

# Deep Learning at 15PF

## Supervised and Semi-Supervised Classification for Scientific Data

Thorsten Kurth  
NERSC, Lawrence Berkeley National  
Lab  
1 Cyclotron Rd, M/S 59R4010A  
Berkeley, California USA 94720

Jian Zhang  
Stanford University  
Gates Computer Science, 353 Serra  
Mall  
Stanford, California USA 94305

Nadathur Satish  
Intel Corporation  
2200 Mission College Blvd  
Santa Clara, California USA 95054

Evan Racah  
NERSC, Lawrence Berkeley National  
Lab  
1 Cyclotron Rd, M/S 59R4010A  
Berkeley, California USA 94720

Ioannis Mitliagkas  
Stanford University  
Gates Computer Science, 353 Serra  
Mall  
Stanford, California USA 94305

Md. Mostofa Ali Patwary  
Intel Corporation  
2200 Mission College Blvd  
Santa Clara, California USA 95054

Tareq Malas  
Intel Corporation  
2111 NE 25th Avenue, JF-1-06  
Hillsboro, Oregon USA 97124

Narayanan Sundaram  
Intel Corporation  
2200 Mission College Blvd  
Santa Clara, California USA 95054

Wahid Bhimji  
NERSC, Lawrence Berkeley National  
Lab  
1 Cyclotron Rd, M/S 59R4010A  
Berkeley, California USA 94720

Mikhail Smorkalov  
Intel Corporation  
Turgeneva Str. 30  
Nizhny Novgorod, Russian Federation  
603024

Jack Deslippe  
NERSC, Lawrence Berkeley National  
Lab  
1 Cyclotron Rd, M/S 59R4010A  
Berkeley, California USA 94720

Mikhail Shiryaev  
Intel Corporation  
Turgeneva Str. 30  
Nizhny Novgorod, Russian Federation  
603024

Srinivas Sridharan  
Intel Corporation  
136 Airport Road  
Bangalore, Karnataka, India 560007

Prabhat  
NERSC, Lawrence Berkeley National  
Lab  
1 Cyclotron Rd, M/S 59R4010A  
Berkeley, California USA 94720

Pradeep Dubey  
Intel Corporation  
2200 Mission College Blvd  
Santa Clara, California USA 95054

### ABSTRACT

This paper presents the first, 15-PetaFLOP Deep Learning system for solving scientific pattern classification problems on contemporary HPC architectures. We develop supervised convolutional architectures for discriminating signals in high-energy physics data as well as semi-supervised architectures for localizing and classifying extreme weather in climate data. Our Intelcaffe-based implementation obtains  $\sim 2$ TFLOP/s on a single Cori Phase-II Xeon-Phi node. We use a hybrid strategy employing synchronous node-groups, while using asynchronous communication across groups. We use this strategy to scale training of a single model to  $\sim 9600$  Xeon-Phi nodes; obtaining peak performance of 11.73-15.07 PFLOP/s and sustained

performance of 11.41-13.27 PFLOP/s. At scale, our HEP architecture produces state-of-the-art classification accuracy on a dataset with 10M images, exceeding that achieved by selections on high-level physics-motivated features. Our semi-supervised architecture successfully extracts weather patterns in a 15TB climate dataset. Our results demonstrate that Deep Learning can be optimized and scaled effectively on many-core, HPC systems.

#### ACM Reference format:

Thorsten Kurth, Jian Zhang, Nadathur Satish, Evan Racah, Ioannis Mitliagkas, Md. Mostofa Ali Patwary, Tareq Malas, Narayanan Sundaram, Wahid Bhimji, Mikhail Smorkalov, Jack Deslippe, Mikhail Shiryaev, Srinivas Sridharan, Prabhat, and Pradeep Dubey. 2017. Deep Learning at 15PF. In *Proceedings of SC17, Denver, CO, USA, November 12–17, 2017*, 11 pages. DOI: 10.1145/3126908.3126916

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

SC17, Denver, CO, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. 978-1-4503-5114-0/17/11...\$15.00  
DOI: 10.1145/3126908.3126916

### 1 DEEP LEARNING FOR SCIENCE

In recent years, Deep Learning (DL) has enabled fundamental breakthroughs in computer vision, speech recognition and control system problems, thereby enabling a number of novel commercial applications. At their core, these applications solve classification and

regression problems, tasks which are shared by numerous scientific domains. For example, problems in identifying galaxies, screening medical images, predicting cosmological constants, material properties and protein structure prediction all involve learning a complex hierarchy of features, and predicting a class label, or regressing a numerical quantity. We assert that that Deep Learning is poised to have a major impact on domain sciences, but there are unique challenges that need to be overcome first.

The primary challenge is in analyzing massive quantities of complex, multi-variate scientific data. Current Deep Learning implementations can take days to converge on O(10) GB datasets; contemporary scientific datasets are TBs-PBs in size. Scientific datasets often contain dozens of channels/variables, which is in contrast to the small number of channels in images or audio data. Scientists need to be able to leverage parallel computational resources to get reasonable turnaround times for training Deep Neural Networks (DNNs). It is therefore imperative that DL software delivers good performance not only on a single node but is also scalable across a large number of nodes. We now elaborate on two scientific drivers that motivate our optimization and scaling efforts.

### 1.1 Supervised Learning for HEP

A major aim of experimental high-energy physics (HEP) is to find rare signals of new particles produced at accelerators such as the Large Hadron Collider (LHC) at CERN, where protons are accelerated to high-energies and collided together to produce resulting particles within highly-instrumented detectors, such as the ATLAS and CMS experiments. Improvements in classifying these collisions could aid discoveries that would overturn our understanding of the universe at the most fundamental level. Neural Networks have been used in HEP for some time [1, 2]. Recently attention has focused on deep learning to tackle the increase in detector resolutions and data rates. Particles produced by LHC collisions (occurring every 25ns) propagate, decay and deposit energy in different detector parts, so creating signals in 100s of millions of channels, with each collision forming an independent ‘event’. Data from the surface of the cylindrical detector can be represented as a sparse 2D image, with data from different layers of instrumentation as channels in that image. We use the energy deposited in the “electromagnetic”, and “hadronic calorimeters”, and the number of “tracks” formed from the “inner detector” in that region as three channels. This is similar to the approach of [3][4] except that we use large images covering the entire detector, and use these directly for classifying entire events rather than individual objects.

The HEP community have simulations of the underlying physics processes and the detector response that can be used for training networks. For this paper, we generate events to match those used for a particular analysis searching for new massive supersymmetric particles in multi-jet final states at the LHC [5]. We use the Pythia event generator [6] interfaced to the Delphes fast detector simulation [7] (with fast jet [8]) to generate events for two classes, corresponding to the new-physics ‘signal’ (6.4M events) and the most prevalent known-physics ‘background’ (64M events). Before training our network we apply some of the physics selections of [5] to filter images to those more challenging to discriminate, resulting

	pixels	channels	#images	Volume
HEP	228x228	3	10M	7.4TB
Climate	768x768	16	0.4M	15TB

**Table 1: Characteristics of datasets used.**

in a training sample of around 10M events. We compare the performance of our deep network to our own implementation of the selections of [5] as a baseline benchmark. We have verified that the samples and baseline selections give performance comparable to that in [5] providing a meaningful benchmark even though those selections were not tuned for these datasets.

### 1.2 Semi-Supervised Learning for Climate

Climate change is one of the most important challenges facing humanity in the 21st century; climate simulations provide a unique approach for understanding the future impact of various carbon emission scenarios and intervention strategies. Modern Climate simulation codes produce massive datasets: a single 30-year run from the CAM5 25-km resolution model produces 100TBs of multi-variate data[9]. In this paper, we are interested in the task of finding extreme weather events in such large datasets. Providing an objective, quantitative tool for finding extreme weather patterns will help climate scientists in understanding trends in such weather patterns in the future (i.e. Do we expect more Category 4/5 hurricanes to make landfall in the 21st century?), and conduct detection and attribution studies (i.e. Is the chance in Tropical Cyclone activity attributable to anthropogenic emissions, as opposed to being an intrinsic property of the climate system?).

The field of climate science typically relies on heuristics, and expert-specified multi-variate threshold conditions for specifying extremes [10–12]. We formulate this task as that of pattern classification, and employ Deep Learning based methods. The problem can be formulated as that of object recognition in images, the difference being that climate images have 16 or more ‘channels’, and their underlying statistics are quite different from natural images. Consequently, we cannot leverage pre-trained weights from contemporary networks such as VGG or AlexNet. Earlier work conducted by [13] demonstrates that convolutional architectures can solve the pattern classification task for cropped, centered image patches. In this work we develop a *unified, semi-supervised* architecture for handling all extreme weather patterns and develop a methodology for predicting bounding boxes. Most importantly, our method provides an opportunity to discover new weather patterns that might have few/no labeled examples.

This paper makes the following contributions:

- We develop Deep Learning models which not only solve the problem at hand to desired precision but are also scalable to a large number of nodes. This includes for example to not use layers with large dense weights such as batch normalization or fully connected units.
- We develop highly optimized Deep Learning software that can process complex scientific datasets on the Intel Xeon Phi architecture

- We build a system based on a hybrid asynchronous approach to scale Deep Learning to the full scale of the Cori supercomputer (~9600 Xeon Phi nodes)
- We demonstrate supervised classification on a 7.4 TB High-Energy Physics dataset
- We develop a novel, semi-supervised architecture, and apply it to detect and learn new patterns on a 15 TB climate dataset
- We obtain a peak performance of 11.73-15.07 PFLOP/s and sustained performance of 11.41-13.27 PFLOP/s for our two problems

While our exploration is conducted in the context of two concrete applications, we believe that our approach, and the resulting lessons learned, can be generalized to a much broader class of data analytics problems in science.

## 2 CURRENT STATE OF THE ART

From a HPC perspective, we can look at deep learning from two dimensions: first, how efficiently can deep learning be mapped to a single compute node; and second, how it scales across a cluster of compute nodes.

### 2.1 Deep Learning on single node

The core computation in deep learning algorithms is dominated by dense linear algebra in the form of matrix multiply and convolution operations. While well-optimized libraries such as implementations of BLAS and LaPACK have long existed for use in HPC applications, the shapes and sizes of the operands differ significantly for deep learning. Hence specific libraries with support for tall-skinny matrix multiplies and convolutions with multiple small filters have been developed for various architectures such as NVIDIA GPUs [14] and CPU architectures [15, 16].

The hardware efficiency of these kernels heavily depends on input data sizes and model parameters (weight matrix dimensions, number of convolutions, convolution strides, padding, etc). DeepBench [17] is a recently developed benchmark from Baidu that captures best known performance of deep learning kernels with varied input sizes and model parameters on NVIDIA GPUs and Intel® Xeon Phi™<sup>1</sup>. Their results show that while performance can be as high as 75-80% of peak flops for some kernels, decreasing minibatch size (dimension 'N' for matrix multiply and convolutions) results in significant efficiency drops to as low as 20-30% (at minibatch sizes of 4-16) on all architectures. As we shall see, this has implications on performance at scale.

### 2.2 Deep Learning on multiple nodes

There have been many attempts to scale deep learning models across a cluster of nodes [18–22]. In this work, we focus on scaling the training of a *single model* across a cluster as opposed to the embarrassingly parallel problem of training independent models [23]. We discuss two common architectures, shown in Figure 1.

**2.2.1 Synchronous-parallel architectures.** Synchronous systems use synchronization barriers and force computational nodes to




Figure 1: Example architectures.

perform every update step in lock-step (See Figure 1). Typically, data parallelism is used where different nodes split a big mini-batch of samples, each processing a chunk of the data. Recent papers that have attempted to scale synchronous deep learning have stopped at a few hundred nodes [20, 21, 24], with the scalability depending on the computation to communication ratio, the speed of the hardware and the quality of the interconnect. Aside from communication there are other factors that limit synchronous scaling:

**Batch size.** Most systems use some variant of SGD with batch sizes that range from 64 to 1024. Large batch sizes have been shown to cause slowdown in convergence [25], and degrade the generalization properties of the trained model [26]. The batch size is a limit on the number of nodes in data-parallel synchronous systems.

**Stragglers.** Since a synchronization barrier is used, the duration of the iteration depends on the slowest node. Variability in the computation needed per sample, OS jitter and, importantly, variations in the throughput and latency in the interconnect leads to significant load imbalance. This effect gets worse with scale.

**2.2.2 Asynchronous and hybrid architectures.** Conceptually, asynchronous architectures [27, 28] remove the synchronization barriers. Each node works on its own iteration (mini-batch) and produces independent updates to the model. Those updates are sent to a central parameter store, the *parameter server* (PS), illustrated in Figure 1. The PS applies the updates to the model in the order they are received, and sends back the updated model to the worker where the update originated. Asynchronous systems do not suffer from straggler effects and are not limited by the total batch size in the same way that synchronous systems are, an important property at scale. Asynchronous methods are known to give significant computational benefits in large-scale systems [29, 30]. Recent work [31] sheds new light on the convergence properties of such systems and shows the importance of momentum tuning for convergence.

**Performance tradeoff.** The main side-effect of asynchrony is the use of out-of-date gradients: each update is computed based on an older version of the model and then sent to the PS to be applied on the latest model. The number of updates that other workers perform between the time a worker reads the model and the time it sends its own update to the PS is called *staleness*. Asynchronous systems may need more iterations to solution, due to staleness: we say they have worse *statistical efficiency* [25, 32]. Synchronous systems typically take longer per iteration due to the straggler effect: they have worse *hardware efficiency*.

<sup>1</sup>Intel, Xeon and Intel Xeon Phi are trademarks of Intel Corporation in the U.S. and/or other countries.

*Hybrid architectures.* The trade-off between statistical efficiency vs. hardware efficiency suggests a third kind of architecture: a *hybrid* system [25]. In this architecture, worker nodes coalesce into separate *compute groups*. Each compute group follows a synchronous architecture: the workers split a mini-batch among themselves and produce a single update to the model. There is no synchronization across compute groups. A parameter server (PS) holds the model and each compute group communicates its updates to the PS asynchronously. Given a cluster of fixed size, the number of compute groups (and their size) is a knob that controls the amount of asynchrony in the system. We can tune the amount of asynchrony along with the other hyper-parameters to find the optimal configuration. We use this hybrid architecture in our paper, as described in Section 3.5.

### 3 INNOVATIONS

#### 3.1 HEP architecture

We formulate the HEP problem as a binary image classification task. We use a Convolutional Neural Net comprised of 5 convolution+pooling units with rectified linear unit (ReLU) activation functions [33, 34]. The kernel sizes used in the convolutional layers are 3x3 pixels with strides 1x1 and 128 filters per layer. In the pooling layers we use 2x2 kernels with strides 2x2. We use max pooling in the first four layers and use global average pooling in the last convolutional layer. The output of the global pooling layer is fed into a single fully connected layer which projects the resulting 128-dimensional vector into a two-dimensional vector on which a softmax function is applied to determine the class probabilities for signal and background. We use softmax with cross-entropy as the loss function. We further employ the ADAM optimizer[35] as the solver. ADAM requires less parameter tuning than Stochastic Gradient Descent and suppresses high norm variability between gradients of different layers by adaptively adjusting the learning rate.

#### 3.2 Climate architecture

We formulate the climate problem as semi-supervised bounding box regression adapted from [36], which is inspired by [37–39]. Essentially, we have a fully supervised convolutional network for bounding box regression and an unsupervised convolutional autoencoder. These two networks share various layers, so the extra unlabelled data input to the autoencoder can help improve the bounding box regression task. We use a series of strided convolutions to learn coarse, downsampled features of the input climate simulations. We call this series of convolutions the encoder of the network. At every location in the features, we compute 4 scores (confidence, class, x and y position of bottom left corner of box, and height and width of box) using a convolution layer for each score. At inference time we keep only the boxes corresponding to confidences greater than 0.8. For the unsupervised part of our architecture, we use the same encoder layers, but use the coarse features as input to a series of deconvolutional layers, which we call the decoder. The decoder attempts to reconstruct the input climate image from the coarse features. The objective function attempts to simultaneously minimize the confidence of areas without a box, maximize those with a box, maximize the the probability of the




Figure 2: Hybrid architecture example.

correct class for areas with a box, minimize the scale and location offset of the predicted box to the real box and minimize the reconstruction error of the autoencoder. As a solver, we use stochastic gradient descent with momentum.

#### 3.3 Single-node performance on manycore architectures

In this work, we used the Intel distribution of Caffe [40] to train our models. This distribution links in the Intel MKL 2017 library [15] with optimized deep learning primitives for Intel Xeon Phi. For our semi-supervised climate network, we needed optimized implementations of deconvolution that were not available. We used the fact that the convolutions in the backward pass can be used to compute the deconvolutions of the forward pass and vice-versa in order to develop optimized deconvolution implementations. These layers perform very similarly to the corresponding convolution layers.

#### 3.4 Multi-node scaling with synchronous approach

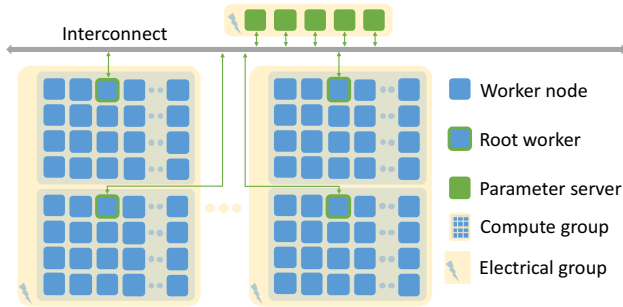
We utilize the new Intel<sup>®</sup> Machine Learning Scalability Library (MLSL) [41] for our multi-node implementation. This handles all communication required to perform training in a synchronous setting, and enables different forms of parallelism - both data and model parallelism - to be applied to different layers of the network without the user/developer worrying about communication details. In this work, we deal with either fully convolutional networks or those with very small fully connected layers, so we only use data parallelism which is well suited for such layers. MLSL also introduces performance improvements over vanilla MPI implementations using endpoints - proxy threads/processes which drive communication on behalf of the MPI rank and enable better utilization of network bandwidth. Results with this library have not been reported at large scales of more than a few hundred nodes; in this work we attempt to scale this out to thousands of nodes.

#### 3.5 Multi-node scaling with hybrid approach

In Section 2.2.2 we outlined the limitations of fully synchronous systems that motivate asynchronous architectures. Asynchronous systems are not limited by the total batch size in the same way that synchronous systems are. Furthermore, asynchrony provides an added layer of resilience to node failures and the straggler effect. In this section we describe the hybrid architecture we use in our system and discuss some of its novel elements.

Architecture	Input	Layer details	Output	Parameters size
Supervised HEP	224x224x3	5xconv-pool,1xfully-connected	class probability	2.3MiB
Semi-supervised Climate	768x768x16	9xconv,5xDeconv	coordinates, class, confidence	302.1 MiB

**Table 2: Specification of DNN architectures used in this study.**



**Figure 3: Topological placement on Cori Phase II.**

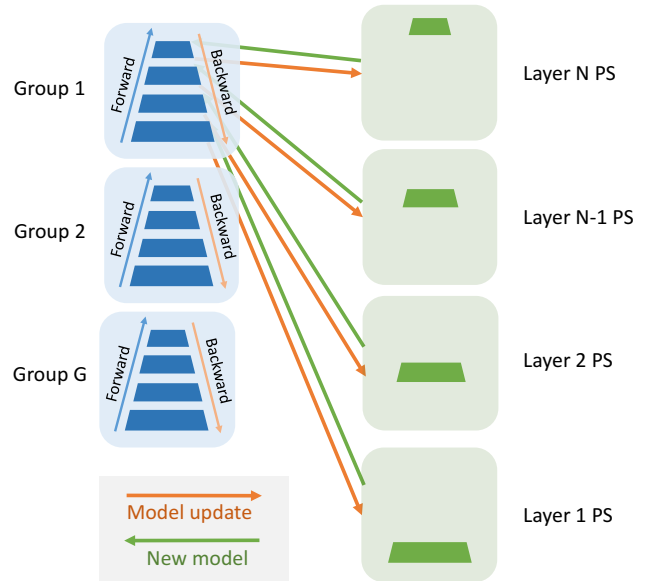
Our architecture is inspired by recently proposed hybrid approaches [25], depicted in Figure 2. Nodes are organized into compute groups. Parallelization is synchronous within (using all-reduce), but asynchronous across groups via a set of parameter servers. The number and size of compute groups, is a knob which controls the level of asynchrony, and allows us to tune asynchrony and momentum jointly, as per recent theoretical guidelines [31]. Figure 3 shows an ideal placement of nodes and compute groups on Cori.<sup>2</sup> All-reduce operations are used to get the aggregate model update from all workers in the group. Then a single node per group, called the *root node* is responsible for communicating the update to the parameter servers, receiving the new model, and broadcasting it back to the group.

*Extreme Scale.* Our work is the first instance of a hybrid architecture that scales to thousands of nodes. Previous implementations were designed (and typically deployed) on dozens or hundreds of commodity machines. For the present work, we deployed our implementation on configurations of up to 9600 nodes on an HPC system.

*Use of MSL library.* MSL does not natively support asynchronous communication. Specifically, all nodes are assumed to communicate with each other and the default library did not allow us to dedicate some subset of nodes for parameter servers. In this work, we extended MSL to enable our hybrid implementation. Specifically, we extended MSL to facilitate node placement into disjoint communication groups and dedicating nodes as parameter servers. Our new MSL primitives allow for efficient overlaying of group communication and endpoint communication with the parameter server.

*Dedicated parameter servers for each layer.* The parameter server needs to be able to handle the volume of network traffic and computation for the updates originating from multiple compute groups

<sup>2</sup>For simplicity PSs are shown in their own electrical group, however this is not typically the case.



**Figure 4: We assign a dedicated parameter server to each trainable layer of the network. Each group exchanges data with the PS for the corresponding layer. For clarity, we only depict the communication patterns for Group 1.**

and for very large models. To reduce the chances of PS saturation, we dedicate a parameter server to each trainable layer in the network (Figure 4). We can consider each compute group as a bigger, more powerful node, that performs the usual forward and backward pass operations on the layers of the network. The backward pass generates a gradient (model update) for each layer of the network. That update is communicated to its *dedicated parameter server*, the update is performed and the model communicated back to the same compute group.

## 4 CORI PHASE II

All experiments reported in this study are conducted on the Cori Phase II system at NERSC. Cori is a Cray XC40 supercomputer comprised of 9,688 self-hosted Intel Xeon Phi™ 7250 (Knight’s Landing, KNL) compute nodes. Each KNL processor includes 68 cores running at 1.4GHz and capable of hosting 4 HyperThreads for a total of 272 threads per node.

The peak performance for single precision can be computed as: (9688 KNLs) x (68 Cores) x (1.4 GHz Clock Speed) x (64 FLOPs / Cycle) = 59 PetaFLOP/s. However, for sustained AVX work, the clock-speed drops to 1.2 GHz, yielding a sustained peak performance of: 50.6 PetaFLOP/s.

Each out-of-order superscalar core has a private 32KiB L1 cache and two 512-bit wide vector processing units (supporting the AVX-512 instruction set<sup>3</sup>). Each pair of cores (a “tile”) shares a 1MiB L2 cache and each node has 96GiB of DDR4 memory and 16GiB of on-package high bandwidth (MCDRAM) memory. The MCDRAM memory can be configured into different modes, where the most interesting being *cache* mode in which the MCDRAM acts as a 16GiB L3 cache on DRAM. Additionally, MCDRAM can be configured in *flat* mode in which the user can address the MCDRAM as a second NUMA node. The on-chip directory can be configured into a number of modes, but in this publication we only consider *quad* mode, i.e. in *quad-cache*, all cores are in a single NUMA domain with MCDRAM acting as a cache on DDR4 main memory. Furthermore, Cori features the Cray Aries low-latency, high-bandwidth interconnect utilizing the dragonfly topology.

## 5 PERFORMANCE MEASUREMENT

We count the executed FLOPs using Intel® Software Development Emulator (SDE) [42]. SDE distinguishes the precision of the FLOP operations and the actual executed FLOPs in the masked SIMD instructions of the code. We use SDE to count the executed single-precision flops in the computational kernels (i.e. the neural network layers) of a single node. Given that all the nodes execute these layers the same number of times and using the same problem size, we compute the total FLOPs by multiplying the single node FLOPs by the number of nodes. The counted FLOPs constitute the vast majority of the application’s FLOP operations. The application time is spent in an iterative training loop, where the computation performed in each training iteration is the same. However, in some iterations, a checkpointing is performed to save the current trained model to the filesystem; this imposes some overhead on runtime. We measure the wall clock time per iteration to obtain the flop rate (i.e. iteration’s measured FLOPS / iteration’s time). The peak flop rate is obtained from the fastest iteration, while the sustained flop rate is computed from the best average iteration time in a contiguous window of iterations.

In the following section, we present the results of training the HEP and climate networks on the Intel Xeon Phi nodes of the Cori supercomputer. All our experiments use 66 of the 68 cores on each node, with 2 being reserved for the OS. All our experiments deal with single precision data and model parameters.

## 6 PERFORMANCE RESULTS

### 6.1 Single node performance

Figures 5a and 5b show the flop rates and time spent in various layers for HEP and Climate networks. For a batch size of 8 images, the overall flop rate of the HEP network stands at 1.90 TFLOP/s, while that of the Climate network stands at 2.09 TFLOP/s. For both networks, most of the runtime is spent in convolutional layers, which can obtain between 3.5 TFLOP/s for layers with many channels, and around 1.25 TFLOP/s on the initial layers with very few channels. As mentioned previously in DeepBench [17], the shapes of the parameters and inputs to a layer can affect performance significantly; we observe that in our experiments.

For the HEP network, about 12.5% of the runtime is spent in the solver update routine which applies the update to the weights and adjusts hyper-parameters for the next iteration. This step spends time in operations like copying models to keep history that do not contribute to flops. The overhead of this step is insignificant (< 2%) in the climate network. For the climate network, time spent in I/O (13%) for loading the data is significant; recall that climate problem consists of high resolution, 16-channel data. In comparison, the I/O time is much lower (2%) for the HEP network, which has low resolution, 3-channel data. We have identified two bottlenecks in our current I/O configuration: first, I/O throughput from a single Xeon Phi core is relatively slow, second, the current HDF5 library is not multi-threaded. We will address these limitations in future work.

### 6.2 Multi-node scaling

We now report on scaling experiments conducted on Cori Phase II.

**6.2.1 Strong Scaling.** The strong scaling configuration (involving keeping the overall batch size per update step fixed while varying the number of nodes) is a natural use-case for deep learning. Figure 6 shows the strong scaling results for HEP and climate networks. We show 3 configurations: 1 synchronous group, 2 and 4 hybrid groups; and show scalability from 1 to 1024 nodes. We use a batch size of 2048 per update. For the synchronous configuration, all nodes split the batch of 2048 images; for hybrid configurations, each compute group independently updates the model and is assigned a complete batch. Figure 6a shows that the synchronous algorithm does not scale past 256 nodes – 1024 node performance is somewhat worse than for 256. The scalability improves moderately for 2 hybrid groups, which saturates at 280x beyond 512 nodes, and more significantly with 4 hybrid groups, with about 580x scaling at 1024 nodes. We observe similar trends for the climate network in Figure 6b - the synchronous algorithm scales only to a maximum of 320x at 512 nodes and stops scaling beyond that point. The 2 and 4 group hybrid groups continue scaling to 1024 nodes; with scalability improving from 580x (on 1024 nodes) for 2 hybrid groups to 780x for 4 hybrid groups. There are two main reasons for this: one, in hybrid algorithms, only a subset of nodes need to synchronize at each time step; this reduces communication costs and straggler effects. Second, the minibatch size per node is higher for the hybrid approaches resulting in better single node performance. Scaling for our hybrid approaches is still not linear due to the single node performance drop from reduced minibatch sizes at scale.

**6.2.2 Weak Scaling.** Figure 7a shows weak scaling for the HEP network, where we keep a constant batch size (8 per node) across all configurations (synchronous and hybrid). On scaling from 1 to 2048 nodes, we find that the performance scales sub-linearly for all configurations: about 575-750x speed-up on 1024 nodes; and about 1150-1250x speed-up on 2048 nodes for asynchronous configurations. We note that the synchronous speed-up on 2048 nodes stands at about 1500x. In contrast, the weak scaling results for the climate network in Figure 7b are near-linear (1750x for synchronous and about 1850x for hybrid configurations). Our analysis indicates significant variability in runtime across iterations for HEP at scale, leading to sublinear scaling. An average convolution layer in HEP

<sup>3</sup>This includes the subsets F, CD, ER, PF but not VL, BW, DQ, IFMA, VBMI.



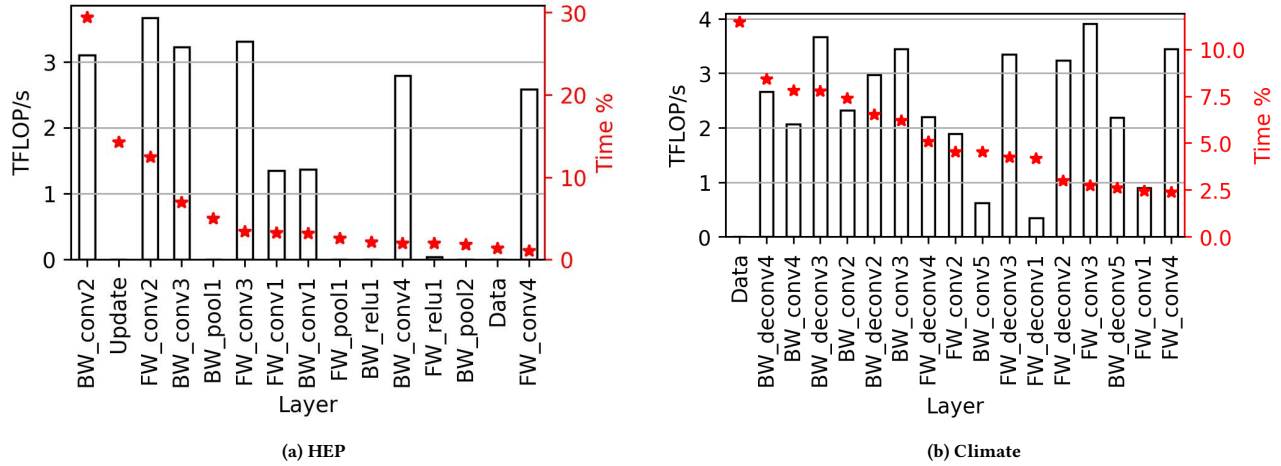


Figure 5: Single node runtime and flop rate of the top time consuming components, with batch size 8




Figure 6: Strong scaling results for synchronous and hybrid approaches (batch size = 2048 per synchronous group).

takes about 12 ms to execute; at the end of which nodes need to synchronize and reduce a small model of ~590 KB. Even a small jitter in communication times can lead to significant variability in this scenario. Hybrid approaches, where we have two additional communication steps (to and from the PS) are more affected by this variability, leading to reduced scaling. Our climate model takes on average over 300 ms per convolution layer, leading to less frequent communication and impact from jitter - we observe slightly better scaling for hybrid over synchronous configurations due to reduced straggler effects.

6.2.3 Overall Performance. For the HEP network, we obtained a **peak** throughput (as described in Section 5) of **11.73 PFLOP/s** for a configuration of 9600 total nodes (9594 compute nodes plus 6 parameter servers) split into 9 groups, with each group using a minibatch of 8528. This corresponds to a speedup of 6173x over single node performance. The **sustained** throughput as measured

over a 100 iteration timespan is **11.41 PFLOP/s**. This corresponds to an average per-iteration runtime of about 106 ms for processing a minibatch.

For the climate network, we obtained a **peak** throughput of **15.07 PFLOP/s** for a configuration of 9622 total nodes (9608 compute nodes plus 14 parameter servers) split into 8 groups, with each group using a minibatch of 9608. This corresponds to a speedup of 7205X over single node performance. The **sustained** throughput as measured over a 10 iteration span is about **13.27 PFLOP/s**, corresponding to a speedup of an average per-iteration runtime of 12.16 seconds. The sustained throughput computed includes the overhead of storing a model snapshot to disk once in 10 iterations, causing slowdowns.

6.2.4 Time to Train. Figure 8 reports the result of different training runs on the HEP network using 1024 worker nodes. We fix the




Figure 7: Weak scaling results for synchronous and hybrid approaches (batch size = 8 per node).

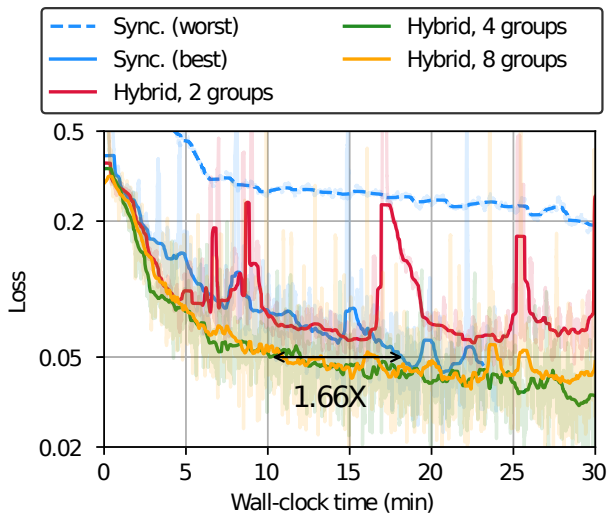


Figure 8: Training losses vs wall clock time for HEP on 1K nodes. Comparing synchronous configuration to 2, 4 and 8 groups.

total batch to 1024 and try a fully synchronous run, and three hybrid runs with 2, 4, 8 groups. We use the Adam update and tune its learning rate in the following range:  $[1e - 4, 1e - 3]$ . For the synchronous setting we fix its momentum to 0.9, but for hybrid runs we tune the momentum on a discrete set of values (0.0, 0.4, 0.7) to account for the momentum contributed by asynchrony [31]. We report the measured training loss over wall-clock for the best configurations. For the synchronous setting, we report (for the same best hyper-parameter configuration) the best and worst run out of 3. We report wall-clock time speedups with respect to a loss of 0.05 that beats the baseline for HEP (as defined in Section 1.1). We establish that the best hybrid configuration achieves the target loss in about 10 minutes, which is about 1.66x faster than the best sync run. The worst sync run is many times slower. We attribute

this, as well as some of the jumps observed in the loss curves of the 2-group case to variability in individual node performance when running on 1K nodes. Note that without additional hyperparameter tuning, we achieve a speedup of 11x in time to convergence for going from 64 to 1024 nodes, which is in line with expectations from weak scaling (cf. Figure 7a).

## 7 SCIENCE RESULTS


### 7.1 HEP Science Result

For the HEP classification problem, it is important to achieve a high signal efficiency at a very low acceptance of the much more prevalent background class. Our benchmark analysis, which is based on selections on high-level physics-derived features, achieves a true-positive rate of 42% at a false-positive rate of 0.02%. To evaluate our results we compare the true-positive rate at this same very low false-positive rate. For the hybrid configuration described in section 6.2.4, we achieve a rate of 72% which represents a **1.7x improvement** over our benchmark. For the full-system runs reported here, even with reduced runtime and without extensive tuning for accuracy, the SGD solver outperforms our benchmark by 1.3X. The capability to achieve high sensitivities to new-physics signals from classification on low-level detector quantities, without the need to design, reconstruct, or tune, high-level features offers considerable potential for enabling new-physics discoveries in future HEP analyses.

### 7.2 Climate Science Result

Figure 9 presents a sample image that illustrates the ability of our semi-supervised architecture to produce bounding boxes and class labels. In the figure, the architecture does a good job of localizing and identifying tropical cyclones. We are working on generating additional metrics for assessing the accuracy of bounding boxes for known classes (including extra-tropical cyclones and atmospheric rivers). More importantly, we are evaluating the ability of the architecture to discover *novel* weather patterns. Since this is fundamentally new approach for pattern detection in the climate





**Figure 9: Results from plotting the network’s most confident (>95%) box predictions on an image for integrated water vapor (TMQ) from the test set for the climate problem. Black bounding boxes show ground truth; Red boxes are predictions by the network.**

science community, we do not have a well-established benchmark to compare our results to.

## 8 IMPLICATIONS

### 8.1 Deep Learning on HPC

To the best of our knowledge, our work is the first successful attempt at scaling Deep Learning on large, many-core HPC systems. We share a number of insights from this unique exercise.

First, at a scale of thousands of nodes, we found significant variability in runtimes across runs, which could be as high as 30%. The probability of one of the thousands of nodes failing or degrading during the run is non-zero. In this work, we report runs where we did not encounter complete node failures. We note that even a single node failure can cause complete failure of synchronous runs; hybrid runs are much more resilient since only one of the compute groups gets affected. However, even in hybrid runs, if model updates from one of the compute groups lags significantly behind others, it can result in “jumps” in the overall loss and accuracy that we have highlighted in Figure 8.

Second, current architectures and software stacks for deep learning are still not as mature as the traditional HPC application stack. Specifically, performance on small batch sizes (essential for scale out) has not been completely optimized in many frameworks. Further, the state of the art in deep learning kernel implementations is rapidly evolving with new algorithms like Winograd [43] and FFT based algorithms. We did not experiment with such algorithms in

this work; studying the impact on per-node performance and scale out behaviour of these algorithms is a direction for future research.

There has been a lot of discussion surrounding training with quantized weights and activations [44, 45]. The statistical implications of low precision training are still being explored [46, 47], with various forms of stochastic rounding being of critical importance in convergence. While supercomputers with architectures supporting low precision computations in hardware are not yet present, we believe that such systems have the potential to further accelerate training time for our applications.

### 8.2 Deep Learning for Science

We believe that science domains that can readily generate vast amounts of representative training data (via simulators) stand to benefit immediately from progress in DL methods. In other scientific domains, unsupervised, and semi-supervised learning are key challenges for the future. In both cases, it is unreasonable to expect scientists to be conversant in the art of hyper-parameter tuning. Hybrid schemes, like the one presented in this paper, add an extra parameter to be tuned, which stresses the need for principled momentum tuning approaches, an active area of research (eg.[25] and recently [48]). With hyper-parameter tuning taken care of, higher-level libraries such as SpearMint [49] can be used for automating the search for network architectures.

We also note that more aggressive optimizations involving computing in low-precision and communicating high-order bits of weight updates are poorly understood with regards to their implications for classification and regression accuracy for *scientific* datasets. A similar story holds with regards to deployment of DL models. Unlike commercial applications where a sparse/compact representation of the model needs to be deployed in-situ, scientific applications will typically utilize DL models within the context of the HPC/Datacenter environment. Nevertheless, the field of Deep Learning is evolving rapidly, and we look forward to adopting advances in the near future.

## 9 CONCLUSIONS

This paper has presented the first 15-PetaFLOP Deep Learning software running on HPC platforms. We have utilized IntelCaffe to obtain ~2 TF on single Xeon Phi nodes. We utilize a hybrid strategy employing synchronous groups, and asynchronous communication among them to scale the training of a single model to ~9600 Cori Phase II nodes. We apply this framework to solve real-world supervised and semi-supervised patterns classification problems in HEP and Climate Science. Our work demonstrates that manycore HPC platforms can be successfully used to accelerate Deep Learning, opening the gateway for broader adoption by the domain science community. Our results are not limited to the specific applications mentioned in this paper, but they extend to other kinds of models such as ResNets [50] and LSTM [51, 52], although the optimal configuration between synchronous and asynchronous is expected to be model dependent. This highlights the importance of a flexible, hybrid architecture in achieving the best performance for a diverse set of problems.

## ACKNOWLEDGMENTS

This research used resources of the National Energy Research Scientific Computing Center (NERSC). This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes. We would like to thank Doug Jacobsen, Brandon Cook, Tina Declerck, David Paul and Rebecca Hartman-Baker for assisting with and troubleshooting Cori reservations. We would like to acknowledge Christopher Beckham, Tegan Maharaj and Christopher Pal, Yunjie Liu and Michael Wehner for help with preparing the climate architecture and dataset. We would like to acknowledge Steve Farrell for assistance with preparing HEP datasets and Ben Nachman and Brian Amadio for physics input on those datasets. Christopher Ré's group at Stanford was the source of valuable advice on asynchrony. We would like to thank Srinivas Sridharan, Mikhail Smorkalov, Mikhail Shiryaev and Dipankar Das for their help in integrating and modifying Intel MLSL.

## REFERENCES

- [1] C. Peterson, "Track finding with neural networks," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 279, no. 3, pp. 537–545, 1989.
- [2] B. Denby, "Neural networks and cellular automata in experimental high energy physics," *Computer Physics Communications*, vol. 49, no. 3, pp. 429–448, 1988.
- [3] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, and A. Schwartzman, "Jet-images – deep learning edition," *JHEP*, vol. 07, p. 069, 2016.
- [4] P. T. Komiske, E. M. Metodiev, and M. D. Schwartz, "Deep learning in color: towards automated quark/gluon jet discrimination," *JHEP*, vol. 01, p. 110, 2017.
- [5] The ATLAS collaboration, "Search for massive supersymmetric particles in multi-jet final states produced in pp collisions at  $\sqrt{s} = 13$  TeV using the ATLAS detector at the LHC," *ATLAS-CONF-2016-057*, 2016.
- [6] T. Sjöstrand, S. Mrenna, and P. Skands, "A brief introduction to PYTHIA 8.1," *Computer Physics Communications*, vol. 178, no. 11, pp. 852–867, 2008.
- [7] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi, "DELPHES 3, A modular framework for fast simulation of a generic collider experiment," *JHEP*, vol. 02, p. 057, 2014.
- [8] M. Cacciari, G. P. Salam, and G. Soyez, "Fastjet user manual," *The European Physical Journal C*, vol. 72, no. 3, p. 1896, 2012.
- [9] M. Wehner, Prabhat, K. A. Reed, D. Stone, W. D. Collins, and J. Baumeister, "Resolution dependence of future tropical cyclone projections of cam5.1 in the u.s. clivar hurricane working group idealized configurations," *Journal of Climate*, vol. 28, no. 10, pp. 3905–3925, 2015.
- [10] T. R. Knutson, J. L. McBride, J. Chan, K. Emanuel, G. Holland, C. Landsea, I. Held, J. P. Kossin, A. Srivastava, and M. Sugi, "Tropical cyclones and climate change," *Nature Geoscience*, vol. 3, no. 3, pp. 157–163, 2010.
- [11] D. A. Lavers, G. Villarini, R. P. Allan, E. F. Wood, and A. J. Wade, "The detection of atmospheric rivers in atmospheric reanalyses and their links to british winter floods and the large-scale climatic circulation," *Journal of Geophysical Research: Atmospheres*, vol. 117, no. D20, 2012.
- [12] U. Neu and et al., "Imilast: A community effort to intercompare extratropical cyclone detection and tracking algorithms," *Bulletin of the American Meteorological Society*, vol. 94, no. 4, pp. 529–547, 2013.
- [13] Y. Liu, E. Racah, Prabhat, J. Correa, A. Khosrowshahi, D. Lavers, K. Kunkel, M. Wehner, and W. D. Collins, "Application of deep convolutional neural networks for detecting extreme weather in climate datasets," *CoRR*, vol. abs/1605.01156, 2016.
- [14] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer, "cuDNN: Efficient primitives for deep learning," *CoRR*, vol. abs/1410.0759, 2014.
- [15] "Introducing DNN primitives in Intel® Math Kernel Library," <https://software.intel.com/en-us/articles/introducing-dnn-primitives-in-intelr-mkl>, 2017.
- [16] A. Heinecke, G. Henry, M. Hutchinson, and H. Pabst, "Libxsmm: Accelerating small matrix multiplications by runtime code generation," in *Proceedings of SC16*. IEEE Press, 2016, pp. 84:1–84:11.
- [17] "Deepbench," [github.com/baidu-research/DeepBench](https://github.com/baidu-research/DeepBench), 2017.
- [18] D. Amodei, R. Anubhai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, J. Chen, M. Chrzanoski, A. Coates, G. Diamos et al., "Deep speech 2 : End-to-end speech recognition in english and mandarin," in *Proceedings of ICML*, 2016, pp. 173–182.
- [19] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le et al., "Large scale distributed deep networks," in *NIPS*, 2012, pp. 1223–1231.
- [20] F. N. Iandola, K. Ashraf, M. W. Moskewicz, and K. Keutzer, "Firecaffe: near-linear acceleration of deep neural network training on compute clusters," *CoRR*, vol. abs/1511.00175, 2015.
- [21] D. Das, S. Avancha, D. Mudigere, K. Vaidyanathan, S. Sridharan, D. D. Kalamkar, B. Kaul, and P. Dubey, "Distributed deep learning using synchronous stochastic gradient descent," *CoRR*, vol. abs/1602.06709, 2016.
- [22] S. Pathak, P. He, and W. Darling, "Scalable deep document / sequence reasoning with cognitive toolkit," in *Proceedings of the 26th International Conference on World Wide Web Companion*, ser. WWW '17 Companion. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2017, pp. 931–934. [Online]. Available: <https://doi.org/10.1145/3041021.3051103>
- [23] "Scaling Deep Learning on 18,000 GPUs," <https://www.nextplatform.com/2017/03/28/scaling-deep-learning-beyond-18000-gpus/>, 2017.
- [24] A. Anandkumar, "Deep Learning at Scale on AWS." [Online]. Available: <https://ml-days-prd.s3.amazonaws.com/slides/speakers/slides/3/Anima-EPFL2017.pdf>
- [25] S. Hadjis, C. Zhang, I. Mitliagkas, D. Iter, and C. Ré, "Omnivore: An optimizer for multi-device deep learning on cpus and gpus," *arXiv:1606.04487*, 2016.
- [26] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, "On large-batch training for deep learning: Generalization gap and sharp minima," *arXiv:1609.04836*, 2016.
- [27] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE transactions on automatic control*, vol. 31, no. 9, pp. 803–812, 1986.
- [28] F. Niu, B. Recht, C. Re, and S. Wright, "Hogwild: A lock-free approach to parallelizing stochastic gradient descent," in *NIPS*, 2011, pp. 693–701.
- [29] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le et al., "Large scale distributed deep networks," in *NIPS*, 2012, pp. 1223–1231.
- [30] T. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman, "Project adam: Building an efficient and scalable deep learning training system," in *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, 2014, pp. 571–582.
- [31] I. Mitliagkas, C. Zhang, S. Hadjis, and C. Ré, "Asynchrony begets momentum, with an application to deep learning," *arXiv:1605.09774*, 2016.
- [32] C. Zhang and C. Re, "Dimmwwitted: A study of main-memory statistical analytics," *PVLDB*, vol. 7, no. 12, pp. 1283–1294, 2014.
- [33] R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung, "Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit," *Nature*, vol. 405, no. 6789, pp. 947–951, 2000.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of ICCV*, 2015, pp. 1026–1034.
- [35] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv:1412.6980*, 2014.
- [36] E. Racah, C. Beckham, T. Maharaj, C. Pal et al., "Semi-supervised detection of extreme weather events in large climate datasets," *arXiv:1612.02095*, 2016.
- [37] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *CVPR*, 2016, pp. 779–788.
- [38] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European Conference on Computer Vision*. Springer, 2016, pp. 21–37.
- [39] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *NIPS*, 2015, pp. 91–99.
- [40] "Intel® distribution of Caffe\*," <https://github.com/intel/caffe>, 2017.
- [41] "Intel® Machine Learning Scaling Library for Linux\* OS," <https://github.com/01org/mlsl>, 2017.
- [42] "Intel® Software Development Emulator," <https://software.intel.com/en-us/articles/intel-software-development-emulator>, 2017.
- [43] A. Lavin and S. Gray, "Fast algorithms for convolutional neural networks," *CoRR*, vol. abs/1509.09308, 2015.
- [44] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Quantized neural networks: Training neural networks with low precision weights and activations," *CoRR*, vol. abs/1609.07061, 2016.
- [45] M. Courbariaux, Y. Bengio, and J. David, "Training deep neural networks with low precision multiplications," *CoRR*, vol. abs/1412.7024, 2014.
- [46] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan, "Deep learning with limited numerical precision," *CoRR*, vol. abs/1502.02551, 2015.

- [47] P. Gysel, M. Motamedi, and S. Ghiasi, “Hardware-oriented approximation of convolutional neural networks,” *CoRR*, vol. abs/1604.03168, 2016.
- [48] J. Zhang, I. Mitliagkas, and C. Ré, “Yellowfin and the art of momentum tuning,” *arXiv preprint arXiv:1706.03471*, 2017.
- [49] J. Snoek, H. Larochelle, and R. P. Adams, “Practical bayesian optimization of machine learning algorithms,” in *NIPS*, 2012, pp. 2951–2959.
- [50] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [51] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997. [Online]. Available: <http://dx.doi.org/10.1162/neco.1997.9.8.1735>
- [52] F. A. Gers, J. Schmidhuber, and F. Cummins, “Learning to forget: Continual prediction with lstm,” *Neural Computation*, vol. 12, no. 10, pp. 2451–2471, 2000. [Online]. Available: <http://dx.doi.org/10.1162/089976600300015015>