


SURVEY PAPER

Open Access



Convergence of artificial intelligence and high performance computing on NSF-supported cyberinfrastructure

E. A. Huerta^{1,2,3,4*†} , Asad Khan^{1,2,4†}, Edward Davis⁹, Colleen Bushell¹, William D. Gropp^{1,2,5}, Daniel S. Katz^{1,5,6,7}, Volodymyr Kindratenko^{1,2,5,6}, Seid Koric^{1,2,8}, William T. C. Kramer^{1,5}, Brendan McGinty^{1,2}, Kenton McHenry¹ and Aaron Saxton^{1,2}

*Correspondence:

eliuhu@illinois.edu

†E. A. Huerta, A. Khan, E. Davis and V. Kindratenko:

Contributed equally

⁴ Department of Physics,

University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA

Full list of author information is available at the end of the article

Abstract

Significant investments to upgrade and construct large-scale scientific facilities demand commensurate investments in R&D to design algorithms and computing approaches to enable scientific and engineering breakthroughs in the big data era. Innovative Artificial Intelligence (AI) applications have powered transformational solutions for big data challenges in industry and technology that now drive a multi-billion dollar industry, and which play an ever increasing role shaping human social patterns. As AI continues to evolve into a computing paradigm endowed with statistical and mathematical rigor, it has become apparent that single-GPU solutions for training, validation, and testing are no longer sufficient for computational grand challenges brought about by scientific facilities that produce data at a rate and volume that outstrip the computing capabilities of available cyberinfrastructure platforms. This realization has been driving the confluence of AI and high performance computing (HPC) to reduce time-to-insight, and to enable a systematic study of domain-inspired AI architectures and optimization schemes to enable data-driven discovery. In this article we present a summary of recent developments in this field, and describe specific advances that authors in this article are spearheading to accelerate and streamline the use of HPC platforms to design and apply accelerated AI algorithms in academia and industry.

Keywords: Artificial intelligence, High performance computing

Introduction

The big data revolution disrupted the digital and computing landscape in the early 2010s [1]. Data torrents produced by corporations such as Google, Amazon, Facebook and YouTube, among others, presented a unique opportunity for innovation. Traditional signal processing tools and computing methodologies were inadequate to turn these big-data challenges into technological breakthroughs. A radical rethinking was urgently needed [2, 3].

Large Scale Visual Recognition Challenges [4] set the scene for the ongoing digital revolution. The quest for novel pattern recognition algorithms [5–7] that sift through large,

high-quality data sets eventually led to a disruptive combination of deep learning and graphics processing units (GPUs) that enabled a rapid succession of advances in computer vision, speech recognition, natural language processing, and robotics, to mention a few [3, 8]. These developments are currently powering the renaissance of AI, which is the engine of a multi-billion dollar industry.

Within just a few years, the curation of high-quality data sets, e.g., ImageNet [9]; GPU-accelerated computing [10]; open source software platforms—TensorFlow [11], PyTorch [12] among others—to design, train, validate and test AI models; improved AI architectures and novel techniques [13, 14] to enhance the performance of deep neural networks, such as robust optimizers [15] and regularization techniques [16], led to the rapid development of AI tools that significantly outperform other signal processing tools on many tasks [17, 18]. Data-driven discovery is now also informing and stirring the design of exascale cyberinfrastructure, in which high performance computing (HPC) and data have become a single entity, namely HPCD [2, 19].

Convergence of AI and HPC

The convergence of AI and HPC is being pursued in earnest across the HPC ecosystem. Recent accomplishments of this program have been reported in plasma physics [20], cosmology [21], gravitational wave astrophysics [22], high energy physics [23], multimessenger astrophysics [24], materials science [25], data management of unstructured datasets [26, 27], and genetic data [28], among others.

These achievements share a common thread, namely, the algorithms developed to accelerate the training of AI models in HPC platforms have a strong experimental component. To date, there is no rigorous framework to constrain the ideal set of hyperparameters that ensures rapid convergence and optimal performance of AI models as the number of GPU nodes is increased to accelerate the training stage. Furthermore, it is customary that distributed training algorithms in HPC platforms are benchmarked using idealized neural network models and datasets, e.g., training a ResNet model [29] using the ImageNet dataset [9]. While this approach provides some guidance about the optimal performance of HPC platforms for deep learning research, it does not impart any insights regarding the actual performance of these facilities when using domain-inspired AI architectures and optimization schemes to do data driven discovery in the context of realistic datasets, which are noisy, incomplete, and heterogenous—vastly different from the ImageNet dataset.

In view of these considerations, some key developments are needed to maximize the potential of AI for data-driven discovery: (i) the development of a rigorous mathematical framework to make informed choices of domain inspired AI architectures and optimization schemes; (ii) the creation of an interdisciplinary effort that brings together domain, information science, AI, data and software experts to inform the collection and curation of experimental and simulation datasets; (iii) the identification of connections between AI data and models, which will facilitate the production of commodity software that may be seamlessly applicable to disparate fields that share common data and computing data challenges; and (iv) the deployment of AI models and data on open source platforms, such as the Data and Learning Hub for

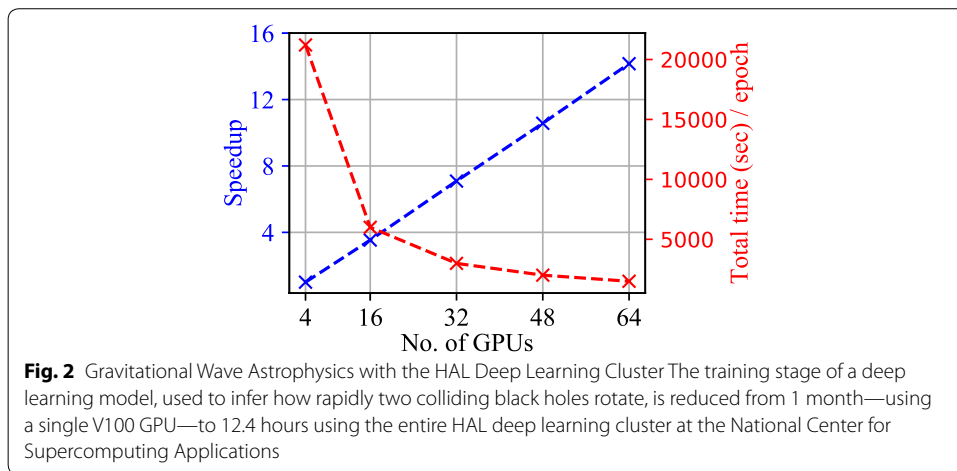
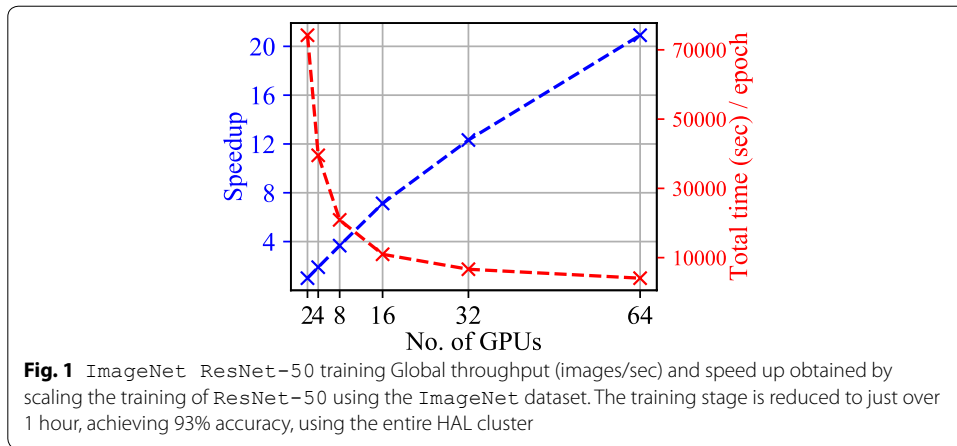
Science [30, 31]. These activities will accelerate the adoption of reproducible and robust AI tools as commodity software across disciplines.

There are several dedicated efforts in the literature to address these timely and relevant challenges, see e.g. [32–34]. In the US, the National Science Foundation (NSF) and the Department of Energy (DOE) are spearheading multi-million dollar programs for the construction of the next generation of HPC platforms to address computational grand challenges at the exascale, and on R&D to accelerate the design, deployment and adoption of innovative AI applications for data-driven discovery in science and engineering, and to translate these innovations into tangible societal benefits, business and industry. The funding of new HPC platforms for innovative AI research such as Bridges-2, Delta, and Neocortex will provide transformative capabilities by introducing new hardware for AI research [35, 36]. The Frontier, Aurora and El Capitan exascale systems will combine simulation, data science, and machine learning to revolutionize how supercomputers are used for scientific discovery and innovation.

In terms of R&D, DOE has launched an initiative to make AI models and data that adhere to FAIR data principles (Findable, Accessible, Interoperable, and Reusable). The goal of this program is to set a standard for the production of data that may be reusable both by researchers and machines, with little human intervention. It is expected that this approach will enable researchers to gain new insights on how AI models abstract knowledge from data, and to quantify how domain-inspired optimization schemes guide AI to the right answer in controlled experiments, while also enabling intuitive AI discovery that is beyond the reach of existing theories that do not fully capture complex phenomena, such as turbulence [37]. This program will maximize the use of exascale HPCD platforms, accelerating the development of AI.

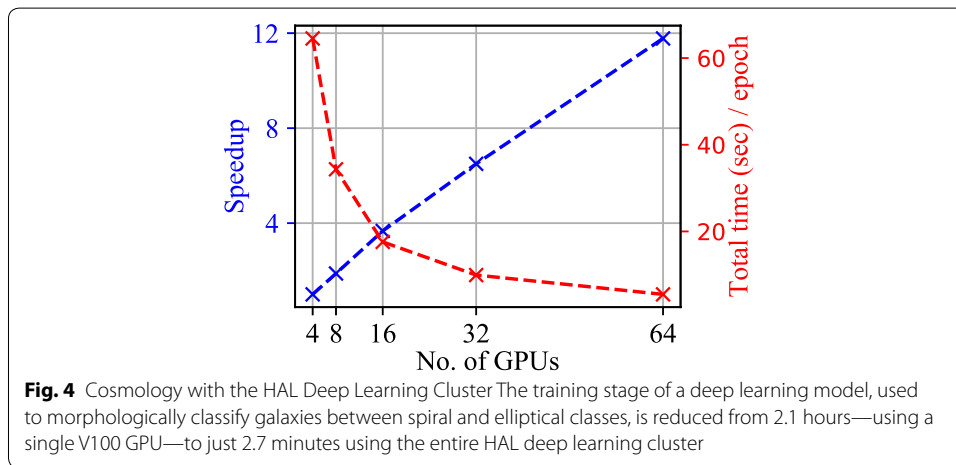
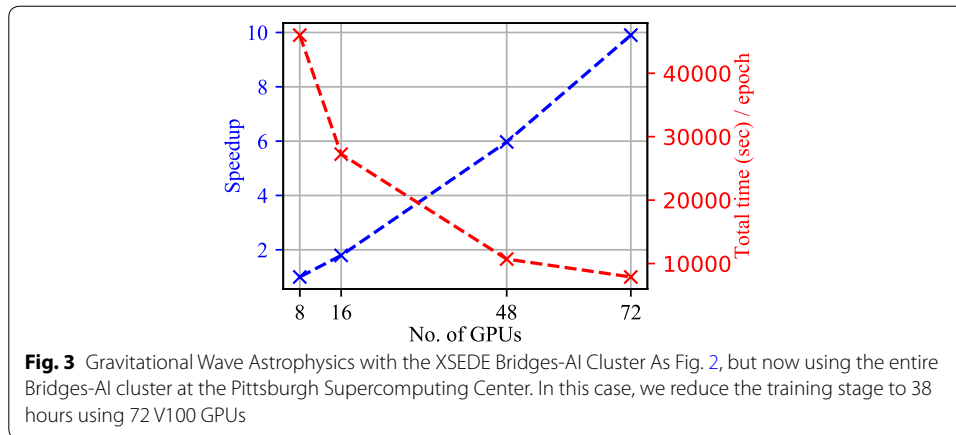
While it is customary to quantify the performance of HPC platforms for distributed training at scale using idealized datasets and vanilla AI models, i.e., ResNet-50 trained with the ImageNet dataset, it is also important to assess the performance of advanced cyberinfrastructure facilities to train more complex, domain-inspired AI models with realistic, experimental datasets. To provide a broad perspective on the state-of-the-art for different domains, we present results for a number of studies that we have conducted on NSF and DOE HPC platforms. The AI models we consider are tailored for image recognition, classification and regression analyses of telescope image datasets, and time-series data that describe the collision of black holes. To showcase the use of these models and datasets, we have used two NSF funded HPC platforms, namely, the Hardware-Accelerated Learning (HAL) cluster [38] at the National Center for Supercomputing Applications (NCSA), and the Bridges-AI system [39] that is part of the Extreme Science and Engineering Discovery Environment (XSEDE) at the Pittsburgh Supercomputing Center (PSC); and the DOE-funded Summit supercomputer at Oak Ridge National Laboratory [40].

HPC Platforms The HAL cluster has 64 NVIDIA V100 GPUs distributed evenly across 16 nodes, and connected by NVLink 2.0 [38] inside the nodes and EDR InfiniBand across the nodes. In Bridges-AI [39] we have used the 9 HPE Apollo 6500 servers, each with 8 NVIDIA Tesla V100 GPUs with 16 GB of GPU memory each, connected by NVLink 2.0.



AI models and datasets We have used three different AI models: (i) ResNet-50; (ii) an AI model to characterize the signal manifold of binary black hole mergers, trained with time-series signals that describe gravitational wave signals [14] (AI-GW); and (iii) an AI model that classifies galaxy images collected by the Sloan Digital Sky Survey (SDSS) [41], and automatically labels galaxy images collected by the Dark Energy Survey (DES) [21] (AI-DES). Our results of these analyses indicate:

- Figure 1 shows that ResNet-50 with ImageNet is trained within 41 hours using 1 V100 GPU in HAL. The training is reduced to just over 1 h, achieving 93% accuracy, using 64 V100 GPUs in HAL.
- Figure 2 shows that AI-GW is fully trained, achieving state-of-the-art accuracy, within 754 hrs using a single V100 GPU in HAL. When scaled to 64 V100 GPUs, the training is reduced to 12.4 h.
- Figure 3 shows that AI-GW is fully trained, achieving state-of-the-art accuracy, within 38 h using 72 V100 GPUs in Bridges-AI.



- Figure 4 shows that AI-DES is trained within 2.1 hrs using a single V100 GPU in HAL. The training is reduced to 2.7 min using 64 V100 GPUs in HAL.

These examples clearly underscore the importance of coupling AI with HPC: (i) it significantly speeds up the training stage, enabling the exploration of domain-inspired architectures and optimization schemes, which are critical for the design of rigorous, trustworthy and interpretable AI solutions; (ii) it enables the use of larger training data sets to boost the accuracy and reliability of AI models while keeping the training stage at a minimum.

Software and hardware challenges

While open source software platforms have played a key role in the swift evolution of AI, they present a number of challenges when used in HPC platforms. This is because open source software platforms such as TensorFlow [11] and PyTorch [12] are updated at a much faster pace than libraries deployed cluster-wide on HPC platforms. For instance, in typical HPC platforms, software updates customarily take place twice per year [42, 43]. In the case of open source AI APIs, releases happen much more often, as can be

seen in the official release timeline of TensorFlow [44]. Furthermore, producing AI models usually requires a unique set of package dependencies. Therefore, the traditional use of modules has limited effectiveness since software dependencies change between projects and sometimes evolve even during a single project. Common solutions to give users more fine-grained control over software environments include containerization, e.g., Singularity [45] or Kubernetes [46], and virtual environments such as Anaconda [47] that is provided in HPC platforms such as Bridges, Bridges-AI, Summit, and HAL. GPUs play a key role in the renaissance of AI because they have unique features to accelerate applications, e.g., they have many cores, provide high throughput, they are good for parallel processing and can perform thousands of operations at once. While these features are particularly relevant for image recognition analysis, gaming and graphics, GPUs are now used extensively in other areas, i.e., autonomous driving and robotics. In the context of HPC and AI, our studies indicate that 5 nodes (each node has 64 Intel KNL 7230 compute cores) in Theta are equivalent to a single V100 GPU. Thus, given how involved it is to optimally scale the training of AI models in HPC platforms, it is apparent the advantage provided by GPU-based HPC platforms for AI research.

We provide below a number of recommendations to streamline the use of HPC resources for AI research:

1. Provide up-to-date documentation and tutorials to set up containers and virtual environments, and adequate help desk support to enable smooth, fast-paced project life-cycles.
2. Maintain a versatile, up-to-date base container image, and base virtual environment that users can easily clone and modify for their specific needs.
3. Distributed training software stacks such as TensorFlow depend on distributed training software stacks, e.g., Horovod [48], which in turn depend on system architecture and specific versions of MPI installed by system and service managers. It is important to have clear up-to-date documentation on system architecture and MPI versions installed, and clear instructions on how to install/update distributed training software packages like Horovod into the user's container/virtual environment.

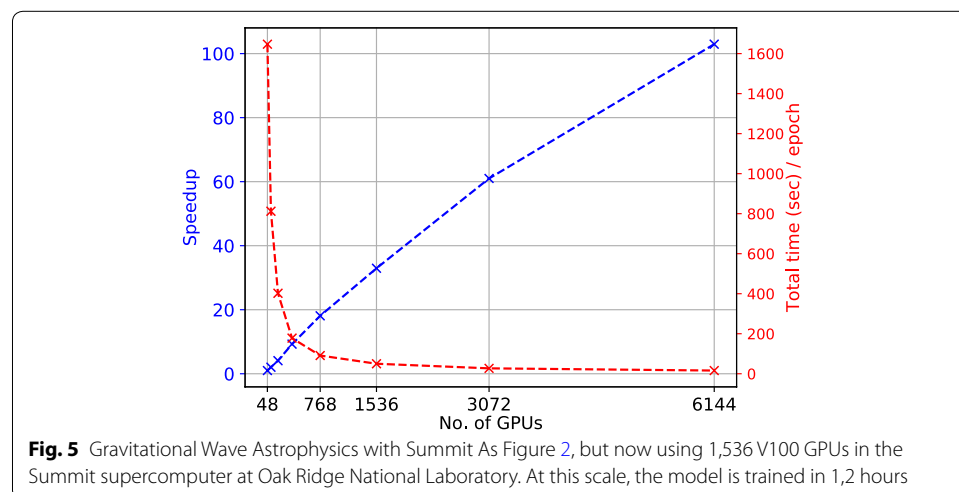
In addition to these considerations, the AI model architecture, dataset, and training optimizer prevent a seamless use of distributed training. Stochastic gradient descent (SGD) [49] and its variants are the workhorse optimizer for AI training. The common way to parallelize training is to use “mini-batches” with SGD. In principle, a larger mini-batch may naively utilize more GPUs (or CPUs). Training time to solution will often scale linearly with small batch size. Figures 2 and 4 show good generalization at 64 GPUs, which amounts to a global batch size of 128 samples. However, it is known that as data sets and number of features grow, naively scaling number of GPUs, and subsequently batch size, will often take more epochs to achieve an acceptable validation error. The state-of-the art in AI training at scale was reported in [50]. Therein, ResNet was trained using a batch size of 64k samples, run across 2048 Tesla P40s. While achieving this level of scaling required a lot of experimental work, this benchmark, and others [51], indicate that scaling AI models to larger data and feature sets is indeed possible. However, it requires a considerable amount of human effort to tune the model and training

pipeline. A mixture of fast human model development cycle mixed with automated hyper-parameter tuning is a candidate solution to tackle this problem.

We have explored whether the methods we have used in the context of HAL and Bridges-AI may work in other HPC platforms optimized for AI research. In Fig. 5 we show that our distributed training algorithms exhibit strong scaling up to 1024 nodes (6144 V100 GPUs) in the Summit supercomputer at Oak Ridge National Lab. The scaling efficiency, i.e., how long it takes to cycle through all of the data once, also known as Total time / epoch (see y-axis label on the right of Fig. 5) can be affected by many factors, e.g., I/O speed, communication, etc., and achieving good efficiency and strong scaling, as shown in this Figure, indicates that we have dealt with properly with these factors.

Furthermore, Fig. 5 shows that using 256 nodes (1,536 V100 GPUs) in the Summit supercomputer we are able to fully train a physics-inspired version of the WaveNet model with time-series data that describes numerical solutions to Einstein's equations that model black hole collisions, attaining state-of-the-art accuracy, within just 1.2 hours. In other words, we can generalize the methods deployed and tested on NSF-funded cyberinfrastructure to HPC platforms that have different scale, hardware and software.

Open challenges A number of challenges remain towards an optimal exploitation of AI and extreme scale computing. For instance, it is recognized that some experimental datasets are not in a suitable format to fully exploit data-driven discovery. To address this pressing issue, DOE has made significant investments to make AI models and data FAIR [52]. Another challenge concerns the design of AI models whose architecture and optimization schemes incorporate domain knowledge, enabling AI models to converge faster while also enabling intuitive, serendipitous discovery that may not be encapsulated by approximate descriptions of complex phenomena [37, 53]. It is also essential to develop a rigorous approach to maximize the use of HPC platforms for distributed training. This requires a systematic approach to select an optimal set of hyperparameters that enables faster convergence, and creative methods to use less training data to achieve state-of-the-art performance. NSF has also funded several institutes to advance



the state-of-the-art in AI, seeking new modes of data-driven discovery in science and engineering. These investments aim to sustain, broaden and accelerate recent breakthroughs in science, technology and industry driven by AI applications [54]. As these projects evolve and mature, it will be essential to facilitate cross-pollination of expertise, avoiding duplication and empowering new AI practitioners to access AI scientific software that is open source, interpretable, reproducible and trustworthy.

Cloud computing and HPC

Cloud computing and containerization became popular for developing customer facing web apps. It allowed a DevOps team—i.e., the team that develops scientific software and manages ongoing operations of a data center—to keep strict control of the customer facing software, while new features and bug fixes were designed, developed, and tested in an environment that “looked the same” as a live one. Depending on the business cycle, companies could dynamically scale their infrastructure with virtually no overhead of purchasing hardware, and then relinquish it when it was no longer needed.

HPC would do well to adopt a DevOps cycle like the ones seen in startup culture. However HPC has some unique challenges that make this difficult. (1) Data storage separated from compute in the form of a shared file system and an instance on maintaining a traditional tree like file system. Cloud computing delivers a unit of compute and storage in tandem as a single instance and isolates distinct resources. A developer using cloud resources treats a compute instance as only the host for their code and must explicitly choose how to move large volumes of data on and off. This is usually done by allocating a specialized cloud instance of a data store, e.g., SQL databases. Improved cloud solutions provide *Kubernetes* (and other cluster manager) recipes to allocate a skeleton of these resources, but it is still up to the developers to choose exactly how data are moved between the resources and to code the specific functions of their app. (2) HPC is a shared resource. That is, many users with different projects see the same file system and compute resource. Each developer must wait their turn to see their code run. In cloud computing, a resource belongs and is billed to the developer on demand. When the resource is released, all of its state-full properties get reset. (3) HPC is very concerned with the compute resources interconnect. To have high bandwidth and low latency between cloud compute instances, one pays a premium.

In the case of distributed training, one needs to ascertain whether the cloud or HPC platforms provide an adequate solution. On-demand, high throughput or cloudbursting of single-node applications are ideally suited for the cloud. For instance, in the case of genetic data analysis, the *KnowEng* platform [28] is implemented as a web application where the compute cluster is managed by *Kubernetes*, and provides an example of a workflow that can be expanded to include methods for intuitively managing library compatibility and cloud bursting. This cloud-based solution includes: (1) the ability to access disparate data; (2) set parameters for complex AI experiments effortlessly; (3) deploy computation in a cloud environment; (4) engage with sophisticated visualization tools to evaluate data and study results; and (5) save results and access parameter settings of prior runs.

However, large distributed training workloads that run for many hours or days will continue to excel on a high-end HPC environment. For instance, the typical utilization

of the HAL cluster at NCSA tends to be well above 70%. Given that the cost of a single V100 GPU node on AWS (p3.2xlarge instance [55]) is \$3.06 per hour, HAL provides over \$141,000 in comparable cloud compute resources every month. This is far higher than the amortized cost of the HAL cluster and its support. A top-tier system like Blue Waters, where a node hour is charged at \$0.60, 4,228 K20 GPUs might have a cloud cost of \$2-3M per month.

Industry applications

The confluence of AI and HPC is a booming enterprise in the private sector. NCSA is spearheading its application to support industry partners from the agriculture, health-care, energy, and financial, sectors to stay competitive on the global market by analyzing bigger and more complex data to uncover hidden patterns, reveal market and cash flow trends, and identify customer preferences [56]. The confluence of modeling, simulation and AI is another area of growing interest among manufacturing and life science partners, promising to significantly accelerate many extremely difficult and computationally expensive methods and workflows in model-based design and analysis [37, 57, 58].

Academic innovation in AI pursues ideas that are exciting and productive, though they may not have immediate, tangible benefits. While academic scholarship is curiosity driven research, innovative AI applications in industry have as a goal to address computational grand challenges at an accelerated pace, and to apply at scale new solutions to profit from them. In brief, while academia and industry pursue distinct goals, it is essential that both spheres of activity maintain a close-knit collaboration [59]. This is a critical endeavor because breakthroughs in industry and technology over the last decade were enabled by basic AI applications. As industrial applications reach new frontiers and computational grand challenges arise, it will be essential to continue leveraging AI innovation, and explore ways to translate it into tangible solutions that may be deployed at scale to produce societal and business benefits. In summary, the training of future AI practitioners demands an interdisciplinary approach that includes a clear vision of industry needs. This approach will ensure that academic AI innovation is readily incorporated and applied, creating a sustainable paradigm that opens up diverse lines of funding for AI researchers.

Conclusion

The convergence of AI and HPC provides the means to address big data challenges in science, engineering and industry, and enables the creation of disruptive approaches for data-driven discovery and innovation. Realizing these goals demands a concerted effort between AI practitioners, HPC and domain experts.

As AI and HPC continue to transform an ever increasing number of disciplines at an accelerated pace, we can only image what the future holds once AI is powered with a rigorous mathematical framework. In that scenario, it will be possible to optimally use oversubscribed HPC platforms, and create intuitive AI solutions that will lead to transformational scientific discoveries, and disruptive solutions in industry and technology

Finally, to contribute to the use of realistic datasets to benchmark HPC platforms, we release two neural network models, along with datasets, that we used to produce Figs. 2, 3, 4 and 5. As the NSF and other funding agencies continue to deploy faster and more

powerful HPC platforms for AI research, it is urgent that we provide guidelines to maximize the use of these resources, and continue training new talent that will catalyze the adoption and best AI practices. This approach was critical in the past to enable the adoption of HPC by industry, and will play a more significant role in the future given the eagerness with which industry is adopting AI solutions.

Abbreviations

AI: Artificial intelligence; GPUs: Graphics processing units; HPC: High performance computing; R&D: Research and development; NCSA: National Center for Supercomputing Applications; NSF: National Science Foundation; HAL: Hardware-Accelerated Learning; XSEDE: Extreme Science and Engineering Discovery Environment; PSC: Pittsburgh Supercomputing Center; GW: Gravitational wave; SDSS: Sloan Digital Sky Survey; DES: Dark Energy Survey.

Acknowledgements

We thank Nicholas A. Nystrom, Paola Buitrago and Julian Uran for their support using Bridges-AI; and Arjun Shankar, Tom Gibbs, Junqi Yin, and Jeff Larking for their support and guidance using the Summit supercomputer. We also thank Ben Blaiszik, Ryan Chard and Logan Ward for their support deploying our neural network models and testing datasets at the Data and Learning Hub for Science hosted by Argonne National Lab.

Authors' contributions

EAH led and coordinated the writing of this article. ED, EAH, AK and VK produced results for Figs. 1 and 2. EAH and AK produced results for Figs. 3, 4 and 5. All authors contributed to developing the ideas, and writing and reviewing this manuscript. All authors read and approved the final manuscript.

Funding

EAH, AK, DSK, and VK gratefully acknowledge National Science Foundation (NSF) award OAC-1931561. EAH and VK also acknowledge NSF award OAC-1934757. This work utilized XSEDE resources through the NSF award TG-PHY160053, and the NSF's Major Research Instrumentation program, award OAC-1725729, as well as the University of Illinois at Urbana-Champaign. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

Availability of data and materials

The neural network models and data used to study characterize black hole mergers, and to classify galaxy images, are readily available at the Deep Learning Hub (DLHub) [30, 31] hosted by Argonne National Laboratory (ANL) [60, 61].

Ethics approval and consent to participate

Not applicable

Consent for publication

The authors approve the publication of this manuscript

Competing interests

The authors declare that they have no competing interests.

Author details

¹ National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA. ² NCSA Center for Artificial Intelligence Innovation, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA. ³ Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA. ⁴ Department of Physics, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA. ⁵ Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA. ⁶ Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA. ⁷ School of Information Sciences, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA. ⁸ Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA. ⁹ University of Queensland, St Lucia, QLD 4072, Australia.

Received: 24 June 2020 Accepted: 27 September 2020

Published online: 16 October 2020

References

- Asch M, Moore T, Badia R, Beck M, Beckman P, Bidot T, Bodin F, Cappello F, Choudhary A, de Supinski B, Deelman E, Dongarra J, Dubey A, Fox G, Fu H, Girona S, Gropp W, Heroux M, Ishikawa Y, Keahey K, Keyes D, Kramer W, Lavignion J-F, Lu Y, Matsuoka S, Mohr B, Reed D, Requena S, Saltz J, Schulthess T, Stevens R, Swamy M, Szalay A, Tang W, Varoquaux G, Vilotte J-P, Wisniewski R, Xu Z, Zacharov I. Big data and extreme-scale computing: Pathways to convergence-toward a shaping strategy for a future software and data ecosystem for scientific inquiry. *Int J High Performance Comput Appl*. 2018;32(4):435–79.
- National Academies of Sciences, Engineering, and Medicine. *Opportunities from the Integration of Simulation Science and Data Science: Proceedings of a Workshop*. The National Academies Press, Washington, DC, 2018.
- Goodfellow Ian, Bengio Yoshua, Courville Aaron. *Deep Learning*. Cambridge: The MIT Press; 2016.

4. Russakovsky Olga, Deng Jia, Hao Su, Krause Jonathan, Satheesh Sanjeev, Ma Sean, Huang Zhiheng, Karpathy Andrej, Khosla Aditya, Bernstein Michael, Berg Alexander C, Fei-Fei Li. ImageNet large scale visual recognition challenge. *Int J Comput Vision*. 2015;115(3):211–52.
5. Lecun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proceed IEEE*. 1998;86(11):2278–324.
6. Lecun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521:436–44.
7. LeCun Y, Boser B, Denker JS, Henderson D, Howard RE, Hubbard W, Jackel LD. Backpropagation applied to handwritten zip code recognition. *Neural Comput*. 1989;1(4):541–51.
8. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778, June 2016.
9. Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L. ImageNet: A Large-Scale Hierarchical Image Database. In: *CVPR09*, 2009.
10. Krizhevsky A, Sutskever I, Hinton G. Imagenet classification with deep convolutional neural networks. *NIPS*, 2012.
11. Abadi Martin, Agarwal Ashish, Barham Paul, Brevdo Eugene, Chen Zhifeng, Citro Craig, Corrado Greg S, Davis Andy, Dean Jeffrey, Devin Matthieu, Ghemawat Sanjay, Goodfellow Ian, Harp Andrew, Irving Geoffrey, Isard Michael, Jia Yangqing, Jozefowicz Rafal, Kaiser Lukasz, Kudlur Manjunath, Levenberg Josh, Mané Dan, Monga Rajat, Moore Sherry, Murray Derek, Olah Chris, Schuster Mike, Shlens Jonathon, Steiner Benoit, Sutskever Ilya, Talwar Kunal, Tucker Paul, Vanhoucke Vincent, Vasudevan Vijay, Viégas Fernanda, Vinyals Oriol, Warden Pete, Wattenberg Martin, Wicke Martin, Yu Yuan, Zheng Xiaoqiang. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
12. Paszke Adam, Gross Sam, Massa Francisco, Lerer Adam, Bradbury James, Chanan Gregory, Killeen Trevor, Lin Zeming, Gímelshein Natalia, Antiga Luca, Desmaison Alban, Kopf Andreas, Yang Edward, DeVito Zachary, Raison Martin, Tejani Alykhan, Chilamkurthy Sasank, Steiner Benoit, Fang Lu, Bai Junjie, Chintala Soumith. Pytorch: An imperative style, high-performance deep learning library. In: *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
13. Raissi Maziar, Perdikaris Paris, Karniadakis George. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Physics*. 2018;378:11.
14. Khan Asad, Huerta EA, Das Arnab. Physics-inspired deep learning to characterize the signal manifold of quasi-circular, spinning, non-precessing binary black hole mergers. *Phys Lett B*. 2020;808:135628.
15. Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization; 2014.
16. Kukačka Jan, Golkov Vladimir. and Daniel Cremers. Regularization for deep learning: A taxonomy; 2017.
17. Schmidhuber Juergen. Deep learning in neural networks: An overview. *Neural Netw*. 2015;61:85–117.
18. Sejnowski Terrence J. The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences*, 2020.
19. National Academies of Sciences, Engineering, and Medicine. *Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017-2020*. The National Academies Press, Washington, DC, 2016.
20. Svyatkovskiy Alexey, Kates-Harbeck Julian, Tang William. Training distributed deep recurrent neural networks with mixed precision on gpu clusters. In: *Proceedings of the Machine Learning on HPC Environments, MLHPC'17*, New York, NY, USA, 2017. Association for Computing Machinery.
21. Khan Asad, Huerta EA, Wang Sibio, Gruendl Robert, Jennings Elise, Zheng Huihuo. Deep learning at scale for the construction of galaxy catalogs in the Dark Energy Survey. *Phy Lett B*. 2019;795:248–58.
22. Shen Hongyu, Huerta E. A., Zhao Zhizhen. Deep Learning at Scale for Gravitational Wave Parameter Estimation of Binary Black Hole Mergers. arXiv e-prints, page [arXiv:1903.01998](https://arxiv.org/abs/1903.01998), Mar 2019.
23. Guest Dan, Cranmer Kyle, Whiteson Daniel. Deep learning and its application to lhc physics. *Annual Rev Nucl Particle Sci*. 2018;68(1):161–81.
24. Huerta EA, et al. Enabling real-time multi-messenger astrophysics discoveries with deep learning. *Nature Rev Phys*. 2019;1:600–8.
25. Ward Logan, Blaiszik Ben, Foster Ian, Assary Rajeev S, Narayanan Badri, Curtiss Larry. Machine learning prediction of accurate atomization energies of organic molecules from low-fidelity quantum chemical calculations. *MRS Commun*. 2019;9(3):891–9.
26. Marini Luigi, Gutierrez-Polo Indira, Kooper Rob, Satheesan Sandeep Puthanveetil, Burnette Maxwell, Lee Jong, Nicholson Todd, Zhao Yan, McHenry Kenton. Clowder: Open source data management for long tail data. In *Proceedings of the Practice and Experience on Advanced Research Computing, PEARC'18*, New York, NY, USA, 2018. Association for Computing Machinery.
27. Padhy S, Jansen G, Alameda J, Black E, Diesendruck L, Dietze M, Kumar P, Kooper R, Lee J, Liu R, Marciano R, Marini L, Mattson D, Minsker B, Navarro C, Slavenas M, Sullivan W, Votava J, Zharnitsky I, McHenry K. Brown dog: Leveraging everything towards autocuration. In *2015 IEEE International Conference on Big Data (Big Data)*, Oct 2015; 493–500.
28. Blatti Charles, Emad Amin, Berry Matthew J, Gatzke Lisa, Epstein Milt, Lanier Daniel, Rizal Pramod, Ge Jing, Liao Xiaoxia, Sobh Omar, Lambert Mike, Post Corey S, Xiao Jinfeng, Groves Peter, Epstein Aidan T, Chen Xi, Srinivasan Subhashini, Lehnert Erik, Kalari Krishna R, Wang Liewei, Weinsillbourn Richard M, Song Jun S, Jongeneel C. Victor, Han Jiawei, Ravaioli Umberto, Sobh Nahil, Bushell Colleen B, Sinha Saurabh Knowledge-guided analysis of 'omics' data using the KnowEnG cloud platform. *PLoS biology*, 2020.
29. He Kaiping, Zhang Xiangyu, Ren Shaoqing, Sun Jian Deep residual learning for image recognition, 2015.
30. Chard R, Li Z, Chard K, Ward L, Babuji Y, Woodard A, Tuecke S, Blaiszik B, Franklin MJ, Foster I. DHub: Model and data serving for science. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2019; 283–292.
31. Blaiszik Ben, Ward Logan, Schwarting Marcus, Gaff Jonathon, Chard Ryan, Pike Daniel, Chard Kyle, Foster Ian. A data ecosystem to support machine learning in materials science. *MRS Commun*. 2019;9(4):1125–33.

32. Balaprakash P, Salim M, Uram TD, Vishwanath V, Wild S. M.. Deephyper: Asynchronous hyperparameter search for deep neural networks. In: 2018 IEEE 25th International Conference on High Performance Computing (HiPC), 2018; 42–51
33. Diaz GI, Fokoue-Nkoutche A, Nannicini G, Samulowitz H. An effective algorithm for hyperparameter optimization of neural networks. *IBM J Res Dev.* 2017;61(4/5):91–911.
34. Frankle, Jonathan, Carbin Michael. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv: Learning*, (2019)
35. NSF Funds Five New XSEDE-Allocated Systems, 2020. <https://www.xsede.org/-/nsf-funds-five-new-xsede-allocated-systems>.
36. Introducing Bridges-2, 2020. <https://www.psc.edu/bridges-2>.
37. Rosofsky Shawn G, Huerta EA. Artificial neural network subgrid models of 2D compressible magnetohydrodynamic turbulence. *Phys Rev D.* 2020;101(8):084024.
38. NCSA. HAL Cluster. <https://wiki.ncsa.illinois.edu/display/SL20/HAL+cluster>.
39. XSEDE. Bridges-AI. <https://portal.xsede.org/psc-bridges>.
40. Oak Ridge National Laboratory. Summit. <https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/>.
41. York Donald G, et al. The Sloan Digital Sky Survey: Technical Summary. *Astron J.* 2000;120:1579–87.
42. What's new with IBM Watson Machine Learning Community Edition, 2020. <https://www.ibm.com/support/pages/get-started-ibm-wml-ce>.
43. IBM Watson Machine Learning Community Edition V1.6.1 helps you get started faster with a software distribution for machine learning running on an enterprise platform for AI, 2019. <https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS219-164&appname=USN>.
44. TensorFlow Release Timeline, 2020. <https://github.com/tensorflow/tensorflow/releases>.
45. Kurtzer Gregory M. Singularity 2.1.2 - Linux application and environment containers for science, August 2016.
46. Kubernetes. <https://kubernetes.io/>.
47. Anaconda. <https://www.anaconda.com/>.
48. Sergeev, A, Del Balso M. Horovod: fast and easy distributed deep learning in TensorFlow. *ArXiv e-prints*, February 2018.
49. Bottou Léon. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *Proceedings of COMPSTAT'2010*, pages 177–186, Heidelberg, 2010. Physica-Verlag HD.
50. Jia Xianyan, Song Shutao, He Wei, Wang Yangzihao, Rong Haidong, Zhou Feihu, Xie Liqiang, Guo Zhenyu, Yang Yuanzhou, Yu Liwei, Chen Tiegang, Hu Guangxiao, Shi Shaohuai, Chu Xiaowen. Highly scalable deep learning training system with mixed-precision: Training imagenet in four minutes. 07 2018.
51. You Y, Zhang Z, Hsieh CJ, Demmel J, Keutzer K. ImageNet Training in Minutes. *ICPP 2018*. Association for Computing Machinery, New York USA, 2018.
52. Department of Energy Announces \$8.5 Million for FAIR Data to Advance Artificial Intelligence for Science, 2020. <https://www.energy.gov/articles/department-energy-announces-85-million-fair-data-advance-artificial-intelligence-science>.
53. van Nieuwenburg Evert P L, Liu Ye-Hua, Huber Sebastian D. Learning phase transitions by confusion. *Nat Phy.* 2017;13(5):435–9.
54. NSF leads federal partners in accelerating the development of transformational, AI-powered innovation, 2020. https://www.nsf.gov/news/news_summ.jsp?cntn_id=299329&org=NSF&from=news.
55. Amazon EC2 P3 Instances, 2020. <https://aws.amazon.com/ec2/instance-types/p3/>.
56. NCSA. NCSA Industry. <http://www.ncsa.illinois.edu/industry>. 2020.
57. Abueidda Diab W., Koric Seid, Sobh Nahil A.. Machine learning accelerated topology optimization of nonlinear structures. *arXiv e-prints*, page [arXiv:2002.01896](https://arxiv.org/abs/2002.01896), Jan 2020.
58. Luo Shirui, Cui Jiahuan, Vellakal Madhu, Liu Jian, Jiang Enyi, Koric Seid, Kindratenko Volodymyr. Review and Examination of Input Feature Preparation Methods and Machine Learning Models for Turbulence Modeling. *arXiv e-prints*, page [arXiv:2001.05485](https://arxiv.org/abs/2001.05485), Jan 2020.
59. Recht Ben, Forsyth David A, Efron Alexei. You Cannot Serve Two Masters: The Harms of Dual Affiliation, 2018. <http://www.argmin.net/2018/08/09/co-employment/>.
60. Asad Khan, Huerta Eliu A., Das Arnab. A deep learning model to characterize the signal manifold of quasi-circular, spinning, non-precessing binary black hole mergers, 2020. <https://doi.org/10.26311/8wnt-3343>.
61. Asad Khan, Huerta Eliu A., Wang Sibbo, Gruendl Robert, Jennings Elise, Zheng Huiho. Deep learning at scale for the construction of galaxy catalogs in the dark energy survey, 2019. <https://doi.org/10.26311/k54a-z689>.
62. HAL at Scale, 2020. <https://github.com/richardkxu/distributed-pytorch>.
63. Kindratenko Volodymyr, Mu Dawei, Zhan Yan, Maloney John, Hashemi Sayed, Rabe Benjamin, Xu Ke, Campbell Roy, Peng Jian, Gropp William. Hal: Computer system for scalable deep learning. 07 2020; 41–48

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.