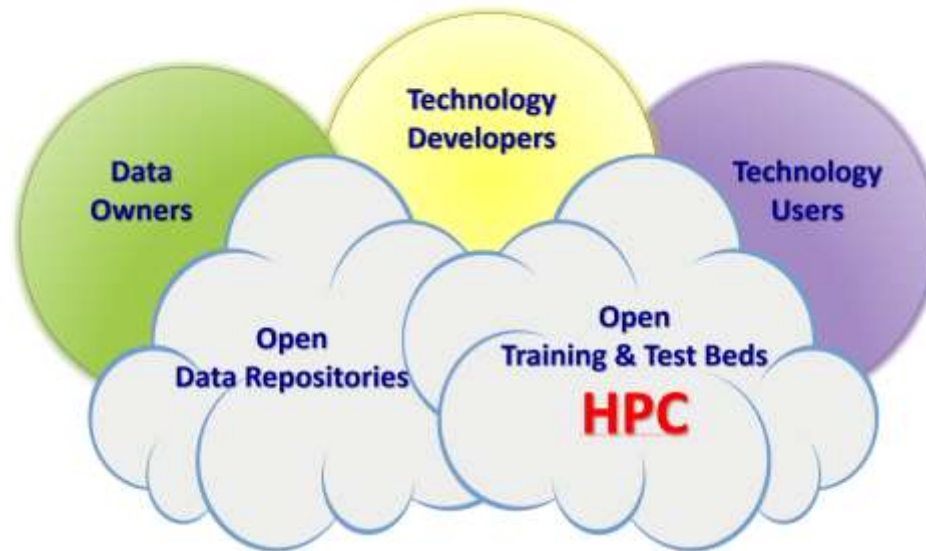
Explainability



Machine Learning



HPC Facilities for Machine Learning



- Reusable / Adaptable
 - Provision / Configure -> stand up and tear down
- Scalable
 - Seamless -> virtual (cloud) with physical (GPUs)
- Public
 - Productization -> creation of intellectual property
 - Small Business



Questions, Comments, & Suggestitons?

Michael Garris
(mgarris@nist.gov)

Senior Scientist / National Institute of Standards and Technology
Co-Chair / National Science and Technology Council,
Subcommittee on Machine Learning and Artificial Intelligence