

Progressive Compressed Records: Taking a Byte out of Deep Learning Data

Michael Kuchnik
Carnegie Mellon University
mkuchnik@cmu.edu

George Amvrosiadis
Carnegie Mellon University
gamvrosi@cmu.edu

Virginia Smith
Carnegie Mellon University
smithv@cmu.edu

ABSTRACT

Deep learning accelerators efficiently train over vast and growing amounts of data, placing a newfound burden on commodity networks and storage devices. A common approach to conserve bandwidth involves resizing or compressing data prior to training. We introduce *Progressive Compressed Records* (PCRs), a data format that uses compression to reduce the overhead of fetching and transporting data, effectively reducing the training time required to achieve a target accuracy. PCRs deviate from previous storage formats by combining progressive compression with an efficient storage layout to view a single dataset at multiple fidelities—all without adding to the total dataset size. We implement PCRs and evaluate them on a range of datasets, training tasks, and hardware architectures. Our work shows that: (i) the amount of compression a dataset can tolerate exceeds 50% of the original encoding for many DL training tasks; (ii) it is possible to automatically and efficiently select appropriate compression levels for a given task; and (iii) PCRs enable tasks to readily access compressed data at runtime—utilizing as little as half the training bandwidth and thus potentially doubling training speed.

PVLDB Reference Format:

Michael Kuchnik, George Amvrosiadis, and Virginia Smith. Progressive Compressed Records: Taking a Byte out of Deep Learning Data. PVLDB, 14(11): 2627 - 2641, 2021.
doi:10.14778/3476249.3476308

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at https://github.com/mkuchnik/PCR_Release.

1 INTRODUCTION

Deep learning training consists of three key components: the data loading pipeline (storage), the training computation (compute), and, in the case of parallel or distributed training, the parameter synchronization (network). A plethora of work has investigated scaling deep learning from a compute- or network-bound perspective [2, 6, 22, 23, 26, 54, 71, 74, 124, 128, 139, 141]. However, little attention has been paid to scaling the storage layer, where training data is sourced.

Current hardware trends point to an increasing divide between compute and the rest of the hardware stack, including network




Figure 1: Three generations of single-node TPU hardware performance on ResNet/ImageNet (batch of 1024). For illustrative purposes, a 1GiB/s limit of data bandwidth is shown (red line) with the corresponding slowdown. ResNet18 on a TPUv3 node can pull over 21k images/second, or over 2.1GiB/s—a challenge for both storage and network.

or storage bandwidth [66, 69, 71] and main memory [67, 103, 130]. Indeed, in the last decade, the amount of compute available to deep learning (DL) has increased exponentially [95]. However, I/O bandwidth has been slower to evolve, potentially resulting in I/O becoming a dominating factor in the overall runtime of deep learning tasks [66, 69, 84, 121]. Recent evidence highlights that, while accelerators are the workhorse of any ML fleet, 30% of resources are spent on the data pipeline [88] in industrial workloads and up to 65% of epoch time is spent in data pipelines in research workloads [87]. While I/O is only a part of the data pipelines, it has the possibility to create bottlenecks and thus lower end-to-end ML training efficiency.

The resource requirements for I/O can be prohibitive, either due to cost, scaling limits of filesystems, or quality of service requirements. Figure 1 shows that this can be problematic even at small scales. We mark the time it takes to fetch one training batch using 1GiB/s of data bandwidth with a dashed red line, since cloud instances are typically limited to 1–4 GiB/s of network [39] and 1GiB/s of disk bandwidth [38]. We have trained ResNet [48] models of varying complexity (ResNet18 being the least complex) using the ImageNet [27] dataset. We find that TPUv2 [53], Google’s second version of their custom AI accelerator, completes computation on a given training batch within the time it takes to fetch the next batch (according to the dashed red line). However, the least complex of the models, ResNet18, is expected to toe that line. TPUv2 is 6 years old, and the third version of Google’s accelerator manages to pack enough compute to speed up batch computations so that two out of the three ResNet models spend more time fetching the next training batch than computing on the current one. This is projected

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 14, No. 11 ISSN 2150-8097.
doi:10.14778/3476249.3476308

to become a problem for even more complex models, according to publicly available per-core performance numbers for TPUv4 [85], which we include in the figure.

These trends highlight that I/O, if left unaltered, stands to dominate training costs. Worse, if the underlying data used in training were to get larger, the problem could become much worse. Contrary to conventional wisdom, datasets like ImageNet consist of *small* images with an average image resolution that is 7× smaller than industry workloads [100], and thus the combination of training rates and data sizes are likely to increase.

To cope with the divide between compute and I/O, architects have turned to hardware/software co-design in an attempt to meet scaling and efficiency goals [53, 65]. Two common, complementary approaches to optimize the storage layer include *caching* [63, 87] and *reducing data volume* [58, 77]. From the caching point of view, I/O pressure can be relieved by keeping a subset of the workload in memory, and optimizing access patterns to hit in the cache and thus avoid I/O. However, for large datasets, one must choose between prohibitively large cache sizes or weaker forms of sampling [81]. From the point of view of data reduction, practitioners can resize images to reduce their size. However, choosing the resize parameters is task-specific, and is subject to error [115], which diminishes task performance. In this work, we show deep neural network training is amenable to a range of JPEG compression; however, unlike resizing, this fact can be exploited as a *mechanism for dynamic data reduction*. Notably, we show that different training tasks—a product of the dataset, model(s), and objective—can tolerate different compression levels (Section 6), and it is non-trivial to determine these levels *a priori*, which motivates a need for dynamic compression.

In this work, we propose *Progressive Compressed Records* (PCRs), a novel data format that reduces the bandwidth cost associated with DL training. Our approach leverages a compression technique that decomposes each data item into a series of *deltas*, each progressively increasing data fidelity. PCRs use deltas to *dynamically* access entire datasets at a fidelity suitable for each task’s needs, while avoiding inflating the dataset’s size. This allows training tasks to control the trade-off between training data size and fidelity. The careful layout of PCRs ensures that data access is efficient at the storage level. Switching between data fidelities is lightweight, enabling adaptation to changing runtime conditions. Using PCRs for a variety of common deep learning models and image datasets, we find that bandwidth (and therefore training time) can be reduced by 2× on average relative to simple JPEG compression without affecting model accuracy. This can allow for a larger fraction of the dataset to be cached in memory, complementing prior work [63]. Overall, we make the following contributions:

- (1) We introduce *Progressive Compressed Records* (PCRs), a novel storage format for training data. PCRs combine progressive compression and careful data placement to enable tasks to *dynamically* choose their data fidelity, increasing the effective training bandwidth.
- (2) We demonstrate that by using PCRs, training speed can be improved by 1.6–2.6× by selecting a lower data fidelity. These speedups are conservative given that the ‘raw’ images we use are in fact already JPEG compressed; speedups are thus likely to be even larger for uncompressed datasets.




Figure 2: The design space of ML image file formats. File-per-Image (Row 1) formats randomly read files in the logical storage address space. Record layouts (Row 2) batch a subset of files into a single, large sequential read, promoting locality in address accesses. PCRs (§3, Row 3) group by image fidelity (3 shown) to maintain the sequential behavior of record layouts while enabling dynamic compression without the need for duplicating data. Reading a full record recovers the full data fidelity for all images. Metadata (not shown) is small and can be kept in memory.

- (3) In experiments with multiple architectures and large-scale image datasets, we show deep neural network training is robust to data compression in terms of test accuracy and training loss.
- (4) We introduce methodology for choosing the particular data fidelity necessary for a task, as well as a tuning heuristic that can be applied automatically. Using PCRs, our method can dynamically switch between multiple data fidelities while training without loss of accuracy.

2 BACKGROUND

Advances in scalable training methods, software, and compute (e.g., accelerators) suggest that the time spent on training computation is decreasing relative to time spent accessing data [63, 66, 69, 71, 87, 121]. Data bandwidth is therefore an increasingly important bottleneck to consider for machine learning pipelines. Two complementary concepts make up the process of storing/loading data: the *data layout*, which helps to utilize the bandwidth of the underlying storage system, and the *data representation*, which can increase bandwidth by reducing the amount of data transferred. In this work, we develop a novel, flexible, and efficient storage format, PCRs, by combining a data representation (progressive compression) with an efficient data layout. Our work serves to lower three fundamental storage costs: storage capacity, storage operations (IOPS), and storage/network bandwidth.

Record Layouts. Learning from data requires sampling points from a training set. In the context of image data, perhaps the simplest way to access data is with a *File-per-Image* layout, such as PyTorch’s ImageFolder, which can cause small, random accesses that are detrimental to storage bandwidth and latency, while also stressing filesystem metadata [80, 98, 127]. *Record layouts*, such as TensorFlow’s TFRecords [112], MXNet’s ImageRecord [99], or even TAR files [98], attempt to alleviate this problem by batching data together into records. Record layouts increase performance (i.e.,

read rate) by exploiting locality (Figure 2). Our experiments indicate a single epoch can take 25× longer with File-per-Image formats compared to reading Record formats—limiting their practicality without caching. To achieve randomness, each Record is read into memory, where it may be shuffled with other Records and broken into minibatches by the data loader. While Record layouts improve over File-per-Image layouts, they are designed to store data at a specific fidelity level, thus requiring multiple copies of each dataset at different fidelities in order to realize efficient training across tasks. In this work, our aim is to combine the efficiency of Record layouts with dynamic compression schemes (described below) to offer quick, easy access to data at multiple fidelity levels.

Image Compression. Compressed forms are commonly used to represent training data. JPEG [119] is one of the most popular formats for image compression and is used ubiquitously in machine learning [27, 33, 73, 101]. Most compression formats (including JPEG) require the compression level to be set at encoding time, which often results in choosing this parameter in an application-agnostic manner. However, as we show in Section 6, it is difficult to set the compression level for deep learning training without over- or under-compressing, as the appropriate level may vary significantly across training tasks. Current approaches resort to storing multiple copies of the dataset at different compression levels, particularly for applications using multiple data fidelities within a single training task [58]. This is infeasible for larger datasets. For example, we find duplicating a 2GiB dataset at 9 resolutions can *amplify the dataset size by 1.5–40× and require hours of extra processing time*. Terabyte-sized datasets rely on distributed frameworks to reduce dataset creation from weeks to days [9].

In Figure 3, we provide a simplified illustration of the JPEG algorithm. First, an image is split into blocks of size 8×8 , which are then converted into the frequency domain. The low frequencies (top left of the matrix) store the bulk of the perceptually relevant content in the image. Quantization, which discards information from the block and results in compression, is used to diminish the high frequency values, compressing the data. *Sequential formats* serialize the image’s blocks from left to right, top to bottom. Decoding the data is simply a matter of inverting this process.

Progressive Image Compression. Progressive compression is an alternative to standard image compression, which—combined with an additional rearrangement of data (Section 3)—forms the basis of the idea behind PCRs. *Progressive* formats allow data to be read at varying degrees of compression without duplication. As an example, over slow internet connections, these formats allow images to be decoded *dynamically* as they are transmitted over the network. With the sequential case, data is ordered by blocks, and thus partially reading the data results in “holes” in the image for unread blocks [119]. Dynamic compression schemes interleave information (*deltas*) from each block, allowing all blocks (and thus the entire image) to be approximated without reading the entire byte stream. As progressive formats are simply a different traversal of the set of quantization matrices, they contain the same information as sequential JPEG [55] and are actually often smaller in practice. As we depict in Figure 3, while non-progressive formats serialize the image matrix in one pass, progressive formats serialize the matrix in disjoint groups of deltas which are called *scans*. Scans are ordered

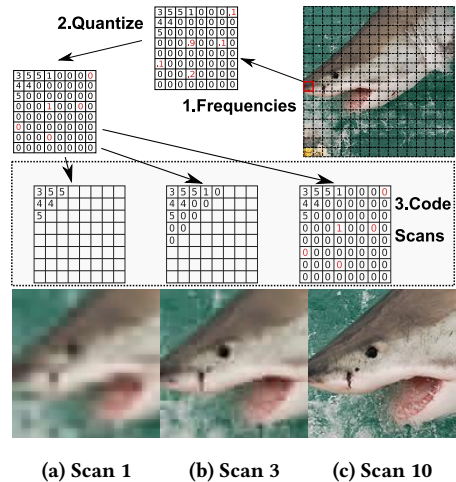


Figure 3: Top: JPEG carves an image into blocks, which are then converted into frequencies, quantized, and serialized. Progressive compression writes out key coefficients from each block before re-visiting the block. Bottom: Higher scans (a→c) have greater fidelity from more frequencies.

by importance (e.g., the first few scans improve quality more than subsequent scans). Thus, any references to images generated from scan n will implicitly assume that the image decoder had access to prior scans. Progressive formats exist not only for images, but also for modalities such as audio [86] and video [102].

3 PROGRESSIVE COMPRESSED RECORDS

We present *Progressive Compressed Records* (PCRs), a novel storage format that reduces data bandwidth for ML training. We specifically explore PCRs for training deep neural networks with image data, but note that the ideas behind PCRs could be readily extended to other modalities (e.g., audio [86] or video [102]), and compression strategies (e.g., cropping [106], interlaced PNG, or neural compression [113]). PCRs define a data layout that ensures bandwidth is fully utilized, and a data representation that permits accessing data at multiple levels of fidelity with minimal overhead.

PCRs are optimized to allow the entire training dataset to be read at a given fidelity. To achieve this, data is rearranged into *scan groups*, i.e., collections of deltas of the same fidelity that are stored together in the address space. To dynamically increase the fidelity of data read and decoded, a task then merely needs to read subsequent scan groups until the desired fidelity level is reached. PCRs differ from other formats (e.g., TFRecord, RecordIO) because PCRs allow these lower fidelity versions of each record to be accessed efficiently (without space/throughput tradeoffs). This efficiency is achieved by using progressive compression and changing the order that data is stored and accessed. Space overhead for PCR conversion is negligible; PCRs are usually 5% smaller than TFRecords. This is because record format size is dominated by the image payload, which is *simply rearranged* with progressive compression. File size differences stem from the efficiency of entropy coding in JPEG, which typically has higher compression ratios over progressive layouts [107]. Since PCRs allow a lower fidelity version of the




Figure 4: PCRs encode label metadata followed by all scan groups. Accessing the dataset at lower quality requires reading up to a given scan group. Reading all scan groups returns the full quality data, and decodes to identical bytes as the conventional JPEG format.

entire dataset to be accessed efficiently, they can drop the effective size and utilized bandwidth of a record by a factor of 2–10×.

Figure 4 depicts the PCR format as it is organized on the storage medium. PCRs logically consist of two parts: *metadata* and *data*. Metadata consists of sample metadata (e.g., labels, bounding boxes) as well as PCR metadata (e.g., mapping of files and scans to storage addresses). Metadata is small in size (e.g., an image label can be represented by a 32 bit int, while an image is 100kiB or more) and can be kept in a database, mapped in memory, or pre-pended to PCRs (for per-sample metadata e.g., labels). The data, which is orders of magnitude larger, consists of the images themselves, organized in terms of increasing levels of fidelity. Each fidelity level for an image is a *scan* and each grouping of images of the same fidelity is a *scan group*. For example, the scan 1 representation of the shark in Figure 3 can be retrieved by reading its data from scan group 1. Likewise, the scan 3 representation will be available once the records up to scan group 3 are read, and the reconstructed representation will be of higher fidelity than that of scan 1. As scan groups consist of scans of the same fidelity, every image contained in a record of the same group offset is available at the same fidelity. Users of PCRs can read data at a given fidelity by simply reading the on-disk byte stream from the start of the PCR to the end of the corresponding scan group. This way of dynamically selecting data fidelity allows for bandwidth savings without re-encoding the data.

4 DESIGN

The key insight behind PCRs is that, for storage or network I/O bound workloads, training tasks can be sped up by reducing data fidelity (and, thus, the amount of data read) to match the available I/O bandwidth. Figure 5 shows the training throughput obtained by using PCRs vs. the traditional TFRecord format. As a fair point of comparison, we use our `tf.data` [88, 110] implementation, and thus only the dataset reader operation has changed. As we describe in Section 4.1, PCRs can reduce the bytes read per image, and thus proportionally increase the throughput of the end-to-end training process (up to the compute limits of the accelerator). However, a speedup is only possible if the CPU overhead introduced by PCRs can be absorbed by the machine (Section 4.2). The final part of the design of PCRs involves choosing the image quality level automatically, which we describe in Section 4.3.

4.1 I/O Speedup Analysis

End-to-End Slowdown. Amdahl’s Law [7] states that if p fraction of a program is waiting for data (see red/blue bars in Figure 1), a $\frac{1}{1-p}$ speedup can be obtained by removing the wait. Recent work

determines possible speedups empirically by finding the gap between the data preparation rate and the I/O rate [87]. However, because PCRs are dynamic, it is important to know what PCR configurations can actually lead to a speedup (i.e., what scan group to select). Although we observe over 500× max/min range on ImageNet, mean size-per-sample, $\mathbb{E}_{x \sim \mathcal{D}}[s(x)]$, is all that is required for an accurate performance model. We tabulate this information in Table 1 and motivate the model below.

Table 1: Image size reduction for various scans and the size of an average image, which can be combined to predict I/O speedups. Scan 10 is approximately the same size as baseline JPEG, and scan 5 is roughly half.

Dataset	Scan 1	Scan 2	Scan 5	Scan 10	$\mathbb{E}_{x \sim \mathcal{D}}[s(x)]$
ImageNet	16×	7×	2×	1×	110kB
HAM10000	30×	15×	3×	1×	250kB
Cars	14×	6×	2×	1×	110kB
CelebAHQ	7×	4×	3×	1×	80kB

Input Pipeline Throughput. Using closed-system Little’s Law [47, 75] and basic assumptions on the characteristics of a storage system (i.e., the cost of large reads is proportional to bytes read), the image throughput, X (e.g., images per second), of an image pipeline

is explained by the equation: $X = \frac{W}{\mathbb{E}_{x \sim \mathcal{D}}[s(x)]}$, where W is the bandwidth and $\mathbb{E}_{x \sim \mathcal{D}}[s(x)]$ is the mean image size (average size of an image sample). The number of bytes in a record (a large read) is, by linearity of expectation, the number of images, n , times the average image size. Thus, the amortized cost per image (dividing by n) is the average image size, and time taken is proportional to W . If a model/accelerator trains at 500 images/second (a function of the resized and cropped input-matrix dimensions [48, 62, 105]), we can conclude that, using ImageNet images, it will use up to $110\text{kB} * 500\text{s}^{-1} = 55\text{MB/s}$ (Table 1), as demonstrated in Figure 5.

Dataset-Level Bounds. To remove dependence on the accelerator’s speed, the equation can be applied on both scans and the original data: Theorem 4.1 presents the asymptotic bounds for the impact of data reduction on training speedups. It is derived by noticing that, when a system with fixed W is bound by the throughput of the I/O subsystem, X , one can calculate the speedup ratio \hat{X} / X , where \hat{X} is using a reduced image size. In sum, reducing the mean data read results in proportionally higher I/O throughput, which results in proportional speedup on I/O bound workloads. For example, using Table 1, which displays the ratios, one can anticipate that a 2× speedup would be seen on ImageNet with scan 5. We defer interested readers to our supplemental material for a more formal discussion of the performance modeling.

THEOREM 4.1. *If a training pipeline is data bound, then the maximum achievable system speedup, S_{max} , for switching from dataset \mathcal{D} to \mathcal{D}' is the ratio of mean sample sizes, $s(x)$:*

$$S_{max}(\mathcal{D}, \mathcal{D}') = \frac{\mathbb{E}_{x \sim \mathcal{D}}[s(x)]}{\mathbb{E}_{x' \sim \mathcal{D}'}[s(x')]}.$$

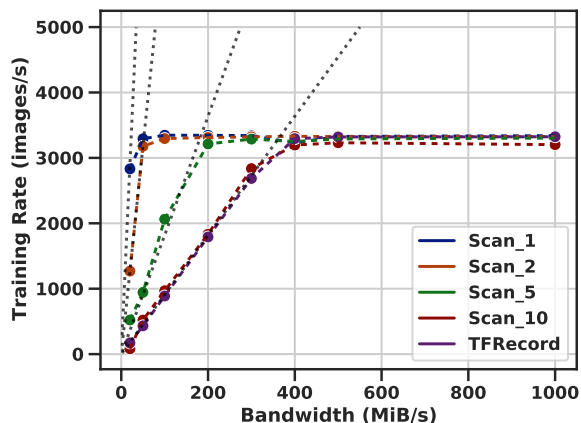


Figure 5: The training rate of a 10-node TitanX GPU cluster with a ResNet-18 workload using PCRs (the scans) and TFRecord. The throughput of the training process is dominated by I/O bandwidth until the compute limit of the GPU is reached. PCRs at scan 10 are approximately the same size as TFRecord, and thus have similar performance. Predicted rates are shown.

4.2 Data Preparation Decoding Overhead

Changing the dataset encoding inevitably changes data-preparation work, which consumes CPU resources and must be managed. The cost of progressive compression is dependent on the image size, scan configuration, and decoder implementation [134]. To analyze the cost of this progressive decoding, we test the rate of decoding 30k images using one thread—these rates can then be multiplied by the number of cores on the machine in a parallel training scenario. The figures are shown in Table 2, indicating that decoding with a subset of scans can be comparable to traditional decoding. On the other hand, using all the scans is over 2× slower than traditional decoding. On a many-core machine (e.g., a 32+ core setup like the one used in Section 6.1), this overhead can be absorbed by idle cores—in Figure 5, we do not see any slowdown by using PCRs relative to baseline JPEG TFRecords because each of the 10 machines can decode 3k images in the worst case. However, for less core-heavy machines, we note three optimization paths to lower decoding overhead. First, excessive and unused scans can be removed. Second, using spectral selection can lower decoding overhead. Third, hardware acceleration can be used (Section 7).

Table 2: The single-core decoding rate (images/s) of various JPEG encodings across the datasets. Progressive decompression can be over 2× more expensive than baseline decoding.

Dataset	Scan 1	Scan 2	Scan 5	Scan 10	Baseline
ImageNet	433	412	340	146	419
HAM10000	465	438	275	96	240
Cars	266	240	225	127	268
CelebAHQ	239	213	195	129	286




Figure 6: The similarity of gradients across epochs for ResNet/HAM10000 (max of 1.0). Legend shows bandwidth savings. Gradient similarity is exact for scan group 10 and decreases for other scans as the model converges. Higher quality scans lead to gradients within 0.1 of the baseline’s gradient, and thus should result in similar final models.

4.3 Autotuning Image Fidelity

Lossy compression of input data creates concerns for the output of model training, and thus creates questions for how to select a tolerable scan group. To analyze the effect of lower image fidelity, we observe that deep learning training is based on stochastic gradient descent, which involves taking a “step” in the *direction* (a vector) that improves the model. If two datasets, \mathcal{D} and \mathcal{D}' , yield the same direction, then they will yield the same model. Therefore, we may intuitively find an alternative dataset \mathcal{D}' , which is close to the original dataset \mathcal{D} in terms of how the model views the gradient direction of the loss function, L . More formally, we want the *angle* between gradient vectors to be small.

To accomplish this, we freeze the model mid-training and empirically measure the gradient direction, $\nabla_{\theta}L$, over the full fidelity dataset, \mathcal{D} , which contains batches of images, X , and labels, y . As the parameters are frozen, we can also measure the gradient on the lower fidelity dataset, \mathcal{D}' , which has alternative images, X' . The angle between the lower fidelity dataset and the original dataset yields the similarity score, which ranges between -1 and 1. Maximizing similarity would yield an *identical* model.

$$\text{score}(\mathcal{D}, \mathcal{D}') = \text{sim}(\nabla_{\theta}L(X, y), \nabla_{\theta}L(X', y))$$

where similarity is the cosine similarity:

$$\text{sim}(\mathbf{A}, \mathbf{B}) = \cos(\theta) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \|\mathbf{B}\|}$$

We evaluate this procedure with HAM10000/ResNet, using 2560 images to estimate the gradient and 3 trials to get 95% confidence intervals. As shown in Figure 6, decreasing image fidelity decreases the similarity with respect to the true gradient. Scan 10 is bit-identical to the baseline dataset, and thus we observe maximum similarity in that case. Meanwhile, scan 1 has the lowest similarity, and the difference increases as the model converges. Given the gradients are well-behaved with respect to fidelity, it is natural to parameterize scan tuning to match some level of gradient similarity.

Using gradient similarity for autotuning quality requires choosing a minimum gradient similarity threshold for scans throughout training, which is the main drawback of this approach. As is shown

in Figure 6, the similarity for high quality data is bounded within a threshold of 0.8—therefore, we find this threshold a good default. The computational cost of evaluation is on the order of tens of gradient steps, and is proportional to the number of scans, epochs, and minibatches used. Our implementation tunes once every 20 epochs and does not tune for the first 5 epochs because models are unstable during this period [42, 48]. As training progresses, low scans become too dissimilar (in terms of gradients) from the higher scans, and therefore are avoided. This, in turn, allows the faster scans to apply a burst of speed to the training process before fine-tuning at a higher fidelity (e.g., when learning rates drop). Compared to static schedules [122], there is only one hyperparameter (the threshold), which is independent of other schedules (e.g., learning rates), and the parameter does not require validation data. We leave tuning using QoS/congestion information [36] to future work.

5 IMPLEMENTATION

There are three fundamental components in the PCR implementation: the encoder, the decoder, and the data loader. The encoder transforms a set of JPEG files into a directory, which contains: a database for PCR metadata and at least one PCR data file. The decoder takes the directory as input and yields a set of JPEG images, efficiently inverting a subset of the encoding. The dataset is split into many records, and, thus, the training process is reading tens to hundreds of PCR data files per epoch. The data loader is where the PCR decoding library interfaces with the inputs provided to deep learning libraries (e.g., TensorFlow [2], MXNet [16], PyTorch [91]). Below, we describe these components in detail.

Encoding. Given a set of images, the PCR encoder breaks images into scans, groups scans into scan groups, and sorts scan groups by fidelity. Once groups are sorted, the PCR encoder can serialize groups while taking note of their offsets (so that subsets may later be decoded). The metadata (e.g., labels) is prepended to the serialized representation, and the resulting byte stream is written to disk. Our implementation uses JPEGTRAN [56] to losslessly transform JPEG images into progressive JPEG images. With the default settings, each JPEG is broken up into 10 scans. The encoder scans the binary representation of the progressive JPEG files, searching for the markers that designate the end of a scan group. The encoder thus has access to all 10 offsets within the JPEG files that can be used to determine the boundaries between scan regions. Forming scan groups requires grouping the scan regions with the same fidelity together, which can be done in one pass over the set of images corresponding to that PCR. An index must be created for ungrouping the scans during decoding; however, serialization libraries, such as Protobuf [96], handle both the packing and unpacking steps transparently. As record format conversion can take hours (§2), PCRs benefit from requiring only a single conversion for multiple tasks. The encoding time is within $2\times$ of conversion to TFRecords in our implementation: for example, converting ImageNet takes 1.4 hours rather than 0.8 hours. When using the widely available TFRecords converter [109], our implementation for PCRs is actually faster due to being parallelized—converters are typically not optimized due to being one-time costs.

Decoding. To decode a PCR file, the file’s scan group offsets have to be located in the PCR metadata. The offsets allow a partial read of the file, i.e., only the bytes of the desired scan group are read. JPEG decoding requires serializing the image deltas of individual scan groups. We terminate the byte stream with an End-of-Image (EOI) JPEG token, which allows most JPEG decoders to render the image with the available subset of scans. The cost of these steps is negligible relative to that of the JPEG decoder, which is the primary challenge facing PCRs.

Loader. We implement PCR loaders using the DALI External-Source operator [89], as well as a C++ version compatible with `tf.data` [88, 110], including a Tensorflow Op [111]. SQLite and RocksDB are supported backing databases, and we support embedding images and metadata in Protobufs or in “raw struct” form. The PCR reader, like most readers, is cheap to evaluate; we can read over 400MiB/s using just a single CPU core. This is because the bulk of the work is not computational, i.e., a file read and a set of `memcpy` operations to re-arrange the PCR images. Serialization libraries can add overhead (e.g., parsing, memory allocations); however, “raw struct” formats avoid these entirely, and flat formats minimize them [35]. Another design point is buffer allocation: in contrast to traditional Record loaders, which can iteratively return individual data samples, PCRs must read (and allocate) buffers for possibly the *entire* record (10MiB+), since later scan groups are used for even the first example (See §7). Thus, an optimized implementation of PCR loaders uses a double-buffer design, where the buffers are re-used and read directly from disk using `O_DIRECT`. Our implementations show that the main bottleneck with using PCRs is the image decode, which is downstream from the loader. For I/O bound workloads, baseline and full-fidelity progressive record readers perform the same (Figure 5) as image size differences are negligible (Section 4.1).

6 EVALUATION

We evaluate the flexibility and efficiency of PCRs using a suite of large-scale image datasets. We begin by describing our experimental setup (§6.1) and present an end-to-end evaluation of PCRs (§6.2), demonstrating their ability to reduce training time. We show that dynamic compression is crucial because the appropriate level of compression varies across models and training tasks (§6.3). We explore metrics that can be used to explain the effectiveness of compression on a training task (§6.4), introduce autotuning heuristics for dynamic training (§6.5), and trace the speedups achieved by PCRs in terms of training time (§6.6). Our supplemental material contains additional experimental details and training plots (e.g., training loss, other datasets).

6.1 Experimental Setup

Our evaluation uses the ImageNet ILSVRC [27, 101], CelebAHQ [58], HAM10000 [116], and Stanford Cars [61] datasets. For CelebAHQ, we classify if the celebrity is smiling or not. A summary of each dataset is given in Table 3. We aimed to select datasets that vary in terms of the image resolution, number of examples, number of classes, and image/scene type. *All of the datasets are in fact already compressed before progressive compression is applied, making the presented speedups conservative estimates of the potential benefit of PCRs.*




Figure 7: Top-1 test performance (with 95% confident intervals) using ResNet and ShuffleNet on ImageNet (a,b) and CelebAHQ (c,d). Lower scan groups (corresponding to further compressed data) can provide faster overall training, often without sacrificing accuracy. However, the appropriate level of compression depends on the model, infrastructure, and data—necessitating the ability to easily access data at multiple fidelities, as with PCRs. We explore these factors in Sections 6.3, 6.4, and 6.5.

Table 3: PCR dataset size and record count information. Datasets vary in terms of number of images, their JPEG quality, and the image sizes. Some datasets, such as HAM10000, have image sizes larger than average. Record sizes concentrate around the dataset size divided by the record count.

Dataset	Records	Images	Size	Quality	Classes
ImageNet	1251	1281167	129GiB	91.7%	1000
HAM10000	125	8012	2GiB	100%	7
Cars	63	8144	887MiB	83.8%	196
CelebAHQ	93	24000	2GiB	75%	2

Specifically, the datasets natively use a JPEG quality level varying from 75% (CelebAHQ) to 100% (HAM10000) (§3). Experiments use resizing, crop, and horizontal-flip augmentations, as is standard for ImageNet training [108, 115]. The sizes of each dataset’s scan groups (used in Table 1) are shown in Figure 8; sizes decrease for lower scan groups. For examples of images under each scan group, see the supplement.

Training Regime. We evaluate two loader implementations of PCRs, comparing PCR scans against themselves and then comparing PCRs against strong baselines (TFRecords). For both setups, we use PyTorch [91] for all model training; the two loaders are using DALI [89] Loader, which was used for initial prototyping, and `tf.data` [88, 110], which we have since made native operator extensions to for maximum performance. In our experiments we use pretrained ImageNet weights for HAM10000 and Cars due to the limited amount of training data. We use standard ImageNet training, starting the learning rate at 0.1 with gradual warmup [42], and dropping it on epochs 30 and 60 by 10×. After augmentations, all inputs are of the same size; thus, a model’s update rates are the same across datasets. The pretrained experiments (HAM10000 and Cars) start at a learning rate of 0.01 to avoid changing the initialization too aggressively. We use mixed-precision training [10, 83] for the DALI runs. We use ResNet-18 [48] and ShuffleNetv2 [78] architectures for our experiments with a batch size of 128 per worker. We run each experiment at least 3 times to obtain confidence intervals. We sample test accuracy every 15 epochs for non-ImageNet datasets. Our

evaluation focuses on the differences obtained by reading various amounts of scan groups. For the DALI runs, we consider reading all the data (up to scan group 10) to be the baseline, as the baseline formats will perform similarly under I/O bounds (Figure 5)—we later provide a direct comparison with baseline TFRecords when using `tf.data`. Our results are conservative as we are already utilizing pre-compressed data and we include evaluation times in our results. For the purpose of evaluation, all scan groups within a dataset were run for the same number of epochs (90 for ImageNet, 150 for HAM10k, 250 for Cars, and 90 for CelebAHQ). We also provide annotated (dashed) lines for subsequent epochs.

System Setup. Our experiments were run on a 16 node Ceph [125] cluster with NVIDIA TitanX GPUs and 4TB 7200RPM Seagate ST4000NM0023 HDD. We use six Ceph nodes: five dedicated Object Storage Device (OSD) nodes, and one Metadata Server (MDS). The remaining 10 nodes are machine learning training workers. This 2:1 ratio between compute and storage nodes results in 400+ MiB/s of peak storage bandwidth; we have also tested a heavily I/O bound 10:1 ratio and found the trends comparable. Ceph is a common production-grade open-source filesystem, but our results would generalize to any setup with a mismatch between compute power and data bandwidth (either storage or network). In addition to microbenchmarks, we evaluate the generalization of PCRs to SSD setups in Section 7. Since state of the art compute is 150× faster than our own setup on a more expensive model (ResNet-50) [136], we focus on models which are fast to train (while still being modern; AlexNet [62] is potentially faster) while limiting read parallelism. The DALI setup uses `O_DIRECT` to ignore caching effects and highlight bandwidth usage. To reflect what PCRs may look like in realistic, heavy-load situations, we provide 20 node experiments in Figure 9 with the same storage system and double the workers, which allows speedups to be seen with full read parallelism per node (over 700MiB/s of peak bandwidth). This setup uses our `tf.data` loader implementation to fairly compare against TFRecords and File-per-Image formats, showing its effectiveness. We use the same setup in Figure 14.




Figure 8: The size in bytes of various levels of scans read. Scan group 0 (not shown) contains only labels and is typically ~ 100 bytes. Each scan adds roughly a constant amount of data (i.e., linear scaling), although certain scans add considerably more than others (i.e., sizes sometimes cluster) due to techniques like chroma subsampling. Using all 10 scans can require over an order of magnitude more bandwidth than 1–2 scans. Interquartile ranges are shown.




Figure 9: Top-1 test performance on ImageNet with ResNet and ShuffleNet, using double the compute (20 workers with same configuration). Doubling the compute forces bottlenecks to appear by approaching hardware limits of aggregate disk throughput. We run this experiment once and terminate at epoch 90, showing a $2\times$ speedup for scan 5.

6.2 Reducing Time to Accuracy via Compression

Observation 1: Training time can be reduced by up to $2\times$ using data compression. PCRs capitalize on this by dynamically reducing training data size, all without adding space overhead.

We begin our empirical study by exploring the effect of data compression on training time and training loss/test accuracy. We provide time-to-accuracy results for ResNet-18 and ShuffleNetv2 training on ImageNet and CelebAHQ (Figure 7), HAM (Figure 10), and Cars (Figure 11). Across these experiments, we find that PCRs can provide a $2\times$ boost on average in time-to-accuracy compared to the baseline, by dynamically providing data at a higher level of compression. We make several observations about these results. First, we note that we tend to see larger speedups for smaller, faster models (e.g. ShuffleNet), than for bigger models (e.g., ResNet). Indeed, the current speedups may in fact become significantly larger with faster compute [e.g., 64, 66, 136]. Such a trend is visible in the heavy ImageNet experiments featured in Figure 9—both TFRecords


and scan 10 are about the same size, and therefore finish simultaneously, but scan 1 and 2 finish nearly an hour faster for ShuffleNet. For this same setup, we observe Files-per-Image take over 2 hours per epoch due to a lack of sequential reads— $25\times$ slower than TFRecords, which scan 5 improves by $2\times$; therefore, we conclude that progressive compression and record formats are both necessary for performance.

Second, while time-to-accuracy is reduced as we move to lower scan groups, there is a *statistical efficiency cost*. Namely, models trained on scans 1 and 2 may not always converge to an acceptable solution, as shown for ImageNet (Figure 7). Certain tasks like CelebAHQ, however, can tolerate this fidelity loss, either because they consist of less compressed images or because the training task is less dependent on high-frequency image features. These results suggest that, while compression saves bandwidth and offers a potential speedup, the ideal amount of compression depends on two factors: (i) the speed of the model and the underlying compute infrastructure, and (ii) the structure of the task and the images in the dataset. We explore these factors in more detail below.

6.3 Task Tolerance to Data Fidelity

Observation 2: Different models can tolerate different fidelities.


Given a fixed dataset, we show that there is variation in the data fidelity/compression level that different models can tolerate for training. This motivates an important use-case of PCRs, as the format allows data to be stored *once* but then accessed at multiple compression levels while models are tuned or various models are applied to the problem at hand. In Figure 10, both ResNet and ShuffleNet are trained with the HAM10000 dataset. While ResNet consistently tolerates low fidelity images, ShuffleNet training tends to degrade with low fidelity data. ShuffleNet reaches its best accuracy at scan 5, but our other results suggest that lowering fidelity results in lower accuracy for the same epoch in nearly all cases (Figure 12). This suggests that different models will experience different speedups for similar accuracy levels, depending on their sensitivity to fine-grained features unavailable in low fidelity data.



(a) HAM10000 ResNet

(b) HAM10000 ShuffleNet


Figure 10: Test accuracy on HAM10000. While ResNet is robust to additional compression, ShuffleNet requires higher fidelity data (at least scan group 5) for higher accuracy. Time is relative to first epoch. 95% confidence intervals are shown.



(a) Original Multiclass

(b) Binary Is-Corvette


Figure 11: Test accuracy with ResNet-18 on the Stanford Cars dataset and a binary variant. The gap between scan groups closes as the task is simplified. Time in x-axis is relative to first epoch. 95% confidence intervals are shown.



(a) Epoch Accuracy

(b) MSSIM

Figure 12: Left: Top-1 test performance vs. epoch on ResNet-18/ImageNet; other models/datasets are similar. Using lower quality scans can only degrade performance; it does not act as a beneficial data augmentation. Right: Corresponding image quality degradation according to MSSIM.



(a) MSSIM Regression

(b) Clustered Convergence

Figure 13: MSSIM vs. accuracy for the Cars dataset with ResNet18 and Shufflenet. We obtain similar results for other datasets. Left: There is a linear relationship between MSSIM and the final test accuracy. Right: Scan groups (ShuffleNet) cluster by MSSIM and accuracy.

Observation 3: Different tasks, e.g., multi-class classification vs. binary classification, can tolerate different levels of data fidelity. The same PCR dataset can service these different tasks.

The difficulty of a task, or training objective of interest, also affects the amount of compression that can be tolerated. Harder tasks, e.g., multi-class classification with a large number of classes, require higher fidelity data. We validate this empirically in Figure 11 (and additional evidence is provided in the supplement). This experiment reduces the number of classes for the classification task, demonstrating that lower scan groups can be used for easier tasks. The full range of classes investigated includes: *Baseline* (i.e., Car Make, Model, Year), *Make-Only* (i.e., car Make only), and *Is-Corvette*, a binary classification task of Corvette detection. Compared to the original task, the coarser tasks reduce the gap between scan groups, decreasing the gap from baseline to the binary case. These results suggest that the optimal image encoding can be dependent on the exact labeling or task complexity. Thus, while static approaches may need one encoding per task, a fixed PCR encoding can support multiple tasks at optimal fidelity by simply changing the scan group depending on how the labels (metadata) are remapped.

6.4 Compression Level Estimation

Observation 4: MSSIM image similarity is a reliable estimator of the accuracy loss between scan groups, and can be used to determine appropriate levels of compression for training with PCRs.

To better explain the effectiveness of compression, we compare how various scans approximate the reference image through MSSIM [123], a standard measure of image similarity. We find a correlation between MSSIM and final test accuracy, especially when comparing scan groups *within* a task. Our preliminary tests show that scan groups with similar MSSIM achieve similar accuracy (Figure 13), which is why only scan groups 1, 2, 5, and the baseline are shown. Due to the way progressive JPEG is coded by default, groups tend to cluster, e.g., scans 2, 3, and 4 are usually similar, while 5 introduces a difference. Such “banding” or clustering is seen in the accuracy trends; the major jumps correlate with the appearance of Y (luminance) AC coefficients in the JPEG encoding. Scan groups of 5 or higher have an MSSIM of 95%+, which is likely why they consistently perform well. MSSIM can therefore be used as a diagnostic for choosing scans, although we acknowledge that changes in perception are hard to predict for large deviations (MSSIM < 95%). For some datasets, linear regression on MSSIM recovers final test accuracy even with different models (Figure 13)




Figure 14: Adaptive tuning on ImageNet for 90 epochs compared to TFRecord training, which is comparable to Scan 10 training. For adaptive tuning, epochs are marked with scatter points. Training is fastest after epoch 5, when autotuning search is enabled. Changing the threshold from 0.8 (shown) to 0.9 results in the last few epochs switching to scan 10.

or augmentations. Test accuracy per epoch degrades with worse image fidelity across our experiments (Figure 12), further highlighting that time-to-accuracy speedups are caused primarily by bandwidth reduction (rather than e.g., a form of regularization induced by lower scans).

6.5 Autotuning Compression Level

Observation 5: *It is possible to automatically determine an appropriate level of compression at runtime by dynamically accessing various data qualities via PCRs.*

In cases where training resolution is not structured into learning [58] or image fidelity heuristics prove too costly to tune (§6.4), automatic tuning of the scan hyperparameter may be desirable. One way of doing this is by tuning with a measurement of the *bias* of the gradient given a lower fidelity image (§4.3), intuitively measuring how the model “sees” the image similarity. As we showed in Section 2, a similarity threshold of 0.8 or higher is sufficient to avoid bad scans throughout training—Figure 6 clusters low-quality scans below that point. We apply this threshold and the rest of the procedure described in Section 4.3 to the ImageNet dataset and observe that such autotuning repeatedly matches accuracy while almost being as fast as a pure scan 5 approach. The main slowdown is due to starting at scan 10 for the first 5 epochs of training, blending the latencies of scan 10 with those of scan 5. We note that, unlike MSSIM, which is statically concentrated above 95% for good quality scans, the gradient similarity changes over training. For example, ResNet18 has a similarity of 0.88 by epoch 85, whereas it had a similarity of 0.95 at epoch 5. We observe that using a higher threshold, approaching 0.9, forces scan 10 to be used for the last few epochs of training when gradient similarity is lower, retaining similar accuracy at slightly longer training times.

6.6 Image Loading Rates

Observation 6: *Image loading rates are directly linked to the compression ratio, i.e., a compression ratio of 2× results in a 2× speedup.*

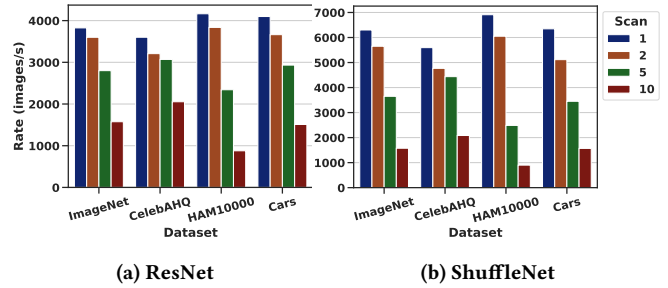


Figure 15: Training rates for ResNet and ShuffleNet. More scans reduce the rate of images/second. From RAM, they can process 4200/7000 images/second, respectively.

Finally, we validate why lower compression levels yield faster training by observing image loading rates. Loading rates for training are shown in Figure 15. Using more scans slows down training significantly, as seen in image rates. Training slowdowns manifest as latency spikes from a data stall, causing rates to fluctuate considerably. Informally, we can perform twice as many read operations if we decrease the data read by each operation by 2× (Section 4.1). The speedup can be calculated through the average PCR size (Figure 5 and Figure 8). Since ShuffleNetv2 is capable of a higher maximum training rate than ResNet-18, it achieves higher speedups. As HAM10000 has the largest images, it is the most bottlenecked by image loading bandwidth—scan 5 is 2.9× smaller than scan 10. For CelebAHQ, scan 2 is roughly the same size as scan 5; as expected, image rates are very similar. For the 20 worker runs, we observe that scan 1 and 2 for ImageNet have a median epoch latency of 100 seconds, while TFRecords and scan 10 have a median epoch latency of 300 seconds—even though the size difference is over 10×, 3× is the best factor that can be achieved before hitting in-memory processing rates. These results indicate that systems with large images, efficient models, and fast compute would be the biggest benefactors of PCRs.

Observation 7: *Image loading rates are directly linked to the underlying model and bandwidth.*

While faster compute hardware can speed up a fixed model (e.g., Figure 1), it is less clear how bandwidth impacts training rates, especially with different hardware/model combinations. To explore this question, we implement the token-bucket algorithm in our `tf.data` implementation. Each second, a node accumulates a fixed amount of tokens, which are traded for bytes read off storage, and thus nodes will block if they use too much bandwidth in a given time. We rate limit the bandwidth of each of the 10 nodes in the cluster across a sweep of aggregate cluster bandwidths in Figure 16, calculate the time per epoch over 7 minutes (using the data shown in Figure 5), and utilize the previous convergence results in Figure 7 to project the accuracy-over-time graphs for 90 epochs. As previously observed, lower scans benefit the most from low bandwidth, and faster models are more bottlenecked. For instance, although ShuffleNet is typically faster than ResNet-18, it still takes a similar amount of time at low bandwidth to finish training—simply because nearly all time is waiting on I/O. We do not see much

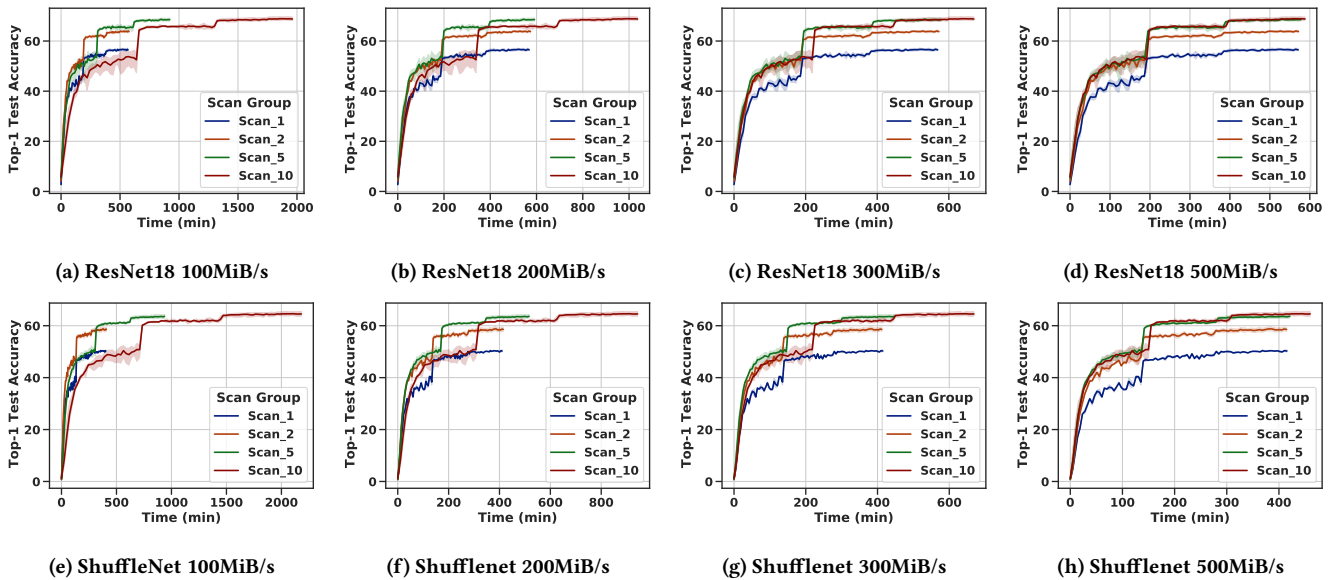


Figure 16: The effect of various amounts of bandwidth on a 10-node cluster running ResNet-18 (top) and ShuffleNet (bottom). At very low bandwidth, all scans provide benefits, while higher bandwidths provide less benefits. Similarly, faster models (e.g., ShuffleNet) or accelerators increase I/O pressure, which enables low scans to obtain higher speedups—scan 1/2 are beneficial for ShuffleNet at 200MiB/s, but not ResNet, and the higher I/O pressure carries into 500MiB/s.

benefit for lower scans at high bandwidth, but low bandwidth (e.g., 20MiB/s) shows gains for even scan 1 over scan 2 (Figure 5).

7 DISCUSSION

Data caching. Data can be partially cached in memory (e.g., OS page cache); however, uniform sampling means that over 75% of the data must be cached to see the majority of speedup—caching does not significantly affect our ImageNet results. Specifically, the expected latency of a read is a convex combination of the disk and memory latency [87], where disk latencies are usually high. PCRs help reduce cache pressure by reducing the number of bytes read (and thus the size of the working set). Further, PCRs facilitate cache sharing in a multi-user environment, as multiple fidelities share common data, eliminating double caching.

Hardware acceleration. Hardware JPEG decoders are popular in mobile phones, and PCRs could take advantage of such hardware support. In fact, NVIDIA’s A100 is the first datacenter GPU to ship with hardware decoding [59]; prior versions used software acceleration [90]. The reasoning behind hardware support is simple: 33% or more of CPU time can be spent on image decoding [87] and 96 cores or more are currently matched with an accelerator [85], which directly cuts into training cost efficiency [53]—empirically, 20% of jobs spend over 30% of their compute budget on data ingest [88]. Frameworks like DALI already offload part of the JPEG decode to the GPU, namely the Discrete Cosine Transform; mapping Huffman decoding to the GPU requires additional work [49, 60, 126]. A different practice is to avoid the CPU by caching *decoded* images in memory [65, 85, 87], though this has limitations as datasets are large, especially when uncompressed.

Non-image datasets and training tasks. Our experiments indicate that PCRs are robust (in terms of accuracy) across a variety of tasks, and we only focus on a subset of tasks due to a limited computation budget. Using ResNet50 [48] on the ImageNet [101] tasks, we obtain 75.14% vs. 75.47% (scan 5/baseline) accuracy, and 38.37/38.80 AP on the FPN-ResNet50/COCO task [72, 73]. While the computational difference between ResNet18 and ResNet50 is only 2–3 years of accelerator progress (Figure 1), we note that detection datasets (like COCO) can be 10-100× slower to compute (albeit with larger images). Apart from image-based datasets, PCRs generalize to other datasets and modalities as long as the encoding is progressive. For example, each component in Principal Component Analysis (PCA) [104] is a progressive approximation of the source dataset, and removing (e.g., compressing) 50% of the components loses 1% of accuracy over YouTube videos [3]. Another general encoding is quantization, which progressively encodes subsets of the higher-order bits (e.g., the first 25%) in a dataset’s features, an approach baked into progressive JPEG and used in YouTube-8m [3]. These results suggest that if one were to implement PCRs via PCA/quantization over videos, they would obtain 8× bandwidth savings total with only a 1% loss in accuracy.

Generalizing across hardware. To test the generalization of our system, we investigate applying PCRs to a Google Cloud [41] n1-instance-16 with a P100 GPU. We attach a 150GB HDD for the operating system and use a SSD for the data loading, which has peak bandwidth of 74MB/s (similar per-GPU load as prior work [87]). On ImageNet/ShuffleNet, we observe 650 images/second for TFRecords, and 680 (scan 10), 1540 (scan 5), 1700 (scan 2), and 1750 (scan 1) for PCRs. The difference between TFRecords and scan 10 can be

explained by progressive compression being 6% smaller in size. Using either as a baseline, scan 5 is over 2× faster. Because the ratio of resources primarily matters, we observe that doubling the CPU, GPU, and SSD resources maintains the same relative performance advantages for PCRs, yielding a 2.2× speedup from TFRecord to scan 5, and a 2.7× speedup from scan 1/2.

8 RELATED WORK

Numerous works have explored methods for decreasing training time with large datasets [42, 52, 64, 66, 133, 136, 137], motivating a need for improved I/O internals [5, 15, 18, 21, 97], formats [98], caching [65], and data pipeline frameworks [88]. We discuss caching, forms of compression, and frequency-domain DL literature below.

Dataset Caching. Caching places data in faster storage tiers to offload the bandwidth burden from slower devices. ML applications lack locality due to uniform sampling [63], requiring either prohibitively large cache sizes or weaker forms of sampling [81]. However, when done correctly, caching can obtain significant speedups [87]. PCRs are designed for datasets which do not fit in caches, and, by virtue of accessing less data, can increase cache hit rates.

Dataset Cardinality Reduction. “Big data” spawned interest in dataset reduction techniques that aim to reduce the *number* of training samples while maintaining model accuracy [11, 24, 34, 57, 70, 79, 129]. Similarly, dataset echoing [20, 138] re-uses subsamples to speedup the data pipeline. PCRs differ in that they reduce I/O burden by modifying data representation and layout.

Dataset Sample Compression. Techniques such as compression [3, 77] or resizing [58] reduce data size by lowering fidelity, reducing I/O pressure. Prior work has shown that resizing as a form of data reduction is particularly effective for DL tasks, as resized data can speed up training, transfer to high fidelity test points, and in some cases, even *increase* accuracy when combined with certain data augmentations [19, 58, 115]. However, resizing parameters are chosen statically and heuristically (therefore suboptimally [115]), and thus may not meet the needs of all applications without duplicated work. PCRs differ in that they provide a *dynamic* mechanism for adjusting I/O load, and thus can adapt to both the system and task at runtime.

Similar to our work, MLWeaving [122] has shown that *transposed layouts* (i.e., column major) can accelerate machine learning training. However, this work differs in that we focus on I/O in the context of deep learning models, whereas MLWeaving focuses on memory bandwidth for general linear models. Additionally, the compression method differs. For image data, the three canonical dimensions of compression are 1) quantization, 2) frequency selection, and 3) spatial selection [119]; MLWeaving uses the first while PCRs use the first and second (via JPEG).

Neural compression [113], which learns custom compression formats using neural networks, is an interesting direction for future work and is compatible with PCRs. However, while neural compression can outperform JPEG in terms of quality [113], it does so at significant cost. Using state-of-the-art neural compression [12, 13, 82], we find decoding to be between 900× and 5000× the cost of baseline JPEG, and thus incompatible with real-time performance.

General Compression. Bandwidth reduction extends to databases, memory hierarchies, and the web [1, 4, 92, 93, 142]. Progressive compression has been used in the context of dynamically saving bandwidth for mobile phone downloads [134]. HippogriffDB [69] uses GPUs to compress data in the context of databases, which lowers I/O bandwidth to get a speedup. Other work has investigated how image degradation affects inference [29, 94, 118, 140]. In contrast, our work is focused on compression for I/O savings in deep learning. Reinforcement learning has been used to choose JPEG parameters for cloud inference [68]; other work has hand-designed JPEG encodings for training [77]. These works are similar to ours in that they tune compression for the model, though they differ in that they are static. Other work investigates compressing models [8, 17, 28, 44–46, 50, 131] or network traffic [6, 71, 74, 124, 128, 139]; these are orthogonal to our work.

Frequency Domain Deep Learning. Prior work modifies models to directly train over compressed representations [37, 43, 114, 117] or with frequency-domain operators [31]; our work does not modify the model. Other work investigates generalization performance from the view of low [14, 132] and high [120] frequencies, which provides insight into our work. JPEG mostly filters low frequency components, and thus prior work has attempted to use JPEG as a defense mechanism against adversarial attacks [25, 32, 76]. Motivated by adversarial attacks exploiting spurious, high-frequency features [40, 51] other work investigates if frequency filters can impact model robustness [30, 135]; our work primarily focuses on retaining test accuracy under non-adversarial conditions.

CONCLUSION

We introduce a novel storage format, *Progressive Compressed Records* (PCRs), to reduce the bandwidth cost of training over large datasets. PCRs use progressive compression to split training data into multiple fidelity levels, while avoiding duplicating space. The format is easy to implement and can be applied to a broad range of tasks dynamically. PCRs provide applications with the ability to trade off data fidelity with storage and network demands, allowing the same model to be trained with 2× less bandwidth while retaining model accuracy. We introduce methodology for choosing the particular data fidelity necessary for a task, as well as a tuning heuristic that can be applied automatically. Using PCRs, our approaches can dynamically switch between multiple data fidelities while training without loss of accuracy. Future directions include alternative compression methods, data modalities, and hardware acceleration.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their help improving the presentation of this paper. We thank the members and companies of the PDL Consortium: Amazon, Facebook, Google, Hewlett Packard Enterprise, Hitachi Ltd., IBM Research, Intel Corporation, Microsoft Research, NetApp, Inc., Oracle Corporation, Pure Storage, Salesforce, Samsung Semiconductor Inc., Seagate Technology, Two Sigma, and Western Digital for their interest, insights, feedback, and support. Michael Kuchnik is supported by a National Defense Science and Engineering Graduate Fellowship. This research was supported with Cloud TPUs from Google’s TPU Research Cloud and research credits from Google Cloud Platform.

REFERENCES

- [1] Daniel Abadi, Samuel Madden, and Miguel Ferreira. 2006. Integrating Compression and Execution in Column-Oriented Database Systems. In *ACM SIGMOD International Conference on Management of Data*. 671–682.
- [2] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems.
- [3] Sami Abu-El-Hajja, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. 2016. YouTube-8M: A Large-Scale Video Classification Benchmark. *arXiv preprint arXiv:1609.08675* (2016).
- [4] Victor Agababov, Michael Buettner, Victor Chudnovsky, Mark Cogan, Ben Greenstein, Shane McDaniel, Michael Piatek, Colin Scott, Matt Welsh, and Bolian Yin. 2015. Flywheel: Google's Data Compression Proxy for the Mobile Web. In *USENIX Networked Systems Design and Implementation*. 367–380.
- [5] Alex Aizman, Gavin Maltby, and Thomas Breuel. 2019. High Performance I/O For Large Scale Deep Learning. In *IEEE International Conference on Big Data*. 5965–5967.
- [6] Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. 2017. QSGD: Communication-efficient SGD via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*, Vol. 30.
- [7] Gene M. Amdahl. 1967. Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities. In *Spring Joint Computer Conference*. 483–485.
- [8] Sajid Anwar, Kyuyeon Hwang, and Wonyong Sung. 2015. Fixed point optimization of deep convolutional neural networks for object recognition. In *International Conference on Acoustics, Speech and Signal Processing*. 1131–1135.
- [9] ApacheBeam [n.d.]. ApacheBeam. https://www.tensorflow.org/datasets/beam_datasets. Accessed: 07-26-2021.
- [10] Apex [n.d.]. NVIDIA Apex. <https://github.com/NVIDIA/apex>. Accessed: 07-26-2021.
- [11] Olivier Bachem, Mario Lucic, and Andreas Krause. 2017. Practical coresets constructions for machine learning. *arXiv preprint arXiv:1703.06476* (2017).
- [12] Johannes Ballé. 2018. Efficient Nonlinear Transforms for Lossy Image Compression. In *Picture Coding Symposium*. 248–252.
- [13] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. 2018. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*.
- [14] Ronen Basri, David Jacobs, Yoni Kasten, and Shira Kritchman. 2019. The Convergence Rate of Neural Networks for Learned Functions of Different Frequencies. In *Advances in Neural Information Processing Systems*, Vol. 32. 4761–4771.
- [15] Denis Baylor, Eric Breck, Heng-Tze Cheng, Noah Fiedel, Chuan Yu Foo, Zakaria Haque, Salem Haykal, Mustafa Ispir, Vihan Jain, Levent Koc, Chiu Yuen Koo, Lukasz Lew, Clemens Mewald, Akshay Naresh Modi, Neoklis Polyzotis, Sukriti Ramesh, Sudip Roy, Steven Euijong Whang, Martin Wicke, Jarek Wilkiewicz, Xin Zhang, and Martin Zinkevich. 2017. TFX: A TensorFlow-Based Production-Scale Machine Learning Platform. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1387–1395.
- [16] 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. In *Neural Information Processing Systems Workshop on Machine Learning Systems*.
- [17] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. 2017. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282* (2017).
- [18] Steven WD Chien, Stefano Markidis, Chaitanya Prasad Sishla, Luis Santos, Pawel Herman, Sai Narasimhamurthy, and Erwin Laure. 2018. Characterizing deep-learning I/O workloads in TensorFlow. In *International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems*. 54–63.
- [19] Ting-Wu Chin, Ruizhou Ding, and Diana Marculescu. 2019. AdaScale: Towards Real-time Video Object Detection using Adaptive Scaling. In *Machine Learning and Systems*, Vol. 1. 431–441.
- [20] Dami Choi, Alexandre Passos, Christopher J Shallue, and George E Dahl. 2019. Faster neural network training with data echoing. *arXiv preprint arXiv:1907.05550* (2019).
- [21] Fahim Chowdhury, Yue Zhu, Todd Heer, Saul Paredes, Adam Moody, Robin Goldstone, Kathryn Mohror, and Weikuan Yu. 2019. I/O characterization and performance evaluation of beegfs for deep learning. In *Proceedings of the 48th International Conference on Parallel Processing*. 1–10.
- [22] Henggang Cui, James Cipar, Qirong Ho, Jin Kyu Kim, Seunghak Lee, Abhimanu Kumar, Jinliang Wei, Wei Dai, Gregory R. Ganger, Phillip B. Gibbons, Garth A. Gibson, and Eric P. Xing. 2014. Exploiting Bounded Staleness to Speed Up Big Data Analytics. In *USENIX Annual Technical Conference*. 37–48.
- [23] Henggang Cui, Hao Zhang, Gregory R. Ganger, Phillip B. Gibbons, and Eric P. Xing. 2016. GeePS: Scalable deep learning on distributed GPUs with a GPU-specialized parameter server. In *European Conference on Computer Systems*. Article 4, 16 pages.
- [24] Amit Daniely, Nevena Lazić, Yoram Singer, and Kunal Talwar. 2017. Short and Deep: Sketching and Neural Networks. *arXiv preprint arXiv:1710.07850* (2017).
- [25] Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Siwei Li, Li Chen, Michael E. Kounavis, and Duen Horng Chau. 2018. SHIELD: Fast, Practical Defense and Vaccination for Deep Learning Using JPEG Compression. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 196–204.
- [26] Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Marc'Aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, Quoc V. Le, and Andrew Y. Ng. 2012. Large Scale Distributed Deep Networks. In *Advances in Neural Information Processing Systems*, Vol. 25.
- [27] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition*. 248–255.
- [28] Emily L. Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. 2014. Exploiting linear structure within convolutional networks for efficient evaluation. In *Advances in Neural Information Processing Systems*, Vol. 1. 1269–1277.
- [29] Samuel Dodge and Lina Karam. 2016. Understanding how image quality affects deep neural networks. In *International Conference on Quality of Multimedia Experience*. 1–6.
- [30] Yinpeng Dong, Qi-An Fu, Xiao Yang, Tianyu Pang, Hang Su, Zihao Xiao, and Jun Zhu. 2020. Benchmarking Adversarial Robustness on Image Classification. In *IEEE Conference on Computer Vision and Pattern Recognition*. 318–328.
- [31] Adam Dziedzic, John Paparrizos, Sanjay Krishnan, Aaron Elmore, and Michael Franklin. 2019. Band-limited Training and Inference for Convolutional Neural Networks. In *International Conference on Machine Learning*, Vol. 97. 1745–1754.
- [32] Gintare Karolina Dziugaite, Zoubin Ghahramani, and Daniel M. Roy. 2016. A study of the effect of JPG compression on adversarial images. *arXiv preprint arXiv:1608.00853* (2016).
- [33] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. 2010. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision* 88 (2010), 303–338.
- [34] Dan Feldman, Melanie Schmidt, and Christian Sohler. 2013. Turning big data into tiny data: Constant-size coresets for k-means, PCA and projective clustering. In *Symposium on Discrete Algorithms*. 1434–1453.
- [35] Flatbuffers [n.d.]. Flatbuffers. https://google.github.io/flatbuffers/flatbuffers_benchmarks.html. Accessed: 07-26-2021.
- [36] Sajjad Fouladi, John Emmons, Emre Orbay, Catherine Wu, Riad S. Wahby, and Keith Winstein. 2018. Salsify: Low-Latency Network Video through Tighter Integration between a Video Codec and a Transport Protocol. In *USENIX Networked Systems Design and Implementation*. 267–282.
- [37] Dan Fu and Gabriel Guimaraes. 2016. *Using Compression to Speed Up Image Classification in Artificial Neural Networks*. Technical Report.
- [38] GCPDiskBandwidth 2021. GCPDiskBandwidth. <https://cloud.google.com/compute/docs/disks/performance>. Accessed: 07-26-2021.
- [39] GCPNetworkBandwidth 2021. GCPNetworkBandwidth. <https://cloud.google.com/compute/docs/network-bandwidth>. Accessed: 07-26-2021.
- [40] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. 2019. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*.
- [41] Google. 2021. Google Cloud. <https://cloud.google.com/>. Accessed: 07-26-2021.
- [42] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2017. Accurate, large minibatch SGD: Training ImageNet in 1 hour. *arXiv preprint arXiv:1706.02677* (2017).
- [43] Lionel Gueguen, Alex Sergeev, Ben Kadlec, Rosanne Liu, and Jason Yosinski. 2018. Faster neural networks straight from JPEG. In *Advances in Neural Information Processing Systems*, Vol. 31.
- [44] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A. Horowitz, and William J. Dally. 2016. EIE: efficient inference engine on compressed deep neural network. In *International Symposium on Computer Architecture*. 243–254.
- [45] Song Han, Huizi Mao, and William J. Dally. 2016. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *International Conference on Learning Representations* (2016).
- [46] Song Han, Jeff Pool, John Tran, and William Dally. 2015. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, Vol. 28.
- [47] Mor Harchol-Balter. 2013. *Performance modeling and design of computer systems: queuing theory in action* (1st ed.). Cambridge University Press.

- [48] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [49] David A Huffman. 1952. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE* 40 (1952), 1098–1101.
- [50] Kyuhyeon Hwang and Wonyong Sung. 2014. Fixed-point feedforward deep neural network design using weights +1, 0, and -1. In *Workshop on Signal Processing Systems*.
- [51] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. 2019. Adversarial Examples Are Not Bugs, They Are Features. In *Advances in Neural Information Processing Systems*, Vol. 32.
- [52] Xianyan Jia, Shutao Song, Wei He, Yangzihao Wang, Haidong Rong, Feihu Zhou, Liqiang Xie, Zhenyu Guo, Yuanzhou Yang, Liwei Yu, Tiegang Chen, Guangxiao Hu, Shaohuai Shi, and Xiaowen Chu. 2018. Highly scalable deep learning training system with mixed-precision: Training ImageNet in four minutes. In *Neural Information Processing Systems Workshop on Systems for ML*.
- [53] Norman P Jouppi, Doe Hyun Yoon, George Kurian, Sheng Li, Nishant Patil, James Laudon, Cliff Young, and David Patterson. 2020. A domain-specific supercomputer for training deep neural networks. *Commun. ACM* 63, 7 (2020), 67–78.
- [54] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, 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 Nix, 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. 2017. In-Datcenter Performance Analysis of a Tensor Processing Unit. In *International Symposium on Computer Architecture*, Vol. 45. 1–12.
- [55] JPEGTran 2015. JPEGTran libjpeg.txt. <https://github.com/cloudflare/jpegtran/blob/master/libjpeg.txt>. Accessed: 07-26-2021.
- [56] JPEGTranManPage [n.d.]. jpegtran(1) - Linux max page. <https://linux.die.net/man/1/jpegtran>. Accessed: 7-26-2021.
- [57] Zohar Karnin and Edo Liberty. 2019. Discrepancy, Coresets, and Sketches in Machine Learning. In *Conference on Learning Theory*, Vol. 99. 1–19.
- [58] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive growing of GANs for improved quality, stability, and variation. *International Conference on Learning Representations* (2018).
- [59] Mahesh Khadatare, Zoheb Khan, and Harun Bayraktar. 2020. Leveraging the Hardware JPEG Decoder and NVIDIA nvJPEG Library on NVIDIA A100 GPUs. <https://developer.nvidia.com/blog/leveraging-hardware-jpeg-decoder-and-nvjpeg-on-a100>. Accessed: 07-26-2021.
- [60] Shmuel Tomi Klein and Yair Wiseman. 2003. Parallel Huffman decoding with applications to JPEG files. *The Computer Journal* 46, 5 (2003), 487–497.
- [61] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 2013. 3D Object Representations for Fine-Grained Categorization. In *Workshop on 3D Representation and Recognition*.
- [62] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*.
- [63] Abhishek Vijaya Kumar and Muthian Sivathanu. 2020. Quiver: An Informed Storage Cache for Deep Learning. In *USENIX Conference on File and Storage Technologies*. 283–296.
- [64] Sameer Kumar, Victor Bitorff, Dehao Chen, Chiachen Chou, Blake Hechtman, HyoukJoong Lee, Naveen Kumar, Peter Mattson, Shibo Wang, Tao Wang, et al. 2019. Scale MLPerf-0.6 models on Google TPU-v3 Pods. *arXiv preprint arXiv:1909.09756* (2019).
- [65] Sameer Kumar, James Bradbury, Cliff Young, Yu Emma Wang, Anselm Levskaya, Blake Hechtman, Dehao Chen, HyoukJoong Lee, Mehmet Deveci, Naveen Kumar, Pankaj Kanwar, Shibo Wang, Skye Wanderman-Milne, Steve Lacy, Tao Wang, Tayo Oguntebi, Yazhou Zu, Yuanzhong Xu, and Andy Swing. 2021. Exploring the limits of Concurrency in ML Training on Google TPUs.
- [66] Thorsten Kurth, Sean Treichler, Joshua Romero, Mayur Mudigonda, Nathan Luehr, Everett Phillips, Ankur Mahesh, Michael Matheson, Jack Deslippe, Massimiliano Fatica, Prabhat, and Michael Houston. 2018. Exascale Deep Learning for Climate Analytics. In *International Conference for High Performance Computing, Networking, Storage, and Analysis*. 1–12.
- [67] Youngeun Kwon and Minsoo Rhu. 2018. Beyond the Memory Wall: A Case for Memory-Centric HPC System for Deep Learning. In *IEEE/ACM International Symposium on Microarchitecture*. 148–161.
- [68] Hongshan Li, Yu Guo, Zhi Wang, Shutao Xia, and Wenwu Zhu. 2019. AdaCompress: Adaptive Compression for Online Computer Vision Services. *ACM International Conference on Multimedia*, 2440–2448.
- [69] Jing Li, Hung-Wei Tseng, Chunbin Lin, Yannis Papakonstantinou, and Steven Swanson. 2016. HippogriffDB: Balancing I/O and GPU bandwidth in big data analytics. *Very Large Data Bases Conference* 9, 14 (2016), 1647–1658.
- [70] Edo Liberty. 2013. Simple and deterministic matrix sketching. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 581–588.
- [71] Hyeontaek Lim, David G. Andersen, and Michael Kaminsky. 2019. 3LC: Lightweight and Effective Traffic Compression for Distributed Machine Learning. In *Conference on Systems and Machine Learning*.
- [72] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision*.
- [73] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *European conference on computer vision*.
- [74] Yujun Lin, Song Han, Huzi Mao, Yu Wang, and William J. Dally. 2018. Deep gradient compression: Reducing the communication bandwidth for distributed training. In *International Conference on Learning Representations*.
- [75] John D. C. Little. 1961. A Proof for the Queuing Formula: $L = \lambda W$. *Operations Research* (1961).
- [76] Zihao Liu, Qi Liu, Tao Liu, Nuo Xu, Xue Lin, Yanzhi Wang, and Wujie Wen. 2019. Feature Distillation: DNN-Oriented JPEG Compression Against Adversarial Examples. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [77] Zihao Liu, Tao Liu, Wujie Wen, Lei Jiang, Jie Xu, Yanzhi Wang, and Gang Quan. 2018. DeepN-JPEG: A Deep Neural Network Favorable JPEG-Based Image Compression Framework. In *Annual Design Automation Conference*.
- [78] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. 2018. Shufflenet v2: Practical guidelines for efficient CNN architecture design. In *European Conference on Computer Vision*.
- [79] Shin Matsushima, S.V.N. Vishwanathan, and Alexander J. Smola. 2012. Linear support vector machines via dual cached loops. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [80] Jim McDonnell. 2012. Large Directory Causes LS to Hang. <http://unixet.co.uk/2012/05/20/large-directory-causes-ls-to-hang/>. Accessed: 07-26-2021.
- [81] Qi Meng, Wei Chen, Yue Wang, Zhi-Ming Ma, and Tie-Yan Liu. 2017. Convergence analysis of distributed stochastic gradient descent with shuffling. In *Advances in Neural Information Processing Systems*, Vol. 31.
- [82] Fabian Mentzer, George D Toderici, Michael Tschannen, and Eirikur Agustsson. 2020. High-Fidelity Generative Image Compression. *Advances in Neural Information Processing Systems* 33 (2020).
- [83] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2018. Mixed precision training. *International Conference on Learning Representations* (2018).
- [84] MLPerfHPCv0.7 2020. MLPerfHPCv0.7 Results. <https://mlcommons.org/en/news/mlperf-hpc-v07>. Accessed: 07-26-2021.
- [85] MLPerfv0.7 2020. MLPerf Training v0.7 Results. <https://mlcommons.org/en/news/mlperf-training-v07/>. Accessed: 07-26-2021.
- [86] Jack Moffitt. 2001. Ogg Vorbis—Open, Free Audio—Set Your Media Free. *Linux Journal* (2001).
- [87] Jayashree Mohan, Amar Phanishayee, Ashish Raniwala, and Vijay Chidambaram. 2021. Analyzing and Mitigating Data Stalls in DNN Training. *Very Large Data Bases Conference* 14, 5 (2021).
- [88] Derek G. Murray, Jiri Simsa, Ana Klimovic, and Ihor Indyk. 2021. tf.data: A Machine Learning Data Processing Framework. *arXiv:2101.12127*
- [89] NVIDIA. 2018. DALI. <https://github.com/NVIDIA/DALI>. Accessed: 07-26-2021.
- [90] nvJPEG [n.d.]. nvJPEG. <https://developer.nvidia.com/nvjpeg>. Accessed: 07-26-2021.
- [91] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*.
- [92] Gennady Pekhimenko, Chuanxiong Guo, Myeongjae Jeon, Peng Huang, and Lidong Zhou. 2018. Tersecades: Efficient data compression in stream processing. In *USENIX Annual Technical Conference*. 307–320.
- [93] Gennady Pekhimenko, Vivek Seshadri, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry. 2012. Base-delta-immediate compression: practical data compression for on-chip caches. In *International Conference on Parallel Architectures and Compilation Techniques*.
- [94] Xingchao Peng, Judy Hoffman, X. Yu Stella, and Kate Saenko. 2016. Fine-to-coarse knowledge transfer for low-res image classification. In *IEEE International Conference on Image Processing*. 3683–3687.
- [95] Raymond Perrault, Yoav Shoham, Erik Brynjolfsson, Jack Clark, John Etchemendy, Barbara Grosz, Terah Lyons, James Manyika, Saurabh Mishra,

- , and Juan Carlos Niebles. 2019. *The AI Index 2019 Annual Report*. Technical Report. Stanford University.
- [96] Protobuf [n.d.]. Protocol Buffers. <https://developers.google.com/protocol-buffers/>. Accessed: 07-26-2021.
- [97] Sarunya Pumma, Min Si, Wu-Chun Feng, and Pavan Balaji. 2019. Scalable Deep Learning via I/O Analysis and Optimization. *Transactions on Parallel Computing* (2019).
- [98] PyTorchWebDataset 2020. [RFC] Add tar-based IterableDataset implementation to PyTorch. <https://github.com/pytorch/pytorch/issues/38419>. Accessed: 07-26-2021.
- [99] RecordIODataset [n.d.]. Create a Dataset Using RecordIO. <https://mxnet.apache.org/api/faq/recordio> and https://gluon-cv.mxnet.io/build/examples_datasets/recordio.html. Accessed: 07-26-2021.
- [100] Vijay Janapa Reddi, Christine Cheng, David Kanter, Peter Mattson, Guenther Schmuelling, Carole-Jou Wu, Brian Anderson, Maximilien Breughe, Mark Charlebois, William Chau, et al. 2020. Mlperf inference benchmark. In *ACM/IEEE International Symposium on Computer Architecture*. IEEE, 446–459.
- [101] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [102] Heiko Schwarz, Detlev Marpe, and Thomas Wiegand. 2007. Overview of the scalable video coding extension of the H. 264/AVC standard. *IEEE Transactions on Circuits and Systems for Video Technology* (2007).
- [103] Abu Sebastian, Manuel Le Gallo, Riduan Khaddam-Aljameh, and Evangelos Eleftheriou. 2020. Memory devices and applications for in-memory computing. *Nature Nanotechnology* (2020).
- [104] Jonathon Shlens. 2014. A tutorial on principal component analysis. *arXiv preprint arXiv:1404.1100* (2014).
- [105] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations* (2015).
- [106] Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. 2001. The JPEG 2000 still image compression standard. *IEEE Signal Processing Magazine* (2001).
- [107] Stoyan Stefanov. 2021. *Book of Speed*.
- [108] 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 *IEEE Conference on Computer Vision and Pattern Recognition*.
- [109] Tensorflow. 2020. `imagenet_to_gcs.py`. https://github.com/tensorflow/tpu/tree/master/tools/datasets/imagenet_to_gcs.py. Accessed: 07-26-2021.
- [110] tf.data 2021. `tf.data`: Build TensorFlow input pipelines. <https://www.tensorflow.org/guide/data>. Accessed: 07-26-2021.
- [111] TFOp 2021. Create an op. https://www.tensorflow.org/guide/create_op. Accessed: 07-26-2021.
- [112] TFRecords 2021. `TFRecord` and `tf.train.Example`. https://www.tensorflow.org/tutorials/load_data/tfrecord. Accessed: 07-26-2021.
- [113] George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, and Michele Covell. 2017. Full resolution image compression with recurrent neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [114] Robert Torfason, Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. 2018. Towards image understanding from deep compression without decoding. *arXiv preprint arXiv:1803.06131* (2018).
- [115] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Herve Jegou. 2019. Fixing the train-test resolution discrepancy. In *Neural Information Processing Systems*.
- [116] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. 2018. The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data* (2018).
- [117] Matej Ulicny and Rozenn Dahyot. 2017. On using CNN with DCT based image data. In *Irish Machine Vision and Image Processing Conference*.
- [118] Igor Vasiljevic, Ayan Chakrabarti, and Gregory Shakhnarovich. 2016. Examining the impact of blur on recognition by convolutional networks. *arXiv preprint arXiv:1611.05760* (2016).
- [119] Gregory K. Wallace. 1992. The JPEG still picture compression standard. *Transactions on Consumer Electronics* (1992).
- [120] Haoan Wang, Xindi Wu, Zeyi Huang, and Eric P. Xing. 2020. High-Frequency Component Helps Explain the Generalization of Convolutional Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [121] Yu Emma Wang, Gu-Yeon Wei, and David Brooks. 2020. A Systematic Methodology for Analysis of Deep Learning Hardware and Software Platforms. In *Machine Learning and Systems*.
- [122] Zeke Wang, Kaan Kara, Hantian Zhang, Gustavo Alonso, Onur Mutlu, and Ce Zhang. 2019. Accelerating Generalized Linear Models with MLWeaving: A One-Size-Fits-All System for Any-Precision Learning. *Very Large Data Bases Conference* 12, 7 (2019), 807–821.
- [123] Zhou Wang, Eero P. Simoncelli, and Alan C. Bovik. 2003. Multiscale structural similarity for image quality assessment. In *Conference on Signals, Systems & Computers*, 2003, Vol. 2. 1398–1402.
- [124] Jianqiao Wangni, Jialei Wang, Ji Liu, and Tong Zhang. 2018. Gradient sparsification for communication-efficient distributed optimization. In *Advances in Neural Information Processing Systems*, Vol. 31.
- [125] Sage A. Weil, Scott A. Brandt, Ethan L. Miller, Darrell D.E. Long, and Carlos Maltzahn. 2006. Ceph: A scalable, high-performance distributed file system. In *USENIX Symposium on Operating Systems Design and Implementation*. 307–320.
- [126] André Weisenberger and Bertil Schmidt. 2018. Massively parallel Huffman decoding on GPUs. In *International Conference on Parallel Processing*.
- [127] Brent Welch, Marc Unangst, Zainul Abbasi, Garth Gibson, Brian Mueller, Jason Small, Jim Zelenka, and Bin Zhou. 2008. Scalable Performance of the Panasas Parallel File System. In *Conference on File and Storage Technologies*. 1–17.
- [128] Wei Wen, Cong Xu, Feng Yan, Chungpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. 2017. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In *Advances in Neural Information Processing Systems*.
- [129] David P. Woodruff. 2014. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science* (2014).
- [130] Wm. A. Wulf and Sally A. McKee. 1995. Hitting the Memory Wall: Implications of the Obvious. *SIGARCH Computer Architecture News* (1995).
- [131] Yuhui Xu, Yongzhuang Wang, Aojun Zhou, Weyao Lin, and Hongkai Xiong. 2018. Deep neural network compression with single and multiple level quantization. In *AAAI Conference on Artificial Intelligence*, Vol. 32.
- [132] Zhi-Qin John Xu, Yaoyu Zhang, and Yanyang Xiao. 2019. Training behavior of deep neural network in frequency domain. In *International Conference on Neural Information Processing*.
- [133] Masafumi Yamazaki, Akihiko Kasagi, Akihiro Tabuchi, Takumi Honda, Masahiro Miwa, Naoto Fukumoto, Tsuguchika Tabaru, Atsushi Ike, and Kohta Nakashima. 2019. Yet Another Accelerated SGD: ResNet-50 Training on ImageNet in 74.7 seconds. *arXiv preprint arXiv:1903.12650* (2019).
- [134] Eddie Yan, Kaiyuan Zhang, Xi Wang, Karin Strauss, and Luis Ceze. 2017. Customizing Progressive JPEG for Efficient Image Storage. In *Workshop on Hot Topics in Storage and File Systems*.
- [135] Dong Yin, Raphael Gontijo Lopes, Jon Shlens, Ekin Dogus Cubuk, and Justin Gilmer. 2019. A Fourier Perspective on Model Robustness in Computer Vision. In *Advances in Neural Information Processing Systems*, Vol. 32.
- [136] Chris Ying, Sameer Kumar, Dehao Chen, Tao Wang, and Youlong Cheng. 2018. Image classification at supercomputer scale. In *Neural Information Processing Systems Workshop on Systems for ML*.
- [137] Yang You, Zhao Zhang, Cho-Jui Hsieh, James Demmel, and Kurt Keutzer. 2018. ImageNet training in minutes. In *International Conference on Parallel Processing*.
- [138] Cyril Zhang, Kunal Talwar, Naman Agarwal, Rohan Anil, and Tomer Koren. 2020. Stochastic Optimization with Laggard Data Pipelines. In *Advances in Neural Information Processing Systems*, Vol. 33. 10282–10293.
- [139] Hao Zhang, Zeyu Zheng, Shizhen Xu, Wei Dai, Qirong Ho, Xiaodan Liang, Zhiting Hu, Jinliang Wei, Pengtao Xie, and Eric P. Xing. 2017. Poseidon: An efficient communication architecture for distributed deep learning on GPU clusters. In *USENIX Annual Technical Conference*. 181–193.
- [140] Stephan Zheng, Yang Song, Thomas Leung, and Ian Goodfellow. 2016. Improving the Robustness of Deep Neural Networks via Stability Training. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4480–4488.
- [141] Hongyu Zhu, Mohamed Akrouf, Bojian Zheng, Andrew Pelegris, Amar Phanshaye, Bianca Schroeder, and Gennady Pekhimenko. 2018. TBD: Benchmarking and analyzing deep neural network training. *arXiv preprint arXiv:1803.06905* (2018).
- [142] Marcin Zukowski, Sandor Heman, Niels Nes, and Peter Boncz. 2006. Super-scalar RAM-CPU cache compression. In *International Conference on Data Engineering*.