

# Lawrence Berkeley National Laboratory

## LBL Publications

### Title

I/O in Machine Learning Applications on HPC Systems: A 360-degree Survey

### Permalink

<https://escholarship.org/uc/item/8tc828xw>

### Authors

Lewis, Noah

Bez, Jean Luca

Byna, Surendra

### Publication Date

2025-03-07

### DOI

10.1145/3722215

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# I/O in Machine Learning Applications on HPC Systems: A 360-degree Survey

NOAH LEWIS, The Ohio State University, and Louisiana State University, USA

JEAN LUCA BEZ, Lawrence Berkeley National Laboratory, USA

SUREN BYNA, The Ohio State University, USA

Growing interest in Artificial Intelligence (AI) has resulted in a surge in demand for faster methods of Machine Learning (ML) model training and inference. This demand for speed has prompted the use of high performance computing (HPC) systems that excel in managing distributed workloads. Because data is the main fuel for AI applications, the performance of the storage and I/O subsystem of HPC systems is critical. In the past, HPC applications accessed large portions of data written by simulations or experiments or ingested data for visualizations or analysis tasks. ML workloads perform small reads spread across a large number of random files. This shift of I/O access patterns poses several challenges to modern parallel storage systems. In this paper, we survey I/O in ML applications on HPC systems, and target literature within a 6-year time window from 2019 to 2024. We define the scope of the survey, provide an overview of the common phases of ML, review available profilers and benchmarks, examine the I/O patterns encountered during offline data preparation, training, and inference, and explore I/O optimizations utilized in modern ML frameworks and proposed in recent literature. Lastly, we seek to expose research gaps that could spawn further R&D.

CCS Concepts: • **Information systems** → **Information storage systems**; *Storage architectures; Hierarchical storage management*; • **Computing methodologies** → *Artificial intelligence; Machine learning*.

Additional Key Words and Phrases: I/O access pattern, HPC I/O, storage, machine learning

## 1 Introduction

Existing High-Performance Computing (HPC) Input and Output (I/O) research has focused on techniques to optimize performance when running traditional HPC application workloads, which typically include simulations and checkpointing partial results [53]. Because of the increased popularity of Machine Learning (ML) workloads, there is a rising demand for I/O systems that can effectively accommodate their distinct I/O access patterns. Write operation bursts commonly dominate traditional workloads; however, ML workloads are usually read-intensive and use many small files [65]. Due to the absence of a well-established consensus on the preferred I/O stack for ML workloads, numerous developers resort to crafting their own ad hoc algorithms and storage systems to accommodate the specific requirements of their applications [82]. This can result in suboptimal application performance due to the under-utilization of the storage system, prompting the necessity for novel I/O optimization methods tailored to the demands of ML workloads.

In Fig. 1, we show the I/O stack used to run ML workloads (on the right side) compared to the traditional HPC I/O stack (on the left side) [53]. The traditional HPC I/O stack has been developed to support massive parallelism. Applications are typically written in C/C++/Fortran languages that use the MPI programming model and I/O middleware, such as HDF5 [22, 163] and netCDF [106], which internally use MPI-IO and POSIX-IO interfaces to move (read or write) data between main memory and parallel file systems (PFS). In contrast, ML workloads are typically written in productivity-oriented languages, such as Python, and use PyTorch and TensorFlow frameworks.

---

Authors' addresses: Noah Lewis, lewis.3621@buckeyemail.osu.edu, The Ohio State University, Baton Rouge, Louisiana, and Louisiana State University, Columbus, Ohio, USA; Jean Luca Bez, jlbez@lbl.gov, Lawrence Berkeley National Laboratory, Berkeley, California, USA; Suren Byna, byna.1@osu.edu, The Ohio State University, Columbus, Ohio, USA.

Although many ML workloads still use the traditional I/O stack on HPC systems, some are shifting to use remote cloud computing systems that offer services such as data warehouses to store and retrieve large datasets. Cloud platforms such as Google Cloud Platform, AWS, and Microsoft Azure offer access to large datasets, including Google Analytics data [87] hosted on BigQuery [88], the 1000 Genomes Project [5] hosted on Amazon Redshift [6], and NYC Taxi & Limousine Commission data [122] hosted on Azure Synapse Analytics [121]. HPC workloads that aim to use these datasets require near-first-class integration with the cloud to enable efficient data transfers to HPC systems. When using the traditional I/O stack, high bandwidth and low latency data access are possible due to the locality of the data and the highly optimized stack. In contrast, access to data stored in a data warehouse is often accomplished through a network interface such as REST [79], SOAP [17], or ODBC [120]. Given the high latency inherent in accessing data stored in a remote storage location, optimizing data retrieval is crucial for

reasonable performance. Despite the lower I/O performance imposed when using a data warehouse, they are still preferred in many scenarios because they provide a centralized data storage location, facilitate data integration from various sources, and enable data governance, which can result in higher-quality data. High-quality data is essential when training ML models because it directly influences their accuracy, reliability, and generalizability [56]. In addition, the monetary cost needed to purchase and maintain an HPC system can far outweigh the cost of using managed cloud services [23, 60]. Furthermore, integration of HPC systems with available cloud resources is necessary due to the cloud being used as a data storage solution from which datasets can be publicly shared and used by all.

The differences listed above in the I/O stacks for ML applications using the traditional PFS compared to cloud computing systems motivate the development of novel storage system designs and I/O performance optimization techniques to adapt to the I/O operations commonly found in ML processes. Ideally, the methods would be transparent to the application developers. Notably, many scientists may lack formal training in software development, as evidenced by Paul et al. [66]. The study revealed that ML workloads are becoming ever more popular and that developers from various scientific domains beyond computer science often overlook the utilization of node-local storage (i.e., burst buffers) within HPC systems. Further evidence for the increasing interest in distributed training is the number of posts on the popular tech website StackOverflow. The cumulative number of posts discussing distributed training increased staggeringly from 257 posts in 2015 to 2,061 posts in 2024 [29]. The increasing number of ML workloads that are running on HPC systems prompts the need for transparent I/O optimization techniques to decrease the complexity required for scientists to write software that takes advantage of the resources offered by HPC systems.

Furthermore, scientists, I/O performance enthusiasts, and HPC practitioners can gain significant insights from benchmarks that accurately emulate ML workloads and profilers that capture critical I/O information from the underlying storage systems. Using these tools, individuals can find potential I/O bottlenecks within their ML applications. Access to

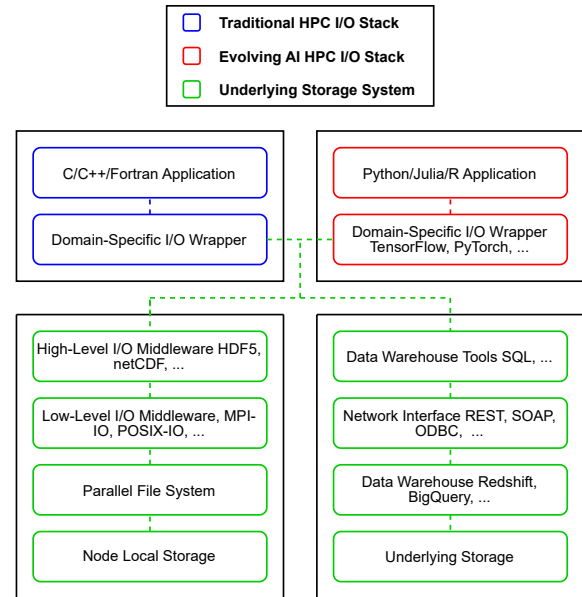


Fig. 1. Traditional HPC I/O stack vs evolving ML I/O HPC stack (adapted from [53]).

realistic benchmarking tools and fine-grained profilers is also invaluable for HPC storage research, allowing researchers to analyze and emulate the diverse range of I/O operations found in ML workloads. Currently, I/O analysis tools often fall short in a number of ways, including the inability to identify I/O requests at the thread level and only capturing traces of specific I/O interfaces (i.e., POSIX and STDIO). This makes it challenging to find a general-purpose tool applicable to a wide variety of workloads, which can reduce the consistency in which I/O traces are captured. In addition, many tools (i.e., Darshan [69] and Scalasca [84]) target MPI applications limiting their compatibility and effectiveness with other communication libraries (i.e., GLOO [77], Spark [11], and Dask [33]). It should be noted that Darshan is capable of capturing I/O traces from applications that do not use MPI through the `NONMPI` environment variable. However, this results in the output of a single DXT log for each process instead of the more common single file output that captures all I/O traces from an application using MPI. In some cases, this leaves developers with the task of manually merging the results from their non-MPI applications in order to conduct further analysis.

The topic of ML I/O and whether it has a significant role in poor performance has been a debate within the HPC community. Many point out that HPC systems often have enough node-local storage to cache the dataset, resulting in minimal time spent waiting on I/O [66, 76, 78, 142]. Others mention that not all HPC systems have large amounts of node local storage or the datasets used are simply too large to be cached [76, 78, 142, 168]. In addition, due to the random sampling that commonly occurs during ML training, modern PFSs can often become an I/O bottleneck [72, 76, 134]. Ultimately, whether I/O is a significant issue is subjective to the workload at hand. It depends on various factors such as dataset size, storage capacity, storage system bandwidth, I/O library implementation and their configurations, and I/O access patterns of applications. Furthermore, many ML application developers are not HPC storage experts; hence, even if their system is capable of providing great performance, transparent solutions such as caching and prefetching are needed to ensure the HPC systems are utilized effectively. Transparent solutions are of particular interest due to the increasing number of heterogeneous processors used within ML applications. These increase the complexity of optimally utilizing available compute resources. For these reasons, we believe that I/O and data management are key areas of interest to the ML community.

There are several common phases in an ML application lifecycle: *data generation*, *data preparation*, *training*, and *inference*. The data management found within each phase varies widely due to different I/O access patterns, data sources, and file formats. We provide a taxonomy of the data management found throughout the ML lifecycle in Fig. 3. The ML lifecycle begins with data generation, which involves collecting data from various sources (e.g., simulation results or web scraping), which are then stored using a chosen file format (e.g., `.tfrecord`, `.csv`, or `.png`) and storage location (e.g., SSD, disk, cloud, or database). This is commonly followed by an offline data preparation phase, which performs a series of transformations and reductions (i.e., filters, aggregations, or generalizations) over the collected data to increase the training efficiency and accuracy of the ML model. The training process involves reading samples from the data sources, which are then fed to computations. The order in which the samples are read largely depends on the chosen optimizer (e.g., Stochastic Gradient Descent and Adaptive Moment Estimation or Adam, etc.) and training distribution strategy (e.g., model or data parallelism). Due to many optimizers requiring random sampling and storage systems being optimized for contiguous sequential reading, training I/O is a potential bottleneck. Once the model has been trained, it is often deployed and fed “new” data during the inference phase. It is common for the online and offline data preparation phases to occur before and during the inference phase as well. More details on each phase of ML can be found in Section 3, but it is important to note the wide range of factors that impact I/O and data management throughout the ML lifecycle, such as varying file formats, data sources, data preparation requirements, and data modalities, among

others. It is impractical to cover all possible combinations of the factors mentioned above, and therefore, in this paper, we aim to discuss the I/O approaches found in the most common ML applications running on HPC systems.

The structure of the paper is as follows. First, in Section 2, we define the scope of the survey and present trends in the cited literature. Then, in Section 3, we describe the data formats and modalities found in ML workloads and provide an overview of the typical phases of ML applications and the popular methods of distributing the model training on HPC systems. Next, in Section 4, we describe the benchmarks and profiling tools scientists and researchers can use to ensure proper system utilization under ML workloads. In Section 5, we describe the I/O access patterns found in the offline data preparation, training phase, and inference phases. The training phase is exemplified by analyzing the currently available representative I/O benchmarks. We then discuss, in Section 6, various I/O optimization techniques used in modern ML frameworks and those proposed in the recent research literature. Lastly, in Section 7, we discuss the gaps in the current research of ML I/O in HPC systems.

## 2 Scope of the Survey

ML applications use various environments, from local machines for small-scale applications to HPC and cloud systems for large-scale modeling. This survey aims to provide a comprehensive review of the I/O patterns in ML running on HPC systems. This means that I/O patterns and other data management solutions that may be common in other environments (e.g., cloud or local) are not relevant. For example, in Sec 5, we discuss only the I/O access patterns found in common HPC scenarios and exclude I/O access patterns that may be common in cloud or local environments.

GPU Mem. Hierarchy	File Format	Sample Caching	Sample Prefetching	Sample Shuffling	Asynchronous Sampling	Profiling	Benchmarking	Model Checkpointing	Model Storage Offloading	Scheduling	ML Analysis	Papers	References
×	×	×	×	×	×	×	×	×	×	×	×	6	[18, 19, 27, 50, 93, 156]
×	×	×	×	×	×	×	×	×	×	×	×	5	[3, 43, 53, 66, 107]
×	×	×	×	×	×	×	×	×	×	×	×	5	[9, 42, 64, 115, 139]
×	×	×	×	×	×	×	×	×	×	×	×	4	[49, 73, 97, 108]
×	×	×	×	×	×	×	×	×	×	×	×	3	[3, 51?]
×	×	×	×	×	×	×	×	×	×	×	×	3	[71, 94, 136]
×	×	×	×	×	×	×	×	×	×	×	×	3	[72, 91, 171]
×	×	×	×	×	×	×	×	×	×	×	×	2	[54, 167]
×	×	×	×	×	×	×	×	×	×	×	×	2	[59, 110]
×	×	×	×	×	×	×	×	×	×	×	×	2	[57, 145]
×	×	×	×	×	×	×	×	×	×	×	×	1	[47]
×	×	×	×	×	×	×	×	×	×	×	×	1	[103]
×	×	×	×	×	×	×	×	×	×	×	×	1	[126]
×	×	×	×	×	×	×	×	×	×	×	×	1	[32]

Table 1. I/O-related topics discussed in the surveyed ML papers.

To ensure a comprehensive review of the current literature, we have selected a 6-year time window from 2019 to 2024 and used ACM Digital Library and IEEE Xplore to filter by papers that mention the following keywords: “I/O Access Patterns”, “I/O pattern”, “HPC”, “Storage performance”, and “I/O performance in AI/ML applications”. We found additional papers through the citations of the work initially discovered using the digital libraries for a total of 39 papers. The various subjects explored in the papers surveyed are shown in Table 1. We selected papers from leading conferences and journals in HPC and AI areas including:

- The IEEE/ACM Int. Conference for High-Performance Computing, Networking, Storage and Analysis (SC)
- The IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID)
- ACM Very Large Database Conference (VLDB)

- ACM Conference on Cloud Computing, Big Data & Emerging Topics (JCC-BD&ET)
- ACM Transactions on Parallel Computing (TOPC)
- IEEE International Symposium on High-Performance Computer Architecture (HPCA)
- IEEE International Conference on Cluster Computing (CLUSTER)
- IEEE International Symposium on High-Performance Computer Architecture (HPCA)
- IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)
- IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC)
- IEEE Conference on High Performance Extreme Computing (HPEC)
- IEEE International Parallel and Distributed Processing Symposium (IPDPS)
- IEEE Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)
- IEEE International Conference on High Performance Computing & Simulation (HPCS)
- IEEE Transactions on Parallel and Distributed Systems (TPDS)
- IEEE Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)
- USENIX Networked Systems Design and Implementation (NSDI)
- NIPS International Conference on Neural Information Processing Systems (NeurIPS)
- SpringerOpen Journal of Big Data




Fig. 2. Trends in cited literature. Each color represents a year. The number to the right of each label represents the total number of papers associated with that label. The publication year and topic each have 1 paper. These two categories each sum to the total number of papers 39. The leftmost and rightmost categories sum to more than 39 (a paper can discuss more than one item).

The parallel coordinates diagram in Fig. 2 illustrates the popularity of various software tools, frameworks, topics, and file formats in the surveyed papers. TensorFlow was discussed in 23 papers, while PyTorch appeared in 30 papers, making these two ML frameworks the most popular by a significant margin. Training emerged as the most common area of focus, with a total of 32 papers, followed by checkpointing, with a total of 5 papers, and lastly, inference, with 2 papers. Furthermore, HDF5 and TFRecord were notably mentioned in 8 and 10 of the surveyed papers, respectively. Based on these results, there has been a strong focus on optimizing I/O for training ML models using TensorFlow and PyTorch over the last couple of years.

In this paper, our contributions are a presentation and analysis of the I/O patterns found when emulating two renowned ML workloads: BERT [37], and Unet3D [173]. In addition, we conduct a review of the currently proposed I/O optimization techniques through R&D efforts and present several areas where future research is needed. Understanding the I/O patterns found in ML applications allows model training speeds to be reduced. This contributes to the rapid

development of ML models, which can be used to perform tasks such as medical imaging analysis and help in the creation of autonomous vehicles. It should be noted that throughout the survey, we use the terms “AI” and “ML” synonymously, as is commonly done in the industry. However, “ML” is the more precise term referring to the concept in which machines extract knowledge from data and learn from it.

#### Summary #1

The analysis of the surveyed papers reveals a significant trend towards optimizing I/O during the training phase. Additionally, TensorFlow and PyTorch were by far the most popular ML frameworks. Lastly, the HDF5 and TFRecord file formats were the most popular.

### 3 ML Data: Data Formats, Modalities, and Common Phases

In this section, we describe the ML lifecycle and the data formats and modalities commonly found throughout this lifecycle. We also discuss the common phases ML applications undergo and the ways of distributing these phases to exploit the resources offered by HPC systems. We then discuss model checkpointing and its impact on training speeds.

ML applications typically undergo four phases during the ML lifecycle: **data generation**, **dataset preparation**, **training**, and **inference** [53]. We show a taxonomy of the data management operations and characteristics throughout this lifecycle in Fig. 3. During data generation, data is obtained through various sources and methods such as simulations, web scraping, and public datasets. During this phase, properties such as the dataset modality, file format, and data storage location are often decided. The data preparation phase is divided into various data preparation steps, including transformation, augmentation, and splitting. The training arm of the taxonomy consists of data I/O operations across common training algorithms and checkpointing tasks. Inference includes a taxonomy similar to that of the training phase. As can be seen, data is a prominent component throughout the entirety of the ML lifecycle.

#### 3.1 Data Formats and Dataset Modalities

Data is at the forefront of the ML lifecycle, and without the correct quantity and quality of data, machine learning models fail to achieve optimal performance and generalization, disabling their ability to provide accurate insights and solutions. Here, we describe the file formats and data modalities commonly used during model training and inference. The file formats described were chosen due to their frequent appearance in the cited literature or their connection with state-of-the-art ML training and preprocessing frameworks (e.g., *.tfrecord* – TensorFlow and *.parquet* – Apache Parquet).

File Format	Features					
	Parallel I/O	Chunking	Partial Compression	Schema	Binary	File Extension
TFRecord [161]	✗	✗	✗	✓ <sup>1</sup>	✓	.tfrecord
HDF5 [163]	✓	✓	✓	✓	✓	.h5 .hdf5 <sup>2</sup>
Apache Parquet [80]	✓	✓	✓	✓	✓	.parquet
RecordIO [10]	✗	✓	✓	✗	✓	.rec
NPZ [131]	✗	✗	✗	✗	✓	.npz
Image Formats	✗	✗	✗	✗	✓	.jpg .png
Video Formats	✗	✗	✗	✗	✓	.mp4 .avi
Audio Formats	✗	✗	✗	✗	✓	.mp3 .wav
Text Formats	✗	✗	✗	✗	✗	.txt .csv

Table 2. A Comparison of commonly used file formats in AI. The partial decompression property indicates that the file format library transparently performs the decompression operation, possibly after it is enabled or disabled. (<sup>1</sup>TFRecords are read and written using Protocol Buffers (protobufs), which require a schema. <sup>2</sup>While .h5 and .hdf5 are used commonly, HDF5 allows any file extension.)





of use and compatibility with common system software such as standard text editors, video players, audio players, and image viewers. However, in addition to their poor performance, these formats lack schema enforcement.

Schema enforcement gives the developer the ability to confine the domain of the file contents. This commonly ensures attempting to write data outside the input domain will fail. Without it, the risk to data integrity can lead to issues such as hard-to-find bugs caused by data corruption. Compression is a common technique used to reduce the size of data by encoding the information with fewer bits. Partial compression enables data loaders to read and write to subsets of a compressed file without having to load the file in its entirety. Due to ML applications commonly reading small random subsets of the data, the inability to read subsets of data without loading the entire file could create a bottleneck. Chunking refers to the ability to divide the data in the file into discrete, manageable “chunks.”

Apache Parquet and HDF5 enable a high level of parallelism by allowing the user to specify the unit (i.e., row group size and data page size for parquet and chunk size for HDF5) in which the data is stored. In addition, Apache Parquet stores data in a columnar format, allowing the file to be split into row groups and processed independently. HDF5, on the other hand, supports features such as chunking and collective I/O, which enables efficient parallel file processing. In many cases, such as when the data is stored in a Lustre PFS [86, 155], this enables parallel reads and writes across independent sections of the file system (e.g., using file striping). Formats such as RecordIO [10], NPZ [131], and TFRecord [161] store data in a more serialized manner. Due to the variety of possible optimizations offered by distributed file formats, making the appropriate selection is essential for optimizing performance on HPC systems.

The selection of file formats model developers can choose from is often constrained by the data modality in use. The dataset modality refers to the mode or type of data being stored (e.g., audio, video, image, etc.). Due to the large variety of ML model goals (e.g., image recognition, optical character recognition (OCR), relationship detection, etc.), modalities vary widely. The dataset modality has a significant impact on the viable file formats available and, therefore, the common I/O operations required to load and use the dataset. In Table 3, we list the dataset modalities commonly used for ML applications and three corresponding publicly available datasets often used for ML research, exemplifying each modality. The provided values are approximate and subject to variation depending on which version of the dataset is selected. For instance, the Charades [149] dataset is available at different video resolutions, while the CIFAR-10 [102] dataset is available as a custom binary format or as a binary containing a Python “Pickled” object. Additionally, in

Modality	Dataset	Total Files	Total Size	Used File Formats
Image	MNIST [36]	4	101.16 MiB	Custom
	CIFAR-100 [102]	60,000	184 MiB	Custom
	ImageNet [35]	1,431,167	150 GiB	.jpg, .txt
Audio	UrbanSound8K [141]	18733	6.60 GiB	.wav, .csv
	LibriSpeech [132]	10,000	60 GiB	.flac, .txt
	AudioSet [85]	2,084,320	2.4 GiB	.tfrecord, .csv
Video	Charades [149]	9,900	55 GiB	.csv, .txt, .mat, .jpg Frames
	UCF101 [164]	13,320	6.69 GiB	.avi
	HMDB [104]	7,000	2 GiB	.txt, .avi
Time Series	NAB [105]	58	9 MiB	.csv
	M4 [113]	13	310 MiB	.csv
	UCR Time Series Archive [34]	416	853 MiB	.tsv
Graph	Cora [116]	1	270 KiB	.csv
	Protein-Protein Int. [2]	1	3,661 KiB	.csv
	Amazon Co-purchasing [170]	1	954,597 KiB	.txt
Text	refinedweb [133]	2,858	2.8 TiB	.parquet
	The Pile [83]	17,103,059	825 GiB	.txt
	The Stack v2 [112]	3,028,000,000	67.5 TiB	Code Files

Table 3. Comparison of different ML dataset modalities.

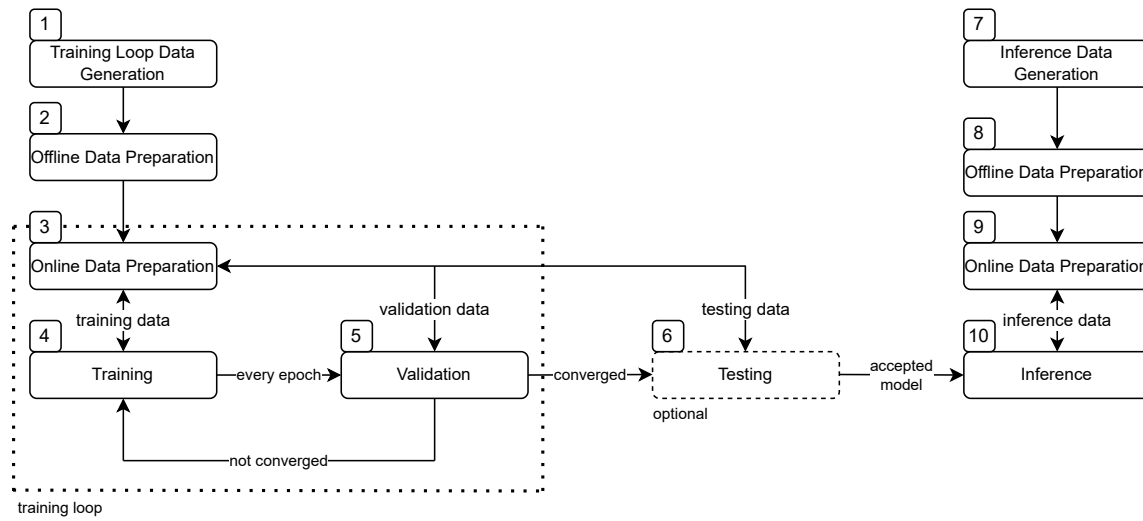


Fig. 4. Common phases during the ML lifecycle. The numbers in the top-left corners of each phase indicate their chronological order. Note that while phases 1 and 2 are sequential, phases 1, 2, and 7, 8 can occur in parallel.

certain cases, like AudioSet, links to the video data are provided rather than the raw video data itself, allowing users to choose the preferred video format.

### 3.2 Common ML Phases

The **dataset preparation** phase is critical in transforming the data into a form that is suitable for model training and inference. It is a critical step in the ML lifecycle and is commonly performed before both the training and inference phases. **Offline dataset preparation** is preprocessing that can be done before the model training or inference phase begins. In contrast, **online dataset preparation** is preprocessing that occurs in real-time during the inference or training phase. There are several major data transformation techniques commonly used during both offline and online dataset preparation to aid in model training and inference. These are listed below:

- **Data Smoothing** is a method used to remove noise from the dataset, ultimately allowing models to identify patterns in the dataset more accurately [154].
- **Feature Engineering** is used to construct new features from existing attributes in the dataset [165].
- **Data Generalization** defines hierarchies within the features of the dataset, allowing the data to be analyzed at varying levels of granularity or abstraction [130].
- **Data Aggregation** is used to summarize or reduce features from the dataset to produce meaningful statistics such as mean, median, or sum among many others [95].
- **Data Discretization** is used to divide continuous values into discrete ones, making analysis easier and improving the efficiency of many algorithms [98].
- **Data Normalization** is used to scale the data changing absolute values while preserving the relative proportions or ratios between data points [44].
- **Data Augmentation** is a strategy used to generate new data from existing data [82].

Augmenting the dataset allows for training models on a larger dataset, which can lead to better performance. The augmentation methods employed often depend on the dataset modality. For text datasets, examples of augmentations include random swap, random insertion, random deletion, and random synonym replacement [148]. In contrast, for image datasets, augmentations include flipping, cropping, and rotation [147]. The data transformation techniques used

affect both data preprocessing speeds and the amount of data needing to move from the preprocessing phase to the training phase. The variety of data transformation techniques that may need to be applied to the dataset causes the file format to be significant. It has a large impact on data transformation speeds as it affects the time needed to apply a given algorithm to the dataset, e.g., when performing data aggregation where the dataset is filtered based on a feature value, processing data in a columnar format (instead of row-based) can be more efficient due to the data locality [41].

The **training** phase that occurs during the ML lifecycle can further be categorized, as seen in Fig. 4, into a **training** phase and a **validation** phase. The figure shows the training loop and the inference phase. The training loop is often run iteratively for a certain number of epochs until the desired accuracy is achieved. An epoch is defined as one complete pass through the entire dataset [21]. It is the responsibility of the model designers to choose the number of epochs that strikes a balance between enabling the model to find meaningful relationships between the input and target datasets while avoiding excessive epochs that lead to over-fitting. Over-fitting occurs when the model becomes excessively entwined with the training data, prioritizing memorization over the ability to generalize and make accurate predictions on new, unseen data. Furthermore, because the entire dataset is read multiple times, model designers need to select the proper data format for their dataset that allows efficient data retrieval.

There are three datasets known as the training, validation, and testing datasets [100] used during the **training**, **validation**, **testing** phases. Although each dataset possesses distinct content, they commonly use the same modality and file formats. All three datasets often require preparation (i.e., cleaning, augmentation, etc.) before they can be used in later stages of the ML lifecycle. Dataset splitting involves dividing the dataset into training, validation, and test sets. Common techniques for this partitioning include random sampling [119],  $k$ -fold cross-validation [8], and stratified sampling [58]. Due to the large amounts of potential I/O and computation required to prepare the data effectively, parallelization techniques are needed to ensure proper system utilization. ML applications typically use the parallelization tools provided by their primary ML framework (e.g., PyTorch [135] and TensorFlow [158]) to parallelize data processing. Once the data has been prepared, it can be fed to the model for training, validation, testing or inference.

During the **training** phase, data from the training dataset is read into memory and fed to the model. The order in which the data is fed to the model is largely impacted by the optimization algorithm of choice. Gradient descent algorithms are a popular set of optimization algorithms used to train ML models. Due to stochastic gradient descent (SGD) being the most popular gradient descent algorithm for large datasets [125], the algorithms we will present are batch gradient descent (BGD) [124], SGD [15] and mini-batch gradient descent (MBGD) [16]. SGD is a specific case of MBGD, and in the future, we will refer to both as SGD unless a distinction is necessary. Deciding which algorithm to use depends on the dataset size, the computational resources, and the underlying model's properties. It is important to note that gradient descent algorithms vary in their mathematical properties [70]. At every epoch, each of the listed algorithms feeds the entirety of the dataset to the model. However, the frequency at which the model is updated, and the order in which the dataset is fed to the model differ.

BGD is the most straightforward gradient descent algorithm. This technique is computationally expensive and is often used only when the entire dataset can fit into memory. The model is trained against the whole dataset before it can perform one update [140]. If the model cannot cache the entire dataset in a fast storage tier, it necessitates utilizing underlying storage tiers, leading to increased I/O operations. MBGD is an algorithm where random subsets or batches are chosen and fed to the model. The batch size varies depending on the specific model being used. However, if the batch size is small, the model becomes susceptible to noise updates (updates that introduce randomness or variability in the parameter updates), while large batch sizes, in contrast, may lead to over-fitting. The SGD algorithm involves updating the model after feeding a single data point. SGD is a specific case of MBGD where the batch size is 1. The random

nature of the data selection when using SGD makes it difficult to cache the same mini-batches or data points from previous epochs for future ones. Each epoch involves a new random sampling of the data, introducing variability in the training process. This commonly leads to an I/O pattern of large amounts of small I/O performed on many random files. When training models on an HPC platform, datasets are located on a shared Parallel File System (PFS), designed for reading large batched files rather than a large number of small files. This poses a significant challenge for PFSs [46].

Finally, after a certain number of epochs, it is common practice to assess the model's performance while tuning hyperparameters using a separate validation dataset [138]. The same validation dataset is often used after each epoch to evaluate the model. As the training and validation phases are time-intensive stages within the ML lifecycle, distribution strategies have been developed to parallelize these tasks. The distribution strategies optimize the utilization of computational resources, thereby expediting the training process and enhancing efficiency in model development. Choosing the appropriate training and validation distribution strategy depends on factors such as model size, data set size, and available computational resources. This decision is pivotal in guaranteeing optimal training performance. It should be noted that after the model has converged, there is sometimes an additional phase called the **testing** phase which evaluates the performance of the model on unseen data using the testing dataset.

### 3.3 Distributing the Training and Validation Phases

Multiple strategies exist to distribute the training and validation phases to take advantage of the parallelism granted by HPC systems. Two common techniques used are **model parallelism** and **data parallelism**, shown in Fig. 5 and Fig. 6, respectively. When using model parallelism, a single model is distributed across  $N$  workers, and each worker trains their partition of the model [28]. It is often used when the model is too large to fit into the memory of a single node. The strategy has two pair synchronization points: one before the validation and another after adjusting each process's model partition. Synchronization occurs among pairs of workers and is completed across multiple steps. This allows pairs of workers who have been updated to continue training without having to synchronize with all other workers. When using data parallelism,  $N$  workers are employed, and the model is replicated  $N$  times across each worker which each train against a different dataset partition. After each iteration, all workers must be synchronized, and the model weights are adjusted according to the results of the backpropagation. This is the parallelization strategy used in Section 5.2 to analyze the I/O patterns observed during training. In most cases, model parallelism's synchronization and communication overhead far exceeds data parallelism's [1]. However, both model parallelism and data parallelism involve inter-device (i.e., TPU, GPU, or CPU) communication to exchange information about model parameters and gradients. This communication overhead can impact training performance. Strategies such as gradient accumulation, where model updates are aggregated over several batches before communication or gradient compression [24] can help reduce the overhead of inter-device communication.

More recent training distribution strategies include hybrid parallelism [4], and pipeline parallelism [38]. Hybrid parallelism combines data and model parallelism, where the data and the model are partitioned across multiple devices. In contrast, pipeline parallelism places different layers of the model on separate devices, allowing for efficient overlapping of computation and communication. Federated learning, introduced by Google in 2016, is a training distribution strategy where the model is partitioned into multiple devices which train their model partition against private data [114, 117]. After some specified amount of training, devices send their model gradients to a central server, which aggregates the model gradients and redistributes the results back to the devices. This is beneficial for use cases in which the data being used to train the model is sensitive, such as in healthcare or finance. The training data is no longer aggregated into one central location resulting in an increase in privacy and security. The sending and receiving of the model gradients

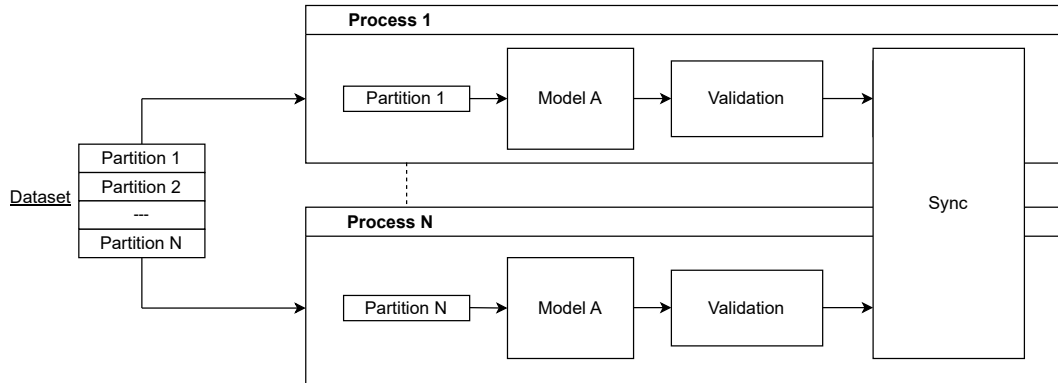


Fig. 5. Data parallelism where the dataset is broken into  $N$  partitions. Each partition is then fed to the model in  $N$  processes. After validation, all processes must synchronize with each other.

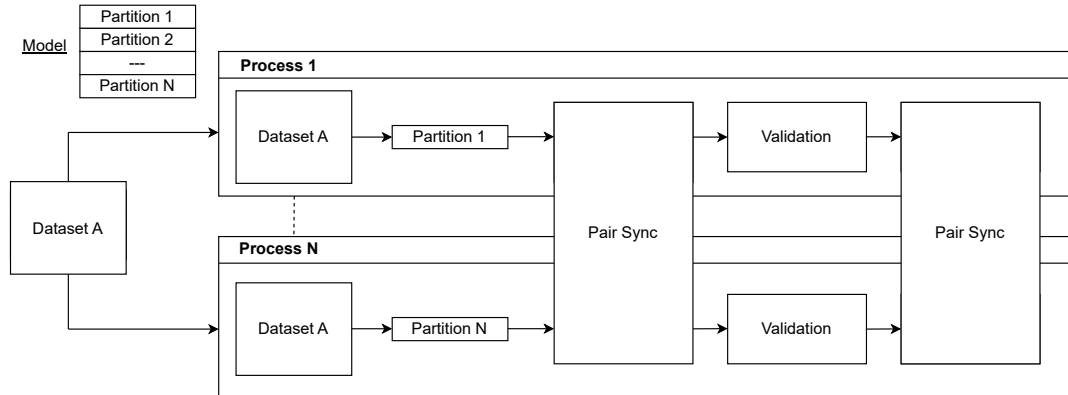


Fig. 6. Model parallelism where the model is broken into  $N$  partitions. Pairs of processes synchronize after both training and validation. Each model partition is trained against the entire dataset at each epoch.

with a central server can result in significant network communication overhead. Various optimizations can reduce communication overhead, such as intelligently selecting the nodes or devices that participate in the training, semi-asynchronous training methods, and clustered federated learning: a strategy in which groups of nodes are prioritized over others during training based on various factors such as latency or physical distance [169].

### 3.4 Model Checkpoints

While large amounts of small reads dominate the I/O operations performed when training ML models, a consistent pattern of writes also occurs [53]. This is due to the need to checkpoint the model as it is being trained. This can significantly impact the training performance, especially if the checkpoints are done frequently. Checkpoints are important as they allow long-running training jobs to pick up where they left off in the event of node failures. In addition, it allows for the analysis of intermediate model features in relation to the training data. Many training jobs begin from an existing model checkpoint from which fine-tuning for their specific task can be accomplished. Starting from a pretrained model can reduce the time and data needed for training. When checkpointing a model, the write size depends largely on the number of parameters the ML model requires [52]. Because of the increasing complexity of ML models [166], checkpoint sizes can be expected to grow proportionally. Most ML frameworks use 32 bit precision to store model parameters [157], the default parameter size on both PyTorch and TensorFlow. Large

language models (LLM)s such as GPT-3 and BERT large have 175 billion [20] and 345 million [37] parameters respectively. The time needed to checkpoint the model depends not only on the number of parameters but also on the file format used [139], making it an important consideration for training performance. In addition, when using frameworks such as TensorFlow, checkpointing is done independently by each worker, or a single worker is assigned the job. This can lead to the stragglers problem due to a single process being responsible for the extra work. When using a single worker to checkpoint the model, the time required does not shrink as more nodes are added to the training. PyTorch only recently (October 4, 2023 in version 2.1) added support for distributed checkpointing, allowing for the saving and loading of models to be accomplished in parallel by multiple ranks. Efforts to optimize model checkpointing include model compression [26] and asynchronous checkpointing. A description of the popular ML frameworks and their corresponding checkpoint APIs can be found in Section 6.1.

### 3.5 BERT Real-world End-to-end ML Workflow

In this section, we present the well-known real-world end-to-end ML workflow of BERT. We discuss BERT in place of GPT because GPT is a closed source, meaning the details about the exact I/O found throughout its lifecycle are not fully disclosed. The authors of BERT mention that it has the most comparable existing pre-training method to OpenAI GPT. The information presented in this section is based on both the original BERT paper and the updated open-source GitHub repository, which specifies important data management attributes such as file formats and pre-training times. A more detailed analysis of the BERT training I/O patterns can be found in Section 5.2 where we simulate the BERT workload using the DLIO [50] benchmark.

*Data Generation* The data used to pre-train the model is the BooksCorpus dataset [172], which consists of 11,000 books with a total of 800 million words, and the English Wikipedia consisting of 2,500 million words.

*Offline Data Preparation* BERT is an unsupervised LLM, meaning it is trained on a large quantity of publicly available plain text data. The Wikipedia data is stripped of all texts other than text passages, ignoring lists, tables, and headers. The text from both datasets is processed into a set of input token sequences, which can then be fed to the model during pre-training. The input datasets are transformed into a single large plain text file, referred to as the corpus, with each sentence being separated by a new line. It is noted that the newline-delimited text sequences contained within the corpus should be actual sentences to enable “next sentence prediction.” Token segmentation is then used to divide the text into individual tokens using a subword tokenizer named WordPiece [152]. A distinguishing factor of BERT is that it uses bidirectional contextual representations, i.e., a given word’s representation is based on the words found to its left and right. This is accomplished by masking out a random 15% of the tokens within each token sequence. The masked-out tokens are then used as the prediction target for the sequence. The token sequence representation used is of the form <Question, Answer>, allowing for both a single sentence and a pair of sentences. This is done by placing special tokens such as [SEP] representing a sentence delimiter and [CLS] representing a token sequence starting point. The [SEP] token delimiter is used to separate different segments of text (e.g., question and answer pairs). The output of the data generation is a set of TFRecord files containing the token sequences.

*Pre-training* The Tensorflow framework is used to train the model during pre-training. This step in the workflow takes around 4-6 days when using 4 to 16 Cloud TPUs. However, it is a one-step procedure from which specific versions of the models can be based on during fine-tuning. Notable hyperparameters used to pre-train BERT in the paper are a batch size of 256 sequences where each sequence is 512 tokens, resulting in 128,000 tokens per batch. During pre-training, random sentences are chosen in pairs A and B in which 50% of the time sentence B is the sentence following A. This is done to enable “next sentence prediction.”

*Fine-tuning* The fine-tuning results presented are relatively inexpensive and can be replicated in at most one hour on a single TPU or a few hours in a single GPU. All results presented in the paper were fine-tuned on a single TPU due to memory constraints. To begin, the model is loaded from a previously generated checkpoint generated from the pre-training phase. The checkpoint consists of three files containing the model weights. The hyperparameters used for fine-tuning are similar to pre-training other than batch size, learning rate, and number of epochs. The paper recommends training for 2, 3, or 4 epochs using a batch size of 16 or 32 and a learning rate of  $5d-5$ ,  $3e-5$ , or  $2e-5$ . Training with long sequence lengths in constrained memory environments can result in out-of-memory errors because the attention method's memory requirement increases by the square of the sequence length.

*Evaluation* After a series of pre-training and training steps, an optional validation or testing phase can be enabled or disabled through the `-do_eval` parameter. This phase is used to track model performance on validation or testing datasets, fine-tune hyperparameters, or select the best model to checkpoint.

*Inference* When using BERT for inference, data preprocessing steps similar to those described above must be done. During either offline or online inference, text input must be tokenized and processed into token sequences. The distinguishing factor of online inference is that the token sequences are fed to the model, and their associated predictions must be processed in real-time. Due to the wide range of possible deployment scenarios, categorizing the data management and I/O that occur in this phase is challenging. However, notable factors distinguishing this phase are the need to store or stream the prediction generated from the model and that there is no longer a need to checkpoint the model.

#### Summary #2

ML applications use a variety of file formats and dataset modalities. The dataset modality impacts what file format is used, affecting how samples are stored and accessed and what features are available. These choices often impact the performance and ease of use in accessing data for ML analysis. Furthermore, during the training phase, SGD is the most popular algorithm in which random batches of samples are fed to the model. The random nature of SGD and the parallelism resulting from training distribution methods (i.e., model and data parallelism) cause I/O operations that are challenging for PFSs. Finally, with the increasing complexity of models, it is crucial to ensure efficient checkpointing methods to achieve fast training speeds.

## 4 I/O Benchmarks and Profiling

Many ML benchmarks have been proposed, such as Fathom [67] and BenchNN [48]. However, these tools focus on identifying the computational requirements of ML workloads and not the I/O requirements [50]. In addition, they target cloud platforms. This leaves a need for a proper benchmark focusing on I/O that simulates ML workloads.

One attempt to address this gap is DLIO [50], a benchmark designed to simulate I/O access patterns commonly found in Deep Learning (DL) workloads. DLIO, which has been released as MLPerf Storage Benchmarks [137], allows for extensive configuration, including the selection of interfaces (HDF5, TFRecord, CSV, NPZ), file access patterns (one file per process versus shared file per process), data access patterns, I/O types, and transfer buffer sizes. Scientists and researchers can take advantage of DLIO by configuring the benchmark to mimic the expected I/O patterns of their DL workloads and using the results to configure their applications effectively. DLIO also supports DALI [129], which is capable of sending batches to the GPU for preprocessing and model training, primarily for image or video datasets. Current DLIO limitations include the inability to store multiple samples per file in certain file formats, the inability to store data/label pairs in separate files, and the assumption that the samples will always be 2D images. The inability to



store multiple samples per file poses a challenge when comparing workloads across ML frameworks. For example, the BERT workload uses the TFRecord format, but currently, there is no PyTorch data loader capable of loading TFRecords. Transitioning to an alternative file format requires configuration modifications, resulting in a workload that diverges from the original and may not faithfully represent it. Thus, while DLIO effectively captures a broad spectrum of DL applications, comparing different frameworks on the same workload while maintaining accuracy proves challenging.

Cheng et al. [27] benchmark two object storage systems (MinIO and Ceph) and three key-value databases (MongoDB, Redis, and Cassandra). Their study attempts to find optimal configurations for parameters such as storage location, storage disaggregation granularity, access pattern, and data format. Notable takeaways are that the object storage systems tested were more sensitive to storage location than key-value databases. In addition, the object storage systems tested were impacted mainly by the storage disaggregation granularity as only one object could be queried at a time.

Many dataset modalities, such as audio, video, and images, have a surplus of public datasets to choose from. In contrast, high-quality, large-scale graph datasets are hard to find; in addition, most commonly used graph datasets are too small to emulate the datasets found in the real world. Hu et al. [93] have created Open Graph Benchmark (OGB), which aims to enhance graph ML research by providing robust and scalable graph datasets. It is capable of producing graph datasets of three different size categories: small, medium, and large, which have 100,000 nodes, 1,000,000 nodes, and 1,000,000,000 nodes, respectively. OGB enables researchers to analyze large-scale graph datasets, facilitating rapid growth in distributed graph machine learning (GML).

iBench [19] is a possible solution for distributed inference simulation and benchmarking and aims to provide several key metrics, including but limited to throughput, latency, ingest rate/bandwidth, pre-processing time, and GPU efficiency. Compared to MLPerf, they note two primary advantages: (1) the ability to measure distributed inference performance and (2) a more realistic performance measure. The simulation places a stream load balancer in front of the HPC inference servers. Client requests are then simulated and scheduled to run via the load balancer. They aim to accurately assess performance by collecting significant factors such as inference time, network response time, payload preprocessing and packing time, and investment time. Due to the limited number of available streaming HPC inference benchmarks, iBench is a promising tool for evaluating network I/O. Current limitations of iBench include the inability to represent many ML models and deployable HPC edge devices. Brewer et al. [18] evaluated iBench, noting several key takeaways, including the impact on performance when using different client batch sizes of 16, 32, or 64, and processing requests asynchronously, the importance of topology-aware GPU-TPU and CPU-GPU communication, and the significance of CPU-GPU in inference compared to training. The evaluation concludes that linear scaling can be achieved if multiple client requests are processed asynchronously and the amount of data sent does not saturate the network.

While benchmarking tools allow ML developers to emulate various ML workload I/O patterns, profiling tools that enable developers to analyze the HPC I/O stack in its entirety are crucial for pinpointing potential bottlenecks. TensorFlow offers the TensorFlowProfiler, released in version 2.2.0, capable of profiling host-side and GPU tracing. However, it only provides information at the TensorFlow level, not capturing lower-level I/O operations such as POSIX or STDIO. Chien et al. [47] created tf-Darshan, a fine-grained I/O profiling tool. It relies on Darshan [69], a highly popular HPC I/O profiler and tracer. Tf-darshan augments Darshan, enabling the runtime extraction and analysis of Darshan's internal data structures. Runtime extraction is accomplished through an augmentation of the Darshan library. It adds extraction functions that return its internal buffers that store information, such as record counters for I/O operations, to individual files. This gives ML applications the ability to perform runtime analysis, enabling I/O optimization techniques such as specialized online I/O schedulers to perform runtime optimizations. Tf-darshan






Fig. 7. Process level accesses for common organizations of samples within files.

also offers a visualization tool that allows users to understand the I/O performance of their ML workloads intuitively. Current limitations include exclusive TensorFlow integration and a moderate overhead of 10% - 20%.

#### Summary #3

Benchmarks play a crucial role in accelerating research within ML I/O. They enable the identification of I/O bottlenecks in ML workloads, thereby enhancing efficiency and productivity. Due to the increase of heterogeneity found in HPC architectures, profilers capable of capturing data movement between a wide range of devices and data sources (e.g., memory, disk, archive, cloud, or database) would be beneficial. Furthermore, existing benchmarks fail to represent the I/O patterns found during all stages of the ML lifecycle, highlighting the need for benchmarks capable of emulating I/O at various stages and data sources.

## 5 Analysis of I/O Access Patterns

In this section, we aim to present the I/O access patterns commonly found during distributed offline data preparation, training, and inference. We searched for benchmarks and found DLIO [50] to be the most promising to simulate the training I/O access patterns of DL applications in HPC systems.

### 5.1 Offline Data Preparation Access Patterns

During the offline data preparation phase, I/O access patterns differ depending on the specific operations needed to prepare the dataset for training or inference. Common operations such as data transformations, normalization, or feature scaling can often be distributed across compute nodes. Due to the high complexity that arises in the programming of distributed systems applications, big data frameworks such as Apache Spark [11], Apache Hadoop [12], or Dask [33] are often utilized to easily distribute data preparation tasks across nodes. Several popular file formats, including .json, .tfrecord, and .npz, do not natively support random I/O access [31]. Instead, data must be read sequentially, enforcing sequential I/O while eliminating parallelism. These file formats are often ideal for streaming scenarios while hindering dataset preparation operations such as selections and filtering. To enable parallel processing for file formats that do not offer random I/O access, the dataset is often sharded, that is, divided into smaller files. Sharding enables distributed processing among compute nodes on each of the "shards" of the dataset. This often results in each process sequentially reading a shard of the dataset, increasing I/O throughput. Although sharding is optional and may require additional processing time, it is often worth the trade-off in reduced I/O times [3]. In general, sharding and the use of big data frameworks provide essential tools to enable ideal I/O access patterns and effective use of the HPC system.

## 5.2 Training I/O Access Patterns

The I/O access patterns found during the training phase are influenced by the gradient descent algorithm used and the training distribution strategy. Here, we describe the I/O access patterns found when using SGD, because SGD is the most popular gradient descent algorithm when training with large datasets. In addition, the presented workloads use data parallelism, where the model is replicated in each worker and the data is partitioned and distributed among them [43]. The organization of the samples within the files, as seen in Fig. 7, has an impact on the I/O operations performed by each worker. Two common organizational strategies are *one sample per file*, often used with file formats such as JPEG or WAV, and *multiple samples per file*, commonly used with file formats such as TFRecord or HDF5. The choice of sample organization plays a large role in the observed I/O operations [107]. A significant amount of additional metadata management overhead can occur when using many files [3]. The Luster PFS used in the workloads presented below has multiple metadata servers (MDS)s which are responsible for providing metadata services such as filename lookup, directory information, file layouts, and access permission. When using *one sample per file* with a dataset with many samples, MDSs can be overwhelmed by concurrent file requests, causing slower training speeds. We show both sample organizations with the Unet3D workload (*one sample per file*) and the BERT workload (*multiple samples per file*).

We obtained the results below by running DLIO and captured the I/O traces using Darshan [151]. All benchmarks were conducted with Perlmutter’s compute nodes. Perlmutter is a supercomputer located at the National Energy Research Scientific Computing Center of the United States Department of Energy. Each of the compute nodes has 64 physical cores and 128 virtual cores, with a total memory capacity of 255 GiB. The PFS is an all-flash Luster file system, containing a total bandwidth of more than 5 TB/sec and supporting 4 million IOPS (4 KiB random). Perlmutter has 16 MDSs and 298 object storage targets (OST)s, and all workloads presented use a 1 MiB stripe size. OSTs are storage devices in which the user’s file data is stored. For visualization, we used DXTE Explorer [45], which is an interactive web-based tool capable of producing graphs from Darshan DXTE logs.




Fig. 8. DLIO TensorFlow reads during the Unet3D workload from the view of an individual file.

In Fig. 8, we used DLIO to simulate the training of Unet3D, which is used for 3D medical image segmentation tasks. The workload consisted of 4 compute nodes, each running 4 Python processes. Python processes were given access to 4 physical CPU cores, and the TensorFlow data loader was configured with 4 I/O threads and a batch size of 4 samples. The dataset for Unet3D consists of one sample per NPZ file where each sample is approximately 146 MiB. There were a total of 168 samples. The model was trained for 10 epochs. Due to the one-to-one correspondence between the number of samples per file, the entire file is requested by the same process at each epoch rather than subsets of the file being requested. In addition, the file was accessed by the same rank (14 in this case) at each epoch.

In Fig. 9, we used DLIO to simulate the training of BERT, which is a large language model (LLM). The computational setup consisted of 8 compute nodes, each running 4 Python processes. Python processes were given access to 2 physical




Fig. 9. DLIO TensorFlow reads during the Bert workload from the view of an individual file.

CPU cores, and the TensorFlow data loader was configured with 1 I/O thread. The dataset size was reduced to increase training speeds; however, we do not expect this to significantly change the overall data access patterns. We reduced the number of files from 500 to 10, and the number of samples per TFRecord file from 313,532 to 31,353. Each sample was 2,500 bytes. Therefore, there were  $10 \times 31,353 = 313,530$  samples in total. The workload was run twice, once with the original batch size of 48 and a second time with a batch size of 96. The model was trained for three epochs.

Fig. 10 is a zoomed-in view of the second epoch of Fig. 9. There are 21 distinct columns of I/O reads that occur at each epoch from every process. This can be expected because the result of  $\frac{\text{total\_samples}}{\text{number\_of\_processes} \cdot \text{samples\_per\_batch} \cdot \text{total\_files}}$  would give us the number of batches that need to be read by every process each epoch with multiple samples per file. Taking into account the 313,530 samples, the 48 samples per batch and the 32 MPI processes, there are  $\frac{313,530}{32 \cdot 48 \cdot 10} = [20.412...] = 21$  batches of samples to read per epoch (rounding up to the nearest integer as the remaining samples need to be read). To ensure this result, we ran the BERT workload again but doubled the batch size to 96 samples. Fig. 11




Fig. 10. Reads during the BERT workload zoomed in on the second epoch of Fig. 9.




Fig. 11. Reads during the BERT workload zoomed in on the second epoch where batch size was doubled to be 96 samples.

is a zoomed-in perspective on the second epoch and can be seen to have 11 distinct columns of I/O reads which agrees with  $\frac{313,530}{32 \cdot 96 \cdot 10} = \lceil 10.206\dots \rceil = 11$  batches of samples to read per epoch. The non-overlapping nature of the columns is expected, because all processes synchronize after every batch is read.

Since no caching is used for the workloads presented, every read request must reach the underlying PFS for every sample needed. The absence of caching poses a potential I/O bottleneck, particularly when dealing with large sample sizes or a large number of small samples. The dividing of the datasets among the ranks and the subsequent I/O reads for batches of samples aligns with the observations of other studies documenting ML training I/O access patterns [27, 50, 126]. The interaction between data organization strategies and the underlying PFS significantly influences the performance, with potential bottlenecks emerging from high volumes of small reads or extensive metadata operations. Understanding and optimizing these patterns is essential for enhancing the efficiency of training ML models on HPC systems.

**5.2.1 Checkpoints** As mentioned in Section 3.4, it is common for a consistent pattern of writes to occur during the training loop to checkpoint the model. Both the Unet3D and BERT workloads use the TensorFlow checkpoint function described in Section 6.1.2, where each checkpoint consists of three files: an index file, a metadata file, and a data file. In both workloads, within the checkpoint directory, a file named “checkpoint” has the most recent checkpoint version. For the Unet3D workload, checkpointing was set to occur after epoch 5, then every 2 epochs thereafter. This resulted in a total of 3 checkpoints after epochs 5, 7, and 9. For BERT, checkpointing was set to occur after every 250 steps. Here, we describe the number of checkpoints when using a batch size of 48. Since the total number of samples was 313,530, the model was trained for 3 epochs, and checkpoints occur every 250 steps, the total number of checkpoints was  $\left\lfloor \frac{3 \cdot 313,530}{48 \cdot 250} \right\rfloor = 78$ . As observed, frequent checkpointing can result in a moderate number of files and, consequently, increased write operations. While checkpointing is essential for fault tolerance and recovery, it also introduces additional I/O overhead that may impact training efficiency. ML developers must strike a balance between ensuring node failure resilience through checkpointing and maintaining reasonable I/O performance to avoid training bottlenecks.

### 5.3 Inference I/O Access Patterns

In this section, we aim to categorize the most common I/O access patterns found during the inference phase within HPC systems. MLPerf [62] developed by MLCommons is a well-known group aiming to provide ML training and inference evaluations. They categorize ML inference into four categories: single-stream, multi-stream, server, and offline [137]. However, these inference categories were not specific to HPC environments, and two of the categories (single-stream and server) were not found in the cited literature or in our own experience. We have therefore excluded these categories from this discussion and have found the two categories described below to be representative of most inference configurations found in HPC systems.

- (1) Single-stream: multiple client devices stream single sample queries to the HPC system in parallel.
- (2) Offline: data is immediately available, and latency and bandwidth are unconstrained, e.g., data is stored in a PFS.

A key difference between streaming and offline inference is the system driving the inference process. In the offline scenario, the HPC system requests data from the storage system, whereas, in streaming scenarios, the HPC system is driven by end users or clients. Latency and bandwidth become more constrained in streaming scenarios, which requires effective load balancing to manage unpredictable query request times to effectively utilize compute nodes [18, 19]. A significant amount of research has been done in this area when ML models are deployed on edge computing devices where constraints such as limited memory, high latency, and limited bandwidth are common [55, 109, 111]. Data reduction techniques such as compression [74] and intelligent grouping of edge computing resources [111] help mitigate

network I/O constraints. Although some of the optimizations mentioned above can be used in offline inference, lower latency and higher bandwidth are often found within HPC systems due to high-speed interconnects.

When inference is compared with training, several distinctions emerge. During inference, data is often not sampled randomly, which can result in an I/O access pattern of large contiguous reads in scenarios such as offline inference. High throughput can be expected during offline inference due to PFSs excelling in large contiguous readings. In a streaming scenario, contiguous reads may not be possible if the data is not readily available. Additionally, checkpointing is less common in inference because the model is no longer updated. It should be noted that there are some scenarios, such as online learning and adaptive inference, that may still involve model updates and checkpointing. In addition to checkpointing, storing the prediction results is another potential source of I/O writes. In a streaming scenario, the prediction results may need to be sent back to the querying client, causing more contention on the network bandwidth.

#### Summary #4

Due to the prevalence of SGD, random batches of samples are read into memory at each iteration during model training. Small random I/O reads can be a bottleneck for PFSs, which motivates the need for I/O optimization techniques such as prefetching and caching. Inference involves a diverse range of scenarios. When streaming over the wide-area network, latency and bandwidth constraints are of greater concern. In an offline scenario, large contiguous reads are possible, which perform favorably on traditional HPC storage systems.

## 6 I/O Optimization Techniques

This section discusses commonly used ML frameworks and the I/O optimization options developers can utilize to distribute their dataset preparation and training phases. Based on the current research, we also discuss potential I/O optimization techniques ML workloads can use to increase training speeds.

I/O Feature	ML Frameworks		
	PyTorch	TensorFlow	Scikit-Learn with Dask-ML
Dataset Streaming	✓	✓	✓
Parallel Dataset Preparation	✓	✓	✓
Sample Prefetching	✓	✓	✗
Sample Caching	✗	✓	✗
Distributed Sample Caching	✗	✗	✓
Sample Shuffling	✓	✓	✗
Asynchronous Checkpointing	✗	✓	✗
Distributed Checkpointing	✓	✗	✗

Table 4. I/O features from popular ML frameworks.

### 6.1 Common I/O Optimization Techniques Used by Current ML Frameworks

Here, we explore the data loading and preprocessing capabilities offered by different ML frameworks. We describe the various classes and methods in state-of-the-art frameworks, including PyTorch, TensorFlow, and Scikit-learn with Dask-ML. Additionally, we touch upon the usage of tools, including Horovod for distributed deep learning across various platforms and direct NVIDIA GPU loading for direct loading from storage and network devices to the GPU. The I/O features supported by each ML framework can be seen in Table 4. It should be noted that ML frameworks often aim to be generally applicable and extendable, meaning that it is often possible to extend one framework to have a feature of another. For example, sample caching can be easily integrated into PyTorch even though it does not offer this option through its API. Table 4 only shows a checkmark for a feature if it is easily enabled through the API without extending the underlying framework significantly.

**6.1.1 PyTorch** PyTorch [135] offers the `torch.utils.data.DataLoader` class, a Python iterable over a dataset, aimed to enable efficient data loading and preprocessing. The data loader supports map-style and iterable-style datasets. Map-style datasets are commonly used when the entire dataset can be indexed efficiently (i.e., the dataset fits into node local memory). In contrast, iterable-style datasets enable the streaming of large datasets where random reads are expensive or when the batch size depends on the fetched data. The user defines the data loading order when using an iterable-style dataset, allowing for implementation optimizations such as chunk reading and dynamic batch sizes. The data loader uses a single process because it is restricted by Python's Global Interpreter Lock (GIL).

To allow asynchronicity, the data loader accepts a `num_workers` parameter to perform multi-process data loading. At each iteration, worker processes are spawned to fetch the data and then destroyed at the end of the iteration or when garbage collection occurs. The data loader also has a `pin_memory` parameter, allowing GPUs to access data faster because it is in page-locked memory [156].

PyTorch uses the Python *pickle* utility, a popular serialization tool, to save and load the model through the `torch.save()` and `torch.load()` methods. PyTorch also has recently (October 4, 2023 in version 2.1) added distributed checkpointing through the `torch.distributed.checkpoint` API. This allows the loading and saving of models to be done by multiple parallel ranks. It produces at least one file per rank and is capable of checkpointing and loading models in topologies of differing world sizes. To allow the distributed checkpoint to work on different topologies, the model checkpoint must be resharded. There is no built-in functionality for asynchronous checkpointing; if necessary, the user must implement it.

**6.1.2 TensorFlow** TensorFlow [158] offers the `tf.data` API [123] to build data loading and preprocessing pipelines. It aims to enable the construction of complex input pipelines from simple reusable pieces. The `tf.data.Dataset` class is used to represent a sequence of samples. It can be constructed from data stored in memory or from one or more files. If the dataset can fit into node local memory, the `Dataset` object can be converted to `tf.Tensor` objects, which can then be fed to the model. Otherwise, the dataset acts as a stream, and data is loaded as needed. The `tf.data.Dataset.cache()` transformation method can be used to cache the entire dataset in memory or in local storage. However, this method should not be used with large datasets that do not fit in memory or local storage as it will lead to cache thrashing.

Asynchronous data loading is possible by prefetching future samples and can be enabled through the `Dataset` class's `prefetch()` transformation method. The number of samples can be specified as a parameter, or users can set the parameter to `tf.data.AUTOTUNE`, causing the runtime to tune the number of samples to fetch dynamically at runtime. Parallel loading of multiple datasets is possible through the dataset classes `interleave()` transformation method. Multiple parameters can be used to configure the `interleave` method, such as the number of datasets to read specified by the `cycle_length` parameter and the level of parallelism to use through the `num_parallel_calls` parameter, which can be set to `tf.data.AUTOTUNE`, allowing TensorFlow to adjust the value at runtime dynamically. To parallelize data preprocessing, the `Dataset` object has a `map()` transformation method that takes a `lambda`, which can then be run across multiple CPU cores concurrently. The `map()` method takes a parameter named `num_parallel_calls`, which can be used to specify the level of parallelism to use. TensorFlow recommends setting this value to the number of cores available as this allows the preprocessing of samples to be run in parallel on multiple cores. If set to one, the preprocessing of samples occurs synchronously which can negatively impact training performance. In addition, users can set the value to `tf.data.AUTOTUNE` allowing TensorFlow to adjust the value dynamically at runtime.

TensorFlow has four checkpointing formats: the `SavedModel` format, `HDF5` format, `.keras` format, and `checkpoint` format. The `SavedModel` format stores the entire model graph in a binary file that is then stored within a directory

containing a TensorFlow checkpoint. The format is useful for model sharing and deployment with other platforms such as TFLite [160], TensorFlow.js [150], and TensorFlowHub [159]. The HDF5 format uses a single HDF5 file and is the legacy choice by TensorFlow when using the higher level `tf.keras` API. The currently suggested format is the `.keras` format. It aims to have an efficient name-based saving scheme that enables easier debugging than the legacy (i.e., HDF5) formats. The checkpoint format uses three files: an index file, a metadata file, and a data file. The index file contains an index of the model data, enabling fast retrieval of model values. The metadata file contains the TensorFlow computation graph, which includes the model architecture required to reconstruct the model and the data file contains the model's actual data (model weights, optimizer states, etc.). It shards the model across multiple files to enable parallelism. Asynchronous checkpointing can be enabled through `tf.train.CheckpointOptions()`. This method takes an `enable_async` parameter, which, when set to `True`, moves the checkpointing of the model off the main thread. The return value can be passed as a parameter to the `tf.train.Checkpoint.restore()` and `tf.train.Checkpoint.save()` methods, which are used to load and save the model.

When performing distributed data parallel training, if the online dataset preparation takes more time than the model training (i.e., time spent performing forward propagation, backpropagation, and sharing of model weights), significant time can be spent waiting for samples to be preprocessed. TensorFlow offers an experimental feature named `tf.data.experimental.service` [40], which disaggregates the online dataset preparation and training workers. This allows for the ratio of preprocessing workers to training workers to be scaled independently, ensuring that training workers do not have to wait for samples to be preprocessed. The number of online dataset preparation workers and training workers and their CPU/RAM resources can be set either manually or auto-scaled by the service. The service is also capable of sharing preprocessed data between jobs as well as executing coordinated reads to avoid stragglers, which are often caused by the varying input data sizes workers use.

**6.1.3 Scikit-learn** Scikit-learn [143] requires workloads with datasets that fit into node-local memory because it does not offer data loading utilities such as PyTorch's `torch.utils.data.DataLoader` or TensorFlow's `tf.data`. Instead, developers often use a separate library named Dask-ML [33] to enable the loading of large datasets. This allows models to be trained against datasets that do not fit in node-local memory. The `Client` class is the primary way of accessing large datasets in parallel. The `Client` class acts as an interface for workers to submit tasks (i.e., Python Futures). The Dask-ML scheduler will then schedule tasks and run them in parallel if possible.

The `Client` class takes a parameter named `asynchronous`, which specifies whether the client acts synchronously or asynchronously. The `dask.dataframe` class has several methods that can be used to load large datasets such as `read_parquet()`, `read_json()`, and `read_hdf5()`, among others. If the requested result of a task is too large to fit into node-local memory, Dask-ML chunks the result and distributes it among the memory of multiple nodes. This allows tasks submitted to the `Client` to be run in parallel by each worker that contains a partition of a distributed result. For distributed data processing, Dask-ML offers methods such as `map()` and `map_partitions()`, which can be used to apply computations to elements or groups of elements in a collection (i.e., arrays, dataframes, and bags). In addition, reduction methods such as `mean`, `max`, and `min` are available and can be applied across partitions.

Scikit-learn model saving and loading is commonly done through the Python *pickle* library. Scikit-learn also recommends using the Python library *joblib* to save objects with large *numpy* arrays, which is common for fitted Scikit-learn estimators (i.e., any object that learns from data). Similarly to PyTorch, no built-in functionality for asynchronous checkpointing exists, and if necessary, the user must implement it.



**6.1.4 Direct NVIDIA GPU Loading [90]** GPUDirect Storage [127] enables direct data loading from storage to GPU memory, and GPUDirect RDMA [128] enables direct data loading from a network device to GPU memory. This avoids loading the data into the CPU memory first, preventing the creation of a bounce buffer. This saves CPU memory and reduces the overall communication overhead between devices.

**6.1.5 TensorFlow/PyTorch with Horovod [144]** Horovod is a popular distributed deep learning framework for TensorFlow, PyTorch, and Apache MXNet [25]. Uber developed Horovod with the goal of making it easier for training scripts written in common ML frameworks to scale to many GPUs. It uses data parallelism to partition the dataset among the workers to enable efficient training.

## 6.2 Proposed I/O Optimization Techniques

In this section, we present the current research towards optimizing ML I/O on HPC systems. We explore a range of strategies aimed at enhancing the efficiency and scalability of ML workflows. The research presented reflects the ongoing efforts to address the unique storage challenges posed by ML.

**Caching** During SGD, samples are commonly shuffled based on a seed using a Pseudo Random Number Generator (PRNG). Each worker is then assigned a partition of the randomly shuffled samples and fetch batches of samples from their assigned dataset partition. The seed can be used to determine the order and frequency at which a given worker will require samples. Dryden et al. [51] propose NoPFS, a machine learning I/O middleware, which capitalizes on these insights to prefetch and cache samples optimally. At the application level, NoPFS can be integrated into the training pipeline of most ML frameworks with minimal code change. NoPFS reduced I/O and improved overall training time by 5.4× on models using the ImageNet-1k [35], ImageNet-22k, and CosmosFlow [61] datasets.

Importance sampling is a DNN acceleration method in which samples are skipped based on their “importance” value. The importance value is based on the influence a given sample has on the model’s accuracy. Chen et al. [49] introduce iCache, a novel caching approach where samples are placed into one of two in-memory caches: the H-cache or the L-cache. The placement of a sample into either cache is determined by its level of importance, with high-importance samples placed in the H-cache and low-importance samples placed in the L-cache. The technique exploits the fact that when using importance sampling, samples with higher importance values should be fed to the model more frequently than samples with lower importance values. The intelligent caching of samples based on their importance values allows I/O operations to reach the underlying storage system less frequently.

**Scheduling** When using data parallel training, the reading and writing of many small files can cause cross-application I/O interference [101] in the underlying PFS. Burst buffers are a reliable technique to absorb large amounts of reading and writing in a short period of time (bursts of I/O); however, when many concurrent applications are performing these operations, congestion may still occur [81]. In addition, not all HPC systems deploy burst buffers, causing scientists to look for alternative solutions. One possible technique to mitigate the congestion is using specialized I/O scheduling algorithms. Adaptive Periodic I/O scheduler (APIO) [171] is a low-profile online I/O scheduler that dynamically adjusts to the periodic and stochastic nature of ML I/O workloads to schedule I/O operations effectively. By controlling the ordering and timing of concurrent processes issuing I/O requests, APIO can mitigate potential congestion.

Due to many concurrent processes performing I/O operations during ML training, it can be difficult for middleware to characterize the status of the underlying storage nodes. This makes it challenging to schedule I/O requests to run optimally. Wang et al. [72] developed a lightweight ML model to predict the status of the underlying storage nodes and present a case study leveraging their model to effectively balance I/O traffic among the OSSs on a Lustre file system.



They tested several models, the most accurate was a support vector machine, which achieved an accuracy of 75% when predicting whether a given OSS was busy or idle.

*Shuffling* Various shuffling techniques have been proposed to shuffle samples between workers before each epoch. Global shuffling involves shuffling all samples across all workers, which can make caching and other I/O optimizations difficult. In contrast, local shuffling involves workers being assigned a subset of the dataset, from which they then reuse and reshuffle before each epoch. Hash-based shuffling involves applying a hash function to the keys of the samples, from which the result determines the partition to which the record will be shuffled. Apache Spark uses this technique as a primary shuffling algorithm as it can help achieve load-balancing across partitions and minimize data movement. The randomness of the shuffling does affect the convergence rate of the model [118]; however, using alternative techniques may allow workers to cache the samples because they no longer read the entire dataset. Partial data shuffling, proposed by Nguyen et al. [126], achieves similar accuracy to global data shuffling while only requiring each worker to store up to 0.03% of the dataset. The partial data shuffling algorithm combines the technique of global and local data shuffling. Their paper also discusses that local shuffling achieves accuracy similar to the more common global shuffling strategy in their experimentation. Integrating their solution into PyTorch requires minor modifications of the training script and a load handler of the dataset. Jie et al. [96] propose ASRDataset, a system aiming to speed up the training of automatic speech recognition (ASR) models by implementing chunk-level and batch-level shuffling. Their system allows the user to specify how many chunks to batch and the number of samples in each chunk. In addition, ASRDataset uses a high-speed audio data processing library implemented in C++ and CUDA, speeding up audio data augmentation and extraction. Lee and Bahn [108] proposed shuffling bundles of samples among workers. This groups samples together to increase locality. To improve buffer caching performance, in each epoch, bundles of samples are accessed in the reverse order of the previous epoch. This improves locality from the buffer cache, allowing the last bundle in the previous epoch to be the first in the next epoch. Liu et al. [110] also propose block shuffling by compacting small files into blocks. They use the ImageNet dataset to test their data-loading pipeline. They cache partially loaded data to avoid redundant reads to the underlying storage system. In the first epoch, all data is cached in memory. All remaining epochs reuse half the loaded data and load the other half from disk to form batches, reducing the underlying storage accesses by half.

*Asynchronicity* Asynchronous I/O enables the overlapping of compute and storage. Sunwoo Lee et al. [59] propose an asynchronous I/O strategy to speed up data parallel model training by creating a dedicated I/O thread to run alongside each process. Their model assigns groups of the dataset to training processes and globally shuffles these groups at each epoch. In addition, each process allocates a buffer large enough to store two or more data groups. Each process's I/O thread reads groups of samples as needed into the buffer, which can occur parallel to model training. The strategy has two drawbacks. The first is the lower degree of randomness due to groups of samples being shuffled, and the second is the increased memory footprint needed to load the additional data groups asynchronously. Serizawa et al. [145] also propose overlapping the copying and reading of the dataset for training. They observed 1.38 to 6.19 times speed up when using their asynchronous strategy compared to using Lustre directly.

*Compression* Compression is a common technique to reduce the I/O needed to load and store data. Kuchnik et al. [103] propose using progressive compression techniques to reduce the overhead of loading data. They present a file format named progressive compressed records (PCRs), which uses progressive compression to store image datasets. Progressive compression stores data samples (images in this case) in a series of deltas, each with more fidelity. When using formats such as JPEG, data fidelity must be chosen at encoding time, often causing image datasets to be stored multiple times at varying compression levels. Using PCR, various models can be trained using the same dataset and accessed at different

fidelities according to their needs. They found that PCR reduced the sizes of images by 5% compared to storing the images with TFRecords. One downside to PCR is that the decoding time can be up to 2x greater than JPEG.

*Dataset distribution* Many supercomputing facilities are not yet equipped with large node-local memory devices, which disables the ability to cache datasets in higher storage. Choi et al. [57] aim to enhance I/O performance by creating a DDStore, an in-memory data store. DDStore evenly distributes the chunks in the memory of compute nodes, after which sample accesses from a particular node are made through in-memory read transactions. To reduce possible congestion due to multiple nodes requesting a batch from the same node, DDStore replicates the dataset and distributes it among groups of workers. This allows the samples to be globally shuffled while reducing communication congestion. This strategy makes a large trade-off between reduced node-local memory and fewer PFS I/O operations. Dantas et al. [31] created Monarch, which aims to exploit storage tiering to enable fast training speeds on single-node workloads. Their solution aims to be ML framework agnostic acting as a transparent middleware which intelligently caches at different storage layers that may not fit into host memory. Aizman et al. [3] created a storage system named AIStore aiming to provide an infinitely scalable namespace. The namespace can be used on an arbitrary number of disks. Metadata overhead is minimized by enabling dataflow directly between compute nodes and storage targets.

*Data redundancy mitigation* Multiple executing workers can often prepare the same sample redundantly during the training phase. Xie et al. [73] introduce JOADER, a data loader aimed at reducing the redundant I/O and data preprocessing of samples. JOADER is capable of sharing data preparation work on overlapping datasets. It accomplishes this by dividing the dataset into intersection and difference sets, correlating the selections, and sharing operations that operate on the intersecting set. They also introduce a domain-specific cache policy with a novel tree data structure aimed at effectively caching samples. This strategy is more effective when the preparation of a sample takes a significant amount of time (e.g., image decoding). Jin et al. [97] created MMDataLoader, a data preprocessing pipeline framework aiming to increase preprocessing throughput and reduce redundant data preparation. It operates at the server level, allowing it to share results across training tasks. The framework enables batches of data to be preprocessed for multiple models at the same time, eliminating the need for each task to recompute batches individually.

*Storing large models* With the increasing complexity of ML models, several methods have been proposed to enable the training of models that do not fit into accelerator or even clustered accelerator memory. Previous methods aiming to reduce memory usage during training include materialization and offloading. Materialization removes some activations during the forward pass. The removed activations are then recomputed during the backward pass. This strategy trades a larger computational cost for less memory. Offloading moves activations from the limited GPU memory to CPU memory during the forward pass trading data movement for memory [14]. The limited HPC bandwidth, in addition to the increased communication caused by the offloading of the model to storage, is a large challenge for these strategies. Rajbhandari et al. [136], propose ZeRO-Infinity to train large models by leveraging GPU, CPU, and NVMe memory. ZeRO-Infinity is fully integrated with PyTorch. Jang et al. [94] propose Smart-Infinity to train large models with storage offloading. Smart-Infinity is also integrated with PyTorch and takes advantage of storage devices to form an extended memory hierarchy. In contrast to ZeRO-Infinity, they use near-storage accelerators to reduce I/O pressure and increase training speeds. In addition, gradient compression/decompression is used to reduce the storage congestion incurred from model offloading. Sun et al. [71] propose STRONGHOLD for large model training. They offload model training from the GPU to the CPU; however, the minimum amount of data to be kept on the GPU to minimize GPU memory usage is dynamically determined at runtime. In addition, they use prefetching to load layers of the model in parallel with GPU computation.

*GPU storage hierarchy* In GPU architectures, storage is typically organized into tiers, comprising high-bandwidth memory (HBM) and on-chip SRAM, facilitating efficient data access and processing [75]. Doa et al. [32] propose FlashAttention which uses an I/O aware exact attention algorithm to reduce the number of memory read and writes that occur between HBM and SRAM on GPU devices. It efficiently moves data between the GPU storage tiers, resulting in faster training speeds while maintaining the same level of accuracy as traditional attention algorithms. FlashAttention has been adopted into several machine learning frameworks such as PyTorch’s nn.Transformer, Microsoft’s DeepSpeed [68], and NVIDIA’s MEgatron-LM [146]. In addition, it yields the fastest in MLPerf [62] BERT training on cloud instances in MLPerf training 2.0 (June 2022) and MLPerf training 2.1 (November 2022).

*Multi-modal pipeline overhaul* FFCV [54] is a library aiming to speed up the data loading and processing pipeline while easily integrating into existing PyTorch applications. FFCV offers a custom file format that can store arbitrary data modalities and enable non-sequential data reading. It stores all samples in a single file but splits the dataset into pages, which trades space for more efficient data reads. If the dataset can fit into memory, it will be cached, and every read of a sample after the first will be read from memory instead of the underlying storage system. In contrast, if the dataset does not fit into memory, FFCV enables efficient data loading by using the PRNG to prefetch data samples. In addition, it uses quasi-random data shuffling in which a permutation of the data samples is loaded into a memory buffer large enough for *batch\_size* pages of the dataset. The model then selects from batches existing within the buffer. This reduces the number of accesses to the underlying storage system. FFCV aims to eliminate unnecessary memory allocations by having operations in the data loading and processing pipeline declare memory requirements so that memory allocation occurs only once before each epoch. It also uses just-in-time (JIT) compilation, where possible, during the data processing pipeline to convert Python code to machine code. Since the machine code is no longer under the constraints of the Python interpreter, it enables multi-threading, which by extension enables collaborative reading and writing to and from memory instead of expensive primitives, i.e., threads can share batches of data. Threads also enable data preparation (data copying and augmentations) to be run in parallel as they share the same CUDA context.

*Image pipeline overhaul* DIESEL+ [167] is an image data loading and preprocessing pipeline that can be integrated with frameworks such as TensorFlow and PyTorch. The pipeline aims to speed up the training of models that use JPG datasets. Diesel+ has a *Fast Mode* in which JPG files are first processed into an intermediate *fast-binary* format and stored in memory for future usage, enabling reduced loading times. There is also a *Compatible Mode* for images that require no manipulation. The time required to load JPG images is further reduced by applying an online region of interest (ROI) technique to decode the images. JPG images are combined into chunks to mitigate the small random I/O reads often performed during training. These chunks are then shuffled amongst the workers, from which random individual samples can be selected in memory. In addition, a per-task distributed cache is employed across the workers to reduce the number of reads reaching the underlying storage system. Finally, to minimize metadata overhead, tasks take metadata snapshots, eliminating the contention over the metadata server.

*Remote storage* Deep Lake [91] is a data lake for deep learning. It can be used with object storage systems such as AWS S3 [7], Google Cloud Storage (GCS) [89], POSIX file systems, or local in-memory storage. Deep Lake has a Tensor Query Language that extends the SQL Parser from Hyrise [39]. This enables multidimensional array operations and deeper integration with the data lake, including indexing, slicing of arrays, and a large set of convenience functions. There is also a data visualizer made available through a web interface, allowing for the visualization of large-scale data. The Deep Lake streaming data loader delegates data fetching to C++ processes to avoid Python’s GIL. Deep Lake also has a version control system that keeps track of different branches of the dataset in the same storage. This allows queries

to be written which specify the dataset version, allowing for more granular and depthful searches. Deep Lake also is capable of storing non-structured data and metadata in a columnar format resulting in fast data streaming speeds.

*Checkpointing* With the growing size of AI models, efficient checkpointing methods are needed to ensure scalability. Nicolae et al. [64] propose DeepFreeze, an asynchronous checkpointing technique built using Very Low Overhead Checkpoint-Restart (Veloc) [63]. Veloc is a lightweight concurrency-optimized checkpointing library aimed at delivering high-performance checkpointing on top of heterogeneous storage hierarchies. DeepFreeze is a transparent checkpointing solution for synchronous data parallel training with layer-wise model parallelism. To increase checkpointing speeds, an asynchronous multilevel checkpointing approach is used in which a local copy of the learning model is captured and persisted in lightweight (i.e., node-local storage of neighboring nodes) and heavy (i.e., PFS) storage to block training for the least amount of time possible. To reduce the overhead of model serialization, a compact binary format is used that leaves out details such as tensor labels. The model's layers are sliced into shards equal to the number of MPI ranks. The ranks then write checkpoints for their partition of the model to local storage. This allows for models to be checkpointed in parallel, avoiding the stragglers problem. Anthony et al. [9] also propose local process model snapshots, which are then stored in node-local storage. Their method uses the SCR-Exa library [30], which is a scalable checkpointing tool for HPC applications. There are two restart methods (cold and hot) available. A cold restart attempts to restart the application using the checkpointing cache within the same allocation of nodes. The hot restart tries to replace faulty nodes with idle spare nodes. Check-N-Run [42] is a checkpointing system for deep-learning recommendation models developed by Facebook. It uses two techniques to address size and bandwidth: differential checkpointing, where only a fraction of the model changes per iteration, and quantization techniques to reduce checkpoint size. The checkpoint system is implemented under the PyTorch framework. Check-N-Run keeps training accuracy as a priority as they state an accuracy loss as small as 0.01% is unacceptable. Maurya et al. [115] introduce DataStates-LLM, an asynchronous multi-level checkpointing approach to speed up LLM checkpointing. By exploiting the periods of time when tensors and optimizer state shards are immutable, the strategy lazily copies shards of the model from the GPU during periods of immutability into a single host and streams the partial checkpoints asynchronously to persistent storage. The copying from GPU to host and the transfer from host to persistent storage can occur in parallel.

#### Summary #5

Modern ML frameworks offer many optimization techniques to enable fast training speeds. There have been many recent advances aiming to optimize ML I/O in HPC systems. From optimized batch shuffling techniques to specialized I/O schedulers, I/O remains a prominent and continually evolving area of research and development.

## 7 Future Research Directions of I/O for ML

While the ML surge is rushing forward driven by rapidly increasing data collection and application use cases, data management and I/O software have been significantly lagging, especially in the HPC domain. In this section, we identify numerous data management and I/O areas that need to be improved to ensure continued progress and innovation in the training and deployment of ML models.

*I/O characterization and benchmarking* While the DLIO benchmarks suite, which is used as the MLPerf Storage Benchmark Suite, focuses on deep learning applications, many more AI models exist and have been evolving whose I/O patterns have not been represented in existing benchmarks. Some significant examples include reinforcement learning, Q-learning, and Deep Q Networks, widely employed for decision-making in dynamic environments. These ML

models may produce unique I/O patterns, motivating the need for additional benchmarks that can mimic these patterns. Benchmarks that can properly emulate the vast number of AI models play a crucial role in assisting researchers in identifying optimization techniques, ensuring that I/O does not become a bottleneck in ML workloads. Development of I/O benchmarks representing ML used specifically in science applications that run on HPC systems is also needed.

*Efficient data preparation strategies* When preparing data from remote storage locations such as data warehouses, it may be necessary to store intermediate forms of the dataset after preprocessing. Further research could involve an I/O middleware that transparently stores intermediate forms of the dataset in clustered memory or the underlying storage system, depending on the dataset size. By extending the memory to the underlying storage system, additional questions arise, such as how data movement between memory and storage can be optimized for high throughput and low latency and what caching strategies should be used to ensure efficient data access.

*Data quality evaluation and monitoring methods* With the increasing amount of data for machine learning, ensuring data quality is essential for model accuracy [56]. A step towards ensuring high-quality data is to use data warehouses because they often provide data governance tools. However, scientists who use HPC systems for high I/O throughput could take a significant performance hit if limited by the bandwidth of the external network. This motivates the need for tools that can aid in obtaining and monitoring high-quality datasets. These tools could offer data management services, such as integrity constraints, versioning, provenance, and quality assessments, while being scalable and supporting multiple data modalities. A survey on metrics for AI data readiness [92] exposes a lack of metrics focusing on quality, bias, and FAIR (Findability, Availability, Interoperability, and Reusability) principle compliance.

*Support for multi-modal data* Multi-modal machine learning models aim to process and relate information from multiple data modalities [13]. I/O optimization strategies targeting multi-modal datasets are needed to ensure reasonable training and inference speeds. Two broad categories for the relationship between the datasets are parallel and non-parallel data. In parallel data, each sample in one modality corresponds to another in a separate modality. In contrast, in non-parallel data, the samples within each modality are independent. Future research could exploit data locality by ensuring optimal data layout grouping related data together depending on whether it is a parallel or non-parallel dataset.

*Support for streaming data* Due to the rise of cloud-based storage being used by ML workloads, methods of efficiently streaming data from multiple data sources are needed. In addition, multi-modal models may need to stream data from several data sources, each with varying data modalities. Future research could involve an I/O middleware capable of concurrently optimizing I/O from multiple data sources. The middleware would take into account the varying latencies between the data sources and use methods such as sample caching and sample prefetching to ensure optimal performance. This would involve a cache eviction strategy in which eviction is based on the dynamic latency of the data sources. In addition, prefetch scheduling could also occur dynamically, ensuring that samples from high-latency data sources have a higher priority than low-latency data sources. It would also take into account the sample size when performing these optimizations to ensure the optimal balance between latency and data size.

*ML application specific I/O optimizations on (1) traditional HPC systems, and (2) heterogeneous systems with GPUs and domain-specific processing units* GPUs have become increasingly popular due to the rise of ML. This is due to their high level of parallelism when doing mathematical computations and, in particular, matrix operations [153]. This is possible because GPUs, unlike CPUs, contain thousands of cores. This has led to the development of ML domain-specific processing units, such as tensor processing units (TPUs) by Google [99]. These are specialized processors specifically built for ML workloads. With the wide range of hardware available, future research is needed to determine how data movement and storage techniques should be used effectively across the diverse range of hardware available. Questions

for future research include where to place and process samples to prioritize processing units and exploit their contrasting abilities effectively, how to efficiently move data between the diverse range of processing units, and how to provide transparent and reasonable solutions for developers facing these challenges.

*HPC I/O inference analysis* The large majority of current ML I/O research is aimed at analyzing the training of ML models. However, the I/O that occurs during model deployment differs widely from training because there is no longer a need to select random samples of data. This causes many of the previously mentioned I/O optimizations in Section 6 to be non-applicable. A thorough analysis of the I/O operations that occur during model deployment should be conducted in order to identify potential I/O bottlenecks. The study could examine a range of model deployment scenarios, including single-model workloads versus multiple-model workloads, as well as different types of models. In addition, while many models no longer need checkpointing, investigation of I/O write operations that occur due to the need to store model predictions could produce valuable insight.

*Integration of HPC I/O and ML programming tools* As shown in Fig. 1, HPC I/O and ML I/O stacks have been organically evolving for faster implementation of the ML algorithms although there is significant scope for higher performance. As the ML algorithms settle, improving the efficiency of data access in running various applications both on HPC systems as well as cloud computing environments need to take priority. ML application development should also target obtaining portability and efficiency on different platforms. Achieving this will require efforts in benchmarking, evaluation of performance with various programming tools, and developing platform-specific optimization strategies.

*Efficient Model Offloading* Model offloading is used in memory-constrained environments to move subsets of the model to lower storage tiers. This has been a prominent area of research in edge computing, where memory constraints are a common limiting factor. With the growth in size of LLMs, many cannot be held in distributed memory. One example of such a model is the DOE’s recently announced one-trillion-parameter (i.e., 4TiB) AI system named “AuroraGPT” [162]. Enabling such a model in a memory-constrained environment requires model offloading. Compounding the issue is that some inference configurations (e.g., ensemble models or pipeline models) require multiple models to be loaded and unloaded from device memory to respond to a single query. Moving the model effectively through the storage tiers to minimize inference response times is a challenging I/O and data management problem where further R&D is needed.

#### Summary #6

The rapid development of AI and its usage in decision support has pushed for quick, easy, and often dirty ways to implement them. As the applications and technologies mature, efficiency will be a priority. We listed various gaps in current research and practice. Addressing the widening gap between computing and storage, additional R&D in I/O and data management are essential to sustain progress and innovation and achieve efficiency in ML model training and development.

## 8 Conclusion

Data plays a pivotal role in machine learning, impacting the learning process and significantly affecting model performance and accuracy. With a wide array of file formats available, each offering distinct I/O optimizations, machine learning applications must carefully select the format that best suits their needs. Additionally, the modality of the dataset influences the choice of file formats, affecting how datasets are accessed and stored throughout the ML lifecycle.

Model training is a critical step in this lifecycle, and due to SGD being the most popular optimization algorithm for training models on large datasets, high amounts of small random I/O reads present a significant challenge for PFSs.



Furthermore, training distribution strategies such as model parallelism and data parallelism, in addition to checkpointing, exacerbate the number of I/O operations PFSs need to handle. In order to find I/O bottlenecks in ML processing, training, and inference, it is essential to have access to appropriate benchmarks and profiling tools. Existing benchmarks and profilers allow I/O performance enthusiasts and HPC practitioners to simulate ML workloads accurately, facilitating the identification of areas of improvement within ML pipelines.

Modern ML frameworks commonly offer several options to optimize data loading and processing, such as parallel data preparation, dataset streaming, and sample prefetching, to optimize the processing and loading of batches of samples. However, caching solutions that work with datasets that do not fit into node local memory are not widely available in ML frameworks such as PyTorch and TensorFlow, often leading to sample accesses always reaching the underlying storage system, which may lead to an I/O bottleneck. Recent advances toward minimizing I/O time include specialized shuffling techniques, I/O schedulers, and domain-specific file formats.

While the rate of ML I/O research has continued to increase alongside the popularity surge of AI, there are still gaps in existing research. All-in-one solutions to existing problems are a challenge for researchers due to the wide variety of data sources where samples are stored, the diverse range of hardware ML training and inference runs on, and the many different types of data used by ML models. An additional challenge is providing solutions to these problems that are either transparent or easily incorporated by ML developers. Ultimately, ongoing R&D in ML I/O is a crucial process to ensure efficient model training and deployment, ultimately aiding in the expanding and ever-growing field of AI.

### Acknowledgments

This research was supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research (ASCR) under contract number DE-AC02-05CH11231 with LBNL. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

### References

- [1] Khalid Abdulaziz Alnowibet, Imran Khan, Karam M. Sallam, and Ali Wagdy Mohamed. 2022. An efficient algorithm for data parallelism based on stochastic optimization. *Alexandria Engineering Journal* 61, 12 (2022), 12005–12017. <https://doi.org/10.1016/j.aej.2022.05.052>
- [2] Monica Agrawal, Marinka Zitnik, and Jure Leskovec. 2018. Large-scale analysis of disease pathways in the human interactome. *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing* 23 (2018), 111–122. [https://doi.org/10.1142/9789813235533\\_0011](https://doi.org/10.1142/9789813235533_0011)
- [3] Alex Aizman, Gavin Maltby, and Thomas Breuel. 2019. High Performance I/O For Large Scale Deep Learning. In *2019 IEEE International Conference on Big Data (Big Data)*. 5965–5967. <https://doi.org/10.1109/BigData47090.2019.9005703>
- [4] Takuya Akiba, Shuji Suzuki, and Keisuke Fukuda. 2017. Extremely Large Minibatch SGD: Training ResNet-50 on ImageNet in 15 Minutes. arXiv:1711.04325 [cs.DC] <https://arxiv.org/abs/1711.04325>
- [5] Amazon Web Services. 2024. 1000 Genomes Project. <https://registry.opendata.aws/1000-genomes/> Accessed: 2024-09-12.
- [6] Amazon Web Services. 2024. AWS Public Datasets. <https://registry.opendata.aws/> Accessed: 2024-09-12.
- [7] Amazon Web Services. n.d. *Amazon Simple Storage Service (S3)*. <https://aws.amazon.com/s3/>
- [8] Davide Anguita, Luca Ghelardoni, Alessandro Ghio, Luca Oneto, Sandro Ridella, et al. 2012. The 'K' in K-fold Cross Validation.. In *ESANN*. University of Genova, Via Opera Pia 11A, I-16145 Genova, Italy, 441–446.
- [9] Quentin Anthony and Donglai Dai. 2021. Evaluating Multi-Level Checkpointing for Distributed Deep Neural Network Training. In *2021 SC Workshops Supplementary Proceedings (SCWS)*. 60–67. <https://doi.org/10.1109/SCWS55283.2021.00018>
- [10] Apache Mesos. 2024. *RecordIO: Apache Mesos Documentation*. <https://mesos.apache.org/documentation/latest/recordio/> Accessed: 2024-02-19.
- [11] Apache Software Foundation. 2024. Apache Spark. <https://spark.apache.org/>.
- [12] Apache Software Foundation. 2024. *Hadoop*. <https://hadoop.apache.org>
- [13] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2019. Multimodal Machine Learning: A Survey and Taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 2 (2019), 423–443. <https://doi.org/10.1109/TPAMI.2018.2798607>
- [14] Olivier Beaumont, Lionel Eyraud-Dubois, Alena Shilova, and Xunyi Zhao. 2022. Weight Offloading Strategies for Training Large DNN Models. (Feb. 2022). <https://inria.hal.science/hal-03580767> working paper or preprint.

- [15] Léon Bottou. 2012. Stochastic Gradient Tricks. In *Neural Networks, Tricks of the Trade, Reloaded*, Grégoire Montavon, Genevieve B. Orr, and Klaus-Robert Müller (Eds.). Springer, 430–445. <http://leon.bottou.org/papers/bottou-tricks-2012>
- [16] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. 2018. Optimization methods for large-scale machine learning. , 223–311 pages. Issue 2. <https://doi.org/10.1137/16M1080173>
- [17] Don Box, David Ehnebuske, Gopal Kakivaya, Andrew Layman, Noah Mendelsohn, Henrik Frystyk Nielsen, Satish Thatte, and Dave Winer. 2000. *Simple Object Access Protocol (SOAP) 1.1*. W3C Note. World Wide Web Consortium. See <http://www.w3.org/TR/SOAP/>.
- [18] Wesley Brewer, Greg Behm, Alan Scheinine, Ben Parsons, Wesley Emeneker, and Robert P. Trevino. 2020. iBench: a Distributed Inference Simulation and Benchmark Suite. In *2020 IEEE High Performance Extreme Computing Conference (HPEC)*. 1–6. <https://doi.org/10.1109/HPEC43674.2020.9286169>
- [19] Wesley Brewer, Greg Behm, Alan Scheinine, Ben Parsons, Wesley Emeneker, and Robert P. Trevino. 2020. Inference Benchmarking on HPC Systems. In *2020 IEEE High Performance Extreme Computing Conference (HPEC)*. 1–9. <https://doi.org/10.1109/HPEC43674.2020.9286138>
- [20] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. , 25 pages.
- [21] Jason Brownlee. 2018. What is the Difference Between a Batch and an Epoch in a Neural Network. *Machine Learning Mastery* 20, 1 (2018), 1–15. <https://lipingyang.org>
- [22] S. Byna, M. Breitenfeld, B. Dong, Q. Koziol, Elena Pourmal, D. Robinson, Jérôme Soumagne, Houjun Tang, Venkatram Vishwanath, and R. Warren. 2020. ExaHDF5: Delivering Efficient Parallel I/O on Exascale Computing Systems. *JCTST* 35 (2020), 145–160.
- [23] Adam G. Carlyle, Stephen L. Harrell, and Preston M. Smith. 2010. Cost-Effective HPC: The Community or the Cloud?. In *2010 IEEE Second International Conference on Cloud Computing Technology and Science*. 169–176. <https://doi.org/10.1109/CloudCom.2010.115>
- [24] Chia-Yu Chen, Jungwook Choi, Daniel Brand, Ankur Agrawal, Wei Zhang, and Kailash Gopalakrishnan. 2017. AdaComp : Adaptive Residual Gradient Compression for Data-Parallel Distributed Training. (2017). arXiv:1712.02679 [cs.LG] <https://arxiv.org/abs/1712.02679>
- [25] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems. arXiv:1512.01274 [cs.DC] <https://arxiv.org/abs/1512.01274>
- [26] Yu Chen, Zhenming Liu, Bin Ren, and Xin Jin. 2020. On efficient constructions of checkpoints. , 10 pages.
- [27] Peng Cheng and Haryadi S. Gunawi. 2021. Storage Benchmarking with Deep Learning Workloads. <https://api.semanticscholar.org/CorpusID:231845927>
- [28] Hyeonseong Choi, Byung Hyun Lee, Se Young Chun, and Jaehwan Lee. 2023. Towards accelerating model parallelism in distributed deep learning systems. *PLoS One* 18, 11 (Nov. 2023), e0293338.
- [29] Stack Overflow Community. 2024. *Search Results for Distributed-Training: Search query for distributed-training up to 2015: distributed training created...2015 and up to 2024: distributed training created...2024*.
- [30] Donglai Dai. 2022. SCR-Exa: Enhanced Scalable Checkpoint Restart (SCR) Library for Next Generation Exascale Computing. (02 2022). <https://www.osti.gov/biblio/1847927>
- [31] Marco Dantas et al. 2022. Accelerating Deep Learning Training Through Transparent Storage Tiering. In *2022 22nd IEEE International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. 21–30. <https://doi.org/10.1109/CCGrid54584.2022.00011>
- [32] Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness. arXiv:2205.14135 [cs.LG] <https://arxiv.org/abs/2205.14135>
- [33] Dask. 2024. *Dask: Scalable Analytics in Python*. <https://www.dask.org/> Accessed on February 5, 2024.
- [34] Hoang Anh Dau, Anthony Bagnall, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatana, and Eamonn Keogh. 2019. The UCR Time Series Archive. arXiv:1810.07758 [cs.LG] <https://arxiv.org/abs/1810.07758>
- [35] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- [36] Li Deng. 2012. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* 29, 6 (2012), 141–142.
- [37] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805 [cs.CL] <https://arxiv.org/abs/1810.04805>
- [38] Jacob Devlin, Myle Ott, Stephen Merity, Michael A. Clark, Kenton Lee, Marketa Zilka, Lukasz Kaiser, and David Cano. 2019. Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism. In *Advances in Neural Information Processing Systems*.
- [39] Markus Dreseler, Jan Kossmann, Martin Boissier, Stefan Klauk, Matthias Uflacker, and Hasso Plattner. 2019. Hyrise Re-engineered: An Extensible Database System for Research in Relational In-Memory Data Management. <https://doi.org/10.5441/002/edbt.2019.28>
- [40] Audibert Andrew et al. 2023. tf.data service: A Case for Disaggregating ML Input Data Processing. In *Proceedings of the 2023 ACM Symposium on Cloud Computing* (Santa Cruz, CA, USA). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3620678.3624666>
- [41] Abadi Daniel J. et al. 2008. Column-stores vs. row-stores: how different are they really?. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/1376616.1376712>
- [42] Assaf Eisenman et al. 2022. Check-N-Run: a Checkpointing System for Training Deep Learning Recommendation Models. In *19th USENIX Symposium on (NSDI) 22*. USENIX Association, Renton, WA, 929–943. <https://www.usenix.org/conference/nsdi22/presentation/eisenman>



- [43] Alonso-Monsalve Saúl et al. 2021. Analyzing the distributed training of deep-learning models via data locality. In *2021 29th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*. 117–121. <https://doi.org/10.1109/PDP52278.2021.00026>
- [44] Ali Peshawa Jamal Muhammad et al. 2014. Data normalization and standardization: a technical report. *Mach Learn Tech Rep* 1, 1 (2014), 1–6.
- [45] Bez Jean Luca et al. 2021. I/O Bottleneck Detection and Tuning: Connecting the Dots using Interactive Log Analysis. In *2021 IEEE/ACM Sixth International Parallel Data Systems Workshop (PDSW)*. 15–22. <https://doi.org/10.1109/PDSW54622.2021.00008>
- [46] Chowdhury Fahim et al. 2019. I/O Characterization and Performance Evaluation of BeeGFS for Deep Learning. In *Proceedings of the 48th International Conference on Parallel Processing (Kyoto, Japan)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3337821.3337902>
- [47] Chien Steven W. D. et al. 2020. tf-Darshan: Understanding Fine-grained I/O Performance in Machine Learning Workloads. In *2020 IEEE International CLUSTER*. IEEE Computer Society, Los Alamitos, CA, USA, 359–370. <https://doi.org/10.1109/CLUSTER49012.2020.00046>
- [48] Chen Tianshi et al. 2012. BenchNN: On the broad potential application scope of hardware neural network accelerators. *Proceedings - 2012 IEEE International Symposium on Workload Characterization, IISWC 2012*, 36–45. <https://doi.org/10.1109/IISWC.2012.6402898>
- [49] Chen Weijian et al. 2023. iCache: An Importance-Sampling-Informed Cache for Accelerating I/O-Bound DNN Model Training. In *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. 220–232. <https://doi.org/10.1109/HPCA56546.2023.10070964>
- [50] Devarajan Hariharan et al. 2021. DLIO: A Data-Centric Benchmark for Scientific Deep Learning Applications. In *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. 81–91. <https://doi.org/10.1109/CCGrid51090.2021.00018>
- [51] Dryden Nikoli et al. 2023. Clairvoyant Prefetching for Distributed Machine Learning I/O. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. Association for Computing Machinery. <https://doi.org/10.1145/3458817.3476181>
- [52] Dey Tonmoy et al. 2020. Optimizing Asynchronous Multi-Level Checkpoint/Restart Configurations with Machine Learning. In *2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. 1036–1043. <https://doi.org/10.1109/IPDPSW50202.2020.00174>
- [53] Edixon Párraga et al. 2021. *Analyzing the I/O Patterns of Deep Learning Applications*. Vol. 1444 CCIS. Springer Science and Business Media Deutschland GmbH, N/A, 3–16. [https://doi.org/10.1007/978-3-030-84825-5\\_1](https://doi.org/10.1007/978-3-030-84825-5_1)
- [54] Guillaume Leclerc et al. 2023. FFCV: Accelerating Training by Removing Data Bottlenecks. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition 2023-June*, 12011–12020. <https://doi.org/10.1109/CVPR52729.2023.01156>
- [55] Hu Yang et al. 2022. PipeEdge: Pipeline Parallelism for Large-Scale Model Inference on Heterogeneous Edge Devices. In *2022 25th Euromicro Conference on Digital System Design (DSD)*. 298–307. <https://doi.org/10.1109/DSD57027.2022.00048>
- [56] Jain Abhinav et al. 2020. Overview and Importance of Data Quality for Machine Learning Tasks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. Association for Computing Machinery. <https://doi.org/10.1145/3394486.3406477>
- [57] Jong Youl Choi et al. 2023. DDStore: Distributed Data Store for Scalable Training of Graph Neural Networks on Large Atomistic Modeling Datasets. *ACM International Conference Proceeding Series*, 941–950. <https://doi.org/10.1145/3624062.3624171>
- [58] Liberty Edo et al. 2016. Stratified Sampling Meets Machine Learning. In *Proceedings of The 33rd International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 48)*. PMLR, New York, New York, USA, 2320–2329. <https://proceedings.mlr.press/v48/liberty16.html>
- [59] Lee Sunwoo et al. 2021. Asynchronous I/O Strategy for Large-Scale Deep Learning Applications. In *2021 IEEE 28th International Conference on High Performance Computing, Data, and Analytics (HiPC)*. 322–331. <https://doi.org/10.1109/HiPC53243.2021.00046>
- [60] Marathe Aniruddha et al. 2013. A comparative study of high-performance computing on the cloud. In *Proceedings of the 22nd International Symposium on HPDC*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/2493123.2462919>
- [61] Mathuriya Amrita et al. 2018. CosmoFlow: Using Deep Learning to Learn the Universe at Scale. In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*. 819–829. <https://doi.org/10.1109/SC.2018.00068>
- [62] Mattson Peter et al. 2020. MLPerf: An Industry Standard Benchmark Suite for Machine Learning Performance. *IEEE Micro* 40, 2 (2020), 8–16. <https://doi.org/10.1109/MM.2020.2974843>
- [63] Nicolae Bogdan et al. 2019. VeloC: Towards High Performance Adaptive Asynchronous Checkpointing at Large Scale. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 911–920. <https://doi.org/10.1109/IPDPS.2019.00099>
- [64] Nicolae Bogdan et al. 2020. DeepFreeze: Towards Scalable Asynchronous Checkpointing of Deep Learning Models. In *2020 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID)*. 172–181. <https://doi.org/10.1109/CCGrid49817.2020.00-76>
- [65] Paul Arnab K. et al. 2020. Understanding HPC Application I/O Behavior Using System Level Statistics. In *2020 IEEE 27th International Conference on High Performance Computing, Data, and Analytics (HiPC)*. IEEE, Pune, India, 202–211. <https://doi.org/10.1109/HiPC50609.2020.00034>
- [66] Paul Arnab K et al. 2021. Characterizing Machine Learning I/O Workloads on Leadership Scale HPC Systems. In *2021 29th International Symposium MASCOTS*. 1–8. <https://doi.org/10.1109/MASCOTS53633.2021.9614303>
- [67] Robert Adolf et al. 2016. Fathom: reference workloads for modern deep learning methods. *2016 IEEE International Symposium on Workload Characterization (IISWC) (2016)*, 1–10. <https://api.semanticscholar.org/CorpusID:1809844>
- [68] Rasley Jeff et al. 2020. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (Virtual Event, CA, USA) (KDD '20)*. Association for Computing Machinery, New York, NY, USA, 3505–3506. <https://doi.org/10.1145/3394486.3406703>
- [69] Snyder Shane et al. 2016. Modular HPC I/O Characterization with Darshan. In *2016 5th Workshop on Extreme-Scale Programming Tools (ESPT)*. 2016 5th workshop on extreme-scale programming tools (ESPT), Salt Lake City, UT, USA, 9–17. <https://doi.org/10.1109/ESPT.2016.006>
- [70] Smith Samuel et al. 2020. On the Generalization Benefit of Noise in Stochastic Gradient Descent. In *Proceedings of the 37th International Conference on Machine Learning*. Hal Daumé III and Aarti Singh (Eds.), Vol. 119. PMLR, Cornell University. <https://proceedings.mlr.press/v119/smith20a.html>

- [71] Sun Xiaoyang et al. 2022. STRONGHOLD: Fast and Affordable Billion-Scale Deep Learning Model Training. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–17. <https://doi.org/10.1109/SC41404.2022.00076>
- [72] Wan Lipeng et al. 2020. I/O Performance Characterization and Prediction through Machine Learning on HPC Systems. In *CUG2020*. CUG, Dallas Convention Center Room A309-A310 650 S Griffin St Dallas, TX 75202, 10.
- [73] Xie Jian et al. 2022. A Deep Learning Dataloader with Shared Data Preparation. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., Virtual, 17146–17156.
- [74] Yao S. et al. 2020. Deep compressive offloading: speeding up neural network inference by trading edge computation for network latency. In *Proc. 18th SenSys (SenSys '20)*. Association for Computing Machinery, New York, NY, USA, 476–488. <https://doi.org/10.1145/3384419.3430898>
- [75] Zhu Maohua et al. 2018. Performance Evaluation and Optimization of HBM-Enabled GPU for Data-Intensive Applications. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 26, 5 (2018), 831–840. <https://doi.org/10.1109/TVLSI.2018.2791442>
- [76] Zhu Yue et al. 2018. Entropy-Aware I/O Pipelining for Large-Scale Deep Learning on HPC Systems. In *2018 IEEE 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*. 145–156. <https://doi.org/10.1109/MASCOTS.2018.00023>
- [77] Facebook Incubator. 2024. Gloo. <https://github.com/facebookincubator/gloo>.
- [78] Steven Farrell, Murali Emani, Jacob Balma, Lukas Drescher, Aleksandr Drozd, Andreas Fink, Geoffrey Fox, David Kanter, Thorsten Kurth, Peter Mattson, Dawei Mu, Amit Ruhela, Kento Sato, Koichi Shirahata, Tsuguchika Tabaru, Aristeidis Tsaris, Jan Balewski, Ben Cumming, Takumi Danjo, Jens Domke, Takaaki Fukai, Naoto Fukumoto, Tatsuya Fukushima, Balazs Gerofi, Takumi Honda, Toshiyuki Imamura, Akihiko Kasagi, Kentaro Kawakami, Shuhei Kudo, Akiyoshi Kuroda, Maxime Martinasso, Satoshi Matsuoka, Henrique Mendonça, Kazuki Minami, Prabhat Ram, Takashi Sawada, Mallikarjun Shankar, Tom St. John, Akihiro Tabuchi, Venkatram Vishwanath, Mohamed Wahib, Masafumi Yamazaki, and Junqi Yin. 2021. MLPerf HPC: A Holistic Benchmark Suite for Scientific Machine Learning on HPC Systems. arXiv:2110.11466 [cs.LG] <https://arxiv.org/abs/2110.11466>
- [79] Roy Fielding. 2000. *Architectural Styles and the Design of Network-based Software Architectures*. Ph.D. Dissertation.
- [80] Apache Software Foundation. 2024. *Apache Parquet*. Apache Software Foundation. <https://parquet.apache.org/> Version 2.10.0.
- [81] Ana Gainaru, Guillaume Aupy, Anne Benoit, Franck Cappello, Yves Robert, and Marc Snir. 2015. Scheduling the I/O of HPC Applications Under Congestion. In *2015 IEEE International Parallel and Distributed Processing Symposium*. 1013–1022. <https://doi.org/10.1109/IPDPS.2015.116>
- [82] Ana Gainaru, Dmitry Ganyushin, Bing Xie, Tahsin Kurc, Joel Saltz, Sarp Oral, Norbert Podhorszki, Franz Poeschel, Axel Huebl, and Scott Klasky. 2022. *Understanding and Leveraging the I/O Patterns of Emerging Machine Learning Analytics*. Springer International Publishing, Cham, 119–138.
- [83] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. The Pile: An 800GB Dataset of Diverse Text for Language Modeling. (2020). arXiv:2101.00027 [cs.CL] <https://arxiv.org/abs/2101.00027>
- [84] Markus Geimer, Felix Wolf, Brian J. N. Wylie, Erika Abraham, Daniel Becker, and Bernd Mohr. 2010. The Scalasca Performance Toolset Architecture. *Concurr. Comput. : Pract. Exper.* 22, 6 (apr 2010), 702–719.
- [85] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio Set: An ontology and human-labeled dataset for audio events. In *Proc. IEEE ICASSP 2017*. New Orleans, LA.
- [86] Anjus George, Rick Mohr, James Simmons, and Sarp Oral. 2021. Understanding Lustre Internals. Second Edition. (09 2021). <https://doi.org/10.2172/1824954>
- [87] Google Cloud. 2024. Google Analytics Data on BigQuery. <https://cloud.google.com/bigquery/public-data/google-analytics> Accessed: 2024-09-12.
- [88] Google Cloud. 2024. Google BigQuery Public Datasets. <https://cloud.google.com/bigquery/public-data> Accessed: 2024-09-12.
- [89] Google Cloud. n.d.. *Google Cloud Storage*. <https://cloud.google.com/storage>
- [90] NVIDIA GPUDirect. 2024. *Enhancing Data Movement and Access for GPUs*. <https://developer.nvidia.com/gpudirect> Accessed: February 26, 2024.
- [91] Sasun Hambardzumyan, Abhinav Tuli, Levon Ghukasyan, Fariz Rahman, Hrnt Topchyan, David Isayan, Mark McQuade, Mikayel Harutyunyan, Tatevik Hakobyan, Ivo Stranic, and Davit Buniatyan. 2022. Deep Lake: a Lakehouse for Deep Learning. arXiv:2209.10785 [cs.DC] <https://arxiv.org/abs/2209.10785>
- [92] Kaveen Hiniduma, Suren Byna, and Jean Luca Bez. 2025. Data Readiness for AI: A 360-Degree Survey. <https://doi.org/10.1145/3722214> Just Accepted.
- [93] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. 2020. Open graph benchmark: datasets for machine learning on graphs. , 16 pages.
- [94] Hongsun Jang, Jaeyong Song, Jaewon Jung, Jaeyoung Park, Youngsok Kim, and Jinho Lee. 2024. Smart-Infinity: Fast Large Language Model Training using Near-Storage Processing on a Real System. arXiv:2403.06664 [cs.AR] <https://arxiv.org/abs/2403.06664>
- [95] Paulo Jesus, Carlos Baquero, and Paulo Sergio Almeida. 2015. A Survey of Distributed Data Aggregation Algorithms. *IEEE Communications Surveys & Tutorials* 17, 1 (2015), 381–404. <https://doi.org/10.1109/COMST.2014.2354398>
- [96] Fei Jie, Haijun Zhang, Jinxiang Wang, and Zhenhua Yu. 2023. ASRDataset: A Multi-granularity Shuffle System for Preparing Large-scale ASR Training Data. In *2023 IEEE International Conference on Knowledge Graph (ICKG)*. 69–76. <https://doi.org/10.1109/ICKG59574.2023.00014>
- [97] Hai Jin, Zhanyang Zhu, Ligang He, Yuhao Li, Yusheng Hua, and Xuanhua Shi. 2024. MMDatLoader: Reusing Preprocessed Data Among Concurrent Model Training Tasks. *IEEE Trans. Comput.* 73, 2 (2024), 510–522. <https://doi.org/10.1109/TC.2023.3336161>
- [98] Ruoming Jin, Yuri Breitbart, and Chibuikue Muoh. 2009. Data discretization unification. *Knowledge and Information Systems* 19, 1 (2009), 1–29. <https://doi.org/10.1007/s10115-008-0142-6>

- [99] 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-Datacenter Performance Analysis of a Tensor Processing Unit. , 12 pages. <https://doi.org/10.1145/3079856.3080246>
- [100] Khalid M. Kahloot and Peter Ekler. 2021. Algorithmic Splitting: A Method for Dataset Preparation. *IEEE Access* 9 (2021), 125229–125237. <https://doi.org/10.1109/ACCESS.2021.3110745>
- [101] Anthony Kougkas, Matthieu Dorier, Rob Latham, Rob Ross, and Xian-He Sun. 2016. Leveraging burst buffer coordination to prevent I/O interference. In *2016 IEEE 12th International Conference on e-Science (e-Science)*. 371–380. <https://doi.org/10.1109/eScience.2016.7870922>
- [102] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images.
- [103] Michael Kuchnik, George Amvrosiadis, and Virginia Smith. 2021. Progressive compressed records: Taking a byte out of deep learning data. *Proceedings of the VLDB Endowment* 14, 2627–2641. Issue 11. <https://doi.org/10.14778/3476249.3476308>
- [104] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. 2011. HMDB: a large video database for human motion recognition. In *Proceedings of the International Conference on Computer Vision (ICCV)*.
- [105] Alexander Lavin and Subutai Ahmad. 2015. Evaluating Real-Time Anomaly Detection Algorithms – The Numenta Anomaly Benchmark. (dec 2015). <https://doi.org/10.1109/icmla.2015.141>
- [106] Choonghan Lee, Muqun Yang, and Ruth A. Aydt. 2008. NetCDF-4 Performance Report. (2008). <https://api.semanticscholar.org/CorpusID:15588867>
- [107] Claire Songhyun et al Lee. 2023. A Case Study of Data Management Challenges Presented in Large-Scale Machine Learning Workflows. In *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. 71–81. <https://doi.org/10.1109/CCGrid57682.2023.00017>
- [108] Jeongha Lee and Hyokyung Bahn. 2023. Analyzing Data Reference Characteristics of Deep Learning Workloads for Improving Buffer Cache Performance. *Applied Sciences* 13, 22 (2023). <https://doi.org/10.3390/app132212102>
- [109] En Li, Liekang Zeng, Zhi Zhou, and Xu Chen. 2020. Edge AI: On-Demand Accelerating Deep Neural Network Inference via Edge Computing. *IEEE Transactions on Wireless Communications* 19, 1 (2020), 447–457. <https://doi.org/10.1109/TWC.2019.2946140>
- [110] Chengjian Liu, Pengcheng Shi, Yihong Li, Lixin Liang, and Lin Lin. 2023. Efficient Data Loading for Deep Neural Network Training. In *2023 9th International Conference on Big Data Computing and Communications (BigCom)*. 211–218. <https://doi.org/10.1109/BIGCOM61073.2023.00036>
- [111] Xiaolan Liu, Jiaodong Yu, Jian Wang, and Yue Gao. 2020. Resource Allocation With Edge Computing in IoT Networks via Machine Learning. *IEEE Internet of Things Journal* 7, 4 (2020), 3415–3426. <https://doi.org/10.1109/JIOT.2020.2970110>
- [112] Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, Tianyang Liu, Max Tian, Denis Kocetkov, Arthur Zucker, Younes Belkada, Zijian Wang, Qian Liu, Dmitry Abulkhanov, Indraneil Paul, Zhuang Li, Wen-Ding Li, Megan Risdal, Jia Li, Jian Zhu, Terry Yue Zhuo, Evgenii Zheltonozhskii, Nii Osae Osae Dade, Wenhao Yu, Lucas Krauß, Naman Jain, Yixuan Su, Xuanli He, Manan Dey, Edoardo Abati, Yekun Chai, Niklas Muennighoff, Xiangru Tang, Muhtasham Oblokulov, Christopher Akiki, Marc Marone, Chenghao Mou, Mayank Mishra, Alex Gu, Binyuan Hui, Tri Dao, Armel Zebaze, Olivier Dehaene, Nicolas Patry, Canwen Xu, Julian McAuley, Han Hu, Torsten Scholak, Sebastien Paquet, Jennifer Robinson, Carolyn Jane Anderson, Nicolas Chapados, Mostofa Patwary, Nima Tajbakhsh, Yacine Jernite, Carlos Muñoz Ferrandis, Lingming Zhang, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2024. StarCoder 2 and The Stack v2: The Next Generation. arXiv:2402.19173 [cs.SE] <https://arxiv.org/abs/2402.19173>
- [113] Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. 2020. The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting* 36, 1 (2020), 54–74. <https://doi.org/10.1016/j.ijforecast.2019.04.014> M4 Competition.
- [114] Priyanka Mary Mammen. 2021. Federated Learning: Opportunities and Challenges. arXiv:2101.05428 [cs.LG] <https://arxiv.org/abs/2101.05428>
- [115] Avinash Maurya et al. 2024. DataStates-LLM: Lazy Asynchronous Checkpointing for Large Language Models. In *Proceedings of the 33rd International Symposium on High-Performance Parallel and Distributed Computing (HPDC '24)*. ACM, 227–239. <https://doi.org/10.1145/3625549.3658685>
- [116] Andrew McCallum. 2017. Cora Dataset. <https://doi.org/10.18738/T8/HUIG48>
- [117] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2023. Communication-Efficient Learning of Deep Networks from Decentralized Data. arXiv:1602.05629 [cs.LG] <https://arxiv.org/abs/1602.05629>
- [118] Qi Meng, Wei Chen, Yue Wang, Zhi-Ming Ma, and Tie-Yan Liu. 2019. Convergence analysis of distributed stochastic gradient descent with shuffling. *Neurocomputing* 337 (2019), 46–57. <https://doi.org/10.1016/j.neucom.2019.01.037>
- [119] Xiangrui Meng. 2013. Scalable Simple Random Sampling and Stratified Sampling. In *Proceedings of the 30th International Conference on Machine Learning*, Sanjoy Dasgupta and David McAllester (Eds.), Vol. 28. PMLR, Atlanta, Georgia, USA. <https://proceedings.mlr.press/v28/meng13a.html>
- [120] Microsoft 2024. *Open Database Connectivity (ODBC)*. Microsoft. <https://docs.microsoft.com/en-us/sql/odbc/reference/odbc>
- [121] Microsoft Azure. 2024. Azure Open Datasets. <https://azure.microsoft.com/en-us/services/open-datasets/> Accessed: 2024-09-12.
- [122] Microsoft Azure. 2024. NYC Taxi & Limousine Commission Data. <https://azure.microsoft.com/en-us/services/open-datasets/nyc-taxi-limousine/> Accessed: 2024-09-12.

- [123] Derek G. Murray, Jiri Simsa, Ana Klimovic, and Ihor Indyk. 2021. *tf.data: A Machine Learning Data Processing Framework*. arXiv:2101.12127 [cs.LG] <https://arxiv.org/abs/2101.12127>
- [124] Takéhiko Nakama. 2009. Theoretical analysis of batch and on-line training for gradient descent learning in neural networks. *Neurocomputing* 73, 1 (2009), 151–159. <https://doi.org/10.1016/j.neucom.2009.05.017>
- [125] David Newton, Raghu Pasupathy, and Farzad Yousefian. 2018. RECENT TRENDS IN STOCHASTIC GRADIENT DESCENT FOR MACHINE LEARNING AND BIG DATA. In *2018 Winter Simulation Conference (WSC)*. 366–380. <https://doi.org/10.1109/WSC.2018.8632351>
- [126] Truong Thao Nguyen, François Trahay, Jens Domke, Aleksandr Drozd, Emil Vatai, Jianwei Liao, Mohamed Wahib, and Balazs Gerofi. 2022. Why Globally Re-shuffle? Revisiting Data Shuffling in Large Scale Deep Learning. (2022), 1085–1096. <https://doi.org/10.1109/IPDPS53621.2022.00109>
- [127] NVIDIA. 2024. *GPUDirect RDMA: Direct Communication between NVIDIA GPUs*. <https://developer.nvidia.com/blog/gpudirect-storage> Accessed: February 26, 2024.
- [128] NVIDIA. 2024. *GPUDirect Storage: A Direct Path Between Storage and GPU Memory*. <https://docs.nvidia.com/cuda/gpudirect-rdma/index.html> Accessed: February 26, 2024.
- [129] NVIDIA Corporation. 2024. *NVIDIA DALI*. <https://docs.nvidia.com/deeplearning/dali/user-guide/docs/index.html>
- [130] Hadeel S. Obaid, Saad Ahmed Dheyab, and Sana Sabah Sabry. 2019. The Impact of Data Pre-Processing Techniques and Dimensionality Reduction on the Accuracy of Machine Learning. In *2019 9th Annual IEMECOM*. 279–283. <https://doi.org/10.1109/IEMECOMX.2019.8877011>
- [131] Travis E. Oliphant. 2006. *NumPy: A guide to NumPy*. *Trelgol Publishing* (2006). <https://numpy.org/>.
- [132] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. LibriSpeech: an ASR corpus based on public domain audio books. <http://www.openslr.org/12>. Accessed: February 26, 2024.
- [133] Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The RefinedWeb Dataset for Falcon LLM: Outperforming Curated Corpora with Web Data, and Web Data Only. arXiv:2306.01116 [cs.CL] <https://arxiv.org/abs/2306.01116>
- [134] Torben Kling Petersen. 2015. *Inside the Lustre File System*. Technical Paper. SEAGATE Technology.
- [135] Pytorch. 2024. *PyTorch 2.2.0 Documentation*. PyTorch. Accessed on February 1, 2024.
- [136] Samyam Rajbhandari, Olatunji Ruwase, Jeff Rasley, Shaden Smith, and Yuxiong He. 2021. ZeRO-Infinity: Breaking the GPU Memory Wall for Extreme Scale Deep Learning. arXiv:2104.07857 [cs.DC] <https://arxiv.org/abs/2104.07857>
- [137] Vijay Janapa Reddi, Christine Cheng, David Kanter, Peter Mattson, Guenther Schmuelling, Carole-Jean Wu, Brian Anderson, Maximilien Breughe, Mark Charlebois, William Chou, Ramesh Chukka, Cody Coleman, Sam Davis, Pan Deng, Greg Diamos, Jared Duke, Dave Fick, J. Scott Gardner, Itay Hubara, Sachin Iidgunji, Thomas B. Jablin, Jeff Jiao, Tom St. John, Pankaj Kanwar, David Lee, Jeffery Liao, Anton Likhomotov, Francisco Massa, Peng Meng, Paulius Micikevicius, Colin Osborne, Gennady Pekhimenko, Arun Tejusve Raghunath Rajan, Dilip Sequeira, Ashish Sirasao, Fei Sun, Hanlin Tang, Michael Thomson, Frank Wei, Ephrem Wu, Lingjie Xu, Koichi Yamada, Bing Yu, George Yuan, Aaron Zhong, Peizhao Zhang, and Yuchen Zhou. 2020. MLPerf Inference Benchmark. arXiv:1911.02549 [cs.LG] <https://arxiv.org/abs/1911.02549>
- [138] Brian D. Ripley. 1996. *Pattern Recognition and Neural Networks*. Cambridge University Press.
- [139] Elvis Rojas, Albert Njoroge Kahira, Esteban Meneses, Leonardo Bautista Gomez, and Rosa M Badia. 2021. A Study of Checkpointing in Large Scale Training of Deep Neural Networks. arXiv:2012.00825 [cs.DC] <https://arxiv.org/abs/2012.00825>
- [140] Sebastian Ruder. 2017. An overview of gradient descent optimization algorithms. arXiv:1609.04747 [cs.LG] <https://arxiv.org/abs/1609.04747>
- [141] J. Salamon, C. Jacoby, and J. P. Bello. 2014. A Dataset and Taxonomy for Urban Sound Research. In *22nd ACM International Conference on Multimedia (ACM-MM'14)*. Orlando, FL, USA, 1041–1044.
- [142] Frederic Schimmelpfennig et al. 2021. Streamlining distributed Deep Learning I/O with ad hoc file systems. In *2021 IEEE International Conference on Cluster Computing (CLUSTER)*. 169–180. <https://doi.org/10.1109/Cluster48925.2021.00062>
- [143] Scikit-learn. 2024. *scikit-learn: Machine Learning in Python*. <https://scikit-learn.org/stable/> Accessed on February 5, 2024.
- [144] Alexander Sergeev and Mike Del Balso. 2018. Horovod: fast and easy distributed deep learning in TensorFlow. arXiv:1802.05799 [cs.LG] <https://arxiv.org/abs/1802.05799>
- [145] Kazuhiro Serizawa et al. 2019. Accelerating Machine Learning I/O by Overlapping Data Staging and Mini-batch Generations. In *Proc. 6th IEEE/ACM BDCAT (Auckland, New Zealand) (BDCAT '19)*. Association for Computing Machinery, 31–34. <https://doi.org/10.1145/3365109.3368768>
- [146] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2020. Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism. arXiv:1909.08053 [cs.CL] <https://arxiv.org/abs/1909.08053>
- [147] Connor Shorten and Taghi M. Khoshgoftaar. 2019. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data* 6, 1 (06 07 2019), 60. <https://doi.org/10.1186/s40537-019-0197-0>
- [148] Connor Shorten, Taghi M. Khoshgoftaar, and Borko Furht. 2021. Text Data Augmentation for Deep Learning. *Journal of Big Data* 8, 1 (19 07 2021), 101. <https://doi.org/10.1186/s40537-021-00492-0>
- [149] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. 2016. Hollywood in Homes: Crowdsourcing Data Collection for Activity Understanding. (2016). arXiv:1604.01753 [cs.CV] <https://arxiv.org/abs/1604.01753>
- [150] Daniel Smilkov, Nikhil Thorat, Yannick Assogba, Charles Nicholson, Nick Kreeger, Ping Yu, Shanqing Cai, Eric Nielsen, David Soegel, Stan Bileschi, Michael Terry, Ann Yuan, Kangyi Zhang, Sandeep Gupta, Sarah Sirajuddin, D Sculley, Rajat Monga, Greg Corrado, Fernanda Viegas, and Martin M Wattenberg. 2019. TensorFlow.js: Machine Learning For The Web and Beyond. , 309–321 pages. [https://proceedings.mlsys.org/paper\\_files/paper/2019/file/acd593d2db87a799a8d3da5a860c028e-Paper.pdf](https://proceedings.mlsys.org/paper_files/paper/2019/file/acd593d2db87a799a8d3da5a860c028e-Paper.pdf)

- [151] Shane Snyder, Philip Carns, Kevin Harms, Robert Ross, Glenn K. Lockwood, and Nicholas J. Wright. 2016. Modular HPC I/O Characterization with Darshan. In *Proceedings of the 5th Workshop on Extreme-Scale Programming Tools* (Salt Lake City, Utah) (ESPT '16). IEEE Press, 9–17.
- [152] Xinying Song, Alex Salcianu, Yang Song, Dave Dopson, and Denny Zhou. 2021. Fast WordPiece Tokenization. arXiv:2012.15524 [cs.CL] <https://arxiv.org/abs/2012.15524>
- [153] D. Steinkraus, I. Buck, and P.Y. Simard. 2005. Using GPUs for machine learning algorithms. In *Eighth International Conference on Document Analysis and Recognition (ICDAR'05)*. 1115–1120 Vol. 2. <https://doi.org/10.1109/ICDAR.2005.251>
- [154] Jonathan J. Stickel. 2010. Data smoothing and numerical differentiation by a regularization method. *Computers & Chemical Engineering* 34, 4 (2010), 467–475. <https://doi.org/10.1016/j.compchemeng.2009.10.007>
- [155] SUN. 2007. *High-Performance Storage Architecture and Scalable Cluster File System*. Technical Report. Sun Microsystems, Inc.
- [156] Ivan Svogor, Christian Eichenberger, Markus Spanring, Moritz Neun, and Michael Kopp. 2022. Profiling and Improving the PyTorch Dataloader for high-latency Storage: A Technical Report. arXiv:2211.04908 [cs.LG] <https://arxiv.org/abs/2211.04908>
- [157] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel Emer. 2017. Efficient Processing of Deep Neural Networks: A Tutorial and Survey. arXiv:1703.09039 [cs.CV] <https://arxiv.org/abs/1703.09039>
- [158] TensorFlow. 2024. *TensorFlow 2.15.0 Documentation*. TensorFlow. [https://www.tensorflow.org/versions/r2.15/api\\_docs](https://www.tensorflow.org/versions/r2.15/api_docs)
- [159] TensorFlow. 2024. TensorFlow Hub. <https://www.tensorflow.org/hub>.
- [160] TensorFlow. 2024. TensorFlow Lite. <https://www.tensorflow.org/lite>.
- [161] TensorFlow. 2024. *TfRecord and tf.train.Example*. TensorFlow. [https://www.tensorflow.org/tutorials/load\\_data/tfrecord](https://www.tensorflow.org/tutorials/load_data/tfrecord) Accessed on 2024-02-19.
- [162] Rajeev Thakur. 2024. AuroraGPT: A Large-Scale Foundation Model for Advancing Science. <https://doi.org/10.5281/zenodo.13345059>
- [163] The HDF Group. 1997-2022. *Hierarchical Data Format, version 5*. The HDF Group. <https://www.hdfgroup.org/HDF5/>.
- [164] THUMOS Challenge Workshop. 2012. UFC-101 Dataset. <http://www.thumos.info/download.html>. Accessed: February 26, 2024.
- [165] C. Reid Turner, Alfonso Fuggetta, Luigi Lavazza, and Alexander L. Wolf. 1999. A conceptual basis for feature engineering. *Journal of Systems and Software* 49, 1 (1999), 3–15. [https://doi.org/10.1016/S0164-1212\(99\)00062-X](https://doi.org/10.1016/S0164-1212(99)00062-X)
- [166] Pablo Villalobos, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, Anson Ho, and Marius Hobbhahn. 2022. Machine Learning Model Sizes and the Parameter Gap. arXiv:2207.02852 [cs.LG] <https://arxiv.org/abs/2207.02852>
- [167] Lipeng Wang, Qiong Luo, and Shengen Yan. 2022. DIESEL+: Accelerating Distributed Deep Learning Tasks on Image Datasets. *IEEE Transactions on Parallel and Distributed Systems* 33 (5 2022), 1173–1184. Issue 5. <https://doi.org/10.1109/TPDS.2021.3104252>
- [168] Cong Xu, Suparna Bhattacharya, Martin Foltin, Suren Byna, and Paolo Faraboschi. 2021. Data-Aware Storage Tiering for Deep Learning. In *2021 IEEE/ACM Sixth International Parallel Data Systems Workshop (PDSW)*. 23–28. <https://doi.org/10.1109/PDSW54622.2021.00009>
- [169] Chenhao Xu, Youyang Qu, Yong Xiang, and Longxiang Gao. 2023. Asynchronous federated learning on heterogeneous devices: A survey. *Computer Science Review* 50 (2023), 100595. <https://doi.org/10.1016/j.cosrev.2023.100595>
- [170] Jaewon Yang and Jure Leskovec. 2012. Defining and evaluating network communities based on ground-truth. , 8 pages. <https://doi.org/10.1145/2350190.2350193>
- [171] Benbo Zha and Hong Shen. 2022. Adaptively Periodic I/O Scheduling for Concurrent HPC Applications. *Electronics* 11, 9 (2022). <https://doi.org/10.3390/electronics11091318>
- [172] Yukun Zhu et al. 2015. Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books. In *2015 IEEE International Conference on Computer Vision (ICCV)*. 19–27. <https://doi.org/10.1109/ICCV.2015.11>
- [173] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 2016. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. arXiv:1606.06650 [cs.CV] <https://arxiv.org/abs/1606.06650>