AI FOR SCIENCE TOWN HALL

AI for Science

Jeff Nichols, ORNL

Rick Stevens, ANL

Kathy Yelick, LBNL

U.S. DEPARTMENT OF ENERGY | Office of Science

Why AI? What can it do for Science?

U.S. DEPARTMENT OF ENERGY | Office of Science

Data Analytics via Supervised Deep Learning

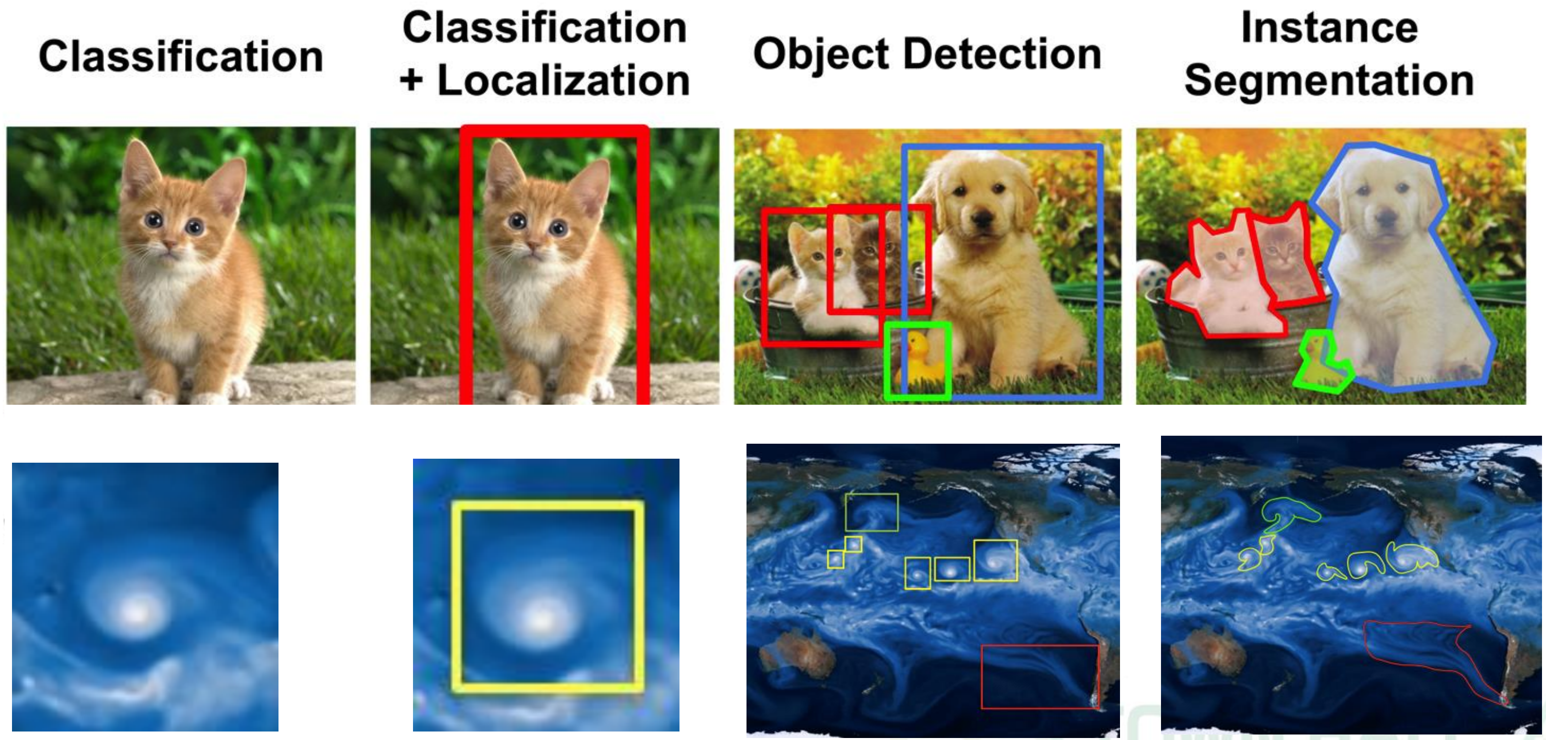

Extending image-based methods to complex, 3D, scientific data sets is non-trivial!

Big Data, Big Model, and Big Iron

Predicted Extreme Weather

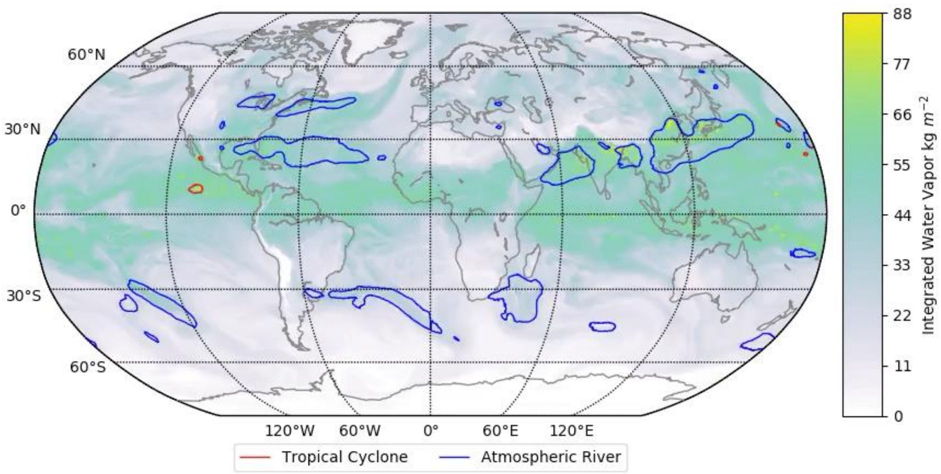

Ground Truth Extreme Weather

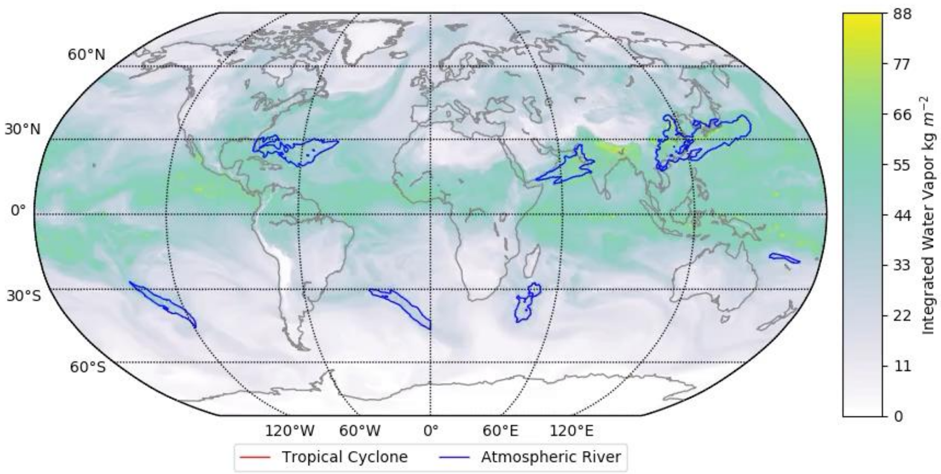

- **Deep learning results are smoother than heuristic labels**
- **Achieved over 1 EF peak on OLCF Summit: Gordon Bell Prize in 2018**

Thorsten Kurth, Sean Treichler, Joshua Romero, Mayur Mudigonda, Nathan Luehr, Everett Phillips, Ankur Mahesh, Michael Matheson, Jack Deslippe, Massimiliano Fatica, Prabhat, Michael Houston

Graph NNs to Classify Neutrinos


17th IEEE International Conference on Machine Learning and Applications
IEEE December 17-20, 2018, Orlando, FL, USA AMLA

BEST PAPER AWARD

Graph Neural Networks for IceCube Signal Classification

Nicholas J Choma (New York University); Federico Monti (Università della Svizzera italiana); Lisa Gerhardt (Lawrence Berkeley National Laboratory); Tomasz Palczewski (University of California, Berkeley); Zahra Ronaghi (Lawrence Berkeley National Laboratory); Prabhat (Lawrence Berkeley National Laboratory); Wahid Bhimji (Lawrence Berkeley National Lab); Michael Bronstein (Imperial College); Spencer R. Klein(University of California, Berkeley); Joan Bruna (Courant Institute of Mathematical Sciences, NYU, USA)

Mehmed Kantardzic, Conference Co-Chair Moamar Sayed-Mouchaweh, Conference Co-Chair

- Apply graph convolutions to irregular, 3D detector grid
- Increase sensitivity of IceCube detector: 6.3x more events
- And improve Signal-to-Noise ratio by 3x

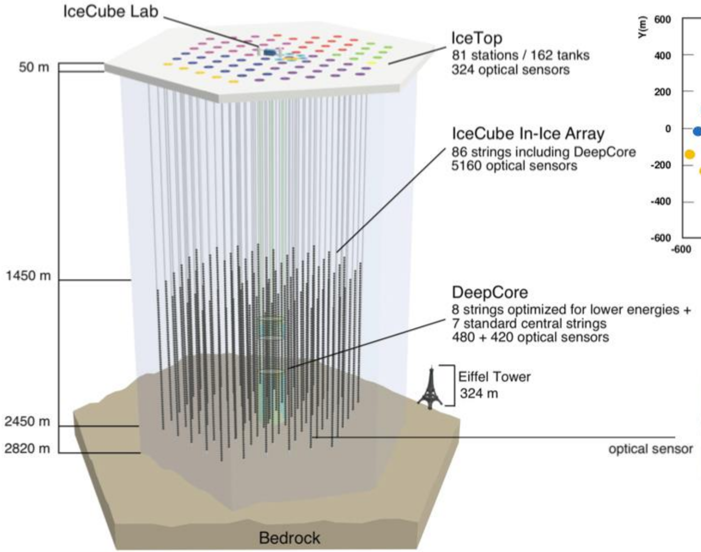

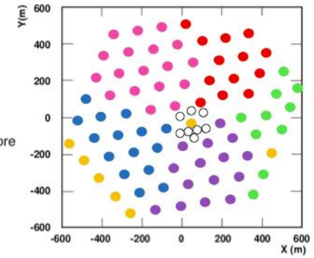

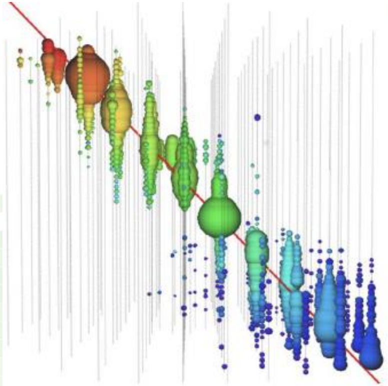

Light deposition for muons bundles

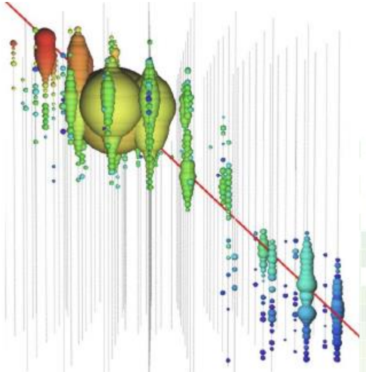

Stochastic light emission from a single high-energy muon

Contributors: Nick Choma, Joan Bruna, Federico Monti, Michael Bronstein, Spencer Klein, Tomasz Palczewski, Lisa Gerhardt, Wahid Bhimji and IceCube collaboration

U.S. DEPARTMENT OF ENERGY | Office of Science

Washington DC Town Hall
October 22-23

Materials Genome Initiative: generating data for learning

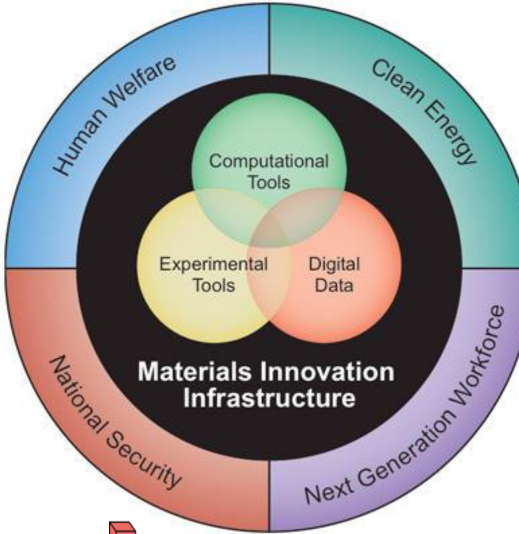

Rotated image **CNN filter output**

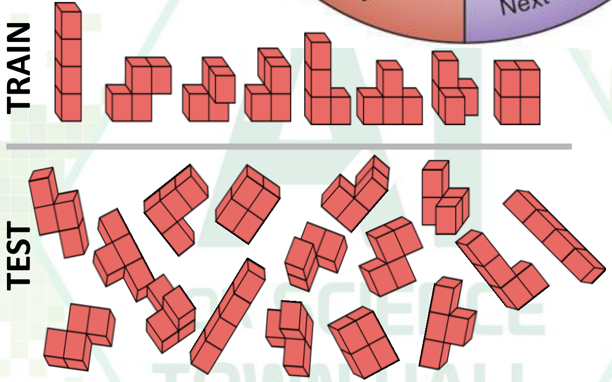

Tensor field networks with 3D translation- and 3D rotation-equivariance

Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, Patrick Riley

How to build next generation instruments and simulations with AI in mind?

(Smaller, faster, cheaper?)

Can we use existing data sets for new discoveries?

How will AI methods change?

Generative Adversarial Neural Nets make "Everybody Dance Now"

Caroline Chan, Shiry Ginosar, Tinghui Zhou, Alexei A. Efros, UC Berkeley

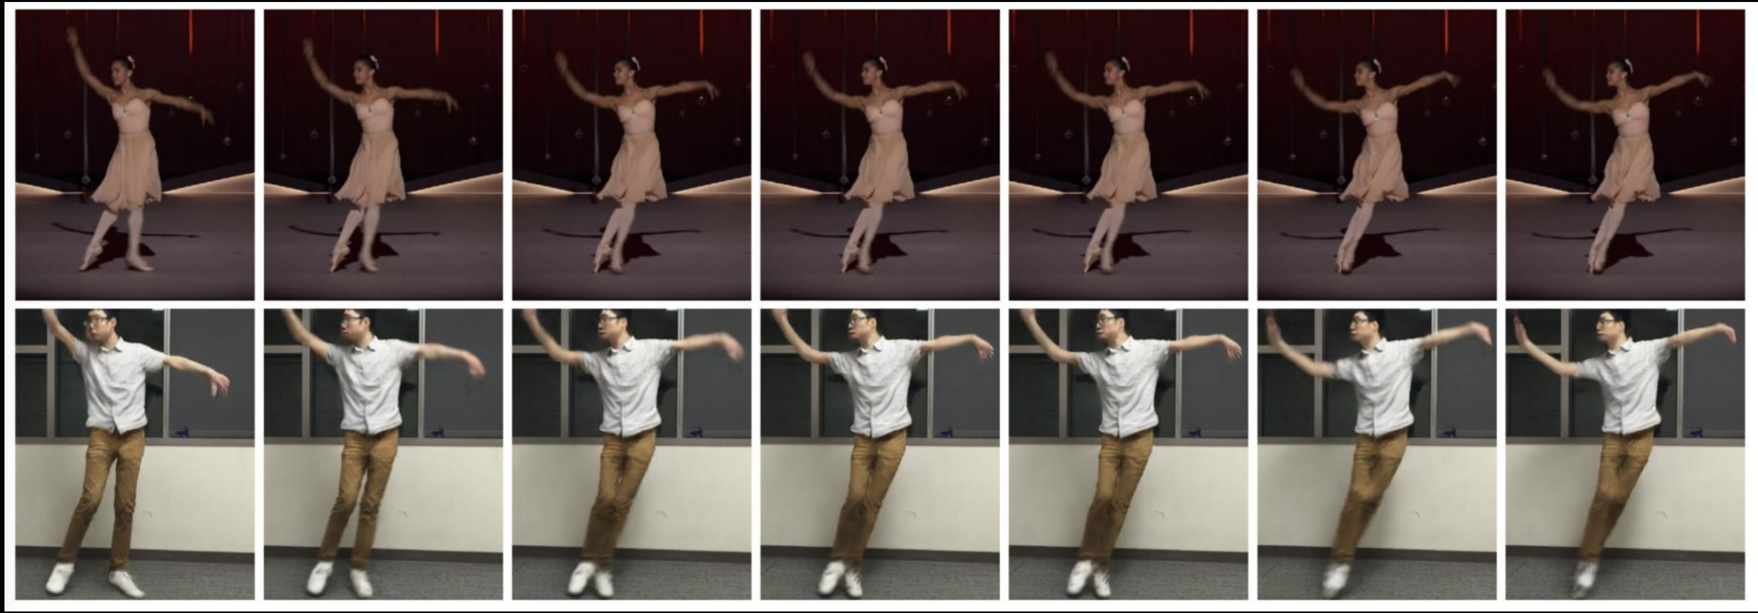

GANs to build convergence maps of weak gravitational lensing

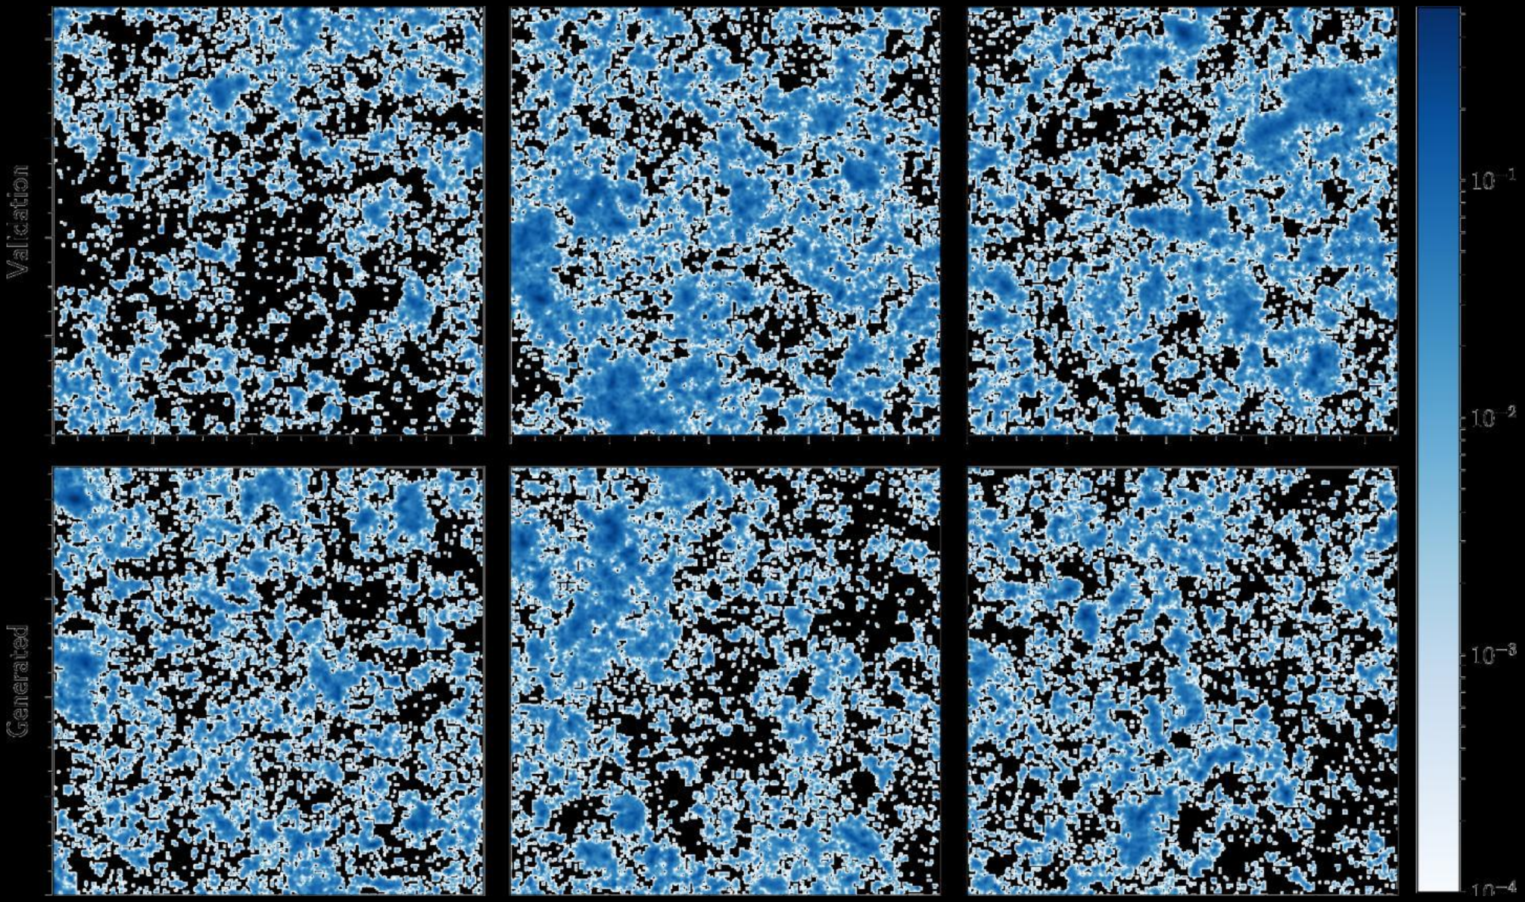

CosmoGAN: Mustafa Mustafa, Deborah Bard, Wahid Bhimji, Zarija Lukić, Rami Al-Rfou, Jan M. Kratochvil

What scientific data sets can AI create?

(Test instruments, evaluate theories, approximate expensive simulations and experiments)

What level of confidence will we need in AI?

Deep Reinforcement Learning to play games of strategy

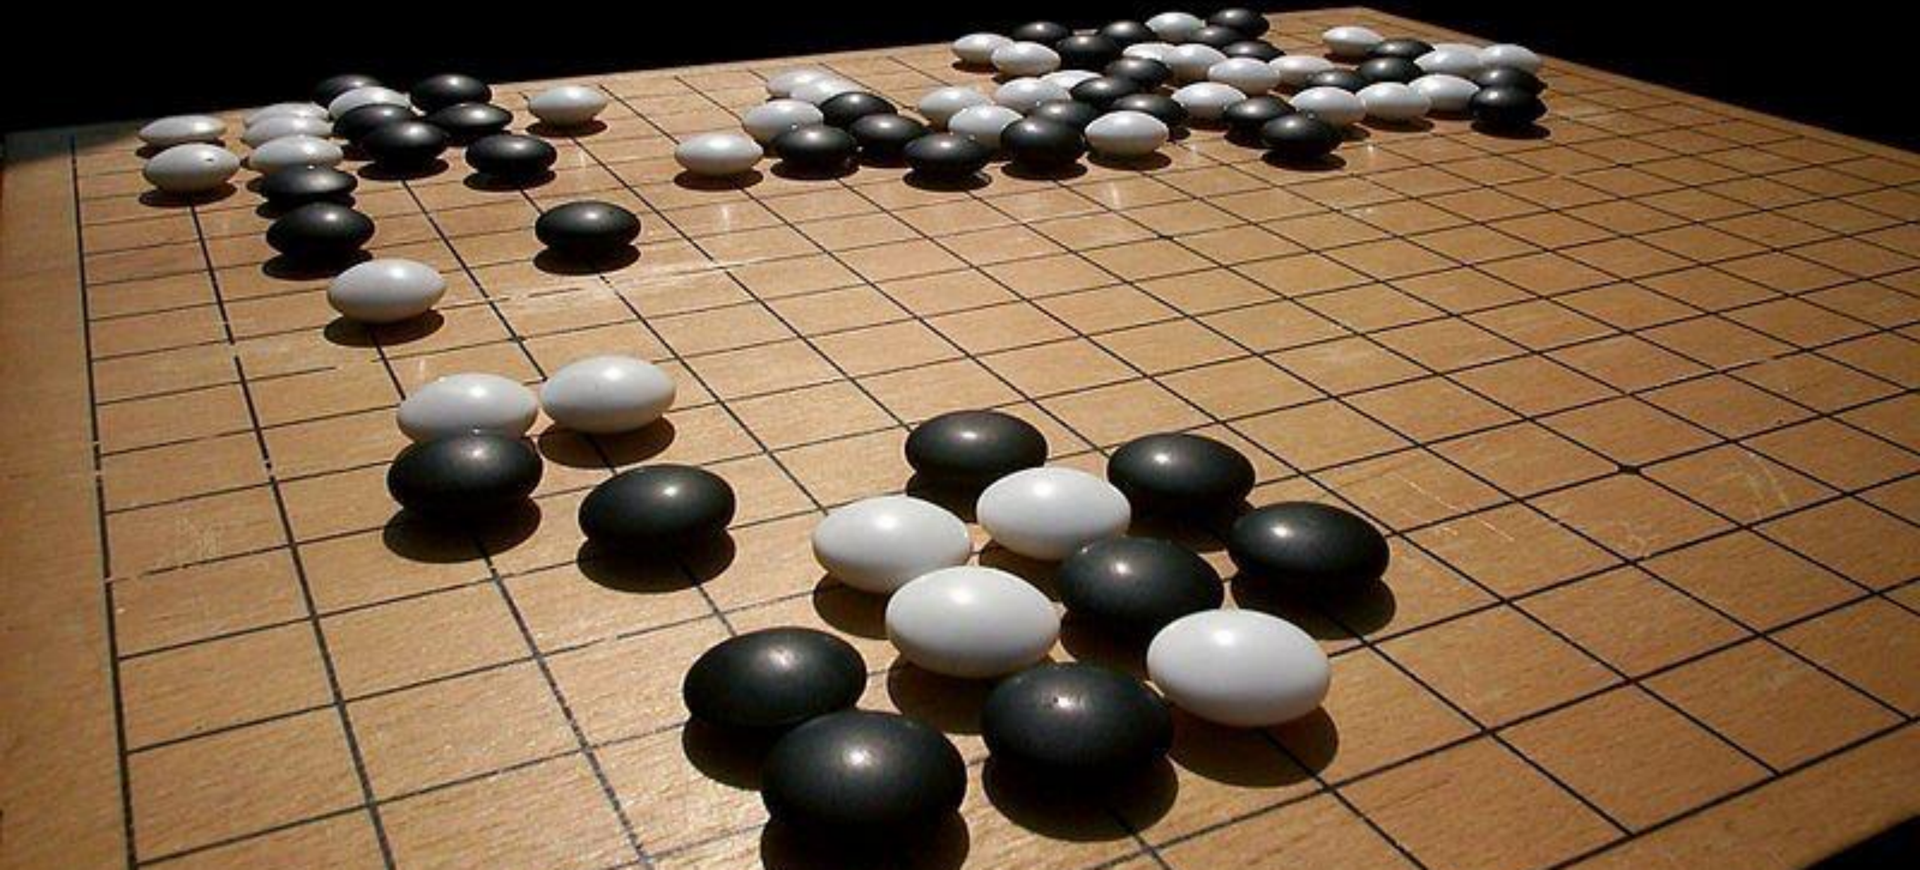

RL to control complex systems like traffic

Baseline Scenario

RL Vehicles

Kathy Jang, Ben Remer, Eugene Vinitsky, Logan Beaver,
Alexandre Bayen, Behdad Chalaki, Andreas A. Malikopoulos

How can AI navigate complex spaces?

(Optimize energy systems, IT systems, find best path through space of experiments)

How do the systems need to change to support automation?

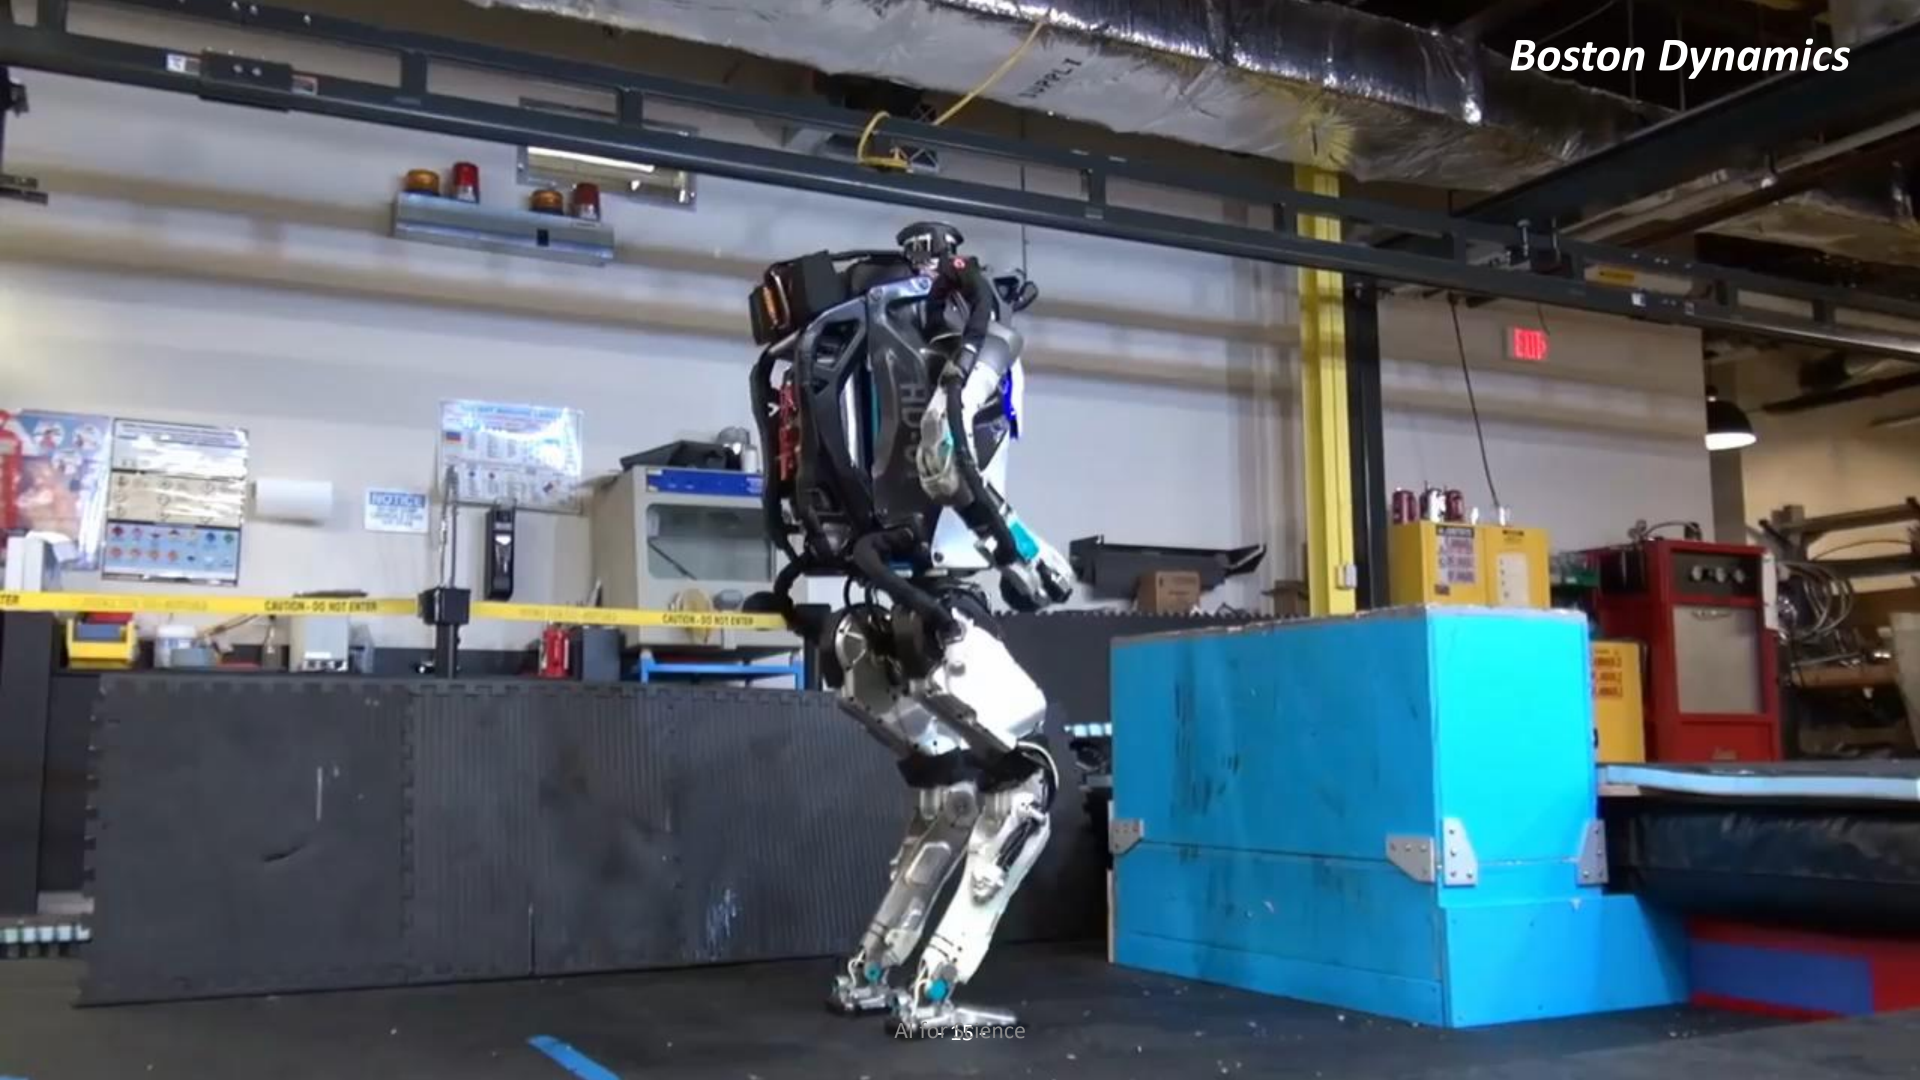
Boston Dynamics
CAUTION - DO NOT ENTER

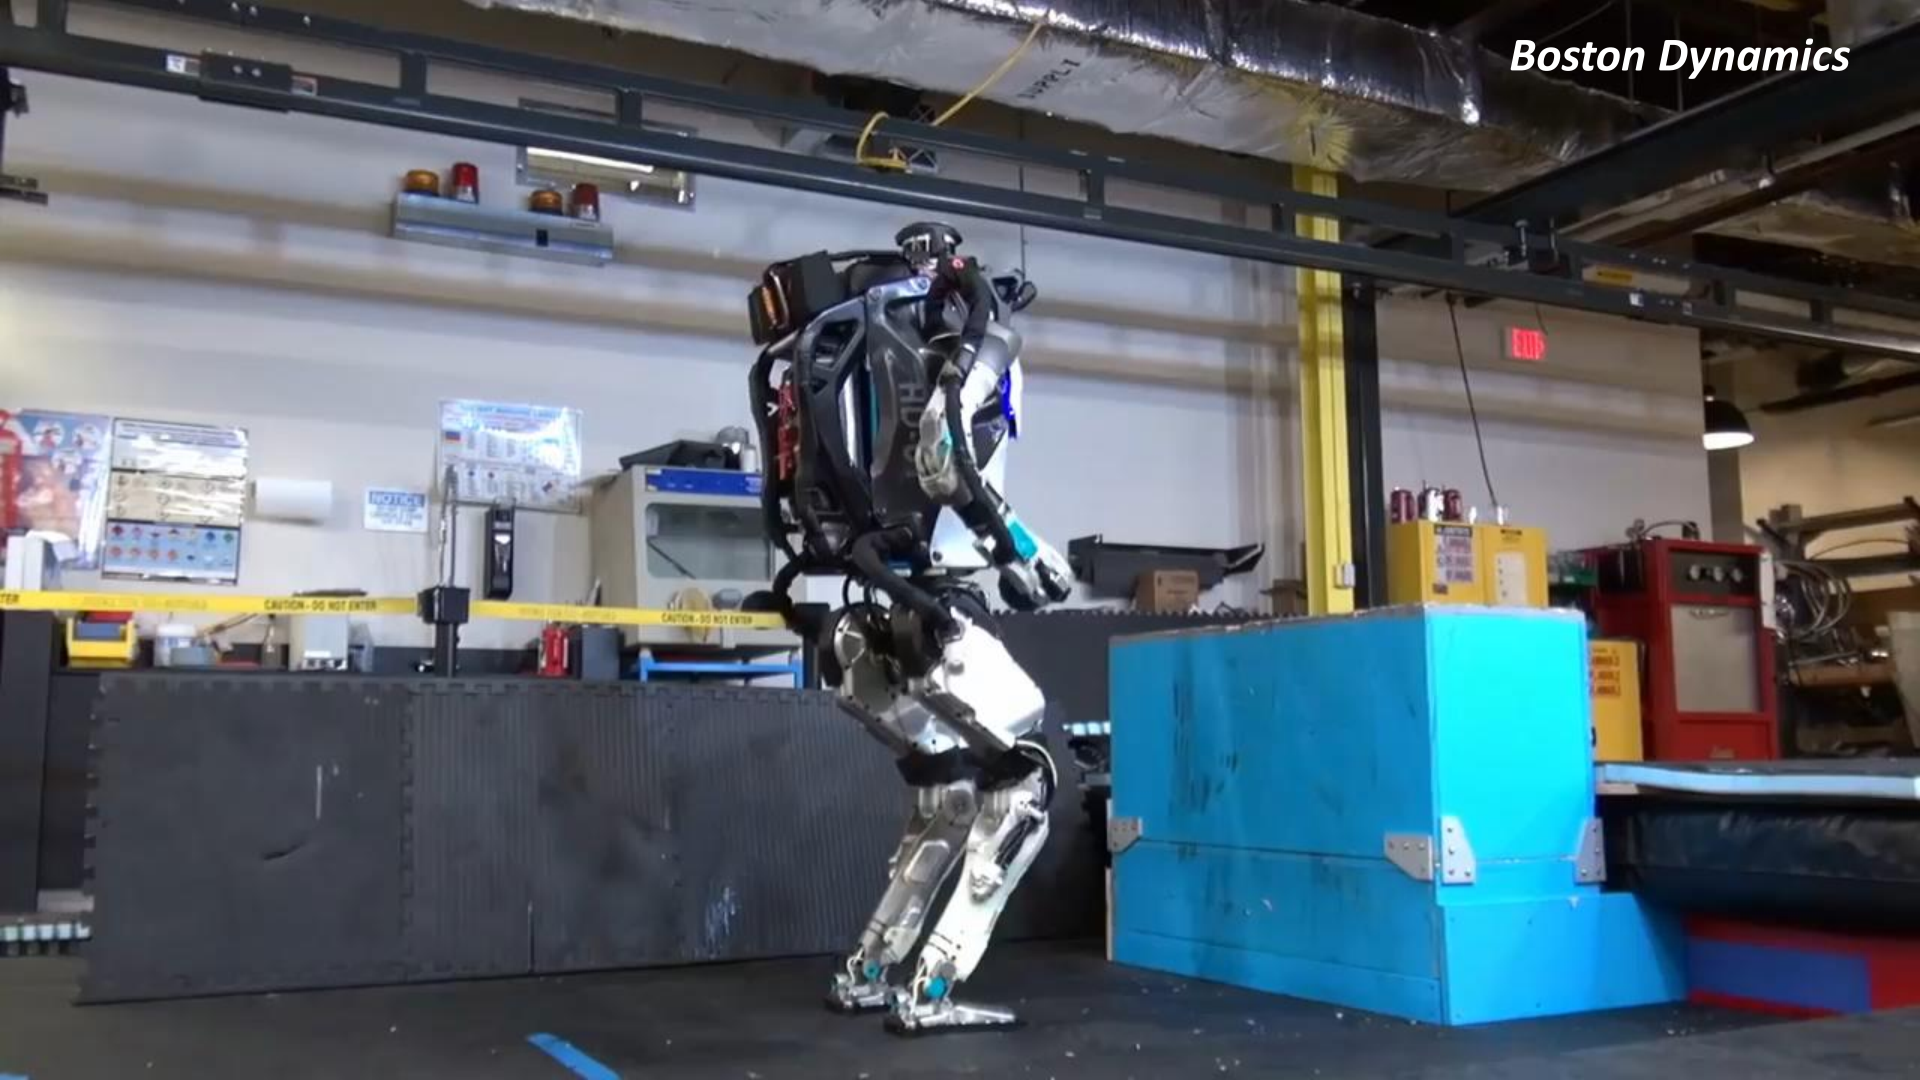
Boston Dynamics
CAUTION - DO NOT ENTER

Robotics and precision control in science

MassSpec robot at JGI

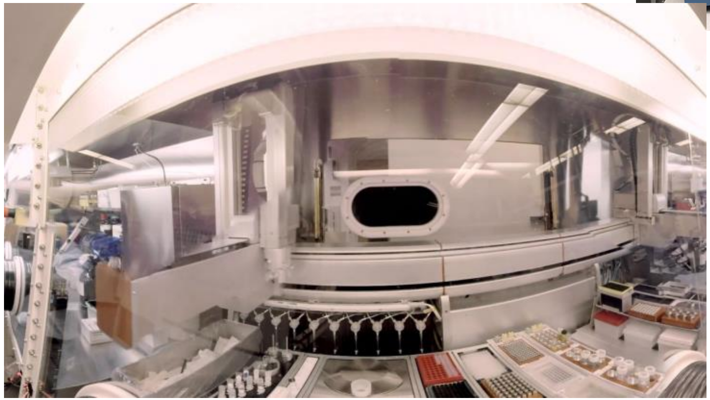

Nanoparticle Robot at the Molecular Foundry

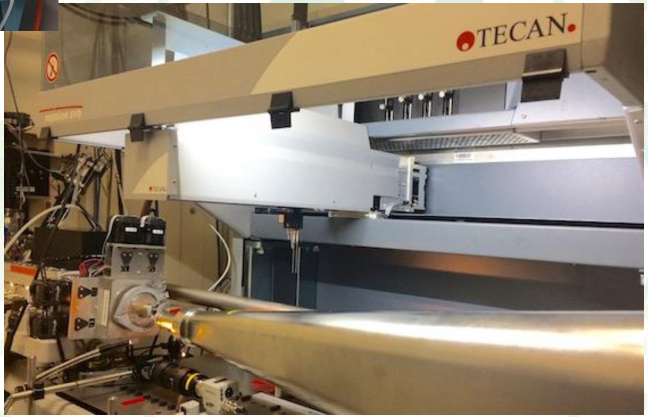

Robot at SYBLIS beamline at ALS

How will robotics and automation transform user user facilities?

(Lab in the cloud, remote access, higher throughput)

How will AI-trained robots enable new instruments?

(Particle accelerators, fusion, etc.)

What level of confidence will we need in AI?

Extracting signals from noisy data: “Visual Microphone”

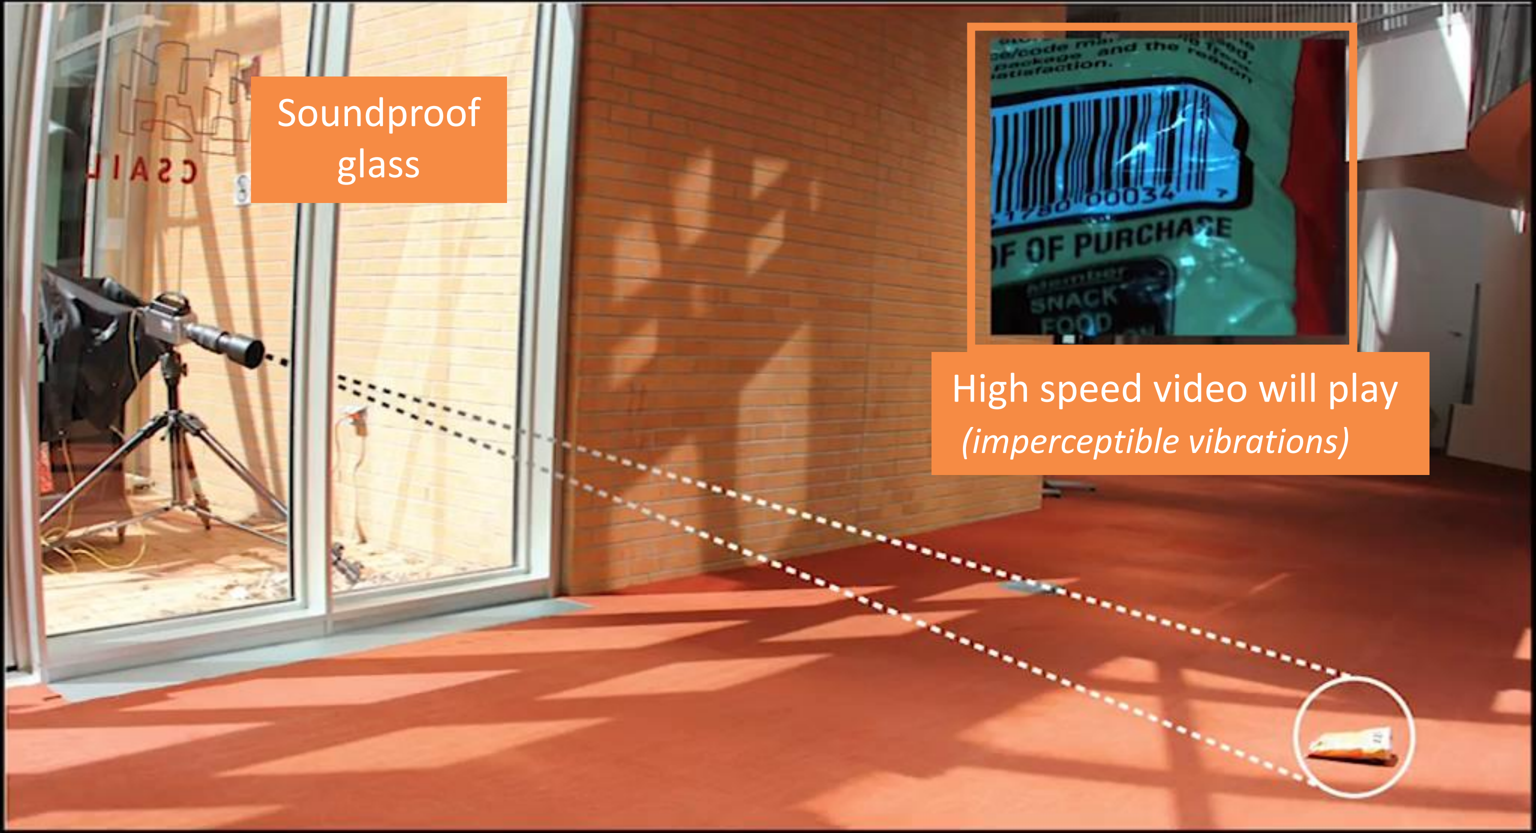

Abe Davis, M Rubinstein, N Wadhwa, GJ Mysore, F Durand, WT Freeman, MIT

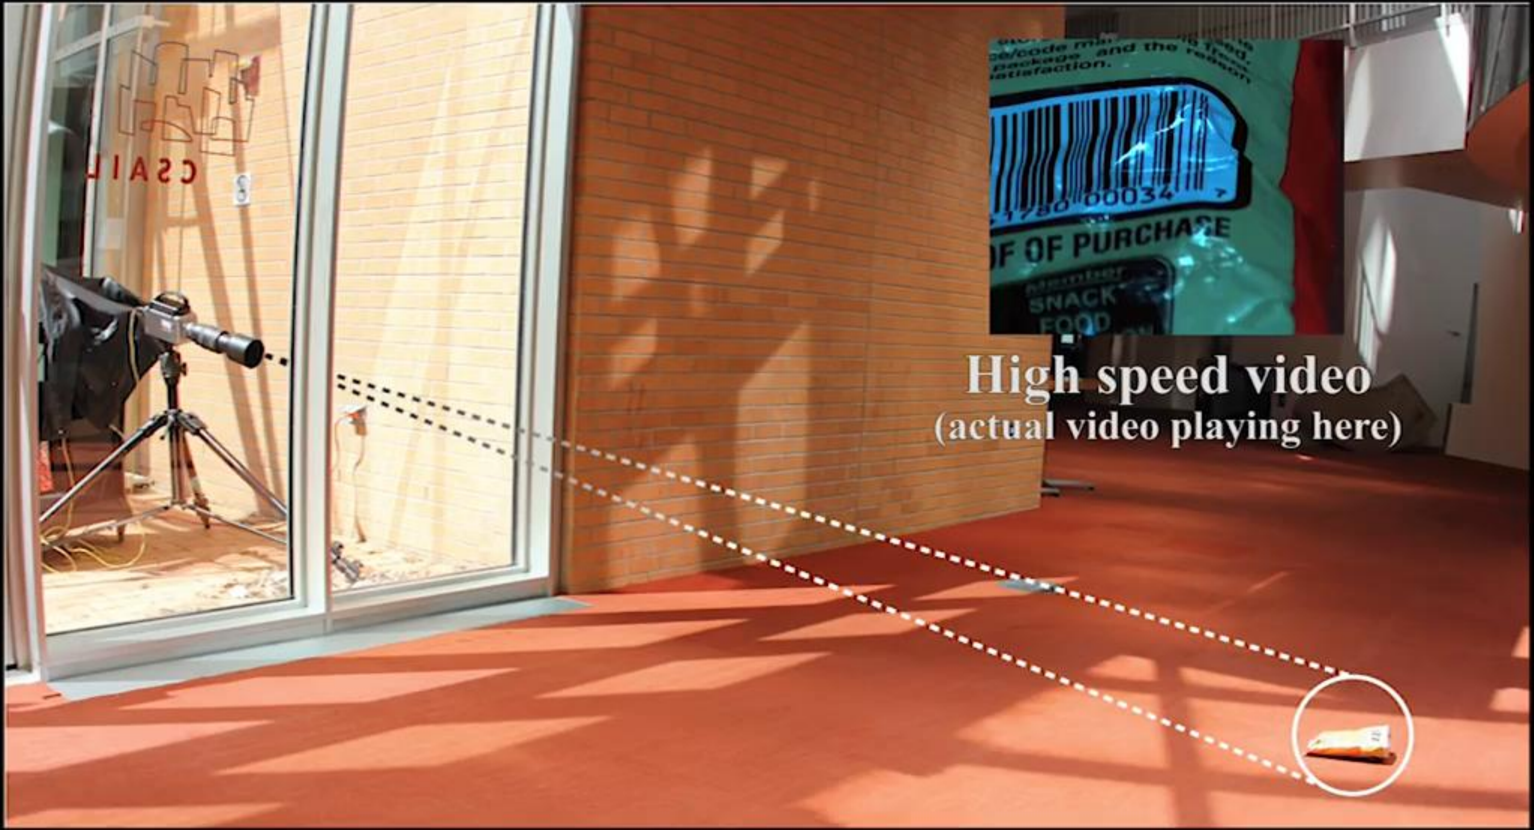
High speed video
(actual video playing here)

Finding Structure in Sparse, Noisy Data

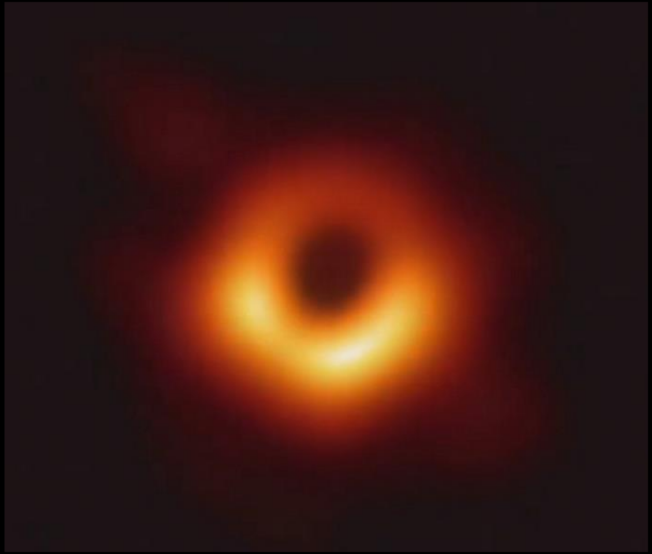

First Image of a Black Hole

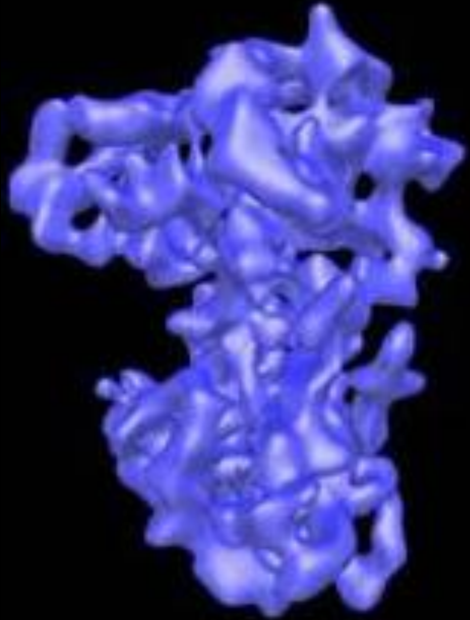

Retinoblastoma Protein

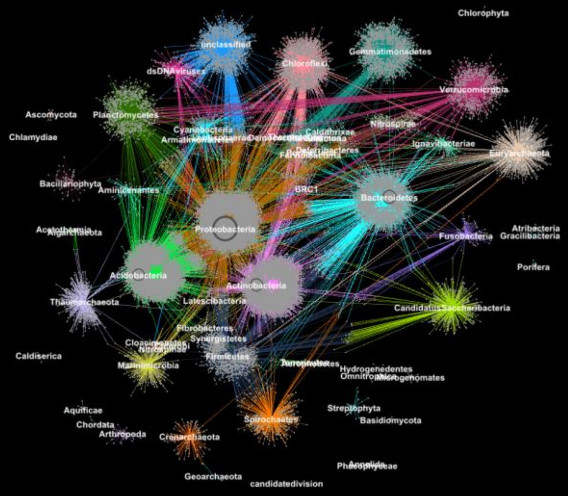

New protein clusters

Where can AI find smaller signals in noisier data?

How can AI learn physically realistic models and design processes?

How can AI help?

Data Analytics

- Classification
- Regression
- Clustering
- Dimensionality Reduction

Inverse problems

- Model reconstruction
- Parameter estimation
- Denoising

Surrogate models

- Approximate expensive simulations
- Approximate experiments
- Fill in missing models in simulations

Design and control

- Optimize design of experiments
- Control instruments
- Navigate state spaces
- Learn from sparse rewards

A Vision for the Future of AI in Science

AI for Science Vision: 2020 ⟹2030

- AI will enable us to attack **new problems**
- AI becomes **equal partners** to modeling and simulation and data analysis
- AI will **enable experimentalists to harness the power of Exascale** computing
- AI will power **automated laboratories** and change the nature of experimental science
- AI will need **new computing architectures**, **new software environments**, **new policies** and create **new user communities** and new ways of dissemination
- AI **will improve how DOE laboratories operate** and how work is done

Things we can do in Science with AI now

Learn predictive models from data without relying upon theory or deep mechanistic understanding

Example: predicting materials and chemistry properties

Learn approximate solutions to inverse problems where we have data and models are not available or are inefficient

Example: phase retrieval in coherent x-ray imaging

Generate large collections of synthetic data that models real data

Example: synthetic sky in cosmology

In Ten Years…

- **Learned Models Begin to Replace Data**
  - queryable, portable, pluggable, chainable, secure
- **Experimental Discovery Processes Dramatically Refactored**
  - models replace experiments, experiments improve models
- **Many Questions Pursued Semi-Autonomously at Scale**
  - searching for materials, molecules and pathways, new physics
- **Simulation and AI Approaches Merge**
  - deep integration of ML, numerical simulation and UQ
- **Theory Becomes Data for Next Generation AI**
  - AI begins to contribute to advancing theory
- **AI Becomes Common Part of Scientific Laboratory Activities**
  - Infuses scientific, engineering and operations

AI Driven Autonomous Laboratory Cluster

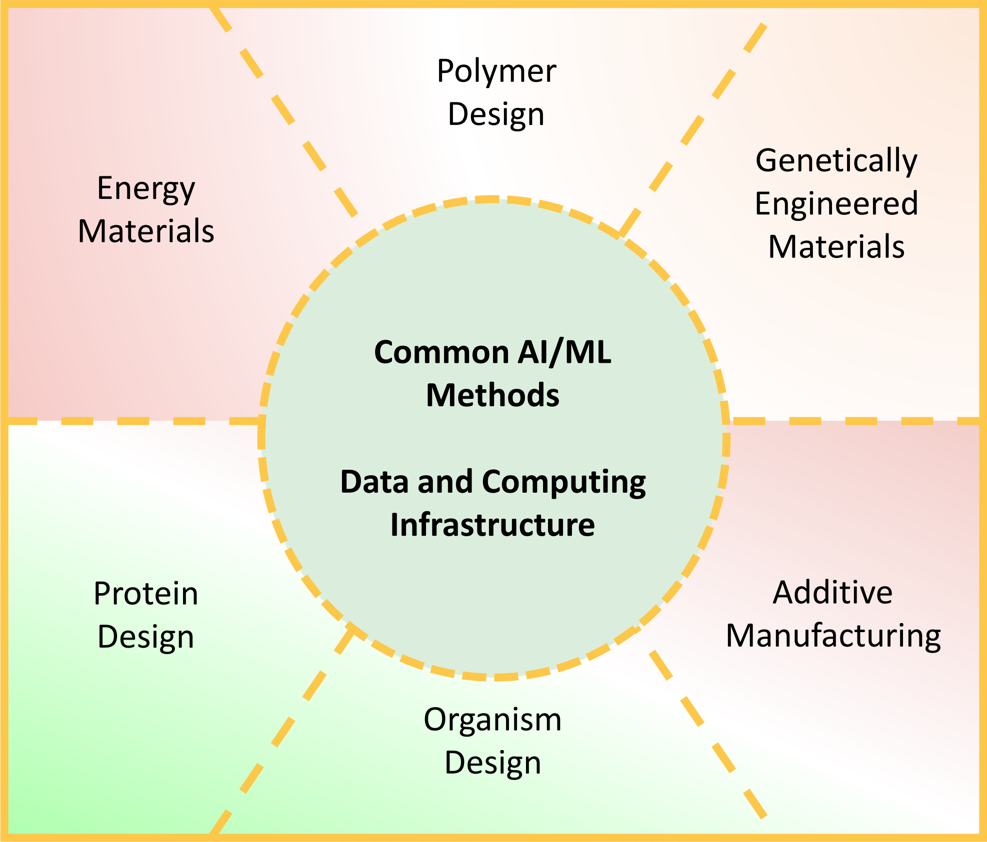

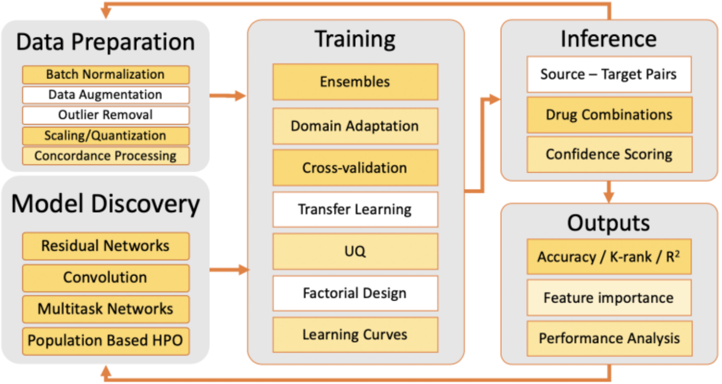

Layered workflow combining AI, HPC and HTS

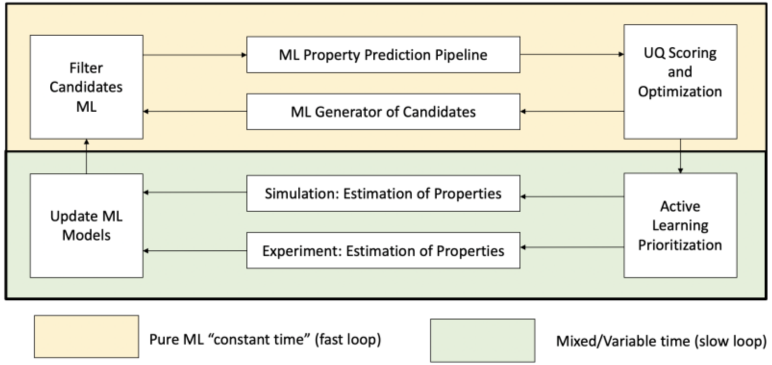

Connecting HPC and AI

In addition to partnerships in AI applications, there are considerable opportunities in foundational methods development, software and software infrastructure for AI workflows and advanced hardware architectures for AI, below we highlight some ideas in the HPC + AI space

- Steering of simulations
- Embedding ML to simulation methods
- Customized computational kernels
- Tuning applications parameters
- Generative models to compare with simulation
- Student (AI) Teacher (Sim) models ⟹learned functions
- Guided search through parameter spaces
- Hybrid architectures HPC + Neuromorphic
- Many, many more

Generative Models

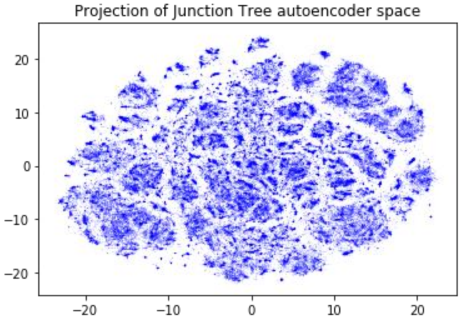

Learned Function Accelerators

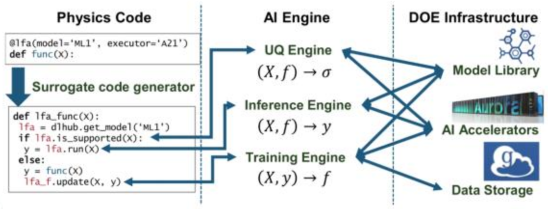

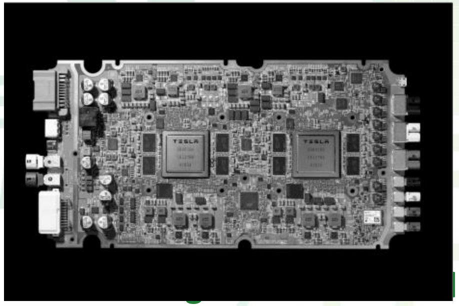

The Landscape of AI Research and Applications

U.S. DEPARTMENT OF ENERGY | Office of Science

AI: a fundamental shift in the economic and military landscape

- Executive Order brings focus to national strategy and government engagement
- Industry focuses on developing AI-based products for business, especially social, financial, health and security
- Universities focus on basic research and education
- DOE has a unique role
  - Mission-driven development and application of AI/ML, i.e., innovation in, for example
    - Science
    - Energy
    - National security
  - Build on its HPC mission
  - Large-scale scientific data for research
  - Talent development

Vision: ***Transform DOE into a world-leading AI enterprise*** by accelerating the research, development, delivery, and adoption of AI.

Observations on the international AI landscape

Of the 35 countries that have AI strategies, only three stand out, the U.S., the E.U. and China.

- The U.S.
  - Leads in research, development and talent (education)
  - Based on historical investments in education, laboratories and the business environment
- China
  - Leads in overall adoption of AI and the collection and use of data
  - Is investing heavily
  - Quality and development is increasing rapidly
- The E.U.
  - Has the most researchers
  - Does not translate this into innovation effectively

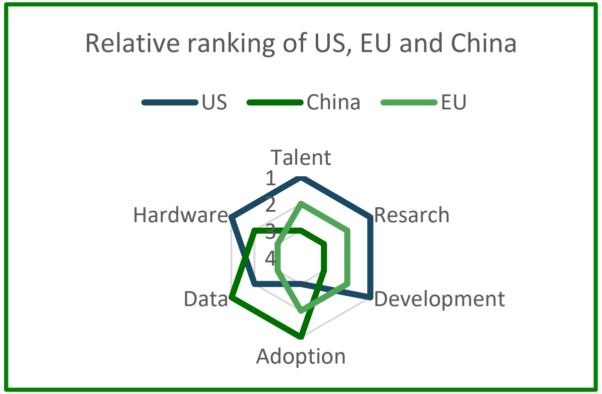

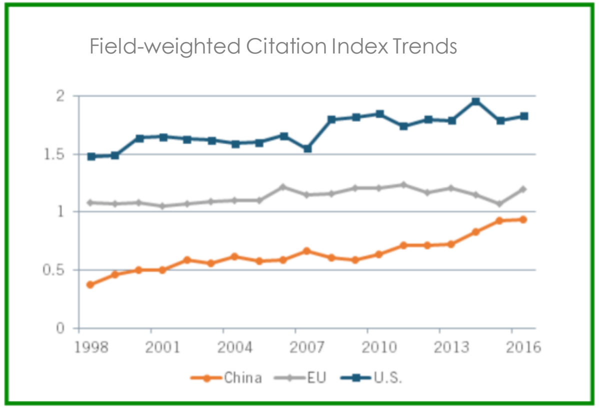

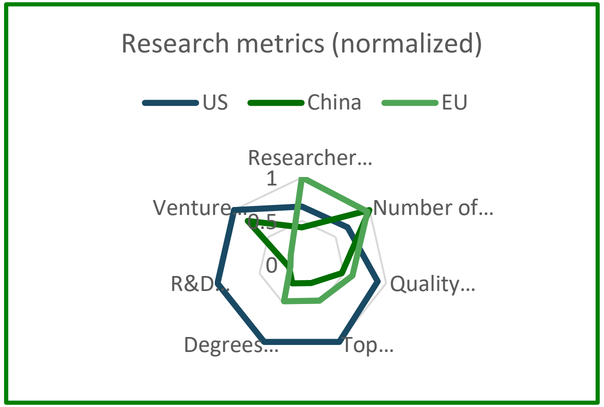

Source: "Who Is Winning the AI Race," Report, Center for Data Innovation

The AI/ML Research Landscape (Measured by Publications)

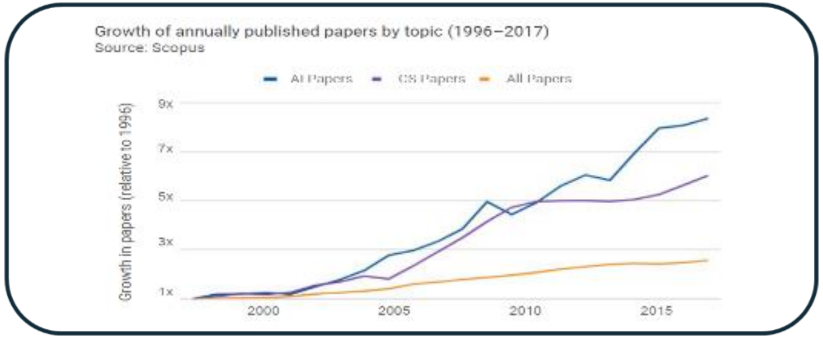

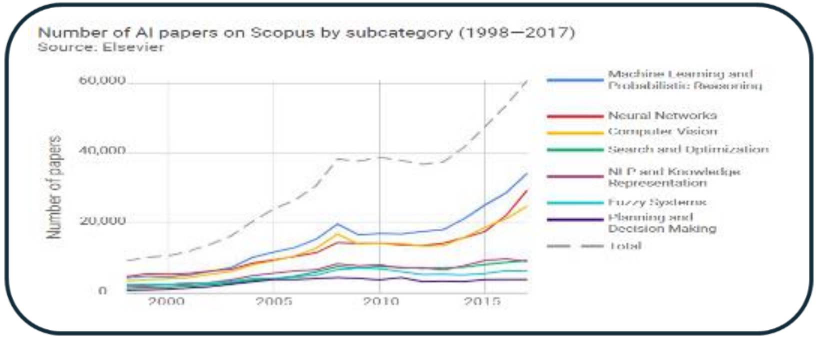

Neural Network Publications
(2016 – Present)

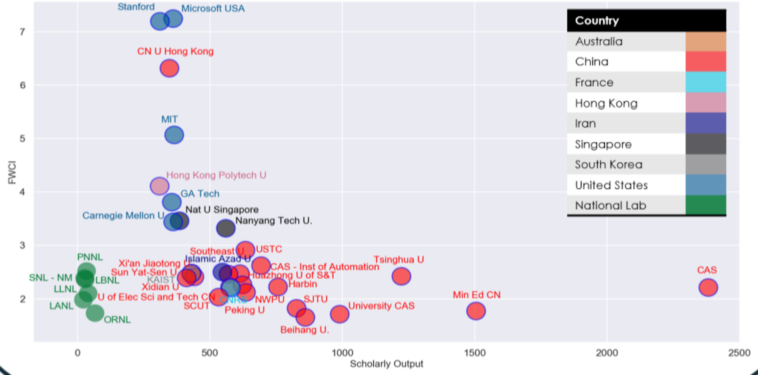

Machine Learning Publications
(2016 – Present)

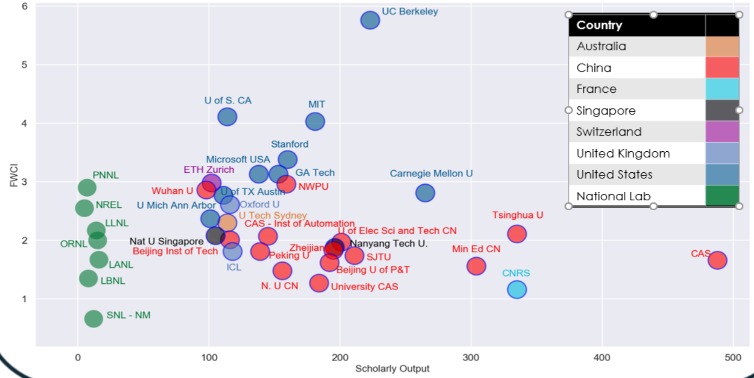

U.S. DEPARTMENT OF ENERGY | Office of Science

The Business Landscape

- Business must incorporate AI
  - The "Big 9" dominate, but don't discount traditional business
  - $7.4B in start-up investments in 488 deals in 2019/Q2 (over $12B in the past 6 months)
  - $803M in "AI for cybersecurity" VC in last six months
- Barriers to insertion
  - Understanding: 37% of executive feel their employees understand the importance of data
  - Trust:
    - 49% of U.S. consumers would trust AI-generated advice for retail,
    - 38% would trust AI-generated advice for hospitality, while only
    - 20% would trust AI-generated advice for healthcare and
    - 19% for financial services
    - Example: 33% of US and 85% of Chinese healthcare professionals have implemented AI into their practice, compared to a 5-country average of 46%.
- Need a consistent approach to regulatory (data and sensitive technologies)

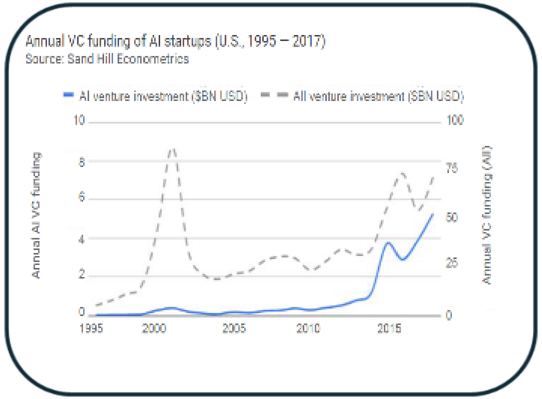

DOE is building on a record of success delivering HPC capabilities

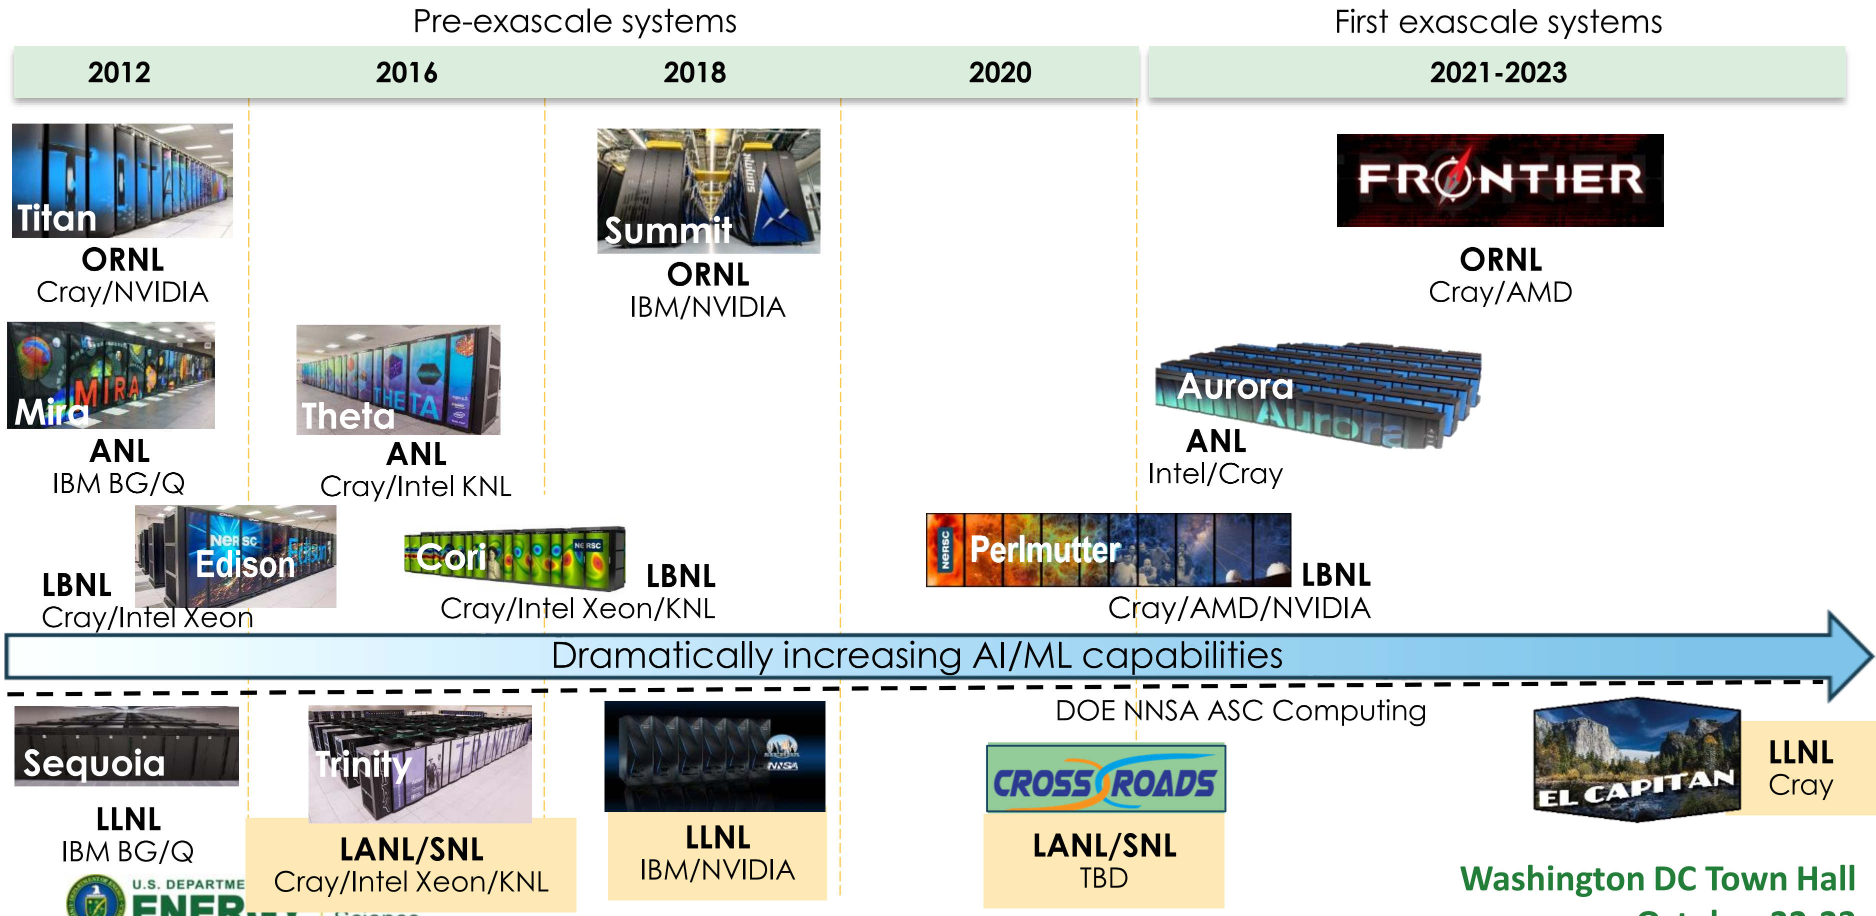

Why Are We Here?

- Exascale Town Halls in 2007
- Led to many other workshops (>10)
- ASCAC engagement
- NSCI
- And the Exascale Computing Initiative
  - ECP: Exascale Computing Project
  - Exascale systems
  - Application efforts across DOE
- Ideas for the next big thing, complementing exascale
- 4 Town Halls organized by the Labs
  - 4th is in DC, October 22-23

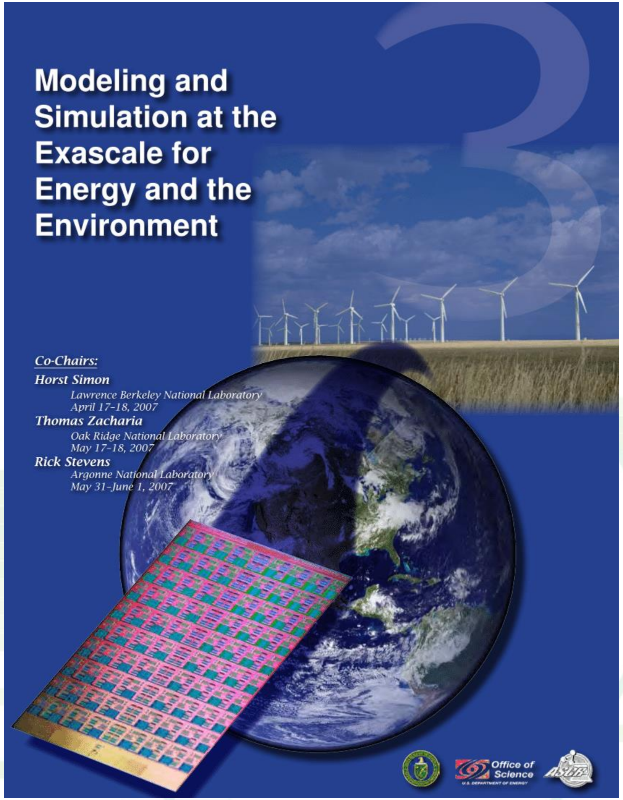

The AI for Science Town Halls so far

- Over 1000 registrations across 4 Town Halls

ANL	357	
ORNL	330	
LBNL	349	+100 online
DC	273	+ ?
Totals	1309	

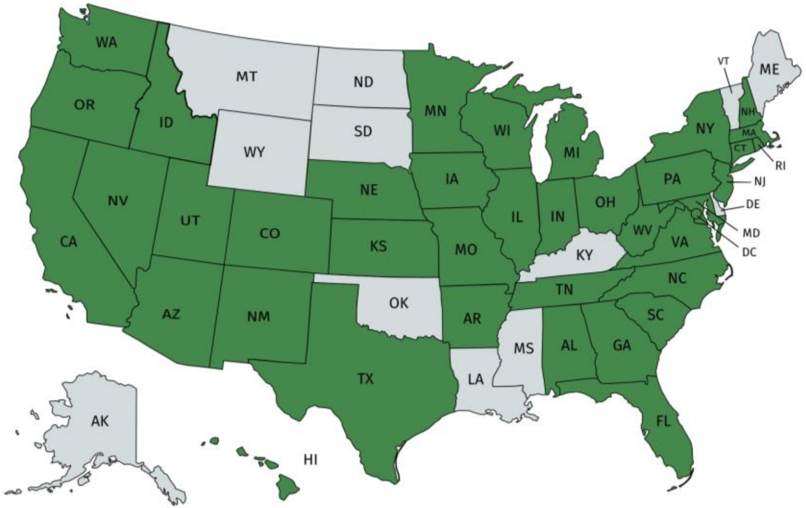

- All 17 DOE National Laboratories
- 39 Companies from large and small
- Over 90 different universities
- 6 DOE/SC Offices + EERE and NNSA

U.S. DEPARTMENT OF ENERGY | Office of Science

The AI for Science Town Halls so far

Argonne

Cosmology

Salman Habib

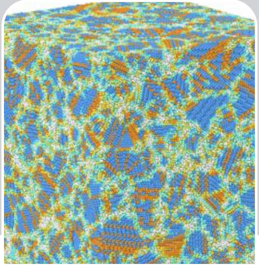

Materials

Ian Foster

Climate

Rao Kotamarthi

Oak Ridge

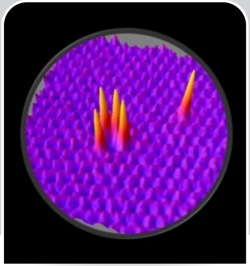

Microscopy

Sergei Kalinin

Manufacturing

Tom Kurfess

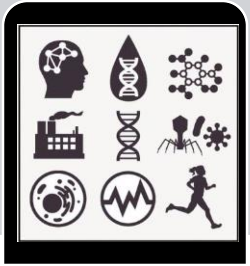

Health

Gina Tourassi

Berkeley

Astrophysics

Josh Bloom

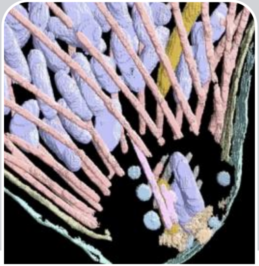

User Facilities

James Sethian

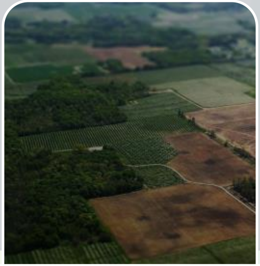

Biology

Ben Brown

U.S. DEPARTMENT OF ENERGY | Office of Science

Washington DC Town Hall
October 22-23

Breakouts and Subtopics

Science and Engineering

High Energy Physics
- Cosmology and Astrophysics
- Particle Physics
- Accelerator Design and Control

Nuclear Physics
- Particle Identification
- Particle Tracking
- Control

Fusion
- Feature Tracking
- Steering
- Design and Control

Energy Sciences
- Materials and Chemistry Modeling
- Photon/Neutron Science
- Electron Microscopy

Earth Sciences
- Climate & Carbon
- Subsurface
- Water

Biological Sciences
- Environmental Biology
- Imaging
- Syn Biology
- Health

Energy Technologies
- Engineering and Manufacturing
- Smart Grid/Urban
- Wind and Solar
- Mobility

Computing Sciences
- AI for hardware / software
- AI for networks / computing facilities

Crosscuts

- Foundations
- Data Lifecycle
- Software
- Hardware
- Computing Facilities

AI Hardware Technology and Industry

Deep Learning Needs High Performance Computing

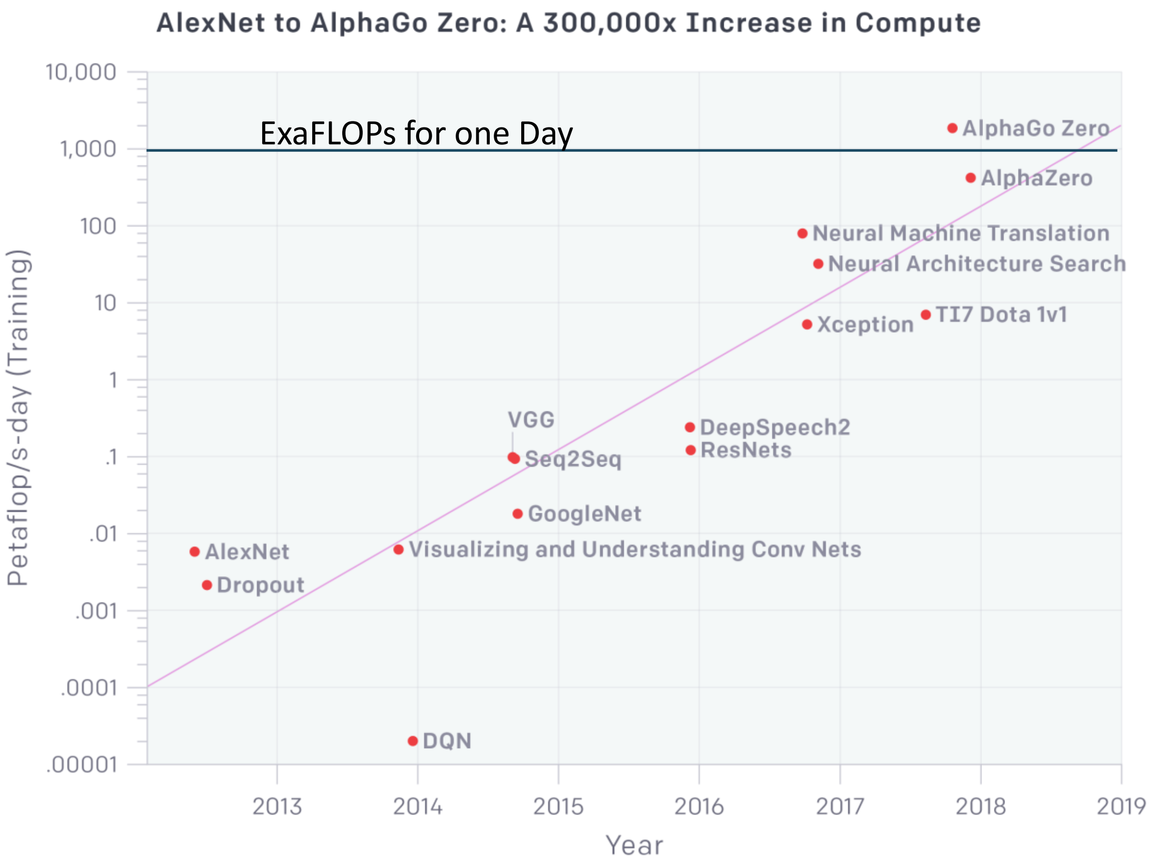

Specialized hardware is emerging that will be 10x – 100x the performance of general purpose CPU and GPU designs for AI

US VCs investing >$4B in startups for AI acceleration

Which platforms will be good for science?

AI Chip Landscape

More on https://basicmi.github.io/AI-Chip/

Tech Giants/Systems

Google
Microsoft
aws
IBM
facebook
HUAWEI
Baidu 百度
Alibaba Group 阿里巴巴集团
FUJITSU
NOKIA
TOSHIBA
Hewlett Packard Enterprise
DELL

IC Vender/Fabless

intel
SAMSUNG
NVIDIA
QUALCOMM
AMD
NXP
ST
XILINX
MEDIATEK
BROADCOM
MARVELL
Rockchip 瑞芯微电子

IP/Design Service

arm
synopsys
Imagination
cadence
CEVA
VeriSilicon
SiFive
ARTERIS IP
alchip
GUC
FARADAY
eSilicon

Startup in China

Cambricon 寒武纪科技
地平线 Horizon Robotics
BITMAIN
intellifusion 云天励飞
ChipIntelli
Think Force
Canaan
云知声 Unisound
Rokid
AISPEECH 思必驰 专注人性化的智能语音
肇观电子 NextVPU
Enflame
亿智科技
清微智能 TSING MICRO

Startup Worldwide

cerebras
WAVE COMPUTING
Graphcore
habana
thinci
SambaNova SYSTEMS
KALRAY
LIGHTELLIGENCE
HAILO Empowering Intelligence
Esperanto TECHNOLOGIES
Tenstorrent
MYTHIC
Preferred Networks
brainchip
PEZY Computing
GREENWAVES TECHNOLOGIES
AIMOTIVE
KONIKU
Tachyum
flexlogix
SYNTIANT
gyrfalcon technology
NOVUMIND
groq

扫码访问AI芯片文章

Compiler

GLOW
tvm
NVIDIA TensorRT
ONNC
nGraph Compiler stack (Beta)
plaidML

Benchmarks

MLPerf
AI - Benchmark
AI Matrix.
AIIA

AI Across Government Agencies

U.S. DEPARTMENT OF ENERGY | Office of Science

Development and Application of AI Critical For All Government Agencies

- Executive Order on AI

Policy Statement: Artificial Intelligence (AI) promises to drive growth of the United States economy, enhance our economic and national security, and improve our quality of life.

… leadership requires a concerted effort to promote advancements in technology and innovation, while protecting American technology, economic and national security, civil liberties, privacy, and American values and enhancing international and industry collaboration with foreign partners and allies.

- Supported by multiple agency strategies and programs

EXECUTIVE ORDERS

Executive Order on Maintaining American Leadership in Artificial Intelligence

INFRASTRUCTURE & TECHNOLOGY

Issued on: February 11, 2019

By the authority vested in me as President by the Constitution and the laws of the United States of America, it is hereby ordered as follows:

Section 1. Policy and Principles. Artificial Intelligence (AI) promises to drive growth of the United States economy, enhance our economic and national security, and improve our quality of life. The United States is the world leader in AI research and development (R&D) and deployment. Continued American leadership in AI is of paramount importance to maintaining the economic and national security of the United States and to shaping the global evolution of AI in a manner consistent with our Nation's values, policies, and priorities. The Federal Government plays an important role in facilitating AI R&D, promoting the trust of the American people in the development and deployment of AI-related technologies, training a workforce capable of using AI in their occupations, and protecting the American AI technology base from attempted acquisition by strategic competitors and adversarial nations. Maintaining American leadership in AI requires a concerted effort to promote advancements in technology and innovation, while protecting American technology, economic and national security, civil liberties, privacy, and American values and enhancing international and industry collaboration with foreign partners and allies. It is the policy of the United States Government to sustain and enhance the scientific, technological, and economic leadership position of the United States in AI R&D and deployment through a coordinated Federal Government strategy,

DOE builds on historical missions and touches all areas

- The U.S. AI strategy includes
  1. **Long-term investment in research**
  2. Effective methods for human-AI collaboration
  3. Address ethical, legal and social implications
  4. **Ensure the safety and security of AI Systems**
  5. **Develop shared datasets and environments**
  6. Standards and benchmarks
  7. Understand the AI workforce
  8. Expand public-private partnerships
- DOE will play a key role in AI for science and engineering
  - AI Technology office
  - Research and talent development
  - Data to support science and engineering research

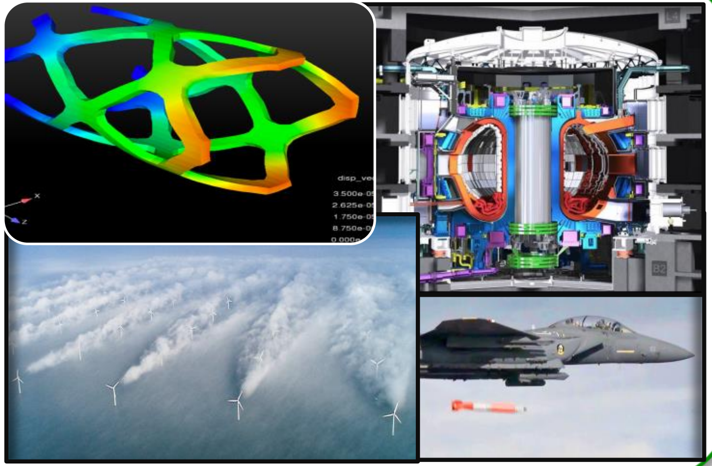

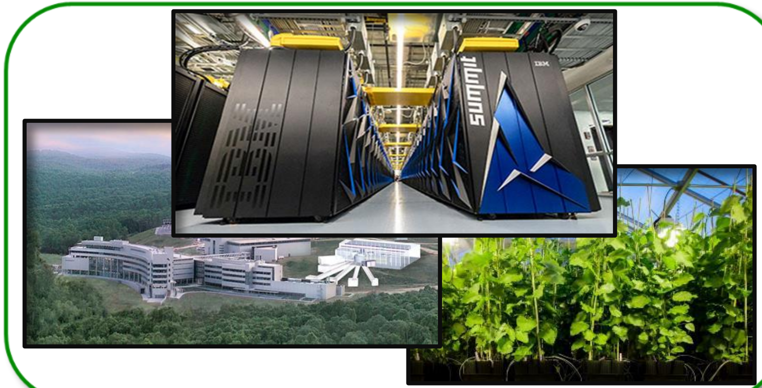

DOE research challenges touch all areas of AI

Data

- Experimental design
- Data curation and validation
- Compressed sensing
- Facilities operation and control

Learning

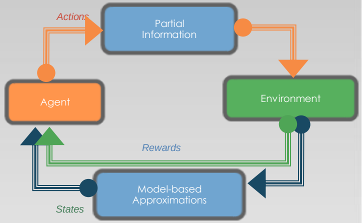

- Physics informed
- Reinforcement learning
- Adversarial networks
- Representation learning and multi-modal data
- “Foundational math” of learning

Scalability

- Algorithms, complexity and convergence
- Levels of parallelization
- Mixed precision arithmetic
- Communication
- Implementations on accelerated-node hardware

Assurance

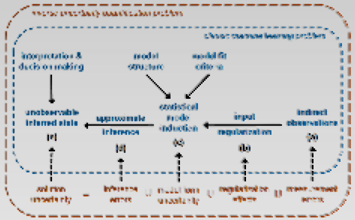

- Uncertainty quantification
- Explainability and interpretability
- Validation and verification
- Causal inference

Workflow

- Edge computing
- Compression
- Online learning
- Federated learning
- Infrastructure
- Augmented intelligence
- Human-computer interface

AI Revolution

- **Learned Models Replace Data**
- **Experimental Discovery Refactored**
- **Questions Pursued Semi-Autonomously**
- **Simulation and AI Merge**
- **Theory Becomes Data**
- **AI Laboratories**