

Integrating Deep Learning in Domain Sciences at Exascale ^{*}

Rick Archibald¹, Edmond Chow², Eduardo D’Azevedo¹, Jack Dongarra^{3,1},
Markus Eisenbach¹, Rocco Febbo³, Florent Lopez³, Daniel Nichols³,
Stanimire Tomov³, Kwai Wong³, and Junqi Yin¹

¹ Oak Ridge National Laboratory
{archbaldrk, dazevedoef, eisenbackm, yinj}@ornl.gov

² Georgia Institute of Technology
echow@cc.gatech.edu

³ University of Tennessee, Knoxville
{dongarra, flopez, tomov}@icl.utk.edu
{rfebbo, dnicho22}@vols.utk.edu
kwong@utk.edu

Abstract. This paper presents some of the current challenges in designing deep learning artificial intelligence (AI) and integrating it with traditional high-performance computing (HPC) simulations. We evaluate existing packages for their ability to run deep learning models and applications on large-scale HPC systems efficiently, identify challenges, and propose new asynchronous parallelization and optimization techniques for current large-scale heterogeneous systems and upcoming exascale systems. These developments, along with existing HPC AI software capabilities, have been integrated into MagmaDNN, an open-source HPC deep learning framework. Many deep learning frameworks are targeted at data scientists and fall short in providing quality integration into existing HPC workflows. This paper discusses the necessities of an HPC deep learning framework and how those needs can be provided (e.g., as in MagmaDNN) through a deep integration with existing HPC libraries, such as MAGMA and its modular memory management, MPI, CuBLAS, CuDNN, MKL, and HIP. Advancements are also illustrated through the use of algorithmic enhancements in reduced- and mixed-precision, as well as asynchronous optimization methods. Finally, we present illustrations and potential solutions for enhancing traditional compute- and data-intensive applications at ORNL and UTK with AI. The approaches and future challenges are illustrated in materials science, imaging, and climate applications.

^{*} This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

1 Background

Deep learning (DL) has become one of the most prevalent technologies today. Applications extend from image recognition to natural language processing. While its commercial applications have garnered the efforts of major technology companies, DL has proven vital to works in numerous scientific domains.

Some of these domain problems are vital to solving many of today’s global issues—for instance, in solving modern, data-dense problems in climate science. This field has seen several advancements with the recent integration of deep neural networks (DNNs). Ise et. al. [1] propose a modern approach towards utilizing neural networks to predict global temperature data. They further conclude that the accuracy of their model increases with the number of input images used. Thus, the model’s accuracy and utility grow with the amount of compute power available. Wang et. al. [2] provide further uses for DNNs in weather forecasting, which compete with the current state-of-the-art, physically based simulations. The neural networks can also reduce the computations necessary, as they possess the ability to extend one trained network to similar environments, or train once, and predict frequently.

Additionally, DL has proven useful in materials science, even leading the field of materials informatics [3]. Feng et. al. [4] were able to utilize deep networks to predict material defects with an accuracy above 98%, even for small data sets. However, these great accuracies require deep and/or pre-trained networks. As with climate modelling, deep learning can be used to replace current physically based simulations, which are expensive and/or do not scale well [5].

Providing high performance hardware and software stacks to facilitate deep learning has become a focus in the high-performance computing (HPC) field. An overview of the parallel algorithms and their challenges for distributed, deep neural networks can be found in Ben-Nun and Toefler [6]. Large-scale modern systems, such as the US DOE’s Summit, are fully loaded with high-performance GPUs to support the compute-intensive nature of deep learning. Additionally, some HPC distributed file systems under-perform in providing I/O for data-hungry DL training [7]. Even in terms of sheer performance, many state-of-the-art deep learning models and libraries fail to utilize a significant portion of the maximum floating-point operations (FLOPs) rate of current top systems [8]. There is significant research in the area of accelerating DL on cloud and other commercial systems, but these methods may not integrate well with varying HPC technologies and modern scientific workflows.

One major area in which HPC can innovate deep learning is model parallelism. In many domain applications, networks can become large, leading to growth beyond the memory of a single GPU. By utilizing existing HPC communication standards, a robust model-parallel training standard needs to be developed in addition to efficient implementations. Currently, several libraries attempt to integrate model parallelism [9], but this requires extensive knowledge of the code base, model, and hardware. There is an active research push into further developing model parallel approaches [10, 11, 12, 13].

R. Stevens presented a vision of the landscape of machine learning as DOE moves towards the exascale era [14], stressing the importance of machine learn-

ing as an integral part of DOE exascale applications. The new paradigm of "HPC + AI" on exascale computing leads towards an "Integrated Sim, Data, Learn Stack." There has been a long history of application code development for compute-intensive applications in the areas of climate, materials, medical, high energy, and urban sciences applications. In light of these needs, we will highlight developments and discuss the necessity of a set of tools and methods particularly relevant to deploying AI on HPC systems.

Realizing the paradigm of "HPC + AI" requires three major components: (1) the HPC system and investments in scaling application-driven capabilities; (2) the AI software achieving the potential performance; and (3) the "+", innovative algorithms and implementations/tools that combine the components of compute- and data-driven applications seamlessly on exascale machines.

2 Deep Learning Software on Modern HPC Systems

There are countless deep learning frameworks in existence today: TensorFlow [15], PyTorch [16], MxNet [17], and many others. Each framework has its own advantages and each falls short in some areas. One advantage of each of these frameworks is their strong corporate backings. Companies pour significant amounts of money into making production-ready libraries, which support various hardware. However, when using these frameworks in an HPC setting, this can cause issues. Much of the research and advancements by the library designers are production focused and target cloud environments. Some of these advancements translate well onto HPC systems, such as the use of hardware accelerators and mixed precision. However, many communication paradigms and implementations differ dramatically between the two. For example, the parameter server method in TensorFlow is more suited for heterogeneous cloud environments, although it is more bandwidth efficient than allreduce; and the socket-based GLOO communication backend built into PyTorch cannot efficiently utilize high-performance interconnects on HPC systems, although it is more flexible.

2.1 Towards a Deep Learning Framework for HPC

We evaluated a number of existing packages for their readiness to efficiently run deep learning models and applications on large-scale HPC systems. As pointed out previously, we found various limitations. Our vision of what is important and needed in deep learning AI packages for HPC is incorporated into the design of the MagmaDNN open source framework [18, 19] (see Figure 1) with current release MagmaDNN v1.2 [20], which we discuss briefly in what follows.

Firstly, MagmaDNN targets HPC systems, rather than cloud systems. It is written in C++ and is open source (available at bitbucket.org/icl/magmadnn).

Many of the core components of MagmaDNN are the same as other frameworks: operations are represented in a compute graph, data is passed around via tensors, gradients are computed using backpropagation, etc. However, each of these was designed with an HPC environment in mind.

MagmaDNN is centered around existing high-performance software packages, enabling easy integration into existing environments and Exascale Computing Project (ECP) software products. First and foremost, MagmaDNN uses the Matrix Algebra on GPU and Multicore Architectures (MAGMA) [21] package for heterogeneous linear algebra, since deep learning is heavily dependent on linear algebra routines. For strictly CPU computations, the package allows the use of any Basic Linear Algebra Subprograms (BLAS)/LAPACK packages such as Intel MKL, OpenBLAS, ATLAS, etc.

To better operate with hardware accelerators, MagmaDNN uses CUDA packages to harness GPU acceleration. Custom CUDA kernels, along with CUDNN’s routines for GPU-accelerated deep learning, allows the package to get maximal performance on GPU-enabled systems.

For distributed computing, MagmaDNN relies on existing HPC tools such as the Message Passing Interface (MPI) for its communication. Most systems have an MPI distribution tuned to that hardware, while some even have custom distributions. Relying on MPI for inter-node communication allows the framework to utilize existing optimized implementations while providing simpler integration with existing scientific codes.

This integration of existing HPC frameworks allows MagmaDNN to compile and link with the fastest libraries for particular hardware. Utilizing existing optimized packages is a consistent design pattern within HPC libraries, and MagmaDNN uses interfaces that are commonly installed on HPC clusters. In terms of integrating with existing systems and scientific codes, this approach offers the package a significant advantage over other frameworks and workflows.

Not only does MagmaDNN integrate with current HPC frameworks, it provides modularity in its design and linking so that users have flexibility in use. Conforming to standard APIs allows linking with any choice of framework; any BLAS/LAPACK/MPI choice implementations can be used. However, it is not just libraries with which MagmaDNN is modular. By utilizing computation graphs and extendable C++ classes, the package allows users to fully customize their neural networks. For instance, to utilize a unique network interface for distributed training, one can define a custom node on the computation graph which handles the transfer of data on that network component. This can be further utilized to define any desired data movement or computational elements. By providing modular components, MagmaDNN provides a structure onto which domain scientists can easily implement desired custom functionality.

Much of this customizability is fully supported through MagmaDNN’s interfaces. For instance, working with GPU and distributed memory is abstracted away in the `MemoryManager` class. Thus, library users do not need to worry about where memory is stored and how it is transferred. The package identifies the memory type, and whenever a memory operation is used (e.g., copy, put,

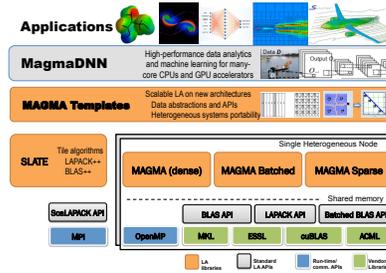


Fig. 1. MagmaDNN software stack.

get, etc.), and the MagmaDNN implementation handles the low-level specifics. Likewise, communication operations between devices or nodes are abstracted into the library as well. However, users can define their own operations, meaning they can redefine those in existing code. This makes it easy, for example, to try out different convolution algorithms and determine which one performs optimally. Another common example is to determine which distribution strategy is optimal for a specific machine. The same training code will run on any machine, and only the communication code might need to be altered.

Looking towards next-generation hardware and exascale systems, MagmaDNN aims to incorporate more technology stacks such as AMD’s HIP (e.g., hip-MAGMA is already available) and Intel’s OneAPI. Support for more technologies will allow researchers at various labs and institutions to focus on domain research—and not on how to integrate their code with HPC stacks. We also hope to provide streamlined and optimal support for model parallelism. As the size of models grows, so does the necessary memory requirements on devices. Soon many cutting-edge networks, which can provide state-of-the-art accuracy in classification, will require model parallelism. Supporting fast, scalable model parallelism on HPC systems is crucial to facilitating next-generation research, and MagmaDNN is moving towards this goal.

2.2 Workflow Software for Modern HPC Systems

Another major component in the integration of AI and HPC is the availability of workflow software: which is critical, for example, in launching, collecting results from, and analyzing various simulations in hyperparameter tuning, DL networks discoveries, etc. Exascale applications largely have their own set of workflow procedures and tools to launch their simulations. Integrating traditional applications seamlessly with existing AI software—either TensorFlow, PyTorch, or MagmaDNN—takes on a new level of challenges. Most workflow frameworks available today focus on cloud deployment, yet a workflow framework tailored to scale on HPC systems is critical to the success to “Automate and Accelerate” the task of “Sim, Data, and Learn Stack”. SWIFT-T [22] is a tool that offers the choice to run on large-scale HPC systems. At The University of Tennessee, Knoxville (UTK), we have developed a parallel workflow framework called openDIEL that aims to give researchers and users of HPC machines an efficient way to coordinate, organize, and interconnect many disparate modules of computation in order to effectively utilize and allocate HPC resources [23, 24]. It provides users an adaptive avenue to run compute-intensive and data science jobs concurrently, allowing specification of DNN architectures, data processing, and hyperparameter tuning. Existing ML tools can be readily used in openDIEL, allowing for easy experimentation with various models and approaches. Most importantly, openDIEL provides a platform to run existing ECP applications in their own workflow procedures. When conducting multi-discipline simulations, it is often required to use a large variety of software, and serial or parallel codes, to answer a research question. Utilizing disparate modules of computation requires the careful coordination of dependencies, com-

munication, and synchronization between them, and there is not always a clear path. This problem is solved by openDIEL: it enables researchers to define complex dependencies between modules and schedule communication between them. OpenDIEL consists of three primary components: the workflow control file, the communication library (COMMLIB), and the Executive, shown in Figure 2. Workflow is defined by users via the control file, which consists of two parts, the definition of functional modules and the interactions and sequence of execution of modules. OpenDIEL provides two communication libraries, one for direct point-to-point data transfer between two modules, and a store-and-move tuple space library for collective and asynchronous data transfers. The Executive is a lightweight manager designed to run and scale on heterogeneous HPC platforms.

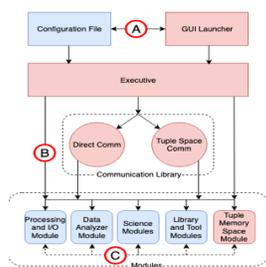


Fig. 2. OpenDIEL Structure.

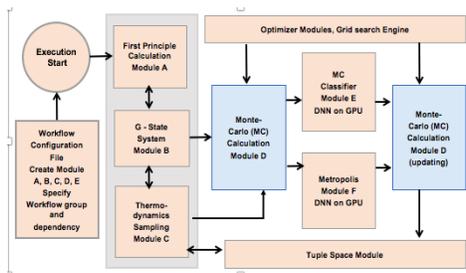


Fig. 3. Workflow in Materials Sciences.

OpenDIEL specifically focuses on unifying modules into a single executable that uses MPI for communication. Climate and materials sciences simulations, with typical workflow shown in Figure 3, are two of the applications that are incorporating deep learning into their simulations.

3 Algorithmic Improvements for DNN AI in HPC

In this section, we select and highlight the development of two technologies that we believe are critical for the scalability and efficient use of HPC systems. The first (Section 3.1) addresses scalability of DL training, stressing the benefits and need for asynchronous methods on today’s highly parallel and heterogeneous systems, featuring powerful nodes and relatively slow interconnects. The second (Section 3.2) addresses hardware changes motivated by DL—in particular powerful hardware acceleration for low-precision arithmetic—and techniques for harnessing this in scientific HPC applications.

3.1 Asynchronous Methods

In a multi-core or many-core shared-memory environment, the speed at which the parameters can be updated by different threads may be a limiting factor for stochastic gradient descent (SGD) performance. This is because threads must lock the parameter vector while they are updating it, preventing other threads

from reading or updating the parameters, so that they must wait. Asynchronous versions of SGD may relax the data dependencies and order of operations compared to the classical SGD algorithm, allowing the algorithm to update parameters more rapidly—but also making the asynchronous algorithm mathematically different and non-deterministic in its execution [25].

In a distributed-memory environment, a common implementation option for SGD is for each compute node to store a copy of the model parameters [26]. When each compute node has computed a gradient for its mini batch of data, all the gradients must be summed and then shared among the compute nodes so that they can update their copies of the model parameters. This leads to a potential bottleneck in the global all-reduce computation at every step of SGD. In addition, nodes that compute their gradients faster than others must wait for the other nodes (sometimes called “stragglers”). This can be a very high synchronization overhead that may be difficult to reduce by load balancing on heterogeneous architectures. An asynchronous approach can address these high overheads, as we demonstrate below.

Another common implementation option for distributed-memory computers is the parameter server approach [27]. Here, one or more compute nodes are dedicated to storing and updating the model parameters. The worker nodes compute gradients for their partition of the data and send these gradients to the parameter server. The parameter server then sends back the updated values of the model parameters. The parameter server model can be used to run SGD synchronously—with the same issues of stragglers slowing down the iteration for all nodes—but it has the advantage that it can also be run asynchronously (i.e., the worker nodes do not have to send gradients at the same time). Workers, however, may be working with different sets of values for the parameters at any instant. This makes the computation asynchronous. In distributed memory implementations, the gradients are typically updated by the parameter server using atomic memory operations, which differs from the shared memory case.

Asynchronous parameter server implementations were proposed around 2012 or earlier, but it was Hogwild! [25], with inconsistent reads and writes of the parameter vector in shared memory, that popularized the asynchronous approach. In particular, Hogwild! showed that, at least in cases of convex problems and gradient sparsity, asynchrony with inconsistent reads and writes does not have to harm convergence asymptotically. Hogwild! was extended in many directions, such as HogBatch [28], designed for efficiently running on multi-core CPU architectures, and Buckwild [29] which exploits low-precision arithmetic to accelerate the training. Improvements on the analysis of Hogwild! showed that convergence can be maintained for the non-convex and non-sparse case [30]. In MagmaDNN, we designed a parallel asynchronous variant of SGD that is similar to a parameter server approach, but on multi-core CPU and GPU architectures [31].

In practice, the accuracy of machine learning models can be dramatically improved by increasing the number of model parameters as well as the sizes of the training sets. As a result, the memory footprint needed for model training often exceeds the memory storage available on a single node. In such cases, models must be trained on large-scale machines where the dataset, and possi-

bly the model as well, are distributed among the compute nodes. From a performance perspective, it is crucial to efficiently parallelize the training in this setup—and this means overcoming the communication costs. It has been shown that when targeting large-scale platforms, asynchronous algorithms can outperform the synchronous ones as in the shared-memory case. In Google’s DistBelief framework [27], for example, the Downpour SGD algorithm uses a centralized parameter server for storing the model, and distributes the dataset to other participating nodes. Asynchrony enables communications to overlap with computations, thus improving resource usage. In addition, the algorithm is particularly well adapted to heterogeneous computing environments where resources may have different processing and communication speeds.

In MagmaDNN, we take a different approach, exploiting remote memory access (RMA) using MPI one-sided communication capabilities for implementing asynchronous SGD in order to maximize the scalability of parallel training. This approach has already been proven efficient in implementing several asynchronous numerical algorithms such as the Jacobi [32] or optimized Schwarz [33] methods for solving sparse systems of linear equations.

3.2 Reduced and Mixed Precision

Deep neural networks can be efficiently trained using low-precision floating-point arithmetic, such as single (fp32) or lower precision, improving the training times on modern hardware [34]. Popular floating-point formats include the 16-bit IEEE half-precision format (fp16) and the 16-bit bfloat16 [35] format. With fewer bits in the mantissa compared to fp16, the bfloat16 format offers less precision but has the advantage of having the same range as fp32—thereby removing the risk of overflow and underflow when converting from fp32 data.

Motivated by the success of low-precision arithmetic in machine learning applications, many specialized hardware accelerators have been designed with reduced- or mixed-precision arithmetic capabilities. For example, the NVIDIA Tensor Cores introduced in the Volta GPU can issue 4×4 matrix multiplication $D = C + A * B$ in one clock cycle, where the input matrices A and B are fp16, whereas D and C can be either fp16 or fp32. The theoretical performance peak on Volta GPU is 125 teraFLOP/s, which is 8 times the performance peak for fp32; and, although the accuracy of tensor core operations is lower than fp32, the ability to accumulate the result of matrix multiplication in fp32 rather than fp16 yields a considerable gain with respect to accuracy over fp16 [36]. In [37], authors use Tensor Cores for computing a reduced-precision LU factorization of a dense linear system, initially generated in double precision (fp64), and use these factors as a preconditioner for a Krylov solver in order to retrieve an fp64 solution. By doing this, they manage to reduce the time to solution up to $4 \times$ fp64 arithmetic for the factorization.

Tensor Cores can also be used in performing FFTs [38]. Note that a matrix in fp32 can be well approximated as the scaled sum of two fp16 matrices

$$A_{32} \approx a_{132} * A_{16} + a_{232} * A_{16} \quad (1)$$

where a_{132} and a_{232} are in fp32. Scaling by $a_{132}(a_{232})$ ensures $A_{16}(A_{216})$ is within the limited dynamic range of fp16. Conceptually, A_{16} roughly captures the 3 most significant decimal digits, and A_{216} the next 3 lower significant decimal digits. If A_{32} is already well scaled, then one can expect $a_{132} = 1$ and $a_{232} = 2^{-11} * a_{132}$. Matrix multiplication of two fp32 matrices can be approximated by 3 matrix multiplications of fp16 matrices on Tensor Cores, which may give a theoretical speedup of $8/3\times$ on Volta GPU. For another perspective, equation (1) suggests the form of operations in DNN are theoretically able to well approximate operations of fp32 using mixed-precision operations in fp16.

4 Applications

4.1 Materials Science and Microscopy

There are multiple opportunities to exploit machine learning techniques in materials science. While many applications have concentrated on materials discovery—as exemplified by the use of ML in conjunction with databases such as The Materials Project [39], AFLOW [40], or OQMD [41]—a tightly coupled integration of traditional simulation techniques has great promise for bridging the accuracy and computational cost divide between first principles calculations and effective models. One promising combination for a tight coupling is for first principles statistical mechanics of materials to calculate the temperature dependence of materials. This requires the calculations of many possible atomic configurations within these materials using a Monte Carlo approach, where the probability of the individual states would be evaluated using an expensive density functional theory calculation [42]. Thus, we will utilize ML to construct surrogate models as an intermediate step to link the first principles calculations to the Monte Carlo simulation. Here, the calculation of ordering phase transitions in alloys might serve as an example [43]. In a solid solution alloy, different chemical species can occupy the sites of an underlying lattice structure (note that the total number of states grows exponentially with the number of atoms in the system). Each of these possible configurations has a different energy that determines the probability of the system being found in this state at a given temperature. To build a model for these interactions, density functional calculations for representative configurations $O(1,000-10,000)$ will be performed to train a surrogate model that can replace the expensive (in the order of multiple node hours per data point) first principles calculation within the Monte Carlo sampling of possible configurations. While this approach can be conducted in a linear workflow, $DFT \rightarrow ML \rightarrow MC$, we envision a tighter coupled workflow, which augments the original training set with new points from important regions of the phase space discovered during the MC simulation, and retrains the model to improve the quantitative predictive power of this approach.

A long-standing inverse problem in atomic imaging is the loss of phase information during measurement (a.k.a., the phase problem). Given the sparse data collection of scanning transmission electron microscopic (STEM) images on different types of materials, a comprehensive database is needed for the community

to study the subject. State-of-the-art electron microscopes produce focused electron beams with atomic dimensions and capture of convergent beam electron diffraction (CBED) patterns. In this dataset [44], we use newly developed electron scattering simulation codes to generate CBED patterns from over 60,000 materials (solid-state materials) from a material project database, representing nearly every known crystal structure. A data sample from this data set is given by a 3D array formed by stacking 3 CBED patterns simulated from the same material at 3 distinct material projections (i.e., crystallographic orientations). Associated with each data sample in the data set is a host of material attributes or properties which are, in principle, retrievable via analysis of this CBED stack. These consists of the crystal space group to which the material belongs, atomic lattice constants and angles, and chemical composition, to name but a few. Of note is the crystal space group attribute (or label).

This dataset could be augmented with experimental data in the future. The generated dataset, based on simulations, will emphasize the scalability aspect of the model with the use of very large images ($\sim 10\text{GB}$ per image).

4.2 Super-Resolution for HPC Simulations

Image compression is a very active field of study, with new methods being constantly generated [45]. The need for improvements in image compression quality is growing in the field of HPC simulations because of the exponential trend in data generation. There exists an untapped potential in this situation due to the nature of simulated data that is not currently exploited. Simulation data from numerical systems of partial differential equations exist on a solution manifold [46]. Thus, the manifold hypothesis in machine learning—which states that real-world, high-dimensional data lie on low-dimensional manifolds embedded within the high-dimensional space—is concrete for simulation data. We can therefore expect that identifying this map to the low-dimensional manifold will provide ideal compression for HPC Simulations. In the next paragraph we describe the basic setup that allows researchers to test in situ machine learning for climate simulations.

The shallow water equations on a sphere is a well-known model used in climate science to simulate basic dynamics of the Earth’s atmosphere. There exist many test case scenarios for this model [47, 48, 49, 50], and in this paper we use the challenging test known as the barotropic instability test [50], which involves perturbing the atmospheric flow with a localized bump to the balanced height field. We simulate the shallow water equations on the sphere by using a fourth-order Runge-Kutta method [51] for time-stepping and a discontinuous Galerkin (DG) spatial discretization with a Legendre basis on a cubed-sphere mesh [52, 53].

The data structure for this simulation consists of a time series of images for each of the six faces of the cube, which for analysis can be converted to a single series lat-lon gridded images. Arguably, the most commonly used form of lossy compression is the discrete cosine transform (DCT) method [54]. ECP has two projects on the development of FFTs and related transformations, one of which

is heFFTe [55], already integrated in the ECP software stack, delivering high-performance (close to 90% of the theoretical roofline peak), and very good weak and strong scaling on systems like Summit at ORNL. This common compression method fills out the setup of the in situ machine learning super resolution method. At each time step of the barotropic instability simulation, in situ machine learning methods are exposed to both the compressed image data and the original image. As the simulation progresses, only the compressed image data is stored to disk, and the machine learning method adaptively learns correct super-resolution transformation of the compressed data. Final stored data contain all compressed image data with the trained machine learning method. It has been demonstrated in [56] that, using this setup, it is possible to train in situ networks to reduce the error in lossy compression—obtaining multiple orders of improvement. Going forward, it will be necessary to further improve in situ compression and analysis in order to maximize discovery with HPC simulations.

5 Meeting Exascale

Exascale systems will appear in 2021. These machines incorporate accelerators from NVIDIA, AMD, and Intel. Efforts in preparing the software stack for exascale systems have been a major focus in the DOE ECP program [57]. We believe a native DNN framework such as MagmaDNN will pave a unique path to meet the challenges of exascale platforms. Assisted by the openDIEL parallel workflow engine, which admits a diverse spectrum of applications as functional modules, we will be able to exploit the full capabilities of exascale machines. In-situ data augmentation will be incorporated with compute-intensive simulations, leading to discovery in multi- and inter-disciplinary, systems-wide, real-time recommendation systems. From instrumentation calibration, experimental and observable results, theoretical simulations, to validation and analysis—exascale computing will be brought to bear on health, transportation, environment, and social sciences. For example, climate and weather recommendation systems will be able to integrate models in storms prediction, rainfall, vegetation, farm growth, pollution, power usage, traffic flow, etc. [58, 59, 60, 61, 62]. As inputs and outputs from many sensor devices become ubiquitous, the importance of a scalable AI framework and an extensible workflow tool will grow.

Challenges will rise from algorithmic approaches as HPC systems continue to expand and evolve. Multiple-precision implementation will be unavoidable. Setting the basis and preparing for product diversity from different vendors will be an important consideration of an AI framework, as well. One challenge in performance is to reduce the impact of communication while maintaining good convergence of the algorithm, such as SGD, in the DNN framework. In the Horovod framework [26], for example, the solution is to distribute the training of neural network on large machines using an efficient gradient reduction strategy referred to as ring-allreduce. Notably, in the 2018 Gordon-Bell award winning paper [63], Horovod was used in the context of TensorFlow to train two DNN models on large systems, including the Summit machine, to detect extreme weather patterns from climate data. Similarly, Horovod was used in [64] to train a

large DNN model for tackling inverse problems in imaging. Algorithmic improvements towards individual or combined synchronous, asynchronous, and pipeline approaches are essential to improve resource usage as well as convergence.

There is a trend in increasingly large model sizes (especially in NLP—the latest model [65] has 17 Billion parameters), and a need in scientific applications (e.g., geospatial imaging) to process larger-dimension inputs. Although there exist exascale deep learning applications [63, 66] on pre-exascale system such as Summit thanks to the Tensor Core technology, those use cases push the limit on large-batch training for data parallelism and are not generally applicable to exascale learning challenges. Early efforts [9, 67] on model parallelism have made progress, but are not yet mature or generic. An AI framework that can efficiently exploit various level of parallelisms (data parallel, model parallel, hyperparameter search, etc) will be in demand.

6 Conclusion

Exascale systems are an important investment in the US. With exascale, we envision a stronger economy and improved quality of life. It will also lead to important scientific discovery and resolution of complex issues related to national security. Development in AI software and tools that scale well on these systems is important, and even more critical for AI frameworks that also work well across the existing spectrum of exascale applications. Although many challenges exist, a primary roadblock is the lack of direct collaborative effort, and a software platform that values performance as the foremost priority. In this paper, we present a unique set of AI tools and algorithms, as well as efforts between collegiate and ORNL researchers, demonstrating that there is a pathway to integrate and deploy machine learning to those ends—with emphasis on two major applications on exascale systems.

7 Acknowledgments

This material is based upon work supported in part by the Laboratory Directed Research and Development program at the Oak Ridge National Laboratory, which is operated by UT-Battelle, LLC., for the U.S. Department of Energy under Contract DE-AC05-00OR22725. The work was partly supported by the Scientific Discovery through Advanced Computing (SciDAC) program funded by U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, with specific thanks to the FASTMath Institutes.

This work was conducted at the Joint Institute for Computational Sciences (JICS) and the Innovative Computing Laboratory (ICL), sponsored by the National Science Foundation (NSF), through NSF REU Award #1659502 and NSF Award #1709069. This work used hardware donations from NVIDIA as well as the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF grant number ACI-1548562. Computational Resources are available through a XSEDE education allocation award TG-ASC170031.

Bibliography

- [1] T. Ise and Y. Oba, “Forecasting Climatic Trends Using Neural Networks: An Experimental Study Using Global Historical Data,” *Frontiers in Robotics and AI*, vol. 6, p. 32, 2019.
- [2] J. Wang, P. Balaprakash, and R. Kotamarthi, “Fast Domain-Aware Neural Network Emulation of a Planetary Boundary Layer Parameterization in a Numerical Weather Forecast Model,” *Geoscientific Model Development*, vol. 12, no. 10, pp. 4261–4274, 2019.
- [3] A. Agrawal and A. Choudhary, “Deep materials Informatics: Applications of Deep Learning in Materials Science,” *MRS Communications*, vol. 9, no. 3, p. 779–792, 2019.
- [4] S. Feng, H. Zhou, and H. Dong, “Using Deep Neural Network with Small Dataset to Predict Material Defects,” *Materials & Design*, vol. 162, p. 300–310, 2019. Citation Key: FENG2019300.
- [5] W. Ye, C. Chen, Z. Wang, I.-H. Chu, and S. P. Ong, “Deep Neural Networks for Accurate Predictions of Crystal Stability,” *Nature Communications*, vol. 9, p. 3800, Sep 2018.
- [6] T. Ben-Nun and T. Hoefler, “Demystifying Parallel and Distributed Deep Learning: an Indepth Concurrency Analysis,” *ACM Computing Survey*, vol. 1, pp. 1–37, January 2019.
- [7] J. Han, L. Xu, M. M. Rafique, A. R. Butt, and S. Lim, “A quantitative study of deep learning training on heterogeneous supercomputers,” in *2019 IEEE International Conference on Cluster Computing (CLUSTER)*, pp. 1–12, 2019.
- [8] Y. You, Z. Zhang, C.-J. Hsieh, J. Demmel, and K. Keutzer, “ImageNet Training in Minutes,” in *Proceedings of the 47th International Conference on Parallel Processing, ICPP 2018*, (New York, NY, USA), Association for Computing Machinery, 2018.
- [9] N. Shazeer, Y. Cheng, N. Parmar, D. Tran, A. Vaswani, P. Koanantakool, P. Hawkins, H. Lee, M. Hong, C. Young, R. Sepassi, and B. Hechtman, “Mesh-TensorFlow: Deep Learning for Supercomputers,” in *Advances in Neural Information Processing Systems 31* (S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds.), pp. 10414–10423, Curran Associates, Inc., 2018.
- [10] J. Geng, D. Li, and S. Wang, “ElasticPipe: An Efficient and Dynamic Model-Parallel Solution to DNN Training,” in *Proceedings of the 10th Workshop on Scientific Cloud Computing, ScienceCloud ’19*, (New York, NY, USA), p. 5–9, Association for Computing Machinery, 2019.
- [11] Y. e. a. Huang, “GPipe: Efficient Training of Giant Neural Networks using Pipeline Parallelism,” in *Advances in neural information processing systems 32*, p. 103–112, Curran Associates, Inc., 2019. Citation Key: NIPS2019_8305.

- [12] A. Gholami, A. Azad, P. Jin, K. Keutzer, and A. Buluc, “Integrated Model, Batch, and Domain Parallelism in Training Neural Networks,” in *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures*, SPAA ’18, (New York, NY, USA), p. 77–86, Association for Computing Machinery, 2018.
- [13] C.-C. Chen, C.-L. Yang, and H.-Y. Cheng, “Efficient and Robust Parallel DNN Training through Model Parallelism on Multi-GPU Platform,” *ArXiv*, vol. abs/1809.02839, 2018.
- [14] R. Stevens, “Exascale Computing : The Coming Integration of Simulation, Data and Machine Learning,” 2018. European Technology Platform for High-Performance Computing (ETP4HPC) ISC’18 Workshop.
- [15] M. A. et. al., “Tensorflow: Large-Scale Machine Learning on Heterogeneous Distributed Systems,” *CoRR*, vol. abs/1603.04467, 2016.
- [16] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshein, L. Antiga, and et al., “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” in *Advances in neural information processing systems 32*, p. 8026–8037, Curran Associates, Inc., 2019. Citation Key: NIPS2019_9015.
- [17] T. Chen, M. Li, Y. Li, M. Lin, N. Wang, M. Wang, T. Xiao, B. Xu, C. Zhang, and Z. Zhang, “MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems,” *CoRR*, vol. abs/1512.01274, 2015.
- [18] D. Nichols, K. Wong, S. Tomov, L. Ng, S. Chen, and A. Gessinger, “MagmaDNN: Accelerated Deep Learning Using MAGMA,” in *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning)*, PEARC ’19, (New York, NY, USA), Association for Computing Machinery, 2019.
- [19] D. Nichols, N.-S. Tomov, F. Betancourt, S. Tomov, K. Wong, and J. Dongarra, “MagmaDNN: Towards High-Performance Data Analytics and Machine Learning for Data-Driven Scientific Computing,” in *ISC High Performance*, (Frankfurt, Germany), Springer International Publishing, Springer International Publishing, 2019-06 2019.
- [20] D. Nichols, R. Febbo, F. Lopez, K. Wong, S. Tomov, and J. Dongarra, “MagmaDNN (Version 1.2),” July 2020. 10.5281/zenodo.3972406.
- [21] S. Tomov, J. Dongarra, and M. Baboulin, “Towards Dense Linear Algebra for Hybrid GPU Accelerated Manycore Systems,” *Parallel Computing*, vol. 36, pp. 232–240, June 2010.
- [22] J. Ozik, N. Collier, J. Wozniak, and C. Spagnuolo, “From Desktop to Large-Scale Model Exploration with Swift/T,” in *Proceedings of the 2016 Winter Simulation Conference*, 2016.
- [23] K. Wong and Z. Trzil, “Tuple Space Implementation in A Parallel Workflow Engine,OpenDIEL,” in *Student Paper,PEARC ’18*, 2018.
- [24] F. Betancourt, K. Wong, E. Asemota, Q. Marshall, D. Nichols, and S. Tomov, “openDIEL: A Parallel Workflow Engine and Data Analytics Framework,” in *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines Learning*, Student Paper,PEARC ’19, (New York, NY, USA), Association for Computing Machinery, 2019.

- [25] F. Niu, B. Recht, C. Re, and S. J. Wright, “HOGWILD! A Lock-Free Approach to Parallelizing Stochastic Gradient Descent,” in *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, (Red Hook, NY, USA), p. 693–701, Curran Associates Inc., 2011.
- [26] A. Sergeev and M. D. Balso, “Horovod: Fast and Easy Distributed Deep Learning in TensorFlow,” *arXiv preprint arXiv:1802.05799*, 2018.
- [27] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, M. aurelio Ranzato, A. Senior, P. Tucker, K. Yang, Q. V. Le, and A. Y. Ng, “Large Scale Distributed Deep Networks,” in *Advances in Neural Information Processing Systems 25* (F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, eds.), pp. 1223–1231, Curran Associates, Inc., 2012.
- [28] S. Sallinen, N. Satish, M. Smelyanskiy, S. S. Sury, and C. Ré, “High performance parallel stochastic gradient descent in shared memory,” in *IPDPS*, pp. 873–882, IEEE Computer Society, 2016.
- [29] C. D. Sa, C. Zhang, K. Olukotun, and C. Ré, “Taming the Wild: A Unified Analysis of HOG WILD! -Style Algorithms,” in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS’15*, (Cambridge, MA, USA), p. 2674–2682, MIT Press, 2015.
- [30] X. Lian, Y. Huang, Y. Li, and J. Liu, “Asynchronous Parallel Stochastic Gradient for Nonconvex Optimization,” in *Advances in Neural Information Processing Systems 28* (C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, eds.), pp. 2737–2745, Curran Associates, Inc., 2015.
- [31] F. Lopez, E. Chow, S. Tomov, and J. Dongarra, “Asynchronous SGD for DNN Training on Shared-Memory Parallel Architectures,” Tech. Rep. ICL-UT-20-04 (To appear in IPDPSW’20 proceedings), Innovative Computing Laboratory, University of Tennessee, 2020-03 2020.
- [32] J. Wolfson-Pou and E. Chow, “Modeling the Asynchronous Jacobi Method without Communication Delays,” *Journal of Parallel and Distributed Computing*, vol. 128, 6 2019.
- [33] I. Yamazaki, E. Chow, A. Bouteiller, and J. Dongarra, “Performance of Asynchronous Optimized Schwarz with One-Sided Communication,” *Parallel Computing*, vol. 86, pp. 66 – 81, 2019.
- [34] M. Courbariaux, Y. Bengio, and J.-P. David, “Training Deep neural Networks with Low Precision Multiplications,” 2014.
- [35] Intel Corporation, “BFLOAT16—Hardware Numerics Definition,” Nov. 2018. White paper. Document number 338302-001US.
- [36] P. Blanchard, N. J. Higham, F. Lopez, T. Mary, and S. Pranesh, “Mixed Precision Block Fused Multiply-Add: Error Analysis and Application to GPU Tensor Cores,” *SIAM Journal on Scientific Computing*, vol. 42, no. 3, pp. C124–C141, 2020.
- [37] A. Haidar, S. Tomov, J. Dongarra, and N. J. Higham, “Harnessing GPU tensor Cores for Fast FP16 Arithmetic to Speed up Mixed-Precision Iterative Refinement Solvers,” in *Proc. of the Intl. Conf. for High Performance Computing, Networking, Storage, and Analysis*, pp. 1–11, 11 2018.
- [38] A. Sorna, X. Cheng, E. D’Azevedo, K. Wong, and S. Tomov, “Optimizing the Fast Fourier Transform Using Mixed Precision on Tensor Core Hard-

- ware,” in *2018 IEEE 25th International Conference on High Performance Computing Workshops (HiPCW)*, pp. 3–7, 2018.
- [39] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, “Commentary: The materials project: A Materials Genome Approach to Accelerating Materials Innovation,” *APL Materials*, vol. 1, no. 1, p. 011002, 2013.
- [40] E. Gossett, C. Toher, C. Oses, O. Isayev, F. Legrain, F. Rose, E. Zurek, J. Carrete, N. Mingo, A. Tropsha, and S. Curtarolo, “Aflow-ml: A restful api for Machine-Learning Predictions of Materials Properties,” *Computational Materials Science*, vol. 152, pp. 134 – 145, 2018.
- [41] S. Kirklin, J. E. Saal, B. Meredig, A. Thompson, J. W. Doak, M. Aykol, S. Rühl, and C. Wolverton, “The Open Quantum Materials Database (OQMD): assessing the accuracy of DFT formation energies,” *npj Comput Mater*, vol. 1, p. 15010, 2015.
- [42] M. Eisenbach, C.-G. Zhou, D. M. Nicholson, G. Brown, J. Larkin, and T. C. Schulthess, “A Scalable Method for Ab Initio Computation of Free Energies in Nanoscale Systems,” in *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, SC '09*, (New York, NY, USA), pp. 64:1–64:8, ACM, 2009.
- [43] M. Eisenbach, Z. Pei, and X. Liu, “First-Principles Study of Order-Disorder Transitions in Multicomponent Solid-Solution Alloys,” *Journal of Physics: Condensed Matter*, vol. 31, p. 273002, apr 2019.
- [44] N. Laanait, A. Borisevich, and J. Yin, “A Database of Convergent Beam Electron Diffraction Patterns for Machine Learning of the Structural Properties of Materials,” *10.13139/OLCF/1510313*, 5 2019.
- [45] K. Sayood, *Introduction to Data Compression*. The Morgan Kaufmann Series in Multimedia Information and Systems, Elsevier Science, 2017.
- [46] W. C. Rheinboldt, “On the Computation of Multi-Dimensional Solution Manifolds of Parametrized Equations,” *Numerische Mathematik*, vol. 53, no. 1, pp. 165–181, 1988.
- [47] D. Williamson, J. Drake, J. Hack, R. Jakob, and P. Swarztrauber, “A Standard Test Set for Numerical Approximations to the Shallow Water Equations in Spherical Geometry,” *Journal of Computational Physics*, vol. 102, pp. 211–224, September 1992.
- [48] R. D. Nair and C. Jablonowski, “Moving Vortices on the Sphere: A Test Case for Horizontal Advection Problems,” *Monthly Weather Review*, vol. 136, no. 2, pp. 699–711, 2008.
- [49] McDonald and Bates, “Semi-lagrangian Integration of a Shallow Water Model on the Sphere,” *Monthly Weather Review*, vol. 117, p. 130, 1989.
- [50] J. Galewsky, R. Scott, and L. Polvani, “An Initial-Value Problem for Testing Numerical Models of the Global Shallow-Water Equations,” *Tellus*, vol. 56A, pp. 429–440, 2004.
- [51] M. Abramowitz and I. Stegun, eds., *Handbook of Mathematical Functions*, ch. 9. Dover Publications, 1972.
- [52] R. Sadourny, “Conservative Finite-Difference Approximations of the Primitive Equations on Quasi-Uniform spherical grids,” *Monthly Weather Review*, vol. 100, no. 2, pp. 136–144, 1972.

- [53] R. Nair, S. Thomas, and R. Loft, “A Discontinuous Galerkin Global Shallow Water Model,” *Monthly Weather Review*, vol. 133, no. 4, pp. 876–888, 2005.
- [54] N. Ahmed, T. Natarajan, and K. R. Rao, “Discrete Cosine Transform,” *IEEE Transactions on Computers*, vol. C-23, no. 1, pp. 90–93, 1974.
- [55] S. Tomov, A. Ayala, A. Haidar, and J. Dongarra, “FFT-ECP API and High-Performance Library Prototype for 2-D and 3-D FFTs on Large-Scale Heterogeneous Systems with GPUs,” ECP WBS 2.3.3.13 Milestone Report FFT-ECP STML13-27, Innovative Computing Laboratory, University of Tennessee, 2020-01 2020. revision 01-2020.
- [56] S. Lee, Q. Kang, S. Madireddy, P. Balaprakash, A. Agrawal, A. Choudhary, R. Archibald, and W. Liao, “Improving Scalability of Parallel CNN Training by Adjusting Mini-Batch Size at Run-Time,” in *2019 IEEE International Conference on Big Data (Big Data)*, pp. 830–839, 2019.
- [57] “Exascale Computing Project.” <https://www.exascaleproject.org>.
- [58] K. Zhou, Y. Zheng, B. W. D. Li, and X. Zhang, “Forecasting Different Types of Convective Weather: A Deep Learning Approach,” *Journal of Meteorological Research*, vol. 33, 2019.
- [59] S. Samsi, C. Mattioli, and V. Mark, “Distributed Deep Learning for Precipitation Nowcasting,” in *2019 IEEE High Performance Extreme Computing Conference (HPEC)*, 2019.
- [60] M. Keane and T. Neal, “Comparing Deep Neural Network and Econometric Approaches to Predicting the Impact of Climate Change on Agricultural Yield,” *UNSW Economics Working Paper*, 2020-02.
- [61] X. Xiuwen Yi, J. Zhang, Z. Wang, T. Li, and Y. Zheng, “Deep Distributed Fusion Network for Air Quality Prediction,” in *Proceedings of KDD’18, London, United Kingdom*, 2018.
- [62] K. Chen, K. Chen, Q. Wang, Z. Ziyu He, Hu, and J. He, “Short-term Load Forecasting with Deep Residual Networks,” *ArXiv*, vol. abs/1805.11956v1, 2018.
- [63] T. Kurth, S. Treichler, J. Romero, M. Mudigonda, N. Luehr, E. Phillips, A. Mahesh, M. Matheson, J. Deslippe, M. Fatica, Prabhat, and M. Houston, “Exascale Deep Learning for Climate Analytics,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, SC ’18*, IEEE Press, 2018.
- [64] N. Laanait, J. Romero, J. Yin, M. T. Young, S. Treichler, V. Starchenko, A. Borisevich, A. Sergeev, and M. Matheson, “Exascale Deep Learning for Scientific Inverse Problems,” 2019.
- [65] S. Rajbhandari, J. Rasley, O. Ruwase, and Y. He, “ZeRO: Memory Optimization Towards Training A Trillion Parameter Models,” *ArXiv*, vol. abs/1910.02054, 2019.
- [66] N. Laanait, J. Romero, J. Yin, M. T. Young, S. Treichler, V. Starchenko, A. Y. Borisevich, A. Sergeev, and M. Matheson, “Exascale Deep Learning for Scientific Inverse Problems,” *ArXiv*, vol. abs/1909.11150, 2019.
- [67] Y. Huang, Y. Cheng, D. Chen, H. Lee, J. Ngiam, Q. V. Le, and Z. Chen, “GPipe: Efficient Training of Giant Neural Networks using Pipeline Parallelism,” *ArXiv*, vol. abs/1811.06965, 2018.