Exploiting Parallelism Opportunities with Deep Learning Frameworks

YU EMMA WANG, Harvard University

CAROLE-JEAN WU, XIAODONG WANG, and KIM HAZELWOOD, Facebook AI

DAVID BROOKS, Harvard University

State-of-the-art machine learning frameworks support a wide variety of design features to enable a flexible machine learning programming interface and to ease the programmability burden on machine learning developers. Identifying and using a performance-optimal setting in feature-rich frameworks, however, involves a non-trivial amount of performance profiling efforts and often relies on domain-specific knowledge. This article takes a deep dive into analyzing the performance impact of key design features in a machine learning framework and quantifies the role of parallelism. The observations and insights distill into a simple set of guidelines that one can use to achieve much higher training and inference speedup. Across a diverse set of real-world deep learning models, the evaluation results show that the proposed performance tuning guidelines outperform the Intel and TensorFlow recommended settings by 1.30× and 1.38×, respectively.

CCS Concepts: • **Computing methodologies** → *Machine learning; Machine learning approaches; Neural networks*;

Additional Key Words and Phrases: Machine learning frameworks, parallel computing, performance analysis

ACM Reference format:

Yu Emma Wang, Carole-Jean Wu, Xiaodong Wang, Kim Hazelwood, and David Brooks. 2020. Exploiting Parallelism Opportunities with Deep Learning Frameworks. *ACM Trans. Archit. Code Optim.* 18, 1, Article 9 (December 2020), 23 pages.

<https://doi.org/10.1145/3431388>

1 INTRODUCTION

The wide adoption of deep learning (DL) tasks has spawned a plethora of domain-specific frameworks to help improve development efficiency and to enable a flexible path from research to production deployment for machine learning (ML). Notable ML frameworks include Caffe/Caffe2 [Jia et al. 2014], PyTorch [Ketkar 2017], TensorFlow [Abadi et al. 2016], and MXNet [Chen et al. 2015]. These frameworks provide high-level APIs as the building blocks of DL models. The ML frameworks can significantly reduce the prototyping cycle because of the substantial (re)use of optimized libraries and greatly improve the productivity of developers building end-to-end DL models. In addition to programmability benefits, the frameworks also provide many DL-specific optimizations

This work was supported by NSF Grant #CCF-1533737.

Authors' addresses: Y. E. Wang, Harvard University, 33 Oxford St, Cambridge, MA; emails: ywang03@g.harvard.edu, dbrooks@eecs.harvard.edu; C.-J. Wu, X. Wang, and K. Hazelwood, Facebook AI, 1 Hacker Way, Menlo Park, CA 94025; emails: {carolejeanwu, xdwang}@fb.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

© 2020 Association for Computing Machinery.

1544-3566/2020/12-ART9

<https://doi.org/10.1145/3431388>

Fig. 1. Time breakdown for Inception v3.

to improve performance and performance portability across software stacks and new hardware systems. The net result is an explosion in the development of ever more complex DL models and a concomitant increase in the performance cost of inference.

Performance is especially important for deep learning inference. When trained DL models are deployed into at-scale production serving, performance optimization for inference can affect data-center resource efficiency, since inference is sensitive to both throughput and latency [Hazelwood et al. 2018; Reddi et al. 2019; Sriraman et al. 2019]. Inference performance also affects whether a model can be deployed to performance- and energy-constrained mobile and edge systems [Wu et al. 2019]. An important research question this article aims to address is *what is the performance cost of high-level abstraction or programmability in modern feature-rich DL frameworks?* A further question examined in the article is *whether and by how much can we reduce this “programmability tax” by tuning the complex set of design choices available in current ML frameworks.* Answering both questions requires a comprehensive understanding of framework features.

Previous works have examined the performance of the modern DL frameworks and focused on cross-framework comparison [Bahrapour et al. 2016; Jain et al. 2019; Shi et al. 2016]. These studies help machine learning model developers and engineers to make informed decisions about framework choices. However, there lacks in-depth, focused performance characterization studies that quantify performance impacts of key framework design features, explaining the underlying performance root cause of each design feature within a framework. These design features include a rich set of parallelism opportunities within and between DL model layers and across input data, for example, across the batch dimension. Furthermore, this article also examines key choices available to users include back-end math kernels, threading libraries, and scheduling policies (Sections 4 to 6).

Figure 1 uses Inception as an example to illustrate the performance impact of important framework features across the programming abstraction.¹ Overall, the performance of Inception can be improved by **3.6×** with properly tuned, optimized framework configurations. Determining performance-optimal framework knobs, however, requires expert knowledge. In this example, tuning the inter-operator parallelism speeds up the performance of native operators by 1.54×, leading to a speedup of **1.47×** for the whole model. Exploiting intra-operator parallelism further speeds up inference performance by another 25×, translating to **2.4×** performance improvement for the entire model. Using the performance tuning guideline described in this work can speed up the native operators by 6.5×, leading to an overall model performance improvement of **3.6×**. Compared to the recommended TensorFlow setting [Google 2019], our tuning guideline speeds up the whole model by 1.15×. The evaluation results across a wide collection of DL workloads demonstrate consistent

¹Performance breakdown for Inception v3 running on a dual-socket CPU platform is presented. Methodological detail is in Section 3.

performance speedup. For some DL workloads, over $2\times$ performance improvement is observed. Selecting the optimal performance setting is not straightforward, even for CPU platforms that serve a large, diverse collection of DL use cases in production datacenter fleets. The significant gap for performance improvement motivates the in-depth studies presented in this article.

This article provides three major contributions:

- We provide a detailed analysis to reason about and quantify performance implications of fundamental framework design features on CPU platforms, including scheduling mechanisms, operator designs, library back-ends, and parallelism mechanisms. We find that the programmability tax ranges from 63% to 1.3%.
- Built upon the insights from this performance analysis, we propose simple guidelines for tuning framework parameters to maximize parallelism and reduce framework overheads.
- We demonstrate the usability of our approach by integrating the proposed methodology into TensorFlow. The code is made available on GitHub.² The proposed settings match the globally optimal performance and outperform the suggested settings from Intel [Anju 2018] and TensorFlow [Google 2019] performance guides by $1.30\times$ and $1.38\times$, respectively, across a set of real-world DL models, including several from the MLPerf suite [Mattson et al. 2020].

2 FRAMEWORK DESIGN OVERVIEW

Deep learning frameworks have lowered the programming effort for DL researchers and decreased prototyping time for new models. Abstraction provided by these frameworks hides some design choices, such as runtime scheduling and actual implementation of operators, which may not be noticeable by users but are important for performance. In this section, we describe how deep learning framework design choices (Section 2.1) exploit parallelism opportunities exposed in deep learning workloads (Section 2.2) and overview our framework parameter tuning methodology (Section 2.3). We also discuss related work. While this work focuses on inference, our analysis and takeaways can be applied for training as well. Thus, we demonstrate both inference and training.

Our performance analysis focuses on CPUs, which are the most widely used platforms serving ML workloads in datacenters [Gupta et al. 2019; Hazelwood et al. 2018]. Performance improvement directly translates into capacity efficiency improvement. As we will see, DL frameworks offer a large set of tradeoffs to exploit parallelism that lead to significant optimization opportunities on CPUs. However, frameworks for accelerators such as GPUs expose less performance knobs for users, because mapping compute kernels to accelerators introduces less variability than mapping to general purpose processors. As developers focusing on accelerating a handful of compute kernels for deep learning, the kernel computation will be increasingly faster and the framework overhead will take larger fraction. At this point, efficient designs of framework is especially important, and the lessons learned from parallelizing CPU frameworks can be transferred to frameworks for accelerators.

2.1 Design Features

Figure 2 presents the stack of DL frameworks and design features that we study. This work focuses on frameworks with opaque operators, i.e., operators or kernels that are manually written and tuned by engineers. It is a design adopted by popular DL frameworks such as TensorFlow, Keras, Pytorch, and Caffe2. Frameworks such as Tensor Comprehensions [Vasilache et al. 2018], TVM [Chen et al. 2018], and Julia [Bezanson et al. 2017] do not belong to this category, since they generate kernels. The configurations they explore are usually not exposed to model developers at

²<https://github.com/Emma926/mcbench>.

Fig. 2. An overview of the framework design features studied in this work.

the framework level such as the ones in Figure 2. Thus, the designs of those frameworks differ in important ways that put them out of the scope of this article.

The design feature study in this work distinguishes itself from previous works that focus on framework-level comparison without revealing the root causes of the performance difference, including Bahrampour et al. [2016], Jain et al. [2019], and Shi et al. [2016]. This work emphasizes that the performance difference of frameworks originates from their different design choices, such as the ones in Figure 2. This applies to different versions of the same framework, such as TensorFlow 1 and 2. The insights from analyzing TensorFlow 1 are applicable to TensorFlow 2, because they share the same design features, and the performance difference is caused by their different choices.

Scheduler. Here “scheduler” refers to the operator scheduler, not the process scheduler of the OS. It takes a computational graph representing the DL workload and schedules its operators based on dependencies and hardware resources. Two common approaches are synchronous and asynchronous scheduling. Synchronous scheduling schedules one operator at a time. Asynchronous scheduling schedules all operators in ready state, such that the operators can execute in parallel if hardware resources are available. In the example of Figure 2, asynchronous scheduling is faster than synchronous scheduling, assuming unlimited hardware units. However, as we will show in Section 4, given limited hardware resources, the optimal mechanism usually falls between the two extremes.

Operator. Frameworks include both native operators and operators based on library kernels. The way operators make use of kernels can have a surprisingly large impact on performance. For example, Figure 2 shows two potential implementations of a MatMul operator based on library kernel MATMUL to implement the MatMul operator at the framework level. The right one passes arguments as is to MATMUL. The left one splits matrix x into smaller blocks and passes each block to a thread in a thread pool. We will show in Section 5.2 that the latter performs better, because it parallelizes data preparation before entering the MATMUL kernel.

Library. Mathematical libraries provide efficient parallel implementations of common kernels. We study three widely used libraries: MKL, MKL-DNN, and Eigen. A thread pool manages a number of threads that execute tasks upon request. Besides the thread pools used by math libraries, DL frameworks use additional thread pools to parallelize computations outside of the math kernels. We study thread pool implementations in the C standard library, Eigen, and Folly.

Beyond One-socket. Parallelism mechanisms need to be applied based on workloads. Common mechanisms include data and model parallelism.

2.2 Parallelism Opportunities

A DL workload can be expressed with a computational graph, where a node represents an operator, and an edge indicates the dataflow dependencies between operators [Abadi et al. 2016]. DL workloads expose the parallelism within an operator (intra-operator), between operators (inter-operator), and among requests. Efficient framework designs should exploit such opportunities.

2.2.1 Parallelism within an Operator. Operators manipulate tensors, i.e., n -dimensional arrays. The parallelism within an operator (intra-operator parallelism) can be exploited with the following techniques:

SIMD. The use of single instruction multiple data (SIMD) architectures, e.g., Intel's AVX instructions [Lomont 2011], is implemented in mathematics libraries. The MKL, MKL-DNN, and Eigen libraries can use AVX2 and AVX512 instructions.

Multi-threading. Multi-threading is implemented at the operator and library levels. For example, MKL uses OpenMP for multi-threading. At the operator level, a framework may have a separate thread pool for further parallelism.

Data Parallelism. Data parallelism splits one batch of data into multiple smaller batches. Thus, it can improve performance of large-batch workloads.

2.2.2 Parallelism between Operators. The parallelism across operators (inter-operator parallelism) can be exploited by scheduling and model parallelism.

Scheduling. Asynchronous scheduling is to place independent operators on different hardware units such as CPU threads, GPU stream multi-processors, or processing elements in accelerators. This way independent operators can execute in parallel.

Model Parallelism. Model parallelism is realized by scheduling different operators (or the same operator after splitting along the model size dimension) on different hardware sockets or nodes, such as distributing large embedding tables across hardware nodes [Eisenman et al. 2018].

Model Pipelining. Model pipelining is a special case of model parallelism [Blog 2019]. One hardware node receives data from a previous node and operates on it while the previous node is computing the next training step. In contrast to data parallelism, model parallelism lowers the memory requirement for large models. This work does not study model pipelining.

2.2.3 Parallelism among Requests. Parallelism among requests can be exploited by batching to transform request-level parallelism to intra-op parallelism. For example, multiple image classification requests can be combined and executed in a single session, such that the number of requests is mapped to the batch size dimension.

2.3 Framework Parameter Tuning

Based on our analysis of design features and parallelism opportunities, we reduce the number of design features needing to be selected from five (scheduling mechanism, operator design, math library, thread pool library, parallelism mechanism) to one, the number of asynchronous scheduling thread pools. Other features, such as operator parallelism, follow from that choice. We propose simple guidelines for tuning framework features based on a model's inter-operator parallelism, as reflected in its computational graph. To demonstrate their usability, we integrate the guidelines with TensorFlow and achieve 1.29 $\times$ and 1.34 $\times$ speedup over Intel [Anju 2018] and TensorFlow [Google 2019] recommended settings, respectively. We also achieve the same average performance as with globally optimal settings and 95% of globally optimal performance in the worst case. We have

Table 1. Hardware Platforms under Study

	SKU	Cores	TFLOPS	Freq	LLC
<i>small</i>	i7-6700k	4	0.423*	4GHz	8MB
<i>large</i>	Platinum 8175M	24	1.64*	2.5GHz	33MB

An additional platform is *large.2* that contains two sockets of *large* with a 120 GB/s bi-directional UPI bandwidth at peak.

*Estimated with GeekBench v4 [Labs 2019].

developed a TensorFlow plugin that sets framework parameters automatically, and it is available on GitHub.³ Details are in Section 8.

Previous work proposed to tune TensorFlow parameters automatically [Hasabnis 2018], which treats the tuning process as a black box and therefore does not explain how parameters affect performance. Our work differs in three ways. First, our tuning method is supported by strong analysis. A deep understanding of framework designs and the root causes of performance difference makes our tuning method intuitive and lightweight. Second, performance with our worst-case settings differs by less than 5% from the *global optimum* obtained by exhaustive search, while previous work [Hasabnis 2018] reports large performance degradation. Finally, the robustness of our guidelines is validated on a different set of real-world DL workloads using a state-of-the-art two-socket system. Our evaluation highlights the importance of framework parameter tuning for state-of-the-art translation and recommendation models.

3 EXPERIMENTAL SETUP

Our scripts and workloads are available on GitHub, together with the TensorFlow plugin.

CPU Platforms. We use three Intel Skylake CPU platforms, *small*, *large*, and *large.2*. *large.2* contains two sockets of *large* with a peak bi-directional bandwidth of 120 GB/s. *large* and *large.2* represent widely used datacenter servers. *small* has fewer cores; we use it to eliminate threading overhead for certain studies. We use *large* and *small* for most of the analysis where the performance tuning guidelines are summarized, and *large.2* for the evaluation of the guidelines. *small* has 32 fused multiply-add (FMA) units per core, while *large* has 64 per core. Although all of the platforms support hyperthreading, each core has only one set of FMA units, which limits the benefits of hyperthreading if both hyperthreads need FMA units. The specifications are summarized in Table 1. *large* and *large.2* are Amazon Web Services m5.metal instances. Unless otherwise specified, we use the *large* platform.

Frameworks. Because TensorFlow supports all features, in most parts of the article, we use TensorFlow v1.13 with the MKL-DNN back-end. Section 4 is an exception, because the Inception architecture is important for this study, and the Caffe2 model zoo makes it more convenient to use Inception. Conducting the same experiments with PyTorch (Caffe2 module) and TensorFlow shows similar trends, and the takeaways in this article apply to both frameworks, because they share the same framework design features as summarized in Figure 2. We set thread affinity to prioritize binding one software thread with one physical core [Anju 2018].

Workloads. We use a set of production-size deep learning models, including three from MLPerf [Mattson et al. 2020] (ResNet-50 [He et al. 2016], Transformer [Vaswani et al. 2017], neural collaborative filtering (NCF) [He et al. 2017]), as well as DenseNet [Huang et al. 2017], SqueezeNet [Iandola et al. 2016], Inception [Szegedy et al. 2016], GoogLeNet [Szegedy et al. 2015], CaffeNet [Jia et al. 2014], ResNext [Xie et al. 2017], and Google’s Wide & Deep Learning model [Cheng n.d.]. To deeply understand the design features, we use micro-benchmarks such

³<https://github.com/Emma926/mcbench>.

Fig. 3. Examples of (a) synchronous scheduling, (b) asynchronous scheduling, and (c) using one and four thread pools with the same total hardware resources.

as matrix multiplication. We use a subset of the aforementioned models to focus our evaluation on the respective design features (Sections 4 and 5). We hold out all the non-vision models for Section 8 to evaluate the proposed method.

Methodology. Our profiling methodology enables thorough analysis. We produce time breakdowns (stack bars) for individual CPU cores using Linux’s `perf` record and profiling one core at a time. Using floating-point performance counters, we measure performance as floating-point operations per second (FLOPS). We trace execution with performance counters by sampling instructions per cycle (IPC) every few milliseconds and ordering the samples by time stamps. The `perf stat` command with the `topdown` option produces the top-down breakdown. LLC misses, memory, and UPI traffic come from the corresponding performance counters.

Terminology. Here are some terms used often later on:

- **MKL Threads:** the threads for MKL and MKL-DNN.
- **Intra-operator Threads:** the threads for an operator at framework level. Abbreviated intra-op threads.
- **Inter-operator Pools:** the independent thread pools in a framework, the size of each set by intra-op threads. Abbreviated inter-op pools, or pools.

4 SCHEDULING MECHANISM

A deep learning model is expressed using a computational graph that represents the data flow between operators. At runtime, operator scheduling exposes optimization opportunities, such as scheduling independent operators simultaneously. In this section, we study the tradeoffs of using such inter-operator parallelism.⁴ We show that not all models benefit from asynchronous scheduling and that the best setting depends on a model’s inter-operator parallelism, quantified by the width of its computational graph.

Figure 3 shows examples of synchronous and asynchronous scheduling of an Inception module [Szegedy et al. 2016] and the use of one and four thread pools. Scheduling of an operator is to submit the job to a thread pool. Sometimes asynchronous scheduling, i.e., running multiple operators simultaneously, can improve performance. For the simple example shown, scheduling one operator at a time takes nine steps to finish (Figure 3(a)); scheduling four operators at a time reduces the steps to five (Figure 3(b)). One simple implementation is to create several thread pools of the same size to share the computing hardware (Figure 3(c)) and to schedule independent operators asynchronously to the thread pools. This design is adopted by popular DL frameworks. In TensorFlow, the number of asynchronous thread pools is called the number of inter-operator

⁴The experiments are conducted with the real-world workloads implemented in Caffe2, because the Inception architecture is important for this study, and the Caffe2 model zoo makes it more convenient to use Inception.

Fig. 4. (Bar Chart) The speedup of using asynchronous scheduling over synchronous. (Table) The maximum computational graph width and best numbers of thread pools. Workloads with more branches benefit from more pools.

threads. Caffe2 calls it the asynchronous thread pool size. In this article, we refer to it as the inter-op pools, as opposed to intra-op threads. We will show that synchronous scheduling is beneficial in both single-socket and multi-socket systems. The best performance is achieved by balancing intra- and inter-operator parallelism.

4.1 Datacenter Platform Performance

We show that the best number of thread pools is no more than the maximum number of parallel operators for a model. We use the *large* platform in Table 1.

Figure 4's bar chart shows the speedup of asynchronous scheduling on different production-size inference and training workloads. The baseline is synchronous scheduling, using one thread pool of size 24. Inference uses three thread pools, each with 8 threads; training uses two pools, each with 12 threads. Workloads benefit from asynchronous scheduling differently. Inception v1 and v2, GoogLeNet, ResNet, and FC-512 speed up more than others.

The performance difference is because of the intrinsic inter-op parallelism of the models, which is quantified by the width, or the number of branches, of their computational graphs. It measures the number of operators that can be scheduled in parallel. The table at the bottom of Figure 4 summarizes the maximum graph width and the best numbers of pools for varying batch sizes. We distinguish between inference and training workloads, because the computational graphs of training workloads contain gradient and sum weight operators, which doubles the number of parallel operators. An intrinsic model limitation is that the best numbers of pools (for varying batch sizes) do not exceed the maximum graph width. For example, Inception modules have a maximum of four branches, as shown in Figure 3. This means the maximum number of operations we can schedule in parallel is four. The optimal number of inter-op pools (2–3) cannot exceed this intrinsic limitation. However, ResNext has only one branch for inference, and two branches for training when counting the backward pass. As a result, the degree of parallelism that can be exploited is constrained by the two inter-op pools of ResNext.

The best number of pools varies based on batch sizes. Large batches increase the best number of pools for inference, but decrease it for training. That is because the parallel operators for training, gradient and sum weight, become imbalanced with large batches. Gradient becomes much compute-intensive than sum weight. Allocating computing resources evenly for the two hurts performance.

4.2 Inception v2 Case Study

To highlight parallelism opportunities at the intra- and inter-op levels, we use Inception v2 as an example. Its model architecture contains operator branches that can execute in parallel. In the

Fig. 5. (a) Inception v2 architecture contains modules with four (b) and three (c) independent branches. Area 1 exhibits inter- and intra-op parallelism and area 2 only intra-op.

Fig. 6. Performance of Inception v2 with different numbers of inter-op pools and MKL threads per pool. The performance is normalized to that of one pool with one MKL thread. Best configuration balances intra- and inter-op parallelism.

baseline implementations that either schedule each branch naively to one CPU core or schedule one operator to all CPU cores, workload imbalance and synchronization overhead significantly reduce performance. We show that synchronization overhead can be mitigated by choosing the number/size of thread pools to better balance intra- and inter-operator parallelism. We use the *small* platform in Table 1, as it enables an exhaustive study of possible cases.

Inception v2 Architecture. To simplify the explanation of later results, we first summarize the Inception v2 architecture [Szegedy et al. 2016] in Figure 5. Figure 5(a) shows the top level architecture, color coded as areas 1 and 2. Area 1 exhibits both intra- and inter-op parallelism, while area 2 only has intra-op parallelism. Area 1 contains two inception modules. Module 4 has four branches (Figure 5(b)) and module 3 has three (Figure 5(c)). The convolution operators are converted to MatMul using `im2col()`, so intensive computation is mostly MatMul.

Performance Scaling with Pools and Threads. Figure 6 shows the relative performance of Inception v2 with a batch size of 16, sweeping inter-op pools and MKL threads per pool. The relative performance is the ratio of the time executing one inference iteration using one pool with

Fig. 7. Execution time breakdown of four cases.

Fig. 8. Execution traces of three cases in Figure 7. Color-coded areas 1 and 2 correspond to the operators in Figure 5.

one MKL thread over that of each configuration. The total number of threads on the system is the product of the two. Hyperthreads are used when more than four threads are created. Exceeding eight threads is labeled over-threading, because there are more software threads than hardware threads. (Scaling is similar with batch sizes from 1 to 128.)

Hyperthreading does not improve performance significantly, such as [4,1] vs. [4,2], and [1,4] vs. [2,4] ([Threads, Pools]). The compute-intensive operators, Convs and MatMuls, are bottlenecked by the fused multiply-accumulate (FMA) units, which are shared between hyperthreads on the same core.

As expected, over-threading, i.e., using more software threads than hardware threads, hurts the performance, because threading overhead increases with more software threads, and computing resources are saturated. As a result, simply setting all framework knobs to the maximum does not yield the best performance.

Performance is best with two pools and two threads per pool. Using four total threads in other ways, such as four pools with one thread each, or one pool with four threads, cannot achieve such performance. Our profiling methodology reveals and visualizes the underlying causes in Figures 7 and 8.

Runtime Breakdown and Execution Traces. We select four cases, a baseline that uses only one thread, and three cases that each use four threads in total. One software thread is bound to one CPU core. Figure 7 shows the aggregate time breakdown, and Figure 8 shows the corresponding execution traces. In Figure 8, one iteration of Inception v2 inference is marked with red bars, and operators in execution are labeled with the corresponding color in Figure 5, where area 1 exhibits intra- and inter-op parallelism, and area 2 has only intra-op parallelism. The fraction of time each core spent executing (rather than synchronizing) is indicated to the right of each trace. It matches the breakdowns in Figure 7.

The first case, using four pools of size one, incurs high synchronization overhead in Figure 7 and Figure 8(a), primarily because the operators in area 2, with only intra-op parallelism, are assigned only one core (thread), and other cores are waiting to synchronize. The trace in area 1 is slightly better, with every core executing one branch of each inception module.

The second case, using one pool with four threads, does not perform well either, because the Caffe2-native operations (the parts labeled as Caffe2 and Caffe2:Math in Figure 7) are single-threaded, and other cores are stalled by core 0, as shown by the long synchronization time in the rightmost three bars of Figure 7 and the traces of cores 1–3 in Figure 8(b).

The third case, using two pools, each with two threads, is a better balance. It reduces synchronization time in both area 1 and area 2 compared to the first case. Area 1 is improved, because the number and size of convolutions in each inception branch is not even, as shown by Figure 5(b), and allocating one core for each branch (the first case) makes the core running the smallest branch wait for a long time. Area 2 is improved, because its operators are sequential and compute-intensive, and having more cores improves performance.

Optimization Opportunity. Operators in a complex model come in different sizes and have different dependencies. Fixing each thread pool size usually incurs synchronization overhead because of work imbalance. Thus, there is an opportunity to implement a global thread pool, allowing the scheduler to determine dynamically how many threads to schedule for each operator. For example, in the traces of Figure 5, providing area 1 with two pools of two threads each and area 2 with one pool of four threads can lead to higher performance.

5 OPERATOR DESIGN

Operators are building blocks of DL frameworks, providing basic semantics through high-level language (like Python) APIs to ease the development process for framework users. As shown in Figure 1, a workload built with a DL framework involves library kernels and framework native computation. In this section, we show that efficient operator design can speed up framework native computation and yields up to 4.2× performance improvement for real-world models.

Implementations of Framework Native Operators. We first describe framework native operators, the operators that do not use library kernels. Native operators handle control flow as well as tensor reshaping, broadcasting, and preprocessing. Some, like those for control flow or input image preprocessing, are necessary. Others can fairly be described as a framework programmability tax. The overhead stems from preparing inputs for library kernels or computing how to parallelize a given workload in the main thread. One example of the latter kind is Eigen::ParallelFor, used by TensorFlow.

Implementations of Compute-intensive Operators. A framework operator must sometimes do more than simply pass arguments to library kernels. Data preparation is often required, for example. Taking matrix multiplication (MatMul) as an example, we list two implementations below. In the context of deep learning, x is an input matrix of size [batch size × number of activations] and w is a weight matrix of size [activations in current layer × activations in next layer]. We

Fig. 9. Speedup from using 24 MKL threads instead of one. TensorFlow exhibits lower speedups than MKL.

assume MatMul is the interface of a framework operator, and MATMUL is the corresponding library kernel, e.g., in MKL.

```
Sec 5.1) MatMul1(x, w):
    data_prep(x, w)
    return MATMUL(x, w)
```

```
Sec 5.2) MatMul2(x, w):
    // Reshape x, w into bx and bw
    for bx, bw in x, w:
        threadpool.run(MatMul1,
            bx, bw)
    return threadpool.join_results()
```

MatMul1 conducts data preparation and passes the matrices to the library kernel. MatMul2 uses an additional thread pool that we call the intra-operator thread pool to distinguish it from the MKL thread pool. The operator splits the matrices into smaller ones and passes those small matrices to the intra-op thread pool. The intra-op thread pool then executes multiple copies of MatMul1 in parallel. This way data_prep() of the whole matrix can be parallelized. Before and after the library call, the data formatting and results gathering work is part of the programmability tax. We will study the two implementations in the following subsections.

5.1 MKL Threads

By analyzing the overhead and scalability of the first operator implementation, MatMul1, we show that both the TensorFlow (TF) operator and the MKL kernel suffer from data preparation overhead, which prevents them from scaling linearly with the number of CPU cores. The results here can also apply when convolution operators are converted to MatMuls using im2col(). We use the *large* platform in Table 1.

Performance Scaling. Both TF and MKL have scaling issues, and TF is slightly worse. Figure 9 shows the speedup of using 24 MKL threads over using one, for both TF operators and MKL kernels. The matrices are squared and represented by one dimension. The total number of floating-point operations is the cube of that number. Figure 9 shows that the speedup of TF is always lower than that of MKL, especially for small matrices. TF speedup is comparable to MKL when matrices are larger than 4K. The maximum speedup achievable is about 16×, which is lower than the number of cores, 24.

Causes of the Poor Scalability. Our profiling methodology reveals that data preparation overhead causes suboptimal performance scaling. We pick two variants of MatMul, MatMul-512 and MatMul-4K, that operate on medium- and large-size matrices, respectively. MatMul-512 represents the fully connected (FC) layers from YouTube [Covington et al. 2016] and Facebook recommendation [Gupta et al. 2019; Naumov et al. 2019] models, while MatMul-4K represents the FC layers in Transformer [Vaswani et al. 2017]. Figure 10 shows the runtime breakdown of all CPU cores

Fig. 10. Runtime breakdown for all CPU cores. Data preparation overhead causes the poor scalability in Figure 9.

running the MatMuls, using 1 and 24 MKL threads. With multiple threads, the thread tasked with lengthy TF data preparation is the main thread, labeled CPU Core 0. The latency of each MatMul workload is normalized to that of using one MKL thread.

The TF parts of Figure 10 show that TF’s scaling issue is caused by framework overhead, due mainly to TF data preparation for MKL kernels. Using one MKL thread, MatMul-512 spends over 10% of its time in TF data preparation; using 24 MKL threads, the overhead exceeds 72%. Overhead is much lower for MatMul-4K: less than 3% in both cases. Without TF overhead, speedup can clearly be much greater. The MKL parts of Figure 10 show that MKL data preparation causes the scaling issue for MKL kernels. MKL kernel execution time is roughly 1/24 of the original runtime for MatMul-512. Speedup drops with time spent in MKL data preparation.

The Role of Framework Design. The Amdahl’s law bottleneck of DL frameworks is non-negligible. The overhead of a MatMul with size $n \times n \times n$ scales linearly with n ($O(n)$), while the number of floating-point operations scales cubically ($O(n^3)$). Thus, the speedup of large MatMuls (e.g., 4K) is closer to ideal speedup. Realistically, however, the most commonly used FC layers are not always large enough. Actually, smaller ones are common in commercial workloads including YouTube’s [Covington et al. 2016] recommendation model (of size 256 to 1K) and Facebook’s [Gupta et al. 2019; Naumov et al. 2019] (of size 64 to 512). Especially when hardware platforms are upgraded to higher floating-point computation capability, we need even larger matrices to amortize the overhead. Thus, it is key to focus optimization efforts on mitigating framework overhead.

5.2 Intra-operator Threads

After decades of optimizing the GEMM kernel, the performance bottleneck has shifted to the overhead of using such kernels, e.g., the data preparation overhead in Figure 10. A natural approach to reducing the overhead in framework design is to parallelize the framework native computation, with an intra-operator thread pool implemented at the framework level, as MatMul2 does. We show that when library kernels are using FMA units, intra-op threads improve performance by utilizing other computational units on the same physical core, thereby benefiting from Intel’s hyperthreading technology. We use the *large* platform from Table 1.

Performance Improvement. We first show how much performance improvement intra-op threads can yield and where it comes from. Figure 11 summarizes speedup and time breakdown when using 1 (left bar) and 24 (right bar) intra-op threads. Both cases use 24 MKL threads. MatMul-512 and MatMul-4K are the same operators as in previous sections. Using 24 intra-op threads reduces the execution time of TF native operators, while that of other parts stays similar. The speedup ranges from 1.05 $\times$ (DenseNet) to 4.21 $\times$ (SqueezeNet).

Workloads bottlenecked by TF native operators benefit more from intra-op threads. Such workloads, including MatMul-512 and SqueezeNet, have small to medium MatMul or convolution

Fig. 11. Runtime breakdown of TensorFlow workloads with 1 (left bar) and 24 (right bar) intra-op threads. Both cases use 24 MKL threads.

Fig. 12. Runtime breakdown of all CPU cores. Intra-op threads parallelize the overhead in cores 24–47.

operations, because TF native operators are likely to consume larger fractions of computation time. For example, SqueezeNet has a small percentage of MKL computation, since it is designed to have fewer parameters than AlexNet by using many small (1×1) convolution kernels [Iandola et al. 2016].

Programmability Tax. Figure 11 also quantifies the framework programmability tax. We estimate the tax using the non-MKL fractions, since they are not compute-intensive and can be largely optimized if written with high-performance language/code as MKL kernels. After optimizing with intra-op threads, the programmability tax ranges from 1.3% (DenseNet) to 63% (MatMul-512). SqueezeNet (47%) is higher than ResNet-50 (26%). MatMul-4K (11%) is slightly smaller. This is the price we are paying for using frameworks. The programmability tax of one framework can be seen as the upper bound of performance difference between itself and better-performed frameworks. Imagine a less programmable framework that does nothing but MKL library calls; the performance difference between it and TensorFlow can be between 1.3% and 63%.

Full-system Profiling. Our profiling methodology visualizes the execution of every core on the CPU platform to expose the reasons for performance improvement. Figure 12 shows the time breakdown for all 48 hyperthreads on the *large* platform. We focus on the two MatMuls, since they are the simplest workloads. Because cores 24–47 are not active when using one intra-op thread, the third bar is omitted for that case. Core 0 of each case is the same as in Figure 11.

Figure 12 shows that with 24 intra-op threads, TF data preparation is distributed to cores 24 through 47. (The bottom of the third bar for MatMul-512 with 24 intra-op threads shows a tiny TF data preparation cost.) That shortens TF data preparation time in core 0, so the TF barrier time of cores 1 to 23 is shortened. With only one intra-op thread, cores 1 to 23 spend about 60% (MatMul-512) and 40% (MatMul-4K) of time waiting in barrier, which is a big waste.

Fig. 13. (a) Cycle breakdown (bottom axis) and IPC (top axis) for three libraries multiplying matrices of various sizes. MKL has the highest retiring ratio and IPC, because it is the least back-end bound. (b) MKL has the lowest LLC miss rate. (c) The memory bandwidth consumption. MKL's prefetching is the most effective: Almost all memory traffic is prefetching.

Hyperthreading. Using intra-op threads takes advantage of Intel's hyperthreading technology by colocating an intra-op thread and an MKL thread on the same physical core. Since they need different hardware resources, they can execute in parallel without contention. The critical path is the MKL thread. The intra-op thread adds no execution time to the overall workload. For example, in Figure 12, logical cores 0 and 24 are on the same physical core. Core 0 executes mostly MKL floating-point operations with FMA units, which core 24 does not need. Even without hyperthreading, the implementation of intra-op threads improves performance through parallelization. In that case, the physical core's critical path combines the intra-op thread with the MKL thread.

6 LIBRARY CHOICE

Thanks to the mathematics and thread pool libraries, deep learning framework developers do not have to implement every basic function from scratch. In this section, we study libraries for machine learning and thread pools. We show that optimization can improve a GEMM kernel's performance by up to 25%, owing to more efficient data prefetching. We also find that robust thread pools such as Eigen and Folly are better able to keep production-critical workloads running with little variation, thus investing in sophisticated implementations is worthwhile for service providers.

6.1 Machine Learning Library

We compare MKL, MKL-DNN, and Eigen with GEMM (general matrix multiplication) microbenchmarks on the *small* platform (Table 1) to expose architectural bottlenecks.

We conduct top-down analysis [Yasin 2014] for single-threaded GEMM kernels with a variety of matrix sizes. Figure 13(a) shows the cycle breakdown (stacked bars) and IPC (dots). The three bars shown for each matrix size are for Eigen, MKL-DNN, and MKL, from top to bottom. For GEMM, MKL performs the best, followed by MKL-DNN. With matrices larger than 4K, about 25% of cycles are back-end bound for Eigen and MKL-DNN.

The back-end bottleneck is caused by last-level-cache (LLC) misses, shown as LLC misses per thousand instructions (MPKI) in Figure 13(b). Eigen and MKL-DNN have much higher LLC MPKI than MKL. The LLC miss rate difference is caused by the aggressiveness and effectiveness of data prefetching, shown by the memory traffic in Figure 13(c), where the right ends of the bars show

Fig. 14. Thread pool overhead, measured as time to run 10K micro tasks. Folly outperforms std::thread and Eigen even in an extreme case with 64 threads on a 4-core CPU.

memory traffic incurred by LLC misses. MKL’s memory traffic is close to that of MKL-DNN, and its much lower LLC miss rate shows that MKL’s software prefetching is more *effective*.

We compare the GEMM kernels to demonstrate how and why kernel performance can vary. MKL-DNN is likely to outperform other libraries for other kernels, because it is a library targeting deep learning with DL-specific optimizations. It reportedly outperforms other frameworks [Bhuiyan et al. n.d.].

We compare three thread pools, one simple implementation using std::thread, and the thread pools in Eigen [Eigen 2019] and Folly [fol 2019]. Our microbenchmark creates a thread pool and starts 10K tasks that increment the value of a globally shared variable. This microbenchmark examines the case where the computation per thread is minimal and the synchronization between threads is the maximum. It is a common way to stress-test thread pools such as in Pytorch⁵ as an isolation study. We benchmark in two scenarios: (1) setting the number of threads to be the same as the number of physical cores and (2) using many more threads than the number of cores. We use the *small* platform from Table 1.

Figure 14 compares the overall latency of running 10K micro tasks. In both cases, Folly outperforms Eigen and Eigen outperforms std::thread. When thread pool size is 64, greatly exceeding the number of CPU cores, Folly and Eigen perform consistently well as with four threads, and over-subscribing the system does not drastically increase the overhead. But the overhead of std::thread increases by over 3×, with every CPU core spending about 60% of its time in synchronization.

7 BEYOND ONE SOCKET

In previous sections, we have explored the framework designs on one-socket CPUs. In this section, we study how those design features can be applied to scale out the workloads beyond one socket. Unsurprisingly, the bottleneck of a two-socket system is the UPI bandwidth between sockets. We study scaling one operator to two sockets and scheduling multiple operators to different sockets as scaling-out versions of intra- and inter-op parallelism studies. The experiments are conducted on the *large* and *large.2* platforms from Table 1.

7.1 Data Parallelism

We study data parallelism by setting the numbers of intra-op and MKL threads to the total number of physical cores and the number of inter-op threads to one.

ResNet Performance. Figure 15 shows the execution time breakdown of ResNet-50 running on one- and two-socket platforms. The latter speeds up ResNet by 1.43×, less than the two-fold hardware increase. The bottleneck is that UPI traffic peaks at 91.4 GB/s, compared to the theoretical maximum of 120 GB/s. UPI saturation increases the latency of TF native operators on a two-socket platform, which now includes both data preparation and transfer time between sockets.

MatMul Performance. To test the limit of the *large.2* platform, we conduct microbenchmarking using TensorFlow MatMul operators. Similarly, the UPI bandwidth is the bottleneck for large

⁵https://github.com/pytorch/pytorch/blob/master/binaries/at_launch_benchmark.cc.

Fig. 15. A two-socket platform speeds up ResNet-50 by 1.43X. The bottleneck is the UPI bandwidth, increasing TF data preparation time as part of the TF native operator time.

Fig. 16. (a) Speedup of a two-socket platform over one socket. (b) Measured peak UPI bandwidth consumption on the two-socket platform is close to 100 GB/s.

Fig. 17. Runtime breakdown of all CPU cores.

MatMuls. Figure 16 shows two-socket speedup and corresponding UPI bandwidth consumption. The speedup and UPI throughput increase with larger MatMul sizes and peak for MatMul-8K. For MatMul-16K, speedup decreases and bandwidth saturates, indicating empirically the maximum UPI bandwidth is around 100 GB/s for such workloads.

The speedup of a workload is determined by its intrinsic parallelism and UPI bandwidth saturation. Figure 17 shows the time breakdown of MatMuls running on one- and two-socket platforms. For medium MatMul sizes like 512, the poor scalability is caused by the limited parallelism of the workload, which cannot hide data preparation overhead. For larger MatMuls, 4K and 8K, the data preparation time of both TF and MKL increases on the two-socket platform because of UPI saturation. MatMul-8K has the best balance of intrinsic parallelism and UPI throughput. It leads to the highest speedup (1.8X, i.e., 44% less execution time than with just one socket), which is close to perfect scaling.

7.2 Model Parallelism

We study model parallelism of a two-socket platform by using two inter-op pools, one per CPU socket. Model parallelism improves performance significantly when the parallel operators are on

critical paths and have similar sizes, as with multiple embedding operators in neural collaborative filtering (NCF). Performance and model parallelism mechanisms will be discussed in the next section.

Model parallelism does not always improve performance. One example is the inter-op parallelism from training workloads, as in Section 4. Assigning gradient and weight sum operators one socket each causes workload imbalance between two sockets when batch size is large. Two-socket platforms are not beneficial when the intra-op thread pool is not implemented at the framework level, since the workload bottleneck is single-threaded operators.

8 FRAMEWORK DESIGN TUNING

At the outset, we identified five design features: scheduling mechanism, operator implementation, math library, thread pool library, and the parallelism mechanism for platforms larger than one socket. Our analysis has shown that the best math library is MKL-DNN with MatMul back-end, and the best thread pool library is Folly. The most effective setting for users to determine the proper number of inter-op pools based on the model architecture. Intra-op parallelism and cross-socket configurations follow from that setting. Based on the analysis in previous sections, we therefore propose a set of simple yet effective guidelines. The guidelines are architecture-irrelevant, because they are summarized based on software design analysis; thus, they can be applied to CPUs with different architectures, including Intel, AMD, and OpenPOWER.

Definitions. To present the guidelines, we first define a few terminologies. Section 4.1 and Figure 4 mentioned the maximum width of a model graph. The *average width* of a model graph quantifies its inter-op parallelism, and we use it to determine the number of inter-op pools for a model. The average width of a model is the floor of the ratio of the total number of (heavy) operators divided by the maximum number of layers. A *heavy operator* is a compute-intensive or embedding operator that usually takes significantly longer execution time than other operators. Examples are the Conv operators in Figure 5, as opposed to lightweight math operators, which are not considered. The average model width of Figure 5(b) is $\lfloor \frac{7}{3} \rfloor = 2$.

Guidelines. The number of inter-op pools (p) is chosen to be the average model width. After p is chosen, we choose the numbers of MKL and intra-op threads such that the entire system is split into p partitions without redundant threads (Section 4.2 and Figure 6). Therefore, the number of MKL threads and the number of intra-op threads for each thread pool should be equal to the total number of physical cores on the system divided by p . That way one MKL thread and one intra-op thread can share the same physical core. MKL threads can use the FMA units and the intra-op threads can use other units via hyperthreading (Section 5.2 and Figure 12).

Evaluation Setup. We integrate our guidelines with TensorFlow v1.13. (TensorFlow refers to inter-op pools as inter-op parallelism threads.) We will open source the TensorFlow plugin to the public to let users apply our guidelines automatically. The size of the design space encompassing the numbers of MKL, intra-, and inter-op threads is the cube of the number of logical cores. For the *large.2* system, that means $96^3 = 884,736$ design points. Our guidelines suggest picking only one of those 884,736 possibilities.

We evaluate our guidelines by applying the rules to a fresh set of workloads and by performing the evaluation on a different platform from that used to develop the guidelines. Our analysis in previous sections uses microbenchmarks and vision models that run with images; for the final evaluation in this section, we add Inception v3 [Szegedy et al. 2016] with ImageNet dataset [Deng et al. 2009], the wide-deep recommendation model [Cheng n.d.] and the neural collaborative filtering model (NCF) [He et al. 2017] with the movielens ml-20m dataset,⁶ and Transformer [Vaswani

⁶<https://grouplens.org/datasets/movielens/>.

Fig. 18. Performance using the recommended TensorFlow settings [Google 2019] (baseline), Intel blog [Anju 2018], our work, and global optimum obtained by exhaustive search. The performance is normalized to the baseline, TensorFlow recommended settings (orange bars). Our work outperforms Intel and Tensorflow suggestions and nearly closes the gap between the state-of-the-art and the global optimum.

et al. 2017] with the WMT 2014 English-to-German dataset,⁷ covering recommendation and translation workloads. The bulk of our previous analysis uses platforms *small* and *large* from Table 1; here, we evaluate the guidelines on the *large.2* platform, the largest AWS bare metal instance.

Speedup. Figure 18 summarizes the speedup of this work over TensorFlow [Google 2019] (baseline) and Intel [Anju 2018] recommended settings. The performance is the inverse of one iteration time over that of the baseline, TensorFlow recommended settings (orange bars). It also compares our settings' performance to the global optimum obtained by exhaustively sweeping the design space. TensorFlow suggests setting the number of MKL and intra-op threads to the physical core count, and inter-op pools to the socket count. Intel suggests setting MKL and intra-op threads to the number of physical cores per socket and inter-op pools to the socket count.

Our analysis shows that TensorFlow suggests exploiting the maximum of intra-op parallelism. This works great for image models, because they are usually compute-intensive and exhibit abundant intra-op parallelism. That is why TensorFlow-recommended settings perform better than Intel's, sometimes perform as well as the global optimum, for image models including DenseNet, ResNet, and SqueezeNet. However, Intel's setting is suitable for models with inter-op parallelism. This benefits sequence and recommendation models, because they are usually memory-intensive and exhibit opportunities for inter-op parallelism. That is why Intel-recommended settings perform better, sometimes achieving the global optimum for models such as NCF and Transformer.

Overall, our performance guidelines perform consistently better than the settings recommended by Intel and TensorFlow for all model types. Our method bridges the performance gap between those state-of-the-art settings and the global optimum for all evaluated workloads except Inception inference and SqueezeNet training. In those two cases, our guidelines achieve 95% of the performance offered by the global optimal setting. On average, this work achieves the same

⁷<http://data.statmt.org/>.

Table 2. Average Model Width, i.e., the Number of Pools Selected for Figure 18 Based on Our Guidelines

Dense	Squeeze	ResNet	IncepV3	W/D	NCF	Trans
1	1	1	2	3	4	4

Intra-op and MKL threads = total physical cores divided by those numbers.

performance as the global optimum and $1.38\times$ and $1.30\times$ better performance than TensorFlow's and Intel's suggestions, respectively.

We summarize the average model width in Table 2. It is the same as the number of inter-op pools in use. The models shown have average width between one and four, which is diverse. The numbers of MKL and intra-op threads for each model is the total number of physical cores (48) divided by the model width. For example, the setting for the W/D (wide and deep) model is 3 inter-op pools, 16 MKL threads, and 16 intra-op threads, which is also the globally optimal setting.

The performance guides from Intel and TensorFlow are general, aiming to make it easy for most users to get reasonable performance, and they perform reasonably well. For vision models, TensorFlow's settings perform as well as the global optima and our guidelines, while Intel's does not perform well for the vision models except for Inception, because Intel's setting favors models with inter-op parallelism that other vision models do not have. Intel's settings perform better than TensorFlow's for recommendation and translation models. The latter have several parallel embedding operators, thus their average width is no less than two. The default TensorFlow setting performs much worse than both the Intel and TensorFlow recommendations. TensorFlow naively sets all parameters—MKL threads, intra-op threads, and inter-op pools—to the number of logical cores. As pointed out in our earlier analysis, this is sub-optimal. Thus, TensorFlow users who run only one model and one session at a time should at least set the number of inter-op pools to one instead of using the default setting.

9 FUTURE WORK

This is the first work that evaluates the performance impact of key design features in machine learning frameworks in depth. As future work, systematic performance analysis can be performed to cover more aspects of machine learning framework design and for different hardware platforms. One example is the choice of compilation styles, just-in-time (JIT) and ahead-of-time (AOT), which provide different tradeoff among compilation overhead, compilation optimization opportunities, and binary sizes.

Frameworks with opaque (implemented) operators are studied. Other framework types provide research opportunities, such as the ones that automatically generate optimized operator implementations including Tensor Comprehensions [Vasilache et al. 2018], TVM [Chen et al. 2018], and Julia [Bezanson et al. 2017]. Such frameworks have unique design features, leading to different performance implications.

As the first work, we focus on CPUs, the most widely used platforms in serving ML workloads in datacenters. For future work, the study and comparison of framework designs on GPUs and accelerators, such as TPUs, open up a lot of interesting questions and directions. For example, whether the programmability tax gets higher on GPUs and TPUs compared to CPUs, given the kernel execution is more highly accelerated. Also, it is known that the frameworks for GPUs and accelerators provide fewer user-level performance knobs; it will be interesting to understand the unique challenges of designing frameworks for such platforms.

10 CONCLUSION

We presented a detailed evaluation and analysis of key design features and the role of parallelism in a machine learning framework, focusing on scheduling, operator implementation, and library back-ends. To maximize parallelism, we proposed simple guidelines for tuning framework parameters, distilled from detailed domain-specific design feature knowledge and analysis. We demonstrated the usability and the additional performance improvement of this approach by integrating and evaluating our methodology with TensorFlow. On average, our method outperformed the suggested settings from Intel and TensorFlow performance guides by 1.30× and 1.38×, respectively, across a set of real-world DL models.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their thoughtful comments and constructive suggestions. The authors would like to thank Udit Gupta, Glenn Holloway, Hsien-Hsin S. Lee, Jhe-Yu Liou, Jongsoo Park, Brandon Reagen, and Xuan “Silvia” Zhang for their feedback.

REFERENCES

- Facebook. 2019. Folly: Facebook Open-source Library. Retrieved from <https://github.com/facebook/folly>.
- Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A system for large-scale machine learning. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI'16)*, Vol. 16. 265–283.
- P. Anju. 2018. Tips to improve performance for popular deep learning frameworks on CPUs. *Intel Dev.Zone* (2018). <https://software.intel.com/content/www/us/en/develop/articles/tips-to-improve-performance-for-popular-deep-learning-frameworks-on-multi-core-cpus.html>.
- Soheil Bahrampour, Naveen Ramakrishnan, Lukas Schott, and Mohak Shah. 2015. Comparative study of Caffe, Neon, Theano, and Torch for deep learning. *arXiv preprint arXiv:1511.06435*.
- Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. 2017. Julia: A fresh approach to numerical computing. *SIAM Rev.* 59, 1 (2017), 65–98.
- Ashraf Bhuiyan, Mahmoud Abuzaina, Niranjana Hasabnis, Niroop Ammbashankar, Faijul Amin, Sheng Fu, and Bhavani Subramanian. [n.d.]. Improving TensorFlow inference performance on Intel Xeon processors. *Intel AI Blog* ([n.d.]). intel.com. <https://www.intel.com/content/www/us/en/artificial-intelligence/posts/improving-tensorflow-inference-performance-on-intel-xeon-processors.html>.
- Google AI Blog. 2019. Introducing GPipe, an open source library for efficiently training large-scale neural network models. Retrieved from <https://ai.googleblog.com/2019/03/introducing-gpipe-open-source-library.html>.
- Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. MXNet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274* (2015).
- Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze et al. 2018. TVM: An Automated end-to-end optimizing compiler for deep learning. In *Proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI'18)*. 578–594.
- Heng-Tze Cheng. 2016. Wide and deep learning: Better together with TensorFlow. *Google AI Blog*. Retrieved from <https://ai.googleblog.com/2016/06/wide-deep-learning-better-together-with.html>.
- Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 191–198.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 248–255.
- Eigen. 2019. Eigen thread pool. (2019). Retrieved from <https://bitbucket.org/eigen/eigen/src/default/unsupported/Eigen/CXX11/src/ThreadPool/>.
- Assaf Eisenman, Maxim Naumov, Darryl Gardner, Misha Smelyanskiy, Sergey Pupyrev, Kim Hazelwood, Asaf Cidon, and Sachin Katti. 2018. Bandana: Using non-volatile memory for storing deep learning models. *arXiv preprint arXiv:1811.05922* (2018).
- Google. 2019. TensorFlow Performance Guide. Retrieved from https://docs.w3cub.com/tensorflow~guide/performance/performance_guide/#general_best_practices.
- Udit Gupta, Xiaodong Wang, Maxim Naumov, Carole-Jean Wu, Brandon Reagen, David Brooks, Bradford Cottel, Kim Hazelwood, Bill Jia, Hsien-Hsin S. Lee, Andrey Malevich, Dheevatsa Mudigere, Mikhail Smelyanskiy, Liang Xiong,

- and Xuan Zhang. 2019. The architectural implications of Facebook’s DNN-based personalized recommendation. *arXiv preprint arXiv:1906.03109* (2019).
- Niranjan Hasabnis. 2018. Auto-tuning TensorFlow threading model for CPU backend. *arXiv preprint arXiv:1812.01665* (2018).
- Kim Hazelwood, Sarah Bird, David Brooks, Soumith Chintala, Utku Diril, Dmytro Dzhulgakov, Mohamed Fawzy, Bill Jia, Yangqing Jia, Aditya Kalro, et al. 2018. Applied machine learning at Facebook: A datacenter infrastructure perspective. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA’18)*. IEEE, 620–629.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 173–182.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4700–4708.
- Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. 2016. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5 MB model size. *arXiv preprint arXiv:1602.07360* (2016).
- Arpan Jain, Ammar Ahmad Awan, Quentin Anthony, Hari Subramoni, and Dhableswar K. D. K. Panda. 2019. Performance characterization of DNN training using Tensorflow and PyTorch on modern clusters. In *Proceedings of the IEEE International Conference on Cluster Computing (CLUSTER’19)*. IEEE, 1–11.
- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM International Conference on Multimedia*. ACM, 675–678.
- Nikhil Ketkar. 2017. Introduction to PyTorch. In *Deep Learning with Python*. Springer, 195–208.
- Primate Labs. 2019. GeekBench v4. Retrieved from <https://www.geekbench.com/>.
- Chris Lomont. 2011. Introduction to Intel advanced vector extensions. *Intel White Paper* (2011), 1–21. <https://software.intel.com/content/dam/develop/external/us/en/documents/intro-to-intel-avx-183287.pdf>.
- P. Mattson, V. J. Reddi, C. Cheng, C. Coleman, G. Diamos, D. Kanter, P. Micikevicius, D. Patterson, G. Schmuelling, H. Tang, G. Wei, and Carole-Jean Wu. 2020. MLPerf: An industry standard benchmark suite for machine learning performance. *IEEE Micro* 40, 2 (2020), 8–16.
- Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Malleevich, Ilia Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy. 2019. Deep learning recommendation model for personalization and recommendation systems. *arXiv preprint arXiv:1906.00091* (2019).
- Vijay Janapa Reddi, Christine Cheng, David Kanter, Peter Mattson, Guenther Schmuelling, Carole-Jean Wu, Brian Anderson, Maximilien Breughe, Mark Charlebois, William Chou, et al. 2019. MLPerf inference benchmark. *arXiv preprint arXiv:1911.02549* (2019).
- Shaohuai Shi, Qiang Wang, Pengfei Xu, and Xiaowen Chu. 2016. Benchmarking state-of-the-art deep learning software tools. In *Proceedings of the 7th International Conference on Cloud Computing and Big Data (CCBD’16)*. IEEE, 99–104.
- Akshitha Sriraman, Abhishek Dhanotia, and Thomas F. Wenisch. 2019. SoftSKU: Optimizing server architectures for microservice diversity at scale. In *Proceedings of the 46th International Symposium on Computer Architecture*. ACM, 513–526.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1–9.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2818–2826.
- Nicolas Vasilache, Oleksandr Zinenko, Theodoros Theodoridis, Priya Goyal, Zachary DeVito, William S. Moses, Sven Verdoolaege, Andrew Adams, and Albert Cohen. 2018. Tensor comprehensions: Framework-agnostic high-performance machine learning abstractions. *arXiv preprint arXiv:1802.04730* (2018).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the International Conference on Advances in Neural Information Processing Systems*. 5998–6008.
- Carole-Jean Wu, David Brooks, Kevin Chen, Douglas Chen, Sy Choudhury, Marat Dukhan, Kim Hazelwood, Eldad Isaac, Yangqing Jia, Bill Jia, et al. 2019. Machine learning at Facebook: Understanding inference at the edge. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA’19)*. IEEE, 331–344.

Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. 2017. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1492–1500.

Ahmad Yasin. 2014. A top-down method for performance analysis and counters architecture. In *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS'14)*. IEEE, 35–44.

Received May 2020; revised September 2020; accepted October 2020