

Deep Learning Inference in Facebook Data Centers: Characterization, Performance Optimizations and Hardware Implications

Jongsoo Park*, Maxim Naumov†, Protonu Basu, Summer Deng, Aravind Kalaiah, Daya Khudia, James Law, Parth Malani, Andrey Malevich, Satish Nadathur, Juan Pino, Martin Schatz, Alexander Sidorov, Viswanath Sivakumar, Andrew Tulloch, Xiaodong Wang, Yiming Wu, Hector Yuen, Utku Diril, Dmytro Dzhulgakov, Kim Hazelwood, Bill Jia, Yangqing Jia, Lin Qiao, Vijay Rao, Nadav Rotem, Sungjoo Yoo and Mikhail Smelyanskiy
Facebook, 1 Hacker Way, Menlo Park, CA

Abstract

The application of deep learning techniques resulted in remarkable improvement of machine learning models. In this paper we provide detailed characterizations of deep learning models used in many Facebook social network services. We present computational characteristics of our models, describe high-performance optimizations targeting existing systems, point out their limitations and make suggestions for the future general-purpose/accelerated inference hardware. Also, we highlight the need for better co-design of algorithms, numerics and computing platforms to address the challenges of workloads often run in data centers.

1. Introduction

Machine learning (ML), deep learning (DL) in particular, is used across many social network services. The high quality visual, speech, and language DL models *must scale to billions of users of Facebook’s social network services* [25].

The power consumption in data centers¹ used to run these models has been rapidly increasing over time. A significant fraction of the future demand is expected to come from workloads corresponding to DL inference, as shown on Figure 1. The higher DL inference demand is due to the expanding range of corresponding applications and the steady improvement in the quality of DL models, which is often associated with the increase in compute and memory requirements [2].

In order to tackle this trend, a lot of research has been done on optimizing computing platforms for DL, including but not limited to [1, 18, 25, 33, 49, 50, 61, 62]. However, a great challenge has been the fast pace of changes in DL applications. For instance, the previously relevant AlexNet [40] is no longer representative of the computation characteristics of today’s computer vision (CV) systems. The rate of change in DL models is so fast that hardware optimized for old models can easily become inefficient for new models.

*jongsoo@fb.com †mnaumov@fb.com

¹The collective power consumption of data centers around the world would be ranked 4th behind only China, US and EU [4]




Figure 1: Server demand for DL inference across data centers

In order to perform a characterizations of the DL models and address aforementioned concerns, we had direct access to the current systems as well as applications projected to drive them in the future. Many inference workloads need flexibility, availability and low latency provided by CPUs. Therefore, our optimizations were mostly targeted for these general purpose processors. However, our characterization suggests the following general requirements for new DL hardware designs:

- High memory bandwidth and capacity for embeddings
- Support for powerful matrix and vector engines
- Large on-chip memory for inference with small batches
- Support for half-precision floating-point computation

These requirements result from the characteristics of DL models important to us now (and projected to be in the future), our experience in optimizing DL applications for current computing platforms as well as their limitations found from our experiences. In particular, we highlight a gap in characteristics between the models commonly studied by the systems community and ones running in our data centers, with implications for future processor design.

2. Characterization of DL Inference

This section highlights characteristics of DL inference workloads that are of interest in our data centers. Section 2.1 describes DL models used in our social network services and discusses trends observed in their evolution over time. Section 2.2 presents detailed characteristics, focusing on aspects related to processor architecture design, and Section 2.3 details their common computational kernels.




Figure 2: A deep learning recommendation model

2.1. Representative Models

We divide inference workloads into three categories. The first provides personalized feed, ranking or recommendations, based on previous user interactions. The second and third are used for content understanding, visual and natural language content, respectively. The latter infer information used for powering recommendations, integrity and security such as detecting objectionable content.

2.1.1. Ranking and Recommendation

Recommendation systems are one of the most common DL workloads in data centers with many applications like ads, feed, and search. Recommendation is usually formulated as an event-probability prediction problem, where an ML model predicts the probability of one or multiple events at the same time. The items associated with the most likely events are ranked higher and shown to the user [28].

Without going into a comprehensive scientific literature review, we point out that over time the ML models and recommendation systems have evolved to incorporate neural networks (NNs). The latter has progressed from matrix and tensor-based factorizations [19, 37] to autoencoder and neural collaborative filtering [27, 41, 59]. Further advances led to the development of more complex models, such as wide and deep as well as deep cross neural networks, which have been successfully applied in practice [13, 26, 70, 76].

These models usually use a combination of signals from dense and sparse features. The former are represented as a vector of real values, while the latter are often represented as indices of an one-hot encoded vector in a high-dimensional space. The sparse features are processed with embedding lookups that project sparse indices to a lower dimensional space. As in Figure 2, the resulting embeddings are combined with the dense features to produce higher order interactions, for example using a set of fully connected layers (FCs) or parameter-less additive and multiplicative mixing [55].

The embedding tables can easily contain billions of parameters, while FCs usually have a modest number of parameters.

The size of these models is often bound by the memory of the system at hand and can easily require a memory capacity exceeding tens of GBs.

These models often have to predict event-probabilities for multiple ad or feed candidates for a single user, usually within 100s ms time constraint. These properties allow us to leverage batching to achieve high performance in FCs. However, the overall model’s execution tends to be memory bandwidth bound and is dominated by the embedding lookups. These look-ups perform a large number of mostly random accesses across table columns, but read an entire column vector for each such random access. For more details, refer to Sparse-LengthsSum operator in Caffe2.

Future Trends:

- **Model Exploration:** recent studies explore explicitly incorporating time into the event-probability models [7, 73]. We believe that such techniques will lead to better models in the future but require more compute demand.
- **Larger Embeddings:** Adding more sparse signals and increasing embedding dimensions tends to improve model quality. Therefore, we expect even larger embeddings to be used. This will further increase the pressure on memory and leads to systems with larger memory capacity, while putting more focus on distributed training and inference.

2.1.2. Computer Vision

CV models were some of the earliest to adopt DL techniques. They rely on convolutions that apply $C_1 \times K_1 \times K_2$ filters on the $B \times [F \times] C_1 \times \text{height} (H) \times \text{width} (W)$ input images with B batch size and C_1 input channels or video clip with F frames and produce a result with C_o output channels.

Image Classification involves matching images to classes. Currently, ResNets are widely used for classification [26]. However, recently much larger ResNeXt models have shown state-of-the-art accuracy even with weakly supervised training on billions of Instagram images [43, 74]. For example, our ResNeXt-101-32x4d model contains 43M parameters and requires 8B multiply-add operations during inference, relying on group convolutions with $G=32$ and $d=4^2$ in its first residual block. The largest configuration ResNeXt-101-32x48d contains 829M parameters and requires 153B multiply-add operations during inference, relying on group convolutions with $G=32$ and $d=48$ in its first residual block. It further improves the Top-1 validation accuracy of ImageNet-1K by 4% to 85.4% [43].

Object Detection involves identifying specific regions that contain objects of interest. One of the largest scale object detection systems running in data centers today is the text detection part of the Rosetta system used to understand text in

²In a group convolution, only the input channels in the same group are used for computing a given output channel. A group convolution with total C_1 input, C_o output channels and G groups is essentially G independent convolutions each with $d=C_1/G$ input and C_o/G output channels. A special case where $C_1=C_o=G$ and consequently group size $d=1$ is called depth-wise convolution.

images [8]. It uses the Faster-RCNN-Shuffle model, that relies on Faster-RCNN architecture [54], with the ResNet trunk replaced with ShuffleNet [75], which uses 3×3 depth-wise convolutions and 1×1 group convolutions with $d=4$. Despite ShuffleNet efficiency, object detection tends to be more time consuming than classification for the following reasons.

First, detection requires high resolution for accuracy. Whereas 224×224 images are sufficient for classification, detection typically resizes images such that the maximum side is 800 while maintaining the aspect ratio. Therefore, a typical input size of dimension $3 \times 800 \times 600$ for object detection is $9.5 \times$ larger than a typical input for classification.

Second, Faster-RCNN employs a region-proposal based approach where the final convolutional block is batched over many proposals of smaller spatial resolution. In Rosetta, the activations tend to be of dimensions [25-100 proposals] \times [544 or 1088 channels] \times [7,14] \times [7,14]. The spatial resolution is typically 7×7 or 14×14 , with large number of channels. Hence the number of proposals is a limiting factor in the number of objects that can be detected and is typically bounded due to computational cost.

Video Understanding historically has taken frame-based approach where sampled video frames are applied through image models. However, recently 3D convolutions gained wide adoption owing to higher accuracies given their ability to model temporal in addition to spatial domain [65]. Extensive studies have been done to analyze the performance vs. accuracy trade-off of vanilla 3D ResNets compared to factorized 3D convolutions as in Res(2+1)D [66], ResNeXt-3D and ShuffleNet-3D [3]. In particular, ResNeXt-3D with depth-wise convolutions, which factorizes the 3D convolution into channel and spatiotemporal dimension, requires $3 \times$ less FLOPs than Res(2+1)D, which factorizes the 3D convolution across spatial and temporal dimension. Further, trading off spatial resolution for increasing clip length shows improved accuracy. In the future, increasing both the temporal and spatial resolution would be important for more complex video understanding tasks, such as object detection or action recognition.

Future Trends:

- **Model Exploration:** There is an increasing trend to fine-tune the last few layers of a model specific to each application (such as adding additional categories for classification) while all applications share a common trunk. This leads to the inference cost per image increasing linearly as a factor of the computational cost of only the final few layers. These final convolutions typically have a large number of channels and work on much smaller spatial resolutions, which can be important optimization targets.
- **Convolution Types:** Group/depth-wise convolutions such as in ResNeXt and ShuffleNet, originally introduced for mobile inference, have increasingly been adopted in the data center due to accuracy and FLOP efficiency. Depth-wise convolutions are memory bandwidth bound, while majority of FLOPs spent elsewhere: e.g. ResNeXt-3D has

97.1% of all FLOPs in $1 \times 1 \times 1$ convolutions.

- **Large Activations:** Image understanding is moving beyond simple classification tasks into more complex domains such as object detection, action recognition, and human pose estimation, which requires larger activations. For instance, tasks like object detection require higher resolution for accuracy, and video inference with more frames per clip demonstrates higher accuracy due to more temporal context. More CV tasks will see this trend, adding pressure on on-chip memory capacity and/or off-chip memory bandwidth.
- **Batch Size:** Although CV inference typically does not have very strict latency requirements, small batches are still preferable in non-classification use cases. Whereas classification tasks perform well when the aspect ratios are distorted into a fixed shape like 224×224 , doing so results in huge accuracy drops in complex tasks like object detection, making batching difficult. Moreover, owing to large activations, increasing batch size puts further pressure on on-chip memory capacity.

2.1.3. Language Models Neural machine translation (NMT) has become the dominant approach to machine translation [5, 34, 47, 64]. It relies on the encoder-decoder approach, also called seq2seq. The former encodes the input sentence, while the latter decodes the encoding into the target output sentence.

This approach has been successfully applied to many natural language processing tasks, such as summarization [17, 46], speech recognition [23], syntactic and semantic parsing [11, 69] as well as question answering and dialog systems [9, 52].

Encoder-decoder approaches vary according to the encoder and decoder implementation. A major challenge with NMT is the dependence of a translation segment on its position in the sentence. This motivated the reliance on recurrent neural networks (RNNs) as one can encode the statement's position in the recurrent network during translation. This approach has shown successful results and is widely used in practice [5, 64]. In this approach, the encoder and the decoder are typically implemented using a Gated Recurrent Unit (GRU) [12] or a Long Short Term Memory (LSTM) cells [29].

Future Trends:

- **Model Exploration:** Results have shown that adding more layers and ensembles improves translation quality, but leads to larger NMT models [64]. Reranking the output of a model is a successful approach that can be used together with ensembles [60]. Also, multilingual models are an attractive way to scale one model to multiple languages but each multilingual model may need more capacity to handle multiple language pairs [32, 58].
- **Parallelism:** While successful, RNN-based approaches impose dependencies on each translated word, making it difficult to utilize highly parallel computing architectures such as GPUs. Recognizing this has motivated NMT models that lift the time dependencies imposed by RNNs. In [20], both the encoder and decoder are implemented as stacked convolutions. In [68], the transformer model is introduced

Category	Model Types	Model Size (# params)	Batch Size (typical)	Max. Live Activations	Arith. intensity (weights)	Arith. intensity (act. & weights)	Latency (constraints)
Recommendation	FCs	1–10M	1–100	>10K	20–200	20–200	10s of ms
	Embeddings	>10 Billion	1–100	>10K	1–2	1–2	10s of ms
Computer Vision	ResNet-50	25M	1 image	2M	avg. 303/min. 100	avg. 164/min. 25	No strict constraints
	ResNeXt-101-32x4-48	43–829M	1 image	2.4–29M	avg. 380/min. 100	avg. 188/min. 28	
	Faster-RCNN-Shuffle	6M	1 image	13.2M	avg. 3.5K/min. 2.5K	avg. 145/min. 4	
	ResNeXt3D-101	21M	1 clip	58M	avg. 22K/min. 2K	avg. 172/min. 6	
Language	seq2seq (GRU/LSTM)	100M-1B	1-8 tokens	>100K	2–20	2–20	10s of ms

Table 1: Resource requirements of representative DL inference workloads implemented on CPU. The batch size can often be increased with more compute throughput, while meeting latency requirements. We point out that 1 clip consists of 8-16 frames.

which removes the need for recurrence or convolution altogether and instead only relies on the attention mechanism to improve achievable hardware parallelism at the expense of additional computation. Results from this work show that NMT training time can be significantly reduced while having the additional benefit of model generality. While these approaches benefit from improved parallelism in both the encoder and the decoder during training and the encoder during inference, a typical inference generates an output sequentially using beam search. A more recent work has attempted to remove the time dependency in the decoder at inference time [24].

- **Batch Size:** Inference with small batches is well suited in the context of instant translation. However, large-scale inference can also be done offline for localization purposes. In that case, using larger batch sizes can be beneficial as throughput becomes more important than latency.

2.2. Compute Characteristics

Let the arithmetic intensity be defined as (# of operations needed to evaluate) / (# of elements incurred in data traffic) during the execution of model. The compute, memory capacity, and memory bandwidth demand of our representative DL inference workloads is shown in Table 1. We report two arithmetic intensities: (i) assuming only weights are incurring the traffic, for example when all activations fit in a level closer to compute in the memory hierarchy, and (ii) assuming that both weights and activations are incurring traffic.

For DL hardware designs, there are notable differences between DL workloads found in our representative sample and those commonly studied in the systems community.

First, *embedding layers stand out with huge model sizes (more than 10s of GBs) and significantly low arithmetic intensities*. Mathematically, the operation we perform on the embedding tables is a sparse-matrix times dense-matrix multiplication. The sparse matrix has >10 rows and >10M columns, each row with >10 non-zeros. This matrix is multiplied with a dense matrix with >10M rows and >10 columns.

The embedding lookup operation can be an interesting opportunity for applying emerging memory technologies and specialized hardware. On one hand, more expensive High-bandwidth memory (HBM) could be useful because it provides higher bandwidth but unfortunately its current capacity

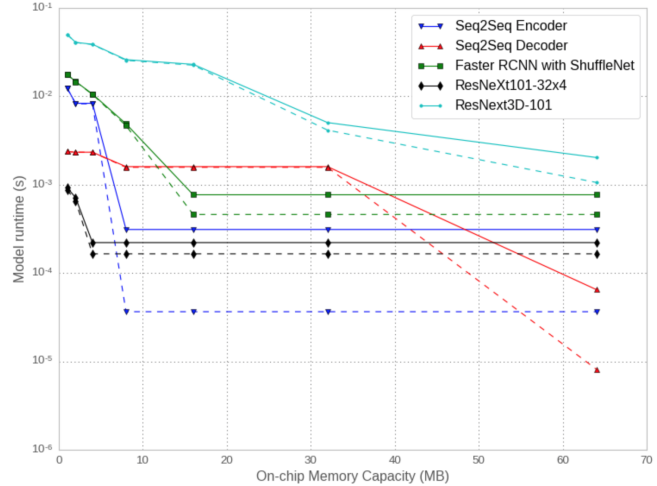


Figure 3: Runtime roofline analysis of different DL models with parameters stored as int8 numbers on a hypothetical accelerator with 100 TOP/s and 100 GB/s DRAM bandwidth. The performance is shown for varying on-chip memory capacity with 1 TB/s (solid) and 10 TB/s (dashed) bandwidth.

is limited. On the other hand, Non-volatile memory (NVM) is an economical alternative to store embeddings compared to DRAM, but the associated bandwidth is too low to be practical out of the box. Further, the memory access pattern to embedding tables has low temporal locality which makes caching challenging, while low spatial locality often results in under-utilization (due to access granularity of 10s of Bytes versus NVM block size). Nonetheless, several techniques have been proposed to mitigate these problems [16].

Second, *recent models can benefit more from larger on-chip memory capacity*. In a hypothetical accelerator with 100 TOP/s and 100 GB/s DRAM bandwidth the performance projected by a roofline model³ improves with larger on-chip memory capacities, as shown in Figure 3. This is not only driven by larger models, such as NMT seq2seq and ResNeXt-101, but also by larger activations, such as 800×600 input images for ShuffleNet and videos for ResNeXt-3D.

Notice that the FC layers in recommendation and NMT

³We assume that the model parameters are stored as int8 numbers. We apply a roofline model for each layer, where each layer differs in whether it reads activations/weights from off- or on-chip memory based on a simple greedy on-chip memory allocation algorithm [72].




Figure 4: Time spent in Caffe2 operators in our data centers.

models use small batch sizes so performance is bound by off-chip memory bandwidth unless parameters can fit on-chip. The batch size can be increased while maintaining latency with higher compute throughput of accelerators [33], but only up to a point due to other application requirements. The number of operations per weight in CV models are generally high, but the number of operations per activation is not as high (some layers in the ShuffleNet and ResNeXt-3D models are as low as 4 or 6). This is why the performance of ShuffleNet and ResNeXt-3D varies considerably with on-chip memory bandwidth as shown in Figure 3. Had we only considered their minimum 2K operations per weight, we would expect that 1 TB/s of on-chip memory is sufficient to saturate the peak 100 TOP/s compute throughput of the hypothetical accelerator. As the application would be compute bound with 1 TB/s of on-chip memory bandwidth, we would expect there to be no performance difference between 1 TB/s and 10 TB/s.

Third, *common primitive operations are not just canonical multiplications of square matrices, but often involve tall-and-skinny matrices or vectors*. These shapes arise from group/depth-wise convolutions that have recently become popular in CV, and from small batch sizes in Recommendation/NMT models due to their latency constraints. Therefore, it is desired to have a combination of matrix-matrix engines to execute the bulk of FLOPs from compute-intensive models in an energy-efficient manner and powerful enough vector engines to handle the rest.

2.3. Computation Kernels

Let us now illustrate the time spent in different computational kernels on CPU in our data centers. Figure 4 shows that FCs are the most time consuming operations, followed by embedding lookups and tensor manipulations⁴.

Following Caffe2 framework convention, the FC operator is defined as \mathbf{XW}^T , with $M \times K$ matrix \mathbf{X} and $N \times K$ matrix \mathbf{W} as inputs, and K being the inner reduction dimension. The convolutions can be logically transformed to matrix multiplications using `im2col`, which results in $M = BHW$, $N = C_o$ and $K = C_i K_1 K_2$ as shown in Figure 5.

We often refer to the number of rows M as effective batch size or batch/spatial dimension, while N as the output feature dimension. If M or N are small (e.g., Δ and \times corresponding to FC and group/depth-wise convolutions), the matrix-matrix multiplication becomes narrow and more closely resembles

⁴ “Tensor Manipulation” refers to concatenation (for combining dense and sparse features in Figure 2), splitting, slicing, and so on, which are good targets for whole graph optimizations discussed in Section 3.3.




Figure 5: Common activation and weight matrix shapes, with Δ : FCs, \times : group and depth-wise convolutions, \circ : other ops

matrix-vector multiplication, with performance deteriorating accordingly from BLAS3 to BLAS2 levels. In this case a matrix-matrix multiplication engine is expected to have a low utilization. This happens for small batch sizes (e.g., recommendation and NMT) and group convolution with few output channels per group (e.g., ShuffleNet).

The number of operations per weight read is proportional to the batch/spatial dimension, and the number of operations per activation read is proportional to the output feature dimension. If either is small, then performance is expected to be bound by memory bandwidth. When an $M \times K$ matrix activation matrix is multiplied with a $K \times N$ weight matrix, we compute $2MKN$ operations while reading KN weights, leading to $2M$ operations per weight. For example, when a batch/spatial dimension (M) is 10, the operations per weight is 20. In this case, if model parameters are stored as `int8` numbers then saturating 100 TOP/s architecture would require 5 TB/s of memory bandwidth for reading weights. Similarly, the number of operations per activation is $2N$. With an output feature dimension of 8, operations per activation of 16 would require 6.25 TB/s for reading input activations.

Note that the overall arithmetic intensity of a DL model can be misleading and we should also look at its individual layers. For example, even though the depth-wise convolutions in ShuffleNet and ResNeXt account for only 2% of total FLOPs, if a hypothetical accelerator can achieve 100 TOP/s for the

other convolutions and only 2 TOP/s for the depth-wise convolutions due to memory bandwidth limitations, time spent in the depth-wise convolutions will be comparable to the others.

Finally, we point out that standard benchmarks, like DeepBench [6], typically give more emphasis on batch sizes larger than what is encountered in our use cases. They do not capture small reduction dimensions in depth-wise convolutions, and big activation tensors in image detection and video models.

3. Performance Optimizations

DL inference workloads running on Facebook’s social network services need to serve billions of users with fluctuating capacity demand over time [25]. Therefore, the availability and flexibility of computing resources is important. In addition, many inference workloads require low latency. These are the reasons why, currently, most inference workloads run on CPUs. Even though accelerators can significantly reduce inference time and improve energy efficiency, CPUs will still be responsible for a good portion of DL inference, especially in cases where tight integration with business logic is needed.

3.1. Data Center Fleet-wide DL Inference Profiling

Our data-centers are running diverse DL inference workloads. Table 1 lists representative models, but by no means covers all of our models and new models with new types of data and varying tensor shapes are always coming online. Therefore, *it is important to continuously monitor DL model performance characteristics fleet wide*. DL operations typically utilize a large fraction of peak compute or memory bandwidth, depending on their arithmetic intensity, and are less limited by memory latency or control overheads compared to other typical data center workloads. They often involve regular compute and memory access patterns, lending themselves as good targets of analytical performance models.

For this purpose we have implemented the observer software design pattern that can be applied to individual operators and are executed at the start and end of the operator. We have developed a number of functions called by observers that track performance metrics for each operator’s execution (refer to Caffe2 operator cost inference functions for more details). When considered in conjunction with the layer’s full specification such as layer type, input/output tensor shapes, and element types, we can understand whether a given layer execution should be memory-bandwidth or compute bound. Viewed differently, we can estimate the benefits of optimizing any specific operator. This is particularly useful as it gives us the necessary data to estimate the priority of a considered optimization.

In order to keep track of the accuracy and identify inefficiencies in the roofline models we maintain detailed per-layer logs that measure execution time, memory bandwidth in GB/s and actual attained FLOP/s that are derived from hardware performance counters for sampled DL operator executions. A telemetry agent running on each host collects and compares

this information with given predictions across all of our data centers. Also, to set realistic goals for our optimization efforts, we developed a number of benchmarks tuned for each potential bottleneck.

3.2. Reduced Precision Inference

The reduced-precision inference has been shown to be effective at improving compute throughput within a power budget, especially in mobile platforms. However, applying reduced-precision inference in data centers is nontrivial.

First, while mobile platforms have widely adopted CV models such as ShuffleNet and MobileNet that trade-off accuracy for significant reduction in compute requirements [30, 75], DL inference in data centers prefers accurate but compute intensive models like ResNet [26] and ResNeXt [74]. In particular, when DL inference is related to core services like feed or integrity/security the accuracy loss should be very small. Usually $<1\%$ change in the accuracy compared with single-precision floating-point results is acceptable.

Also, while general purpose CPUs have high availability in data-centers, they have not yet adapted to rapidly increasing compute demand of DL inference and hence lack good support for high-performance reduced-precision inference. This is exacerbated by less mature high-performance and high-accuracy reduced-precision linear algebra libraries for CPUs compared to their higher precision counter parts.

3.2.1. Performance Challenges Current generations of x86 processors [31] provide conversion instructions between half- and single-precision floating point numbers (`vcvtps2ps` and `vcvtps2ph`), but without native half-float (`fp16`) computation. They also require a sequence of instructions (`vpmaddubsw + vpmaddwd + vpadd`) to implement 8-bit integer multiplications with 32-bit accumulation with marginally higher ($\sim 33\%$) compute throughput than that of single-precision floating point (`fp32`) [56]. The compute throughput of 8-bit integer multiplications with 16-bit accumulation can be about twice higher than `fp32`, but this often results in significant accuracy drops unless combined with outlier-aware quantization that will be described shortly. On the other hand, VNNI instructions provide higher throughput `int8` multiplications with 32-bit accumulation but they are not available in current x86 microarchitectures [71]. As a result, we had to tailor optimization strategies based on the performance bottleneck.

If the performance is memory-bandwidth bound, then using `fp16` when storing weights or using 8-bit multiplications with 32-bit accumulation (`i8-acc32`) can increase the arithmetic intensity by up to a factor of $2\times$ and $4\times$, respectively. In this case, we can obtain speedups proportional to the memory bandwidth saving, even when we save nothing with respect to the number of instructions. For example this happens in FCs with small batch sizes and group convolutions with a small number of channels per group (the extreme case being depth-wise convolution with just one channel per group).

We have designed and implemented a reduced-precision lin-

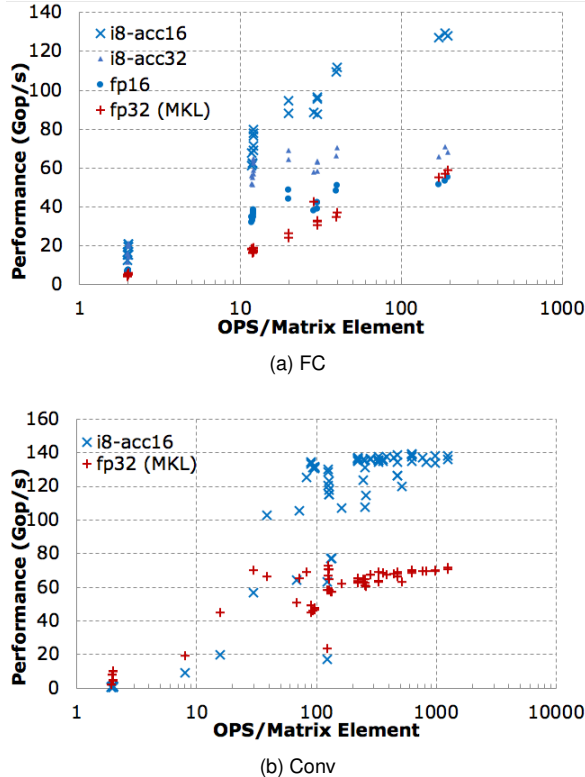


Figure 6: Performance of FBGEMM in Gop/s vs. arithmetic intensity ($2NMK/(NK + MK)$) for multiplications of $M \times K$ and $K \times N$ matrices, compared with MKL GEMM in fp32.

ear algebra library for DL inference called FBGEMM [35, 36]. Figure 6(a) plots the performance of our optimized fp16 and i8-acc32 matrix multiplication (GEMM) in FBGEMM compared with Intel MKL’s fp32 GEMM. The experiments are performed on a single thread running on Intel Xeon E5-2680 v4 with turbo mode off using Intel MKL version 2017 update 3. Notice that for cases with low arithmetic intensity our fp16 and i8-acc32 GEMMs obtain up to $2\times$ and $4\times$ speedups over MKL’s fp32 GEMM, respectively. For instance, applying our fp16 GEMM, we obtain up to $2\times$ speedup in FC layers in a recommendation model with 15% overall latency reduction. Also, applying our i8-acc32 GEMM, we obtain overall $2.4\times$ speedup in the Faster-RCNN-Shuffle used for our optical character recognition application.

If the performance is bound by the instruction throughput, then we use 8-bit multiplications with 16-bit accumulation and periodic spills to 32-bit accumulators (i8-acc16), which can provide $\sim 2\times$ compute throughput over fp32. To avoid saturation and accuracy drops, we employ outlier-aware quantization that separates out weights with bigger magnitude as outliers [50]. Here, we consider a typical threshold for outliers, where a weight is not an outlier if representable with 7 bits (i.e. the value of weight is between -64 and 63). We split the weight matrix into two parts, $\mathbf{W} = \mathbf{W}_{main} + \mathbf{W}_{outlier}$, where \mathbf{W}_{main} is in 7 bits and $\mathbf{W}_{outlier}$ contains the residual. The matrix multiplication, \mathbf{XW}^T , is calculated in two stages, where \mathbf{XW}_{main}^T

uses 16-bit accumulation, and $\mathbf{XW}_{outlier}^T$ uses 32-bit accumulation. We find that $\mathbf{W}_{outlier}$ becomes a sparse matrix, often with density less than 0.1%, especially when combined with symmetric quantization [39]. Since $\mathbf{W}_{outlier}$ is sparse, $\mathbf{XW}_{outlier}^T$ accounts for a small fraction of the total time. Figure 6(b) plots the performance of our i8-acc16 GEMM compared with MKL GEMM in fp32, which achieves up to $2\times$ speedup for matrix shapes with high arithmetic intensity. In particular, applying our i8-acc16 GEMM to ResNet-50, we obtain $1.7\times$ speedup over MKL in fp32.

Even though some of the applied optimizations are done to work around limitations of current x86 processors, they provide insight for future DL hardware optimizations. Our optimizations show it is useful to apply different quantization techniques depending on where the performance bottleneck lies. For example, quantization techniques that are primarily for saving storage and bandwidth should be tested with embedding layers, FCs with small batch size, and depth-wise convolutions. Our experience with outlier-aware quantization shows that a high-performance sparse linear algebra engine will be helpful not only for pruned models but also for reducing required precision of non-pruned models. For example, 6-bit quantized models can be computed in 4-bit for main values while the outlier part is computed with the 6-bit sparse engine.

3.2.2. Accuracy Challenges: Impressive progress has been made in low-precision DL inference, some of which consider even ternary or binary quantization [53, 77]. However, even 8-bit quantization has presented its own set of challenges to meet the accuracy requirements of our DL workloads in data centers. The following five techniques were effective at meeting the accuracy requirements:

1. *Fine-grain Quantization.* Instead of having a single quantization parameter per tensor, applying quantization in a finer granularity is often required. Examples are per output feature quantization in FCs, per output channel quantization in convolutions, per group quantization in group convolutions, or per-entry quantization in embedding tables.
2. *Quantization-aware Training.* We found that quantization-aware training for example using fake quantization is important for meeting the accuracy requirements. This aligns with a recent white paper [39] that shows the importance of per-channel quantization and quantization-aware training in quantizing CNNs for mobile platforms.
3. *Selective Quantization.* Unlike mobile platforms which can highly prefer end-to-end quantization, DL inference in data centers should be able to fall back to floating-point in accuracy-sensitive parts of DL models. We systematically profile errors introduced by quantization per layer and skip quantization when the error is too high. Examples include the first and last few layers of CNNs.
4. *Outlier-aware Quantization.* In addition to the outlier-aware quantization technique described previously for 16-bit accumulation, we can take advantage of the fact that the

range of values can be confined much more once outliers are ignored. For example, instead of quantizing a tensor \mathbf{W} for the range $[\min(\mathbf{W}), \max(\mathbf{W})]$, we can quantize for a smaller range, such that the **L2 norm of quantization error** is minimized with respect to the distribution of A values. Unlike weight tensors, activation tensors are not constant, so we collect distribution of activation tensors by running with calibration inputs from the training data.

5. *Net-aware Quantization*. We can often further reduce the range we’re quantizing for based on neighboring operators. For example, if an operator is only followed by ReLU, we can narrow down the range by excluding negative values.

For instance, using these techniques, **a ResNet-50 model with int8 quantization** (except softmax) achieves 75.6% Top-1 and 92.8% Top-5 accuracy for ImageNet-1K validation set [15], which corresponds to only 0.3% Top-1 and 0.1% Top-5 accuracy drop compared to the baseline fp32 model [22].

3.2.3. Software Challenges: Linear algebra operations for machine learning inference require optimizations that are quite different from those for high-performance scientific computing (i.e. HPC). The standard BLAS interface cannot provide the desired performance for the matrix shapes that are common in DL inference. Since the compute requirement in DL is rapidly changing, it can be premature to attempt to standardize a new linear algebra interface for DL, but it worthwhile to discuss the associated requirements and challenges.

As shown in Figure 5, typical matrix shapes in DL inference are smaller and often tall and skinny, compared to those in typical HPC applications. High-performance matrix multiplications often “pack” a block of input matrices into a format friendly for vectorization and cache locality. For large enough square matrices, the overhead of packing can be amortized inside a single matrix multiplication adhering to the standard BLAS interface. However, for tall-skinny matrices, we need to amortize the packing overhead across multiple matrix multiplications for constant weight matrices which requires a new interface that accepts a custom pre-packed matrix.

A significant fraction of DL computation is not strictly matrix multiplication. For example, the convolution operator in CNNs can be expressed as `im2col` followed by matrix multiplication, but this often does not lead to the highest performance due to the duplication of input activations and the overhead of `im2col`. Therefore, it is important to promote convolution as a first-class citizen of the interface to enable the computation of direct convolutions without `im2col`. This will also enable algorithmic optimizations such as Winograd or FFT-based convolution as in cuDNN with automatic choice of the best algorithm for given tensor shapes. The native convolution interface is particularly important for group convolution with only a few channels per group. If we individually apply `im2col` followed by GEMM for each group, the reduction dimension and the output feature dimension are too small for efficient vectorization and parallelization. Note that even the FC layer cannot be implemented strictly with only a GEMM

operation as it involves a bias term which should be fused with the main GEMM to save memory bandwidth. It is also desirable to fuse other common operations such as ReLU.

Reduced-precision fixed-point computation requires additional steps such as handling non-zero offsets used in asymmetric quantization and rescaling 32-bit intermediate results of matrix multiplication, which should be fused with the main GEMM to save bandwidth. Google’s `gemmlowp` library [21] provides a well-designed interface of fusing “output pipeline” with the main GEMM. However, `gemmlowp` doesn’t provide native convolution interface and is mostly optimized for ARM Neon and Intel x86 SSE4.1, not for AVX2 and AVX-512.

Intel MKL-DNN is another library that provides high performance implementations of DL primitives on CPU. MKL-DNN implements advanced features such as Winograd convolution in int8. On the other hand, FBGEMM has features such as outlier-aware quantization. Therefore, some of our DL inference applications use FBGEMM and some others use MKL-DNN, depending on compute characteristics and operator availability. Low-precision linear algebra for DL inference is still a young field, and we believe it is healthy to have multiple complementary solutions that can experiment different optimizations while adopting proven techniques from each other.

The below code snippet shows an example of our FBGEMM library interface. In this example, a templated C++ function that performs a matrix multiplication for different data types is shown. The salient features of this interface are the way it accepts a packed B matrix (usually weights that can be packed once and used multiple times) and also a parameter for packing matrix A. The packing of matrix A can be specialized and fused with memory bandwidth bound operations such as `im2col`, row-wise sum for asymmetric quantization, or depth-wise convolution. `outProcess` parameter is templated to support various types of data processing on output once the main GEMM part is finished (similar to `gemmlowp`’s output pipeline). As previously mentioned, many matrices in DL inference are tall-skinny so the main kernels of matrix multiplication are dynamically generated just-in-time to take advantage of matrix size specific optimizations. The FBGEMM library is **open source** and integrated with Caffe2 deep learning framework. For more complete examples, refer to the tests and benchmarks in our open source project.

```
template<typename T_PACK_A, typename T_PACK_B,
        typename T_C, typename OUT_FUNCTOR>
void gemmPacked(
    // packed inputs
    T_PACK_A& packA, T_PACK_B& packedB,
    // output
    T_C* C, uint32_t ldc,
    // post-processing functor, e.g. Relu
    OUT_FUNCTOR& outProcess);
```


3.3. Whole Graph Optimization

While it is important to optimize the performance of individual operators as outlined in the previous subsections, we can get additional significant performance improvements by looking at the DL graph as a whole and performing cross-operation optimizations. A few different optimizations fall into this category, including operator fusion, data movement elimination, operator scheduling, and threading for both inter- and intra-parallelism. This section focuses on operator fusion, specifically quantifying potential speedups of operator fusion. The realized speedup from operator fusion will heavily depend on the efficiency of underlying fused kernel. Automatic generation of fused kernels is an active area of research and early productization efforts are underway [10, 42, 57, 67]. However, it is still often necessary to write fused kernels manually. For this reason, we focus our efforts in two directions: 1) to find the top few opportunities where we will get the most gains from fusion for our models that can be worth manual attention and 2) to find a broader set of opportunities for compiler generated kernels.

Our approach to identify fusion opportunities for both cases is similar. We aim at identifying subgraphs that occur commonly in our workloads across the entire fleet; and are expected to have high speedup potentials. We log the complete graphs annotated with operator dependencies, frequency, and input/output tensor shapes. We then run a frequent subgraph mining algorithm on the nets captured. The idea here is to find all subgraphs that are executed frequent enough and order them on the basis of speedup potential from fusion. To perform the ordering, we use the input/output dimensions for the operators to compute a simple roofline model for the subgraph being considered. Specifically, we compute performance projected by the roofline model before and after fusion, and use the difference to estimate speedup potential. Note that we filter out some subgraphs based on specific operator pattern rules. For example, we rule out subgraphs with operators that are not data parallel and hence challenging to fuse. Finally, we run a top-k algorithm on the ordered subgraphs to return the top opportunities.

With this analysis, we were able to find several opportunities for merging batched matrix multiplies with tensor manipulation operations. As analyzed in Figure 4, these tensor manipulation operations comprise about 17% of the overall DL inference CPU time. Most of these operations are memory bandwidth limited; merging them with compute bound operations resulted in a total of over 10% savings in run time.

4. Application Driven HW Co-design Directions

This section discusses implications of the DL model characteristics and their optimization for software and hardware co-design. We believe that the server-side DL workload optimizations should be considered as a co-design problem along three axes: DL models, numerics (quantization, Winograd/FFT con-

volution, and sparsity), and hardware platforms. Also, the process should be driven by DL models because of their rapid changes and diversity. We highlight a few relevant observations in this regard next.

Workload Diversity: DL is a fast moving field while the design space of inference hardware is huge. Therefore, one needs a fast turn-around loop with performance modeling capability to predict benefits of various hardware and software co-optimizations based on workload characteristics captured from a wide range of up-to-date DL models. This study reveals the following characteristics of DL models running in our data centers. First, they have diverse compute patterns where matrices do not necessarily have “nice” square shapes. There are also many “long-tail” operators other than FC and convolutional layers. Therefore, in addition to matrix multiplication engines, hardware designers should consider general and powerful vector engines. Second, DL models in our data centers have diverse and sometimes conflicting demands to memory subsystem. For example, due to larger activation matrices or matrices with tall-and-skinny shapes, recent CV and NMT models need bigger on-chip memory capacity to sustain high compute throughput without being bottlenecked by off-chip memory bandwidth. However, we should not solely rely on on-chip capacity to fit the entire model because it is difficult to project on-chip memory capacity demand of future models. Some of our current recommendation models are already too big to fit on-chip memory. Recommendation models not only require a huge memory capacity but also high bandwidth.

Data Center Requirements: When co-designing inference hardware for data centers, it is important to keep in mind that data center server-side DL inference has different requirements from mobile/embedded/IoT devices. For example, some quantization and pruning techniques report 2–3% accuracy drops but that is often too high for data center environment and they are often not deployed. If quantization drops the accuracy of say 32x32d model by more than 1% with less than 2x speedup, it can be more advantageous to just use the 32x16d model without quantization. In order to minimize accuracy drops from quantization, inference hardware for data centers should support per-channel quantization. They also should support fp16 or bfloat16 compute as a fallback in accuracy sensitive parts such as the last layer of some DL models.

Service Dis-aggregation: DL applications have distinctive compute and memory characteristics, compared to other typical data center workloads. Specifically, DL inference often utilizes a higher fraction of peak FLOPs, memory capacity, and bandwidth. As a result, other jobs on the same machine can suffer memory capacity and bandwidth pressure, or power limitation, e.g. reduction in turbo frequency when AVX2 or AVX-512 is used by deep learning workload [38]. This reduces the performance of other important components such as business logic and has detrimental effect on full system performance. Hence, a natural decision is to dis-aggregate DL inference into a separate tier (accelerated or not). Dis-aggregation

can also allow to pool requests from many front-end servers, increasing the batch size and hence compute efficiency. A challenge is that inference queries and results need to be transferred between the tiers over the network. Thus, the tier design, network bandwidth and latency, and compression techniques need to be carefully considered. For example, a hypothetical accelerator with 100 TOP/s compute throughput would require a few GB/s PCIe and/or network bandwidth for the DL models listed in Table 1, unless image decompression can be done within the accelerator or on the same host.

DL Model and Hardware Co-design: It is important to co-design DL models to be aware of the cost associated with the required hardware resources. While power/energy budget, on-chip memory capacity, and off-chip memory bandwidth are typically more scarce resources, research on efficient DL model architectures often only optimizes the number of floating-point operations. When the performance is bandwidth bound, adding more FLOPs without increasing the bandwidth consumption can be a good way to improve the accuracy while maintaining the performance. If adding $2\times$ FLOPs to the FC part of a recommendation model and increasing the embedding dimension of its embedding table by $2\times$ provide similar accuracy improvements, we would expect adding FLOPs will be the more economical direction. Recovering accuracy losses from quantization by making DL models wider is an example of hardware cost aware trade-offs: int8 multiplication consumes more than $5\times$ less energy and area compared to fp16 multiplication, hence there is a big room to recover the accuracy while maintaining the energy savings [14, 44]. NMT models with higher arithmetic intensity and parallelism such as the transformer architecture also illustrate hardware cost aware trade-offs.

5. Related Work

Recently, Hazelwood et al. presented a holistic characterization of ML workloads in data centers, covering inference, training, data acquisition and including a discussion of their diversity, huge data and compute capacity demands [25]. In contrast, our paper aims to provide insights that are useful for software/hardware co-design of DL applications, focusing on DL inference characteristics.

Hardware accelerators for server ML workloads have been actively studied by academia and industry, with NVIDIA GPUs, Google TPUs and Microsoft Brainwave computing platforms being successfully used in data centers [18, 33, 48]. In particular, Google TPU relies on a systolic array accelerator mainly targeted for 8-bit matrix-matrix multiplication which is challenging to utilize for small batches and group/depth-wise convolutions. On the other hand, Microsoft Brainwave is a matrix-vector accelerator for low latency AI applications in data centers. It consists of dot-product engines which perform, with the broadcast vector and its local matrix weights, dot-product operations in parallel. The salient features of Brainwave are model pinning and block floating point representa-

tion. The large on-chip memory of FPGA is exploited to store weights on chip, avoiding off-chip DRAM accesses. Block floating point offers low precision computation by enabling 4- or 5-bit multiplications of mantissa and 5-bit additions of shared exponents. However, it is not clear if architectures like Brainwave are general enough to efficiently target our diverse DL inference workloads

Moreover, a number of techniques has been proposed to improve energy efficiency of DL inference by taking advantage of reduced precision and sparsity. NVIDIA’s SCNN skips computation with zero input in matrix multiplications, thereby offering significant improvements in energy efficiency [49]. Akhlaghi et al. propose early stopping of convolution when the output is expected to be non-positive followed by ReLU [1]. Sharma et al. present a systolic array accelerator called Bit-Fusion which supports variable bit precision [61]. Park et al. present a zero-aware 4-bit accelerator called OLAcel which applies reduced precision to the majority of data while keeping a small fraction of large value data in high precision [50], a technique also used in the optimizations described in this paper. Fleischer et al. propose a multi-TOP/s AI core supporting a wide range of precision from fp16 (for training) to 1- or 2-bit (for inference) [62]. Our paper shows that, while low-precision and sparse computation can significantly improve energy efficiency, they should meet the accuracy requirements of server-side DL inference to be widely used in data centers.

Finally, we point out that a number of DL benchmarks are actively being developed [45, 63]. A benchmark framework has been presented where a model zoo of benchmark neural networks is provided and the performance of neural networks, optimized by users, is measured on real mobile devices remotely [63]. MLPerf aims at providing benchmarks for both server and mobile devices [45]. These benchmarks will facilitate system-level performance measurements and comparisons on diverse software platforms like TensorFlow [42] and PyTorch [51] as well as hardware architectures.

6. Conclusion

In the face of rapid innovation in deep learning and the increase of their computation and memory requirements, co-designing DL inference hardware for current and future DL models is an important but challenging problem. We believe our DL inference characterization and optimization experience can provide useful insights for DL inference hardware designs. We hope our paper can also contribute to discussion around software ecosystem such as benchmarking suites, linear algebra interface optimized for DL, and the compiler for optimizing and scheduling the whole graph of DL models, which are important parts of co-design process.

7. Acknowledgements

We would like to thank AML, Caffe2 and Glow team members for help with collecting information and reviewing this study.

References

- [1] Vahideh Aklaghi, Amir Yazdanbakhsh, Kambiz Samadi, Hadi Esmaeilzadeh, and Rajesh K. Gupta. Snapea: Predictive early activation for reducing computation in deep convolutional neural networks. In *ACM/IEEE Int. Symposium on Computer Architecture*, 2018.
- [2] Dario Amodei and Danny Hernandez. AI and Compute. <https://blog.openai.com/ai-and-compute/>, OpenAI, 2018.
- [3] Anonymous. An empirical study of groupwise convolution for video classification networks. In *Under review*, 2018.
- [4] Maria Avgerinou, Paolo Bertoldi, and Luca Castellazzi. Trends in data centre energy consumption under the european code of conduct for data centre energy efficiency. *Energies*, 10(10):1470, 2017.
- [5] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv:1409.0473*, 2014.
- [6] Baidu. DeepBench. <https://svail.github.io/DeepBench/>.
- [7] Alex Beutel, Paul Covington, Sagar Jain, Can Xu, Jia Li, Vince Gatto, and Ed H Chi. Latent Cross: Making Use of Context in Recurrent Recommender Systems. In *Proc. 11th ACM Int. Conf. Web Search and Data Mining*, pages 46–54, 2018.
- [8] Fedor Borisjuk, Albert Gordo, and Viswanath Sivakumar. Rosetta: Large scale system for text detection and recognition in images. In *Proc. 24th ACM SIGKDD Int. Conf. Knowledge Discovery & Data Mining*, pages 71–79, 2018.
- [9] Daniel Braun, Adrian Hernandez-Mendez, Florian Matthes, and Manfred Langen. Evaluating Natural Language Understanding Services for Conversational Question Answering Systems. In *Proc. 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 174–185, 2017.
- [10] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Haichen Shen, Eddie Yan, Leyuan Wang, Yuwei Hu, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. TVM: end-to-end compilation stack for deep learning. In *Proc. SysML*, 2018.
- [11] Jianpeng Cheng, Siva Reddy, Vijay Saraswat, and Mirella Lapata. Learning structured natural language representations for semantic parsing. In *Proc. 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 44–55, July 2017.
- [12] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv:1406.1078*, 2014.
- [13] Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for youtube recommendations. In *Proc. 10th ACM Conf. Recommender Systems*, pages 191–198, 2016.
- [14] William Dally. High-performance hardware for machine learning. *NIPS Tutorial*, 2015.
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [16] Assaf Eisenman, Maxim Naumov, Darryl Gardner, Misha Smelyanskiy, Sergey Pupyrev, Kim Hazelwood, Asaf Cidon, and Sachin Katti. Bandana: Using non-volatile memory for storing deep learning models. *arXiv:1811.05922*, 2018.
- [17] Angela Fan, David Grangier, and Michael Auli. Controllable abstractive summarization. *arXiv:1711.05217*, 2017.
- [18] Jeremy Fowers, Kalin Ovtcharov, Michael Papamichael, Todd Massengill, Ming Liu, Daniel Lo, Shlomi Alkalay, Michael Haselman, Logan Adams, Mahdi Ghandi Stephen Heil, Prerak Patel, Adam Sapek, Gabriel Weisz, Lisa Woods, Sitararam Lanka, Steven K. Reinhardt, Adrian M., Caulfield Eric, S. Chung, and Doug Burger. A Configurable Cloud-Scale DNN Processor for Real-Time AI. In *ACM/IEEE Int. Symposium on Computer Architecture*, 2018.
- [19] Evgeny Frolov and Ivan Oseledets. Tensor methods and recommender systems. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 7(3):e1201, 2017.
- [20] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. *arXiv:1705.03122*, 2017.
- [21] Google. Low-precision matrix multiplication. <https://github.com/google/gemmlowp/>.
- [22] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyröla, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch SGD: training imagenet in 1 hour. *arXiv:1706.02677*, 2017.
- [23] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *IEEE Int. Conf. Acoustics, Speech and Signal Processing*, pages 6645–6649, 2013.
- [24] Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. Non-autoregressive neural machine translation. *arXiv:1711.02281*, 2017.
- [25] Kim Hazelwood, Sarah Bird, David Brooks, Soumith Chintala, Utku Diril, Dmytro Dzhulgakov, Mohamed Fawzy, Bill Jia, Yangqing Jia, Aditya Kalro, James Law, Kevin Lee, Jason Lu, Pieter Noordhuis, Misha Smelyanskiy, Liang Xiong, and Xiaodong Wang. Applied Machine Learning at Facebook: A Datacenter Infrastructure Perspective. In *IEEE Int. Symposium on High Performance Computer Architecture*, pages 620–629, 2018.
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [27] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *Proc. 26th Int. Conf. World Wide Web*, pages 173–182, 2017.
- [28] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, and Joaquin Quiñero Candela. Practical lessons from predicting clicks on ads at facebook. In *Proc. Eighth Int. Workshop on Data Mining for Online Advertising*, pages 1–9, 2014.
- [29] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [30] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv:1704.04861*, 2017.
- [31] Intel 64 and IA-32 Architectures Optimization Manual. <https://software.intel.com/en-us/download/intel-64-and/ia-32-architectures-optimization-reference-manual>.
- [32] Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. Google’s multilingual neural machine translation system: enabling zero-shot translation. *arXiv:1611.04558*, 2016.
- [33] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre Iuc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gullett, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Ni, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. In-datacenter performance analysis of a tensor processing unit. In *ACM/IEEE Int. Symposium on Computer Architecture*, pages 1–12, 2017.
- [34] Nal Kalchbrenner and Phil Blunsom. Recurrent continuous translation models. In *Proc. 2013 Conf. Empirical Methods in Natural Language Processing*, pages 1700–1709, 2013.
- [35] Daya Khudia, Protonu Basu, and Summer Deng. Open-sourcing FBGEMM for state-of-the-art server-side inference. <https://code>.

- fb.com/ml-applications/fbgemm/, Facebook, 2018.
- [36] Daya Khudia, Protonu Basu, Summer Deng, Jianyu Huang, Haixian Liu, Jongsoo Park, and Mikhail Smelyanskiy. Facebook GEMM library. <https://github.com/pytorch/fbgemm>, Facebook, 2018.
- [37] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, (8):30–37, 2009.
- [38] Vlad Krasnov. On the danger of Intel’s frequency scaling. <https://blog.cloudflare.com/on-the-dangers-of-intels-frequency-scaling/>.
- [39] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv:1806.08342*, 2018.
- [40] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [41] Oleksii Kuchaiev and Boris Ginsburg. Training deep autoencoders for collaborative filtering. *arXiv:1708.01715*, 2017.
- [42] Chris Leary and Todd Wang. XLA: TensorFlow, compiled. *TensorFlow Dev Summit*, 2017.
- [43] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. *arXiv:1805.00932*, 2018.
- [44] Asit Mishra, Eriko Nurvitadhi, Jeffrey J Cook, and Debbie Marr. WRPN: wide reduced-precision networks. *arXiv:1709.01134*, 2017.
- [45] MLPerf. <https://mlperf.org>.
- [46] Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, and Bing Xiang. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv:1602.06023*, 2016.
- [47] Graham Neubig. Neural machine translation and sequence-to-sequence models: A tutorial. *ArXiv:1703.01619*, 2017.
- [48] NVIDIA. Turing architecture whitepaper. 2018.
- [49] Angshuman Parashar, Minsoo Rhu, Anurag Mukkara, Antonio Puglielli, Rangharajan Venkatesan, Bruce Khailany, Joel Emer, Stephen W Keckler, and William J Dally. SCNN: An accelerator for compressed-sparse convolutional neural networks. In *ACM/IEEE Int. Symposium on Computer Architecture*, 2017.
- [50] Eunhyuk Park, Dongyoung Kim, and Sungjoo Yoo. An Energy-Efficient Neural Network Accelerator based on Outlier-Aware Low Precision Computation. In *ACM/IEEE Int. Symposium on Computer Architecture*, 2018.
- [51] Adam Paszke, Sam Gross, Soumith Chintala, and Gregory Chanan. Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration, 2017.
- [52] Holger Quast and Robert Bosch. *Speech Dialogue Systems and Natural Language Processing*, pages 67–106. Springer Berlin Heidelberg, 2004.
- [53] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. XNOR-Net: Imagenet classification using binary convolutional neural networks. In *European Conf. Computer Vision*, pages 525–542, 2016.
- [54] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015.
- [55] Steffen Rendle. Factorization Machines. In *IEEE Int. Conf. Data Mining*, 2010.
- [56] Andres Rodriguez, Eden Segal, Etay Meiri, Evarist Fomenko, Young Jin Kim, Haihao Shen, and Barukh Ziv. Lower Numerical Precision Deep Learning Inference and Training. <https://software.intel.com/en-us/articles/lower-numerical-precision-deep-learning-inference-and-training>.
- [57] Nadav Rotem, Jordan Fix, Saleem Abdulrasool, Summer Deng, Roman Dzhabarov, James Hegeman, Roman Levenstein, Bert Maher, Satish Nadathur, Jakob Olesen, Jongsoo Park, Artem Rakhov, and Misha Smelyanskiy. Glow: Graph lowering compiler techniques for neural networks. *arXiv:1805.00907*, 2018.
- [58] Holger Schwenk and Matthijs Douze. Learning joint multi-lingual sentence representations with neural machine translation. *arXiv:1704.04154*, 2017.
- [59] Suvash Sedhain, Aditya Krishna Menon, Scott Sanner, and Lexing Xie. Autorec: Autoencoders meet collaborative filtering. In *Proc. 24th Int. Conf. World Wide Web*, pages 111–112, 2015.
- [60] Rico Sennrich, Barry Haddow, and Alexandra Birch. Edinburgh neural machine translation systems for WMT 16. In *Proc. 1st Conf. Machine Translation*, pages 371–376, August 2016.
- [61] Hardik Sharma, Jongse Park, Naveen Suda, Liangzhen Lai, Benson Chau, Joon Kyung Kim, Vikas Chandra, and Hadi Esmaeilzadeh. Bit Fusion: Bit-Level Dynamically Composable Architecture for Accelerating Deep Neural Networks. *arXiv:1712.01507*, 2017.
- [62] Vijayalakshmi Srinivasan, Bruce Fleischer, Sunil Shukla, Matthew Ziegler, Joel Silberman, Jinwook Oh, Jungwook Choi, Silvia Mueller, Ankur Agrawal, Tina Babinsky, Nianzheng Cao, Chia-Yu Chen, Pierce Chuang, Thomas Fox, George Gristede, Michael Guillorn, Howard Haynie, Michael Klaiber, Dongsoo Lee, Shih-Hsien Lo, Gary Maier, Michael Scheuermann, Swagath Venkataramani, Christos Vezyrtzis, Naigang Wang, Fanchieh Yee, Ching Zhou, Pong-Fei Lu, Brian Curran, Leland Chang, and Kailash Gopalakrishnan. A Scalable Multi-TeraOPS Deep Learning Processor Core for AI Training and Inference. In *Proc. VLSI Symposium*, 2018.
- [63] Fei Sun. Benchmark Data Driven Co-design for Neural Network Applications. In *Design Automation Conf.*, 2018.
- [64] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, pages 3104–3112, 2014.
- [65] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proc. IEEE Int. Conf. Computer Vision*, pages 4489–4497, 2015.
- [66] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A Closer Look at Spatiotemporal Convolutions for Action Recognition. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 6450–6459, 2018.
- [67] Nicolas Vasilache, Oleksandr Zinenko, Theodoros Theodoridis, Priya Goyal, Zachary DeVito, William S Moses, Sven Verdoolaege, Andrew Adams, and Albert Cohen. Tensor comprehensions: Framework-agnostic high-performance machine learning abstractions. *arXiv:1802.04730*, 2018.
- [68] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [69] Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. Grammar as a foreign language. In *Advances in Neural Information Processing Systems*, pages 2773–2781, 2015.
- [70] Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. Deep & cross network for ad click predictions. In *Proc. ADKDD*, page 12, 2017.
- [71] WikiChip. Cascade Lake Microarchitecture. https://en.wikichip.org/wiki/intel/microarchitectures/cascade_lake.
- [72] Samuel Williams, Andrew Waterman, and David Patterson. Roofline: an insightful visual performance model for multicore architectures. *Communications of the ACM*, 52(4):65–76, 2009.
- [73] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J Smola, and How Jing. Recurrent recommender networks. In *Proc. 10th ACM Int. Conf. Web Search and Data Mining*, pages 495–503, 2017.
- [74] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *IEEE Computer Vision and Pattern Recognition*, pages 5987–5995, 2017.
- [75] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices. *arXiv:1707.01083*, 2017.
- [76] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. Deep interest network for click-through rate prediction. In *Proc. 24th ACM SIGKDD Int. Conf. Knowledge Discovery & Data Mining*, pages 1059–1068, 2018.
- [77] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv:1606.06160*, 2016.