

AI for Science




Rick Stevens
Argonne National Laboratory
The University of Chicago

Crescat scientia; vita excolatur

AI for Science Townhalls

Organized by Argonne, Oak Ridge and Berkeley with participation from all the laboratories..

- Four “Townhalls” aimed at getting input from the DOE community on opportunities and requirements for the next 5-10 years in computing with a focus on convergence between HPC and AI
- July (Argonne), August (Denver), September (Berkeley), October (Washington)
- Modeled after the 2007 Townhalls that launched the Exascale Computing Initiative
- Each meeting covers roughly the same ground, geographically distributed to enable local participation
- Applications in science, energy and technology
- Software, math and methods, hardware, data management, computing facilities, infrastructure, integration with experimental facilities, etc.
- Expect ~200 people per meeting
- Output will be a report to guide strategic planning at Labs and DOE



Innovation XLab Artificial Intelligence Summit

The next in the series of Innovation XLab events will be hosted by Argonne in Chicago on October 2-3, 2019

- Event date confirmed for Oct 2-3, 2019
- 11 of 17 national labs actively involved in planning: ANL, LLNL, ORNL, LBNL, BNL, LANL, SNL, NETL, FNAL, PNNL, SLAC
- Industry focus areas: Energy, Manufacturing, Healthcare, Risk
- Set up Steering and Program Committees consisting of PIs and tech transfer participants from all the labs; regular calls held to coordinate input and ensure broad participation
- Initial list of industry attendees generated with ~650 names
- Initial list of speakers and panel participants generated with ~75 names
- Target agenda draft by June 21
- DOE-OTT weekly call with the Organizing Committee kicked off on June 10

In symbols one observes an advantage in discovery which is greatest when they express the exact nature of a thing briefly and, as it were, picture it; then indeed the labor of thought is wonderfully diminished.

— Gottfried Wilhelm Leibniz


DOE/Argonne was the home to a leading symbolic AI group from the 1960's to the mid 2000's working on Automated Theorem Proving

Alan Bundy	University of Edinburgh
Edmund Clarke	Carnegie Mellon University
Tammi Henry	University of Tennessee
Larry Hines	University of Texas
Deepak Kapur	State University of New York at Albany
Matt Kaufmann	Computational Logic
Ken Kunen	University of Wisconsin
Vladimir Lifschitz	Stanford University
Ewing Lusk	Argonne
William McCune	Argonne
Ross Overbeek	Argonne
Dana Scott	Carnegie Mellon University
Mark Stickel	SRI
Rick Stevens	Argonne
Robert Veroff	University of New Mexico
Richard Waldinger	SRI
Steve Winker	Argonne
Larry Wos	Argonne
Hantao Zhang	The University of Iowa

Attendees at an Argonne
ATP "theory institute" in 1990.

- [1. Algebraic Geometry](#)
 - [o Cancellative Semigroups on a Cubic Curve](#)
 - [o Uniqueness of the 5-ary Steiner Law](#)
- [2. Cancellative Semigroups](#)
- [3. Lattice Theory](#)
 - [o A Simpler Absorptive Basis for Lattice Theory](#)
 - [o A New Schema for Single Axioms](#)
 - [o A Shorter Single Axiom for Lattice Theory](#)
 - [o A Single Axiom for Weakly Associative Lattices](#)
- [4. Quasilattice Theory](#)
- [5. Uniqueness of Operations in Lattice-like Algebras](#)
- [6. Self-dual Bases for Boolean Algebra](#)
- [7. Self-dual 2-Basis for Group Theory](#)
- [8. Self-dual Bases for Group Varieties](#)
- [9. Quasigroup Theory](#)
- [10. Quasigroup Design Problems](#)
- [11. Single Axioms for Ternary Boolean Algebra](#)
- [12. Single Axioms for Groups](#)
 - [o Ordinary Groups](#)
 - [o Abelian Groups](#)
 - [o Exponent Groups](#)
 - [o Some Permutative Varieties](#)
 - [o Ordinary Groups \(Kunen\)](#)
 - [o Groups of Exponent 4 \(Kunen\)](#)
 - [o Odd Exponent Groups \(Hart and Kunen\)](#)
- [13. Simple Bases for Moufang Loops](#)
- [14. Single Axioms for Inverse Loops and Subvarieties](#)
- [15. Left Group and Right Group Calculi](#)
- [16. Fixed Point Combinators](#)
- [17. Semigroup Structure \(F3B2\)](#)
- [18. Illative Combinatory Logic \(Jech\)](#)
- [19. Robbins Algebra and Boolean Algebra](#)
- [20. Equivalential Calculus Single Axioms](#)
- [21. Semigroups, Antiautomorphisms, and Involutions](#)
- [22. Independence of Ternary Boolean Algebra Axioms](#)
- [23. Two-valued Sentential Calculus](#)
- [24. Many-valued Sentential Calculus](#)
- [25. Short Proofs in Various Logic Calculi](#)
- [26. Pure Proofs in Equivalential Calculus](#)

Machine Learning Arxiv Papers per Year





What is possible?

Things we can do 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

Things We Want To Do With AI In The Future

- Develop methods that can learn from both encoded symbolic theory (e.g. QM/GR) and large-scale data so we can leverage the vast theoretical knowledge we have accumulated over hundreds of years
- Automate and accelerate discovery from planning, to conjecture, to experiment, to confirmation and analysis \Rightarrow end-to-end automated science
- Create an ability to use AI for generating new theories that address the problematical areas of existing theories

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

The background features a complex, glowing network of blue and purple lines and nodes, resembling a digital circuit board or data flow. The lines are interconnected, creating a sense of depth and movement. The overall color palette is dominated by cool blues and purples, with some warmer yellow and orange highlights where the lines intersect or glow.


A Sampling of Science Opportunities

Materials and Chemistry

- Design of materials and molecules
- AI-guided synthesis
 - automated design of chemical pathways
 - mapping metastable phases
 - extracting mechanisms
- Predictive interfacial transport of ions and charge
- AI-accelerated ab Initio molecular dynamics
- Quantification of energy drivers for separations
- Describing multiscale charge, spin, lattice correlations
- Exploring energy landscapes in ultrafast, nonequilibrium, and driven systems and processes
- Inverse design, bandstructure engineering



Tian et al. PRL 2016



Nonequilibrium superconductivity

Table 1: The opportunity areas described in this report, with observations on AI-related requirements and challenges.

§	Area	AI requirements and challenges
5.1	AI-Accelerated Ab Initio Molecular Dynamics for Catalysis	Methods development to enable application of ML/AI methods to extremely large collections of samples obtained from simulation studies, and for efficient coupling of simulation and AI components.
5.2	Ultra-Fast Simulations of Complex Materials	Processing billions of DFT energy evaluations is likely to require extremely large neural networks. Handling data from multiple sources is also a key need.
5.3	Designing New Chemical Pathways Automatically	Tight integration with experiment. Reinforcement learning and active learning algorithms to guide experimental campaigns. Representation and update of kinetic table and associated uncertainties.
5.4	Real-time Inversion of Multi-modal Characterization Data	Requires methods for integrating physical constraints into neural networks (NNs). May also build up large enough NNs to require specialized AI accelerators.
5.5	Panoramic Synthesis for Discovery and Deployment of New Materials	Would benefit from symbolic AI to create human-interpretable (and, ideally, scientifically testable) design rules for panoramic synthesis.
5.6	AI-Driven Material Discovery for Energy Storage	Tight integration with computational simulation. Reinforcement learning and active learning algorithms to guide computational campaigns.
5.7	Discovery and Design of Magnetic Topological Materials and Magnetic Order	Learning from small data. Transfer learning between different classes of materials. Integration of experimental and simulation data.
5.8	AI-Generated Designs of Unconventional Structures	Requires method development for generative models for networks/paths and supervised learning methods on graphs/path data.
5.9	Comprehensive Atlas of Phase Diagrams of All (Meta)Stable Materials	Requires advances in natural language processing (NLP) and in methods for propagating uncertainty through many different supervised learning and physical models.
5.10	Optimizing Gas-phase Chemistry for Scale-up of Complex Materials	AI-based surrogate models for manufacturing processes are needed that can enable near-real-time feedback; current multi-scale simulation methods take days or weeks.

Advanced Photon Source Upgrade

AI can drive the scientific and measurement motifs enabled by APS-U

Detect rare events/features in large volumes with nanoscale resolution

Capture dynamic processes

Enable multidimensional inquiry, exploring spaces of higher dimension and size

Metal fatigue, solid-state batteries, brain circuitry

Catalyst coarsening, precision synthesis, additive manufacturing



High-entropy alloys, metal fatigue, catalysts

APS-U's 2-3 orders of magnitude increased brightness and coherent flux, will lead to:

- Massive data, too much for humans to handle
- Data rates too fast for human management

Control: real-time autonomous execution

Analyze: reconstruct, feature extraction, viz, optimized photon dose





Advanced accelerator control:
100's of control points, 1000's of inputs

AI at the edge:
autonomous data reduction near the source

Climate and Biology

- Accelerated Climate Models (PDE/ML hybrids)
- Improved integration of remote sensing and ground truthing into Climate Models (cloud/precipitation, land cover/biogeochem, sea ice/calibration, etc.)
- Improvement in ARM data pipelines, automated model extraction from data, smart data fusion

- Vast applications in genomics and metagenomics ($G \Rightarrow P$)
- Automation of bioinformatics methods (improved productivity)
- Automating hypothesis formation in biology (causal analysis)
- Forward design of novel pathways, proteins, regulons, operons, organisms, etc. for secure biodesign
- Anomaly detection (discovery in sequencing, biosecurity, etc.)




High Energy Physics

Energy/Intensity Frontier:

- Search for Beyond the Standard Model (BSM) physics through AI-driven anomaly detection
- AI-reduced uncertainties to enable precision electroweak measurements for BSM clues
- Generative Adversarial Networks (GANs) for large-scale Large Hadron Collider detector simulation

Cosmic Frontier – AI in end-to-end application:


- Precision Cosmic Microwave Background emulation – AI simulation speed-up of a factor of 1000
- Search for strong lensing of galactic sources for precision cosmology measurements using AI classification, regression, and GANs for image generations
- AI-based Photometric Redshift Estimation
- Combination of AI methods to enable searches for hidden space variables



AI applications in an “end-to-end” Cosmic Frontier application: 1) GANs for image emulation, 2) GP and DL-based emulators for summary statistics, 3) CNN-based image classification, 4) AI-based photometric redshift estimation, 5) Likelihood-free methods for inference [Work performed under the Argonne-led SciDAC-4 project: “Inference and Machine Learning at Extreme Scales”]

Nuclear Physics

- AI- and deep learning-guided insight to unravel new physics in quantum chromodynamics
 - Active Learning and Generative Adversarial Networks (GANs) to discover new sum rules and violations of constraints
- AI and deep learning for ATLAS and Electron-Ion Collider to probe fundamental questions: How do mass and spin of nucleons arise, how do nucleosynthesis and stellar evolution produce current abundances?
 - Deep neural network for detector and accelerator design optimization
 - GANs for self-tuning performance-maximizing detector configurations and time-saving online accelerator tuning in multi-beam/multi-detector experiments
 - AI-assisted data analysis of many-body break-up and dynamics: tag recoil spectators to isolate struck nucleon
- AI-driven data analysis of neutrino-less double beta decay
 - Sparse neural network with scalable machine learning techniques accelerate computations and extend range of experiments
 - GANs and segmentation networks improve detector understanding and resolution



Transverse momentum slices of u and \bar{d} quarks in a longitudinally polarized proton.


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

Generative Models

Projection of Junction Tree autoencoder space




Learned Function Accelerators




AI Accelerators


AI at Argonne: Broad Span of Scientific Targets




Reduced order modeling of laser sintering




Nowcasting with convolutional LSTMs




Prediction of radiation stopping power




Strong and weak lensing in sky survey data




Prediction of antimicrobial resistance phenotypes



Identification and tracking of storms




Efficient climate model emulators




Defect-level prediction in semiconductors




Structure-property-process triangle in additive manufact.



Parameter extraction in atom probe tomography




Learning for dynamic sampling in spectroscopy




Vehicle energy consumption prediction




Flying object detector for edge deployment




Discovery of new energy storage materials




Cosmic Microwave Background emulation



Photometric red shift estimation



New materials for efficient solar cells




Enhancement of noisy tomographic images




Example from Cancer Research

Personalized Cancer Therapy



Modeling Cancer Drug Response



Drug (s)


descriptors

fingerprints


structures

SMILES

dose




$$\mathcal{R} = f(\mathcal{T}, \mathcal{D}_1, \mathcal{D}_2)$$






IC50
AUC
GI50
% growth
Z-score
Response

gene expression levels
SNPs
protein abundance
microRNA
methylation
Tumor




Why deep learning




“Uno” Model Predictions with Dropout UQ (trained on ALMANAC)

All Samples colored by Sample ID



RTS subset of drug pairs

Samples found in Cluster 1 or Cluster 2




- x NCIPDM.237351~077-R~AL-IR0
- x NCIPDM.994434~217-R~AK3YH7
- x NCIPDM.CN0446~F447~M12M52




Example from Traumatic Brain Injury


Anatomical Segmentation




Connectomics



CT ↔ MRI "Super-resolution"



Diagnosis/Prognosis




Training with diverse data modalities and phenotypes

4RTNI 129 APPLY BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	A4 0 APPLY BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	ABIDE 1,112 GO BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	ABVIB 554 APPLY BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER
ACE 449 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	ADNI 3,051 GO BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	ADNID 131 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	ADNIDOD 278 GO BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER
AIBL 852 APPLY BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	ALLOP1 27 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	BIOFIND 232 APPLY BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	FTLD 0 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER
GSP 1,639 PENDING BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	HCP 40 GO BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	HDNI 369 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	ICBM 853 GO BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER
LEFFTDS 736 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	MAPP 344 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	NAPLS 849 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	NAPLS3 681 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER
NIAD 229 PRIVATE BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	NIFD 346 APPLY BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	PAD 3 GO BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER	PPMI 2,263 GO BIO CLINIC CT dMRI fMRI GENE sMRI SPECT PET OTHER



Enhance CT imaging and exploit labels from other modalities

Generative Adversarial Networks



GAN Model trained on TBI patient data



Real CT



Real MR




Fake MR


Diverse brain disease MRI data for identifying abnormal CT

CNN Model trained on normal/abnormal MRI slices

Normal MRI




Tumor



Meningioma


Glioma

Stroke Lesion




Knowledge transfer for CT

Normal CT



TBI Lesion/Midline Shift





Building the AI Environment for Science

AI for Science Requires New Research and Infrastructure

Applications

AI applications across science and engineering. Transformative approaches to simulation and experimental science.

Learning systems

AI software. Software infrastructure for managing data, models, workflows etc., and for delivering AI capabilities to 1,000s of scientists and engineers.


Foundations

Mathematics, algorithms; general AI, reinforcement learning, uncertainty quantification, explainability, etc.


Hardware

Advanced hardware to support AI. Evaluation of new architectures and systems; exploration of neuromorphic and quantum as long term accelerators for AI.

Infrastructure for AI-enabled Science



Infrastructure for AI-enabled Science



DLHub: Organizing and Serving Models

<https://www.dlhub.org>



- Collect, publish, categorize models
- Serve models via API with access controls to simplify sharing, consumption, and access
- Leverage ALCF resources and prepare for Exascale ML
- Deploy and scale automatically
- Provide citable DOI for reproducible science

Argonne Advanced Computing LDRD


Models and Processing Logic as a Service




X-Ray Science

```
from dlhub_sdk.client import DLHubClient
dl = DLHubClient()



struct = dl.run("cherukara_structure", X)
```



Energy Storage



Tomography





TomogAN: Liu et al.


CANDLE: Exascale Deep Learning Tools

Deep Learning Needs Exascale

- Automated model discovery
- Hyper parameter optimization
- Uncertainty quantification
- Flexible ensembles
- Cross-Study model transfer
- Data augmentation
- Synthetic data generation
- Reinforcement learning




<https://github.com/ECP-CANDLE>





Future Directions in Foundations

- Leverage DOE expertise in automatic differentiation, symbolic computing and optimization to ensure that machine learning for science is forward looking, methods are robust and models interpretable
- Many facets relevant to science
 - Integration of symbolic computing with machine learning
 - Prediction and inference of spatio-temporal processes
 - Derivatives for training, sensitivity analysis, optimization, and UQ
 - Rapid data analysis to reduce volume or identify features of interest
 - Variety of new approaches to inference and UQ
 - Identify and account for uncertainty in data sources and computations



Aurora: HPC and AI

>> Exaops/s for AI




CRAY[®]

Architecture supports three types of computing

- Large-scale Simulation (PDEs, traditional HPC)
- Data Intensive Applications (scalable science pipelines)
- Deep Learning and Emerging Science AI (training and inferencing)



Robust Learned Function Accelerators



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

**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



IC Vender/Fabless



IP/Design Service



Startup in China



Startup Worldwide



扫码访问AI芯片文章

Compiler



Benchmarks



AI Accelerator Testbed

Engaging the community to understand and improve specialized AI hardware for science

Dozens of proposed AI accelerators promise 10x - 1000x acceleration for AI workloads. AI testbed will:

1. Provide an **open and unbiased environment** for evaluation of AI accelerator technologies
2. **Disseminate information** about use cases, software, performance on test problems
3. **Support collaborations** with AI technology developers, academics, commercial AI, DOE labs



IC Vendors	Intel, Qualcomm, Nvidia, Samsung, AMD, Xilinx, IBM, STMicroelectronics, NXP, Marvell, MediaTek, HiSilicon, Rockchip	13
Tech Giants & HPC Vendors	Google, Amazon_AWS, Microsoft, Apple, Aliyun, Alibaba Group, Tencent Cloud, Baidu, Baidu Cloud, HUAWEI Cloud, Fujitsu, Nokia, Facebook, HPE, Tesla	12
IP Vendors	ARM, Synopsys, Imagination, CEVA, Cadence, VeriSilicon, Videantis	7
Startups in China	Cambricon, Horizon Robotics, Bitmain, Chipintelli, Thinkforce, Unisound, AISpeech, Rokid, NextVPU, Canaan, Enflame, Eesay Tech	12
Startups Worldwide	Cerebras, Wave Computing, Graphcore, PEZY, Tenstorrent, ThinCI, Koniku, Adapteva, Knownm, Mythic, Kalray, BrainChip, Almotive, DeepScale, Leepmind, Krtkl, NovuMind, REM, TERADEEP, DEEP VISION, Groq, KAIST DNPU, Kneron, Esperanto Technologies, Gyrfalcon Technology, SambaNova Systems, GreenWaves Technology, Lightelligence, Lightmatter, ThinkSilicon, Innogrit, Kortiq, Hailo, Tachyum, AlphalCs, Syntiant, Habana, aiCTX, Flex Logix, Preferred Network, Cornami, Anaflash, Optaylsys, Eta Compute	44

Staged evaluation enables identification of most promising systems for science



Argonne is developing AI infrastructure

- Argonne is partnering with Cerebras to develop and deploy an AI computing platform
- Scientific AI models from Cancer, cosmology, brain imaging and materials science are the first examples that will be deployed
- Our goal is to accelerate relevant AI model types for problems in materials, biomedical, cosmology, high-energy physics, energy systems, synthetic biology, climate, software optimization, architecture research etc.



CONCLUSION: **FUTURE**



- ❖ Massive multi-core engines that **enable model parallelism**
- ❖ Orders of magnitude **greater memory and communication BW**
- ❖ Unconstrained methods, e.g., large and **small mini-batch**
- ❖ Capture weight and activation **sparsity** for higher performance
- ❖ Support research and execution of **emergent model architectures** (not just those of today)

The background features a complex, glowing circuit board pattern in shades of blue, purple, and teal, set against a dark, starry space background. A central, semi-transparent dark square is positioned behind the text.


AI Driven Experimental Science

autonomous molecular discovery system with multiple feedback loops

Tanja Dimitrov, Christoph Kreisbeck, Jill S. Becker, Alán Aspuru-Guzik, and Semion K. Saikin


ACS Applied Materials & Interfaces Article ASAP

DOI: 10.1021/acsami.9b01226




The ATOM Platform


Active Learning Drug Discovery Framework



Layered workflow combining AI, HPC and HTS



Pure ML "constant time" (fast loop)



Mixed/Variable time (slow loop)







Come to a Townhall and tell us what you need!



Chicago AI for Science Town Hall
Argonne National Laboratory
July 22-23, 2019
To register for Chicago, [click here](#)
DRAFT Agenda: [Click here](#)



Denver AI for Science Town Hall
LOCATION
August 20-21, 2019
Registration link here
DRAFT Agenda: [Click here](#)



San Francisco AI for Science Town Hall
Lawrence Berkeley National Laboratory
September 11-12, 2019
Registration link here
DRAFT Agenda: [Click here](#)



Washington DC AI for Science Town Hall
LOCATION
October 22-23, 2019
Registration link here
DRAFT Agenda: [Click here](#)

