

Al for Science @ Argonne



IAN FOSTER

Director, Data Science and Learning Division, Argonne National Laboratory Professor of Computer Science, University of Chicago foster@anl.gov

Fermilab, January 14, 2020

Briefly: What is AI?

Artificial Intelligence

Any technique that enables computers to mimic human intelligence: e.g., logic, if-then rules, decision trees, and machine learning (including deep learning)

Briefly: What is AI?

Machine Learning

A subset of AI that includes statistical techniques that enable computers to improve at tasks with experience. Includes deep learning. Any technique that enables computers to mimic human intelligence: e.g., logic, if-then rules, decision trees, and machine learning (including deep learning)

Artificial Intelligence

Briefly: What is AI?

Machine Learning

Deep Learning

The subset of machine learning composed of algorithms that permit software to train itself to perform tasks, like speech and image recognition, by exposing multilayered neural networks to vast amounts of data A subset of AI that includes statistical techniques that enable computers to improve at tasks with experience. Includes deep learning.

Artificial Intelligence

Any technique that enables computers to mimic human intelligence: e.g., logic, if-then rules, decision trees, and machine learning (including deep learning)





Predicting formation enthalpies of crystalline materials



Predicting formation enthalpies of crystalline materials



Best conventional machine learning method,

Predicting formation enthalpies of crystalline materials

Best conventional machine learning method,



Jha, Ward, et al., 2018.

New methods have limitations: E.g., ImageNet Roulette



judge, justice, jurist: a public official authorized to decide questions brought before a court of justice

- person, individual, someone, somebody, mortal, <u>soul > adjudicator</u> > judge, justice, jurist
- person, individual, someone, somebody, mortal, soul > worker > skilled worker, trained worker, skilled workman > official, functionary > judge, justice, jurist

https://www.excavating.ai



orphan: a child who has lost both parents

> person, individual, ٠ someone, somebody, mortal, soul > juvenile, juvenile person > child, kid, youngster, minor, shaver, nipper, small fry, tiddler, tike, tyke, fry, nestling > orphan







- Queen of England: the sovereign ruler of England
 - person, individual, someone, somebody, mortal, soul > leader > aristocrat, blue blood, patrician > female aristocrat > <u>queen, queen</u> regnant, female monarch > Queen of England



- religious teacher; used as a title of respect
- person, individual, • someone, somebody, mortal,
 - soul > religious person > Hindu, Hindoo > swami

Why are we excited about "AI for science"?

Push

- Step changes in AI/ML methods, notably deep neural networks
- Major advances in areas like machine translation, speech recognition, image processing
- New hardware specialized for deep neural networks

Pull

- Exploding volumes of data due to new sensors and instrumentation exceed human capabilities
- End of Moore's Law puts hard problems out of reach
- Growing complexity of science and engineering problems slowing rate of discovery

Argonne's major Initiatives offer compelling targets

Five major initiatives build on our core capabilities to deliver cutting-edge science and enable future energy technologies

Hard x-ray sciences

Transform understanding of materials and chemical systems through 3D microscopy

Advanced computing

Deploy exascale computer and advance machine learning and quantum and neuromorphic computing

Materials and chemistry

Discover emergent phenomena and synthesize novel materials and chemical systems

The universe as our laboratory

Make leading contributions to physics experiments that explore the early universe and its dynamics

Energy manufacturing science and engineering Create science-based approaches to speed scaling of manufacturing processes for energy technology





"AI" is an opportunity in all five areas

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 model 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 problematical areas of existing theories

In 10 years ...

Learned models begin to replace data

-Queryable, portable, pluggable, chainable, secure

- Experimental discovery processes are dramatically refactored
 - Models replace experiments, experiments improve models
- Many questions are 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
- Al becomes common part of scientific laboratory activities

 Infuses scientific, engineering, and operations

http://bit.ly/2FOnJi3

Al for Science @ Argonne targets both research and infrastructure

Applications	AI applications across science and engineering. Transformative approaches to simulation and experimental science.
Learning systems	Al software. Software infrastructure for managing data, models, workflows etc., and for delivering AI capabilities to 1000s 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.

Al for Science @ Argonne targets both research and infrastructure

Applications	Al applications across science and engineering. Transformative approaches to simulation and experimental science.
Learning systems	Al software. Software infrastructure for managing data, models, workflows etc., and for delivering AI capabilities to 1000s 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.

Al at Argonne: Dozens of applications projects



Al for Science: Advanced Photon Source Upgrade

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



Beyond the depth of focus limit: APS and ALCF

Depth of focus goes like DOF \simeq 5.4(transverse resolution)²/ λ .

- Example: DOF=4 μm for 10 nm resolution at 10 keV
- We can't treat images as simple projections of an object when going beyond DOF!
- How to combine penetrating power of X rays with tomorrow's resolution limit?
- We can describe the forward problem using multislice wave propagation [1,2]
- We can then "learn" how to adjust a 3D object to agree with microscope measurements.

- First demo required 200,000 core hours at Argonne's LCRC [3]

- Now using Automatic Differentiation (part of AI toolkits) to "learn" how to adjust a 3D object with additional info included
 - AD applied to parallelized 2D ptychography [4]
 - Presently using Argonne's ALCF Data Science Program allocation, plus support from LDRD and also NIH
 - Ming Du, Saugat Kandel, Sajid Ali, Northwestern; ANL APS; ANL MCS; ANL MSD.

Simulations: 12 DOF thick

DOF effects ignored: reduced resolution



DOF effects included with multislice/AD: full resolution 3D imaging!



[1] Cowley, Acta Cryst. 10, 609 (1957). [2] Li et al., Optics Express 25, 1831 (2017). [3] Gilles et al., Optica 5, 1078 (2018). [4] Nashed et al., Procedia Computer Science 108, 404 (2017).

Al in materials and chemistry

AI will accelerate discovery through autonomy in data collection and physics-based model building

Discover New Phenomena in Extracting rules from **Heterogeneous Data Streams** multimodal data Supervised and unsupervised learning will uncover pathways and intermediates encoded in multidimensional data **Achieve Deterministic Control** Predicting of Multiscale Transformations transformations across Active learning and deep neural networks will navigate reaction scales networks and multicomponent phase behavior **Deliver on the Promise of** Translating design to **Predictive Design of Matter** creation Reinforcement learning and decision trees will autonomously explore and control multiscale synthesis Multimodal Streams at APS (10⁴-10⁵/day) Advanced sampling for sparse data **Data Needs** AI Needs High Throughput Synthesis (10³/day) Rare-event detection and interpretation

- Simulation/Synthetic Data (10¹-10⁴/day)
- Domain specific models



Interrogate organization, dynamics, and reactivity of multiscale systems



Derive rules of dynamical organization among charges, spins, and atoms



Design and synthesize functional atomic and defect landscapes

Opportunities for Al in Materials and Chemistry: Findings of an Argonne internal task force

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

AI for materials and chemistry: Examples









23

Liu et al. (2019)

Al for science: High Energy Physics at Argonne

Energy/Intensity Frontier:

- GANs for detector simulation
- CNN-based event classification
- Learn detector effects that impact physics
- Hyperparameter scans to reduce reconstruction time
- Anomaly detection with autoencoders
- Neuromorphic models/accelerators for use in triggers

Cosmic Frontier (SciDAC-4):

- GANs for object catalog emulation
- GP/DL-based emulators for diverse cosmic probe summary statistics
- CNN-based image classification
- AI/ML-based methods for large-scale training data production and estimation methodology for photometric redshifts
- Likelihood emulation and likelihood-free methods for statistical inference



AI applications in an "end-to-end" Cosmic Frontier application:

GANs for image emulation, 2) GP and DL-based emulators for summary statistics,
 CNN-based image classification, 4) AI-based photometric reshift estimation,
 Likelihood-free methods for inference [Work performed under the Argonne-led SciDAC-4 project: "Inference and Machine Learning at Extreme Scales"]

Al for transformative manufacturing science

Science Impact	Available Data	Al Role	Manufacturing Impact
Mimic Pt-group metal catalysis using earth-abundant elements	High throughput catalyst performance testing	Data simulation and prediction	Predict economical catalytic conversions
Expand thermochemical tables	Rotational and mass spectra of gas-phase reaction products	On-the-fly spectral deconvolution, queryable models for pathway prediction	Predict new synthesis pathways via flame spray pyrolysis
Improve models of ion transport in materials	Real-time battery performance and lifetime data	Anomaly detection, correlation modeling	Improve battery performance

Al for Science @ Argonne targets both research and infrastructure

Applications	AI applications across science and engineering. Transformative approaches to simulation and experimental science.
Learning systems	Al 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.
	Advanced hardware to support AI. Evaluation of new
Hardware	architectures and systems; exploration of neuromorphic

New applications require new methods



Deep learning uses massive computing



Relationships between AI/ML and HPC



- HPC for AI/ML: HPC Technologies are applied to learning tasks to accelerate computation and/or solve larger problems
- AI/ML for HPC: Learning Technologies are applied to HPC computations to improve their performance in some way, e.g., by choosing the next simulation(s) to perform

Robust Learned Function Accelerators (RLFAs) Fluidity between simulations and learned models



Logan Ward, Ben Blaiszik, 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











A rapidly evolving computing/data continuum



Sources: http://bit.ly/2SDGHzT, https://doi.org/10.1007/978-3-319-31903-2_8, Pete Beckman

New methods require new mechanisms



Project Celerity* (My "Office of Science Distinguished Scientists Fellow" project)

Identity new mechanisms needed to bridge the gap between new (especially "AI-first") scientific applications and the emerging data/computing continuum

Via a process of experimentation, discussion, and debate

* Celerity (n) rapidity; swiftness; speed – from Latin celeritas, from which also c for speed of light in vacuum

Bridging the gap

What mechanisms will facilitate the programming of this distributed, heterogeneous, dynamically evolving AI-first continuum?

Some things that we surely need:

- Function: Compute wherever is fastest, cheapest, closest, most accurate, ...
- Data: Access where fastest, cheapest, closest, most accurate, ...
- Trust: Balance certainty vs. cost
- **Cost**: Useful estimates of the state of this dynamic system



Coding the continuum: Elements of an open solution

Thanks to colleagues, especially:



Rachana



Yadu Babuji Ananthakrishnan













Kyle Chard





Ryan Chard





Zhuozhao Li









Logan Ward

Coding the continuum: Elements of an open solution





Coding the continuum: Elements of an open solution

In [1]:	<pre>from funcx_sdk.client</pre>	: import FuncXClient	
	<pre>fxc = FuncXClient()</pre>		
In [2]:	<pre>func = """ def add(data): sum_val = sum(dat return sum_val """</pre>	ta['data'])	
In [3]:	fxc.register_function	n("add_func", func, description	="Sum a list of numbers.")
In [4]:	input_data = [1, 2, 3	3]	
	res = fxc.run(input_c	data, "user#laptop", "add_func"))
In [5]:	<pre>print(res)</pre>		
	6		
Ро	rtable code	Any access	Any compute
Do	Python	SSH, Globus,	Clusters,

scheduler

accelerators

Singularity



*func*X: Transform clouds, clusters, and supercomputers into high-performance function serving systems



*func*X: Transform clouds, clusters, and supercomputers into high-performance function serving systems





DLHub: Organizing and Serving Models

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
- Citable DOIs for reproducible science



Ward et al.

TomoGAN: Liu et al.

Cherukara et al.

Argonne Advanced Computing LDRD

DLHul

https://www.dlhub.org

Data and Learning Hub for Science





We must also rethink other technologies



We must also rethink other technologies



Al for Science @ Argonne targets both 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.
	Mathematics, algorithms; general AI, reinforcement
Foundations	Mathematics, algorithms; general AI, reinforcement learning, uncertainty quantification, explainability, etc.

Future directions in foundations

- Leverage 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

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

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

Attendees at an Argonne

ATP "theory institute" in 1990.

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 Yo	rk at Albany
Matt Kaufmann	Computational Logic	
Ken Kunen	University of Wisconsin	
Vladimir Lifschitz	Stanford University	Attandaas
Ewing Lusk	Argonne	ATD "theor
William McCune	Argonne	AIP theor
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	

1. <u>Algebraic Geometry</u>

- Cancellative Semigroups on a Cubic Curve
- Uniqueness of the 5-ary Steiner Law
- 2. Cancellative Semigroups
- 3. Lattice Theory
 - A Simpler Absorptive Basis for Lattice Theory
 - A New Schema for Single Axioms
 - <u>A Shorter Single Axiom for Lattice Theory</u>
 - 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
 - Ordinary Groups
 - <u>Abelian Groups</u>
 - Exponent Groups
 - Some Permutative Varieties
 - Ordinary Groups (Kunen)
 - Groups of Exponent 4 (Kunen)
 - 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. <u>Two-valued Sentential Calculus</u>
- 24. Many-valued Sentential Calculus
- 25. Short Proofs in Various Logic Calculi
- 26. Pure Proofs in Equivalential Calculus



Credit: Lloyd DeGrane / The New York Times

Dr. William McCune at Argonne Labs, Illinois in his office with computer. The "Proof of Robbins Conjecture" problem is on the screen.

1996

8 Days on an IBM RS/6000 30 megabytes of memory

Powerful satisfiability modulo theory (SMT) solvers like Microsoft's Z3 adapt strategies from Otter

- Bounded model-checking of model programs
- Termination
- Security protocols
 - Business application modeling
- Cryptography Mag
- Model Based Testing (SQL-Server)
- ZB
- Your killer-application here

Sources: https://nyti.ms/37X8fUM, http://bit.ly/35J3lcE

Al research has strong symbolic as well as connectionist roots



https://neurovenge.antonomase.fr

Co-citation network of 100 most influential authors in publications mentioning AI

Evolution of academic influence of connectionist and symbolic approaches to AI





Al for Science @ Argonne targets both research and infrastructure

Applications	AI applications across science and engineering. Transformative approaches to simulation and experimental science.
Learning systems	Al 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.

Aurora: HPC and AI

>> Exaops/s for AI





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)



intel

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 Al acceleration

Which platforms will be good for science?

Al Chip Landscape

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



An 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, AlSpeech, Rokid, NextVPU, Canaan, Enflame, Eesay Tech	12
Startups Worldwide	Cerebras, Wave Computing, Graphcore, PEZY, Tenstorrent, ThinCl, Koniku, Adapteva, Knowm, 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, <u>Tachyum</u> ,AlphalCs,Syntiant, Habana, aiCTX, Flex Logix, Preferred Network, Cornami, Anaflash, Optaylsys, Eta Compute	44





https://github.com/basicmi/AI-Chip

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 Al model types for problems in materials, biomedical, cosmology, high-energy physics, energy systems, synthetic biology, climate, software optimization, architecture research etc.







- * 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)

DOE's AI for Science Townhalls

Organized by Argonne, Oak Ridge and Berkeley with participation from all DOE labs

- 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 (Oak Ridge), 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.
- ~300 people per meeting
- Output will be a report to guide strategic planning at Labs and DOE



I have covered just a few of the many activities underway at Argonne on AI in science

We look forward to collaboration with Fermilab and Uchicago in applications, learning systems, foundations, and hardware

lan Foster foster@anl.gov A 20-Year Community Roadmap for Artificial Intelligence Research in the US http://bit.ly/2JK8OZ9



Computing Community Consortium

