



The government seeks individual input; attendees/participants may provide individual advice only.

Middleware and Grid Interagency Coordination (MAGIC) Meeting Minutes¹
April 1, 2020, 12-2 pm ET

Virtual

Participants

Valentine Anantharaj (ORNL)	Alex Lackpour (NOAA)
Bradford Benbow (DOE/HQ)	Eric Lancon (BNL)
Doug Benjamin (ANL)	Ronit Langer (MIT)
Tekin Bicer (ANL)	Joyce Lee (NCO)
Jessica Breet (ORNL)	Zhengchun Liu (ANL)
Ben Brown (DOE/SC)	Glenn Lockwood (LBL)
David Brown (LBL)	Arthur Barney Maccabe (ORNL)
Eric Burger (OSTP)	David Martin (ANL)
Richard Carlson (DOE/SC)	Linden Mercer (NRL)
Joaquin Chung (ANL)	Michael Roy Nelson (Carnegie)
Dhruva Chakravorty (TAMU)	Mark Neubauer (UI)
Vipin Chaudhary (NSF)	Robinson Pino (DOE/SC)
Ewa Deeknab (ISI)	Yuriy Posherstnik (DISA)
Martin Doczkat (FCC)	Sonia Sachs (DOE/SC)
Katherine Evans (ORNL)	Arjun Shankar (ORNL)
Carlos Fernando Gamboa (BNL)	Alan Sill (TTU)
Geoffrey Fox (IU)	Suhas Somnath (ORNL)
Rob Gardner (UChicago)	Bill Spatz (DOE/SC)
Cheryl Ingstad (DOE/HQ)	Martin Swany (IU)
Shantenu Jha (Rutgers)	Valerie Taylor (ANL)
Raj Kettimuthu (ANL)	Aristeidis Tsaris (ORNL)
Kathryn Knight (ORNLO)	Cong Wang (Renci)
Vangelis Kourlitis (ANL)	Sean Wilkinson (ORNL)
Patrycja Krawczuk (USC)	

Proceedings

This meeting was chaired by Richard Carlson (DOE/SC) and Vipin Chaudhary (NSF).

Guest Speaker: Valerie Taylor, Division Director of the Mathematics and Computer Science Division, Argonne National Laboratory, presenting on the [AI for Science Report](#).

AI for Science

¹ Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program.

Looking beyond exascale at research opportunities in AI (10 years out) (Slide 2):

- Applying AI to different science domains (e.g., biology, chemistry, materials, climate, physics) and also computer science- what's needed from computer science and applied math to drive AI.
- ASCAC, looked at next steps (Early version of report due out May 2020)

Town Halls and feedback on report (Slides 3-4)

- Began with plenary looking at science domain – and use of AI facilitate direction in scientific discovery in those domains (10 years out). Took advantage of lab expertise
 - ANL: cosmology, materials and climate
 - ORNL: microscopy, manufacturing and health
 - LBNL: astrophysics, user facilities and biology
- Post-plenary, chapter discussions and reports.

DOE's Role in AI for Science (Slide 5)

Different space within DOE – mission and science driven. Unique aspects:

Data sets (generated by experiments, simulation) – many data sets available for different science domains.

- HPC (CY2021, DOE will have most powerful exascale machine)
- Match and CS Research: research in applied math and computer science needed to drive scientific discovery
- User facilities available for generating data: light sources (ALS); ARM, etc.
- Team science: end-to-end science solutions; math and computer science research and data generated

Expected AI for Science Impacts (Slide 6)

Acceleration of discovery rates

- E.g., new materials, physics, engineering, climate modeling
- Targeted search, optimization

From simple automation to goal directed systems

- E.g., generative models generating potential materials meeting particular design– having labs explore developing materials
- Automation: AI driving science impact

Simulation + AI hybrids

- E.g., restructuring Is simulation – replacing functions with AI-driven models
- Using techniques (optimization of parameters)

Accessible and integrated knowledge bases

- Access to knowledge bases (may also consider different sensitivities around information in different knowledge bases)
- Looking at new interfaces - extract information from literature and data

Comprehensive transformation of science support and operations

- Moving towards AI everywhere; smart processes

Bold Vision for AI at DOE (10 years out) (Slide 7)

- Learned models may begin to replace data- available to be queried, pluggable, chainable, secure
- Experimental discovery processes dramatically refactored -Models replace experiments
- Questions pursued semi-autonomously at scale - seek materials; self-driving labs, new physics
- Simulation and AI approaches merge -Deep use/integration of AI in simulation methods- hybrid methods; numerical simulation and UQ
- Theory becomes the next data driving next-generation AI -AI used to generate hypothesis, advance theory
- AI becomes part of scientific laboratory activities - AI integrated into science, engineering and operations

DOE building on record of success delivering HPC capabilities (Slide 8)

Timeline of pre-exascale and exascale systems (CY2021: Aurora and Frontier -first deployed exascale systems in U.S.)

Chapter Outcomes (Sampling of 3 out of 16 chapters)

AI Driving Science Domains: Biology and Life Sciences (Slide 9)

Biological systems – predict, control and understand systems, often molecular scale

Challenges:

- building capacity to design custom, biological systems (important for global health)
- systematically manage and engineer global environmental systems through predictive understanding; looking at full ecosystem
- Develop AI-enabled self-driving labs

Accelerating Development – what’s needed?

- Improving scalability of data set (quality, provenance, quantity)
- Establish Infrastructure – needed to make communal use of data – develop foundational technology.
- Develop foundational technologies to promote statistical framework to integrate knowledge across disciplines, including data-efficient learning
- Understand biases and inaccuracy threaten model performance on subgroups in heterogeneous settings

Outcomes:

- Delivery accuracy to precision medicine
- Discover controls of massively multi-scale, dynamic biosystems
- Build life to spec (what’s needed in biological systems)
- Engineer our troposphere

High Energy Physics (Slide 11-12)

Discovering ultimate constituents of matter and uncovering nature of space time

Challenges:

- Reconstructing history of universe using AI
- Advance knowledge of cosmic structure formation with AI-driven Automated Cosmology Experiment (ACE)
- Zettascale AI to uncover new fundamental physics

Accelerating Development

- Usable tools
- Training methodologies – to detect rare features and want to be robust against systematic effects
- Tools to quantify these impacts

Outcomes

- Enable data exploration from next-gen surveys
- Make movie of universe from its earliest moments until today
- New era of precision physics in energy and intensity

Applied math and Computer science driving AI (Slides 13-14)

AI Foundations and Open Problems: looking at mathematical, statistical and information –theoretic foundations needed for AI; and advancing these foundations

Challenges:

- Incorporating domain knowledge in AI/ML – arose in science domain (incorporating physics principles)
- Establishing assurance for AI (important re: trust of models, robustness of models; understanding uncertainty among models)
- Recognizing need to achieve efficient learning of AI systems – need efficient algorithms, esp. when working with large masses of data

Accelerating Development

- Use scientific principles: incorporate in modeling and simulation; inform and advance AI
- Address robustness, interpretability of AI system – leads to trust
- Learning for inverse problems
- Reinforcement and active learning to develop AI – important for control needed for self-driving labs and data acquisition system

Outcomes:

- Increase trust in ML and AI as scientific techniques
- Provide efficient computational algorithms for ML and AI
- Maximize the understanding realized from science-informed AI

AI Science Applications – 3 components (AI-enabled) (Slide 15):

- Design workflows (what to make): comes up in materials, polymers, organisms
- Experimental (How to make): self-driving labs, synthesis search
- Comprehension (meaning): insights from data sets, literature, science “news”, strategies

AI for Science: AI Building blocks (examples) – arose in chapters (Slide 16)

Examples of building blocks needed to drive scientific discovery across domains

- Computer science and applied math to drive AI

Blocks involving:

- Design (e.g., materials, chemicals, drugs)
- control (e.g., experiments, simulations, accelerators, reactors)
- augmented simulations (surrogate models, optimize around parameters, search)
- science and math comprehension (understanding more re: physics and mathematics)
- Generative models (cosmology, bio design)
- Inverse problems (waveform to source, spectra to structures)

- Multimodal Learning: getting data in different modalities (e.g., images, waveforms)
- Decision-making (risk assessment)

AI: a fundamental shift in economic and military landscape (Slide 17)

DOE set up AI and technology office:

- DOE's unique role (mission driven development and application of AI/ML; i.e., innovation in science, energy, national security)
- Build on wide expertise, HPC mission
- Large-scale scientific data for research
- Talent development

Discussion

AITO: Actively developing strategy and developing AI with focus on national and economic security

- AI sensor working with joint intelligence center that DoD stood up recently.
- external face to government, Industry and industry

International collaboration?

AITO dialogued with 13 allies last month and following up currently

Geoffrey Fox, Digital Science Center, Indiana University Bloomington, Issues at the intersection of AI, Streaming, HPC, Data-centers and the Edge

Outline: Discuss working with industry (e.g., MLPerf activity- Science Research benchmarks) and software issues

DOE and NSF funded streaming workshops (2015-2016)

- Since then, software models improved, but essentially same, but algorithms have changed.
- New algorithms to be developed
- Deep Learning for geospatial time series – applied to streaming data; relevant to virus data (very promising and relevant to Covid-19 studies)

ML Perf Consortium Deep Learning Benchmarks (Cos, universities, DOE labs) (Slide 3-5)

- Mission: build benchmarks for measuring training and inference performance of ML hardware, software and services; what users care about. Benchmarks online: framing.
- Working groups – (e.g., HPC, Deep learning for time series, science data proposed by Fox)
- Benchmarks (Slides 4-5) learn about efficiency of scaling, Any member can contribute results
- Science Data and MLPerf I – would like more scientists involved (Slide 6)
 - Science Research Data – proposed working group proposed by Fox and Tony Hay-
 - Collect science data and establish benchmarks
 - Example benchmarks (Slide 7)

Linkage of Deep Learning (AI) and HPC (Slide 8)

- Working in HPC systems – using GPUs or CPUs which are similar in performance and customization
- Need to explore new architectures
- Many Programs in using HPC for ML and using ML to make HPC run fast

- Industry very important – 2 dominant frameworks, PyTorch and TensorFlow come from Facebook and Google. Need to make these systems run very well
- Hyper-parameter search needs to be deployed broadly: Industry always runs. Not possible at most universities because insufficient GPUs
- Need to advance tools for time series (e.g., LSTM, GRU) – for streaming data, edge data
 - Image streams: overlap between industry and science; others: industry logistics, ride-hailing, speech – are different
- Industry- switch to Deep Learning for all ML; clustering, dimension reduction and other classic ML problems being replaced by deep learning algorithms

Diagrams (Slides 9-11)

- HPC for ML: Integration Challenges (Slide 9, includes ML/Deep learning)
- Classic Streaming Software solution: classic architecture (e.g., Apache storm) unchanged since streaming workshops
- Google Cloud AI Platform – state of the art edge to cloud workflow with google cloud data flow open source runs streaming analysis programs;

MLSys Edge Deep Learning Papers on the Edge: Top conference on interface of systems and AI (List of papers, Slide 12)

- Over half of attendees from industry.
- Papers focused on, edge computing, Deep Learning and the intersection of systems.

Examples

Deep Learning Prediction of Fusion Tokamak Instabilities (Slide 13)

- Great example of what can be currently done with Deep Learning.
- Using Deep Learning to tackle real-time data.
- Standard approach: mixture of convolutional and current networks.
- Good job predicting disruptive instabilities in controlled fusion

Ride Hailing community- Time Series Analysis (most advanced work in this area) (slide 14)

- Deep Learning Time series based analyses: span minutes- predict load and traffic across days and weeks.
- Multi-graph convolutional neural network

Indy 500 real-time anomaly detection and ranking prediction (Slides 15-16, graph)

- Analyzed in real time by convolutional network
- Train on previous races; used to predict winner

Denosing Images from Light Sources (Slide 17, graph)

Rapid and real-time denoise images of light sources. Use edge devices (Edge TPU from Google, Nvidia); Intel CPU

Learn Newton's laws with recurrent neural networks (Slide 18 graph)

Using time series of recurrent neural nets, which if set up with 65k parameters, can build new operator that is highly accurate (up to 4k times normal time set).

- Learned Newton's laws from 5k training simulations
 - Neural networks better at extrapolating than classical molecular dynamics algorithms
- Surrogates replace classic algorithms of HPC simulations
- Promising area relating edge and disk computing to simulations as they are all time series

Hidden Theories and Instances - Earthquake (Slide 19, simulation; predictions)

- Time series from earthquakes.
- Deep Learning Networks (DNN): hidden variables that learn the theories
- Model learning simulated earthquake and predictions compared to observations for annual earthquake activity in S. California
- Observed data contains many uncertainties

Deep Learning in particle physics data analysis (Slides 20-21)

- Deep Learning network learns to classify events from their energy structures; replaces classic ideas (Fox Wolfram moments).
- Shows improvement of past methods – design system with the right architecture and have sufficient training data: good chance to learn

Conclusions

- MLPerf- initiative to develop performance measures for ML; consider science research benchmarks in MLPerf and enhance collaboration between industry and research, HPC and MLPerf/MLSys communities
- Software issues: not change drastically in last 5 years
- Edge to cloud –
 - switching from trad algorithms to Deep Learning will revolutionize geospatial time series work - relevant to Covid-19 studies
 - convolutional recurrent networks are so powerful
- Timely vs. Real-time – sometimes real-time is important. other times, care more about throughput (value may be more important than latency)
- Looking forward to collaborating to help advance science.

Discussion

Opportunities for involvement

- [MLPerf](#): new activity with MLCommons. Currently free and open membership. .
- TinyML – system working group doing benchmarks for TinyML (ultra-lower-power machine learning)

DOE/Lab Contacts: Look at Computing Research Leadership Council

Roundtable

DOE/SC: Rich Carlson

COI workshop for DOE labs: vision of laboratory structure in relation to computing. Website will be circulated. Needs to be in-person (moving April)

NSF: Vipin Chaudhary

[NSF Principles and Practice of Scalable Systems](#) proposal deadline moved to April 6th.

Meetings:

April 1 – April 3, 2020: [CASC meeting](#), Westin Crystal City, Virginia

April 29- May 1, 2020: [Women in HPC Summit](#), Vancouver, BC

July 26 – 30, 2020, [PEARC20](#) Meeting, Portland, OR (February 17, 2020 deadline for submissions)

Next Meeting: May 6, 2020 (12 noon ET)