

MINATOLoader: Accelerating Machine Learning Training Through Efficient Data Preprocessing

Rahma Nouaji
McGill University
Canada

Stella Bitchebe
McGill University
Canada

Ricardo Macedo
INESC TEC & U. Minho
Portugal

Oana Balmau
McGill University
Canada

Abstract

Machine learning (ML) frameworks, such as PyTorch and TensorFlow, rely on data loaders to preprocess data before feeding it to accelerators. When preprocessing is inefficiently pipelined, GPUs can remain idle over long periods of time, leading to substantial training delays. For example, PyTorch’s default data loaders can cause up to 76% GPU idleness. A key bottleneck is the variability in preprocessing time across samples within the same dataset. Existing data loaders are oblivious to this variability, training all samples uniformly. In this case, a single slow sample can stall the entire batch, causing head-of-line blocking.

We present MINATOLoader, a general-purpose data loader for PyTorch that accelerates training and improves GPU utilization under single-server, multi-GPU settings. It continuously prepares data in background and constructs batches by prioritizing fast-to-process samples, while slower samples are processed in parallel.


Experiments conducted over NVIDIA V100 and A100 GPUs show that MINATOLoader accelerates training by up to 7.5× (3.6× on average) over PyTorch DataLoader and Pecan, and up to 3× (2.2× on average) over DALI. It also increases average GPU utilization from 46% with PyTorch to 90%, while preserving model accuracy and enabling faster convergence.

CCS Concepts: • Computing methodologies → Machine learning; Parallel algorithms; • Computer systems organization → Data flow architectures.


Keywords: ML training, Data loader, Data preprocessing

ACM Reference Format:

Rahma Nouaji, Stella Bitchebe, Ricardo Macedo, and Oana Balmau. 2026. MINATOLoader: Accelerating Machine Learning Training Through Efficient Data Preprocessing. In *European Conference on*



(a) Data preprocessing and training pipeline of PyTorch DataLoader.



(b) CPU and GPU usage of PyTorch DataLoader during 3D-UNet training.

Figure 1. Inefficient PyTorch DataLoader pipeline. Slow data samples delay the batch construction process, resulting in GPU under-utilization and poor training performance.

Computer Systems (EUROSYS '26), April 27–30, 2026, Edinburgh, Scotland Uk. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3767295.3769376>

1 Introduction

The efficacy of Machine Learning (ML) deployments relies on both high-quality algorithms and high-quality data, the latter being shaped through data preprocessing. Although considerable research focused on improving the training algorithms – leading to advances in techniques [22, 28, 40], software libraries [3, 5, 6, 31], and hardware accelerators (e.g., GPUs, TPUs, DPUs)– the data preprocessing step has been relatively underexplored [16, 37, 41]. Yet, data preprocessing plays an important role in training efficiency and model generalization. Transformations such as cropping, resizing, shuffling, and random augmentations are commonly applied on the fly during training to increase model robustness and accuracy [19, 32, 42, 46]. These operations expose the model to diverse inputs, improving generalization and training speed.

To pipeline the data preprocessing with training, modern ML frameworks rely on *data loaders*. These components act as intermediaries between the dataset and the model, streamlining and optimizing the process of accessing data



This work is licensed under a Creative Commons Attribution 4.0 International License.

EUROSYS '26, Edinburgh, Scotland Uk

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2212-7/26/04

<https://doi.org/10.1145/3767295.3769376>

for preprocessing, training, and inference. The data loader plays a critical role in the training pipeline, and inefficiencies in this stage can bottleneck the overall training performance.

In most modern applications, this functionality is handled by the PyTorch DataLoader [9], which is widely used due to PyTorch’s popularity. In the PyTorch DataLoader, data preprocessing runs on the CPU, and model training happens on the GPU, as depicted in Figure 1a. Data samples are loaded from storage into main memory, transformed in parallel, assembled into batches, and finally transferred to the GPU for training. To reduce loading latency, the data loader exposes several tuning knobs that enable parallel data loading and prefetching, allowing multiple workers to load and prepare data samples concurrently and in advance of training.

However, when preprocessing times vary across samples, tuning these parameters alone is not enough to avoid training delays and GPU stalls. This is because PyTorch constructs each batch synchronously, by waiting for all individual samples to be ready before sending it to the GPU. As a result, a single slow sample can block the entire batch and cause head-of-line blocking, where a single sample delays the batch creation to undergo training despite other samples being ready. This inefficiency is evident in the PyTorch DataLoader pipeline (Figure 1a), and results in significant GPU underutilization (Figure 1b). For instance, during heavy data transformations, the GPU often remains idle. Our experiments reveal that data preprocessing times can vary significantly across samples, even within the same dataset (§3.1). For example, in an image segmentation workload using the 3D-UNet model [15], preprocessing times range from 0.01 to 2.2 seconds per sample, with an average of approximately 0.5 seconds (Figure 2a). Importantly, it is not trivial to predict which samples will take longer to process, as no simple heuristic (e.g., sample size, number of transformations) reliably predicts preprocessing time across workloads (§3.2). This unpredictability makes it difficult to schedule data loading efficiently and exacerbates the impact of head-of-line blocking.

In this paper, we address the following question: *How can the data loader deliver preprocessed data to the GPU as fast as possible to minimize idle periods?*

To achieve this, we propose MINATOLoader, a general-purpose, drop-in replacement for the PyTorch DataLoader that improves training time and GPU utilization without requiring prior knowledge of the dataset or preprocessing pipeline. Contrary to prior work [4, 9, 18], which treats all data samples uniformly and blocks batch creation on the slowest sample, MINATOLoader introduces a sample-aware scheduling strategy that dynamically adapts to per-sample preprocessing variability. The key idea behind MINATOLoader is to prioritize fast-to-process samples and defer slow ones for later use. This is a simple yet powerful approach that improves training performance, is generalizable across workloads, and maintains virtually the same model accuracy despite sample reordering inside the batches.

In a nutshell, MINATOLoader provides a dynamic load balancer that classifies samples on the fly during training. All samples are optimistically assumed to be fast at first. After a brief profiling phase to gather lightweight heuristics, MINATOLoader applies a per-sample timeout – samples that exceed a given threshold are flagged as slow and processed in the background, while the rest are immediately used to construct training batches. Fast and slow samples are managed through separate queues, enabling MINATOLoader to bypass slow samples and prevent stalling batch construction.

To adapt to changing workload demands, MINATOLoader dynamically adjusts the number of CPU worker threads involved in preprocessing during training. It starts with a default number of CPU workers per GPU, selected based on factors such as dataset size, preprocessing complexity, and I/O intensity. During execution, MINATOLoader continuously monitors queue occupancy and CPU utilization to detect underutilization or bottlenecks. Based on these, it adaptively scales the number of workers to sustain high throughput.

We validate the efficiency of MINATOLoader through a comprehensive experimental evaluation across multiple testing scenarios on a single-server with multiple GPUs setup, over three widely-used data preprocessing workloads from the MLPerf benchmarking suite [31]: image segmentation (3D-UNet) [15], object detection (Mask R-CNN) [30], and speech recognition (RNN-T) [29]. We compare MINATOLoader against three baselines: the native PyTorch DataLoader [9]; DALI [4], a data loader framework from NVIDIA that delegates preprocessing to the GPU; and Pecan [18], a state-of-the-art data loader for TensorFlow that automates data preprocessing worker placement and transformation reordering decisions. Across all experiments, MINATOLoader significantly outperforms prior work, reducing total training time by up to 7.5× compared to PyTorch DataLoader and Pecan, and up to 3× compared to DALI. It also improves GPU utilization, increasing the average usage from 46.4% with PyTorch DataLoader to 90.5% with MINATOLoader. Furthermore, MINATOLoader scales effectively with the number of GPUs, and even when using a single GPU, it achieves comparable or better training performance than the remainder of data loaders configured with 4 GPUs (outperforming them by up to 60.6%). In memory-constrained settings, it maintains high throughput and avoids input stalls, outperforming existing data loaders up to 2×. Finally, we show that MINATOLoader has almost no impact on model accuracy while substantially accelerating convergence.

In summary, this paper makes the following contributions:

- We extensively study the causes of GPU idleness during training stemming from inefficient data preprocessing. Notably, we identify that sample preprocessing time variability is a major bottleneck, causing head-of-line blocking in state-of-the-art data loaders, including PyTorch DataLoader and DALI.

- We design a *dynamic, sample-aware load balancer* that classifies data samples as fast or slow at runtime, enabling efficient batch construction that avoids pipeline stalls and GPU idleness.
- We combine the above-mentioned observations and techniques into *MINATOLOADER*, a *general-purpose data loader* that proactively prepares training data without requiring any prior knowledge of the dataset or preprocessing pipeline. We open source MINATOLOADER at <https://github.com/Rahm-no/MinatoLoader> and Zenodo [36].

2 Background

This section first provides background on popular data loaders used throughout this paper. It then outlines the datasets we use in the paper and their data preprocessing pipelines.

2.1 Data Loaders

Data loaders iterate over datasets, serving data to the model. They fetch data from persistent storage, to system memory (DRAM), to GPU memory, enabling on-the-fly data augmentation and tasks like batching, shuffling, and parallelization.

PyTorch DataLoader [9] uses a standard iterator to access the dataset and supports parallel data loading with worker threads, where each worker can prefetch several samples (controlled by `prefetch_factor`). While these settings help increase throughput, they are limited when preprocessing becomes the bottleneck. The issue lies in how PyTorch constructs and processes batches. It first defines the order in which data indices are drawn from the dataset, and then groups these into batches. However, by default, the data loader is unaware of differences in preprocessing cost. Consequently, a batch may contain both fast and slow samples, leading to head-of-line blocking where an entire batch is delayed by its slowest sample, as illustrated in Figure 1a.

DALI [4] is a GPU-accelerated data preprocessing library from NVIDIA that allows users to define mixed CPU-GPU pipelines for preparing data during training. Its design leverages pipelining and asynchronous execution to improve data throughput. While these features can reduce preprocessing latency, we found that tuning DALI’s execution parameters had little effect on overall training time. In practice, DALI’s reliance on the GPU comes with trade-offs: it increases memory usage on the accelerator and limits flexibility, as most preprocessing operations are implemented as low-level CUDA kernels that are hard to extend or debug. Most importantly, GPUs are expensive and are best reserved for training. Using them for preprocessing is generally only feasible in environments with abundant GPU capacity, where dedicated GPUs can be allocated to preprocessing tasks.

Pecan [18] is a recent research state-of-the-art solution, built on `tf.data` [35] (TensorFlow’s data loader), which aims to reduce ML training costs by optimizing data preprocessing. Pecan automates data preprocessing, worker placement,

and transformation reordering decisions through two main policies: *AutoPlacement* and *AutoOrder*. The *AutoPlacement* policy dynamically schedules data preprocessing workers across CPU resources. *AutoOrder* reorders transformations in the input pipeline to increase per-worker preprocessing throughput. Pecan categorizes transformations as *inflationary*, if they increase data volume (e.g., image padding, one-hot encoding), or *deflationary* if they reduce data volume (e.g., sampling, filtering, cropping). Following this categorization, the *AutoOrder* policy moves deflationary transformations earlier in the pipeline and postpones inflationary transformations. To maintain correctness, Pecan divides the input pipeline into sections at specific transformations that act as barriers to ensure that reordering is restricted within these sections and does not cross barrier boundaries.

2.2 Data Preprocessing Pipelines

Throughout this paper, we use three representative workloads from the MLPerf Training Benchmark suite [31, 33], selected for their diversity in data modalities (3D image, 2D image, and audio), preprocessing pipelines, and dataset sizes. This variety allows us to test MINATOLOADER under different levels of preprocessing complexity. Table 1 depicts the preprocessing transformations applied over each workload.

Image Segmentation. This workload performs 3D segmentation of kidney tumors using the 29GB KiTS19 dataset [1, 20], and is trained with the 3D-UNet model [15]. The data pipeline applies five sequential preprocessing steps (Table 1). Before preprocessing, the size of input samples ranges from 30MB to 375MB, with an average of 136MB. After preprocessing, all samples are standardized to a uniform size of 10MB. We include this workload because it presents heavy and variable preprocessing, primarily due to the complexity and size of volumetric 3D medical images.

Object Detection. This workload detects objects in 2D images using the Mask R-CNN model [30] with a ResNet50 backbone. It operates on the 58GB COCO dataset [27], and applies three common preprocessing steps (Table 1). Before preprocessing, sample sizes range from 0.1MB to 1MB, with an average of 0.8MB. After preprocessing, the sample size increases to between 4MB and 12MB, averaging 7MB. We select this task as a representative, widely used computer vision benchmark with relatively lightweight preprocessing. Compared to image segmentation, its processing time is lower, which helps evaluate whether MINATOLOADER can still yield benefits when preprocessing is less accentuated.

Speech Recognition. This workload performs speech-to-text transcription using the RNN-T model [29], trained on the 228GB LibriSpeech dataset [38], which contains approximately 1,000 hours of English speech. As summarized in Table 1, the preprocessing pipeline transforms raw audio inputs, ranging from 0.06MB to 0.34MB (average 0.2MB), into spectrograms of size 0.4MB to 9MB (average 4MB).

Table 1. Preprocessing pipelines for object detection, image segmentation, and speech recognition (*Speech-X*) workloads. The most time-consuming steps are highlighted in **bold**. *LightStep* simulates lightweight preprocessing (e.g., volume normalization, frame splicing). *HeavyStep* simulates compute-intensive steps that may include complex filtering or augmentation (e.g., long-context time-stretching, multi-pass spectrogram enhancement) and is performed every 5 samples.

Workload	Preprocessing Pipeline
Object Detection	Resize → RandomHorizontalFlip → ToTensor → Normalize
Image Segmentation	RandomCrop → RandomFlip → RandomBrightness → GaussianNoise → Cast
Speech-3s	Pad → SpecAugment → FilterBank → FrameSplicing → PermuteAudio → LightStep (0.5s) → HeavyStep (3s)
Speech-10s	Pad → SpecAugment → FilterBank → FrameSplicing → PermuteAudio → LightStep (0.5s) → HeavyStep (10s)

We include this workload to evaluate MINATOLoader on a different data modality, namely audio, and to demonstrate that its benefits generalize beyond image-based tasks. Additionally, we design this workload as a microbenchmark, in order to showcase MINATOLoader’s performance under heavy compute: all samples undergo a *LightStep* transformation that takes 0.5 seconds and simulates lightweight preprocessing, such as volume normalization or frame splicing. Every fifth sample is additionally subjected to a *HeavyStep* transformation that simulates more compute-intensive steps, taking either 3 seconds (Speech-3s) or 10 seconds (Speech-10s). These may include complex filtering or augmentation, such as applying long-context time-stretching or multi-pass spectrogram enhancement.

3 Issues in the Data Preprocessing Pipeline

Modern accelerators can ingest data at TB/s rates [7], yet existing data loaders often fail to feed data to the accelerators at a matching pace. In this section, we present a comprehensive experimental study to uncover the bottlenecks introduced by the data preprocessing pipeline during training.


Hardware and OS configurations. We run experiments in two hardware configurations: *Config. A* is a server with two 64-core AMD EPYC processors, 512 GB of memory, 4× A100 40 GB NVIDIA GPUs, connected to a shared Lustre file system via a 200 Gb/s interconnect, using Rocky Linux 8; *Config. B* corresponds to a server with two 40-core Intel Xeon processors, 512GB of memory, 8× V100 32GB NVIDIA GPUs, and a 7TB NVMe SSD, using Ubuntu 20.04.

Workloads. We use three representative workloads from the MLPerf Training Benchmark, with different models and preprocessing pipelines, as described in §2.2 and Table 1.

Data loading frameworks. The experiments were conducted over three ML data loading frameworks presented in §2.1: PyTorch DataLoader [9], DALI [4], and Pecan [18].

3.1 Preprocessing Time Variability

We begin by exploring the variability in preprocessing times across individual data samples with the PyTorch DataLoader.



(a) Image segmentation (3D-UNet).

(b) Object detection (R-CNN).

Figure 2. Variability in per-sample preprocessing time for image segmentation and object detection workloads. The red dashed lines depict the average preprocessing time across all samples – 0.5s in (a) and 35ms in (b).

Table 2. Preprocessing time (in *ms*) for each workload.

Workload	Avg	Med.	P75	P90	Min–Max–Std
Obj. Det.	31	28	30	35	11–176–19
Img. Seg.	500	470	630	750	10–2230–197
Speech-3s	998	508	509	3008	502–3017–992
Speech-10s	2351	508	509	10008	502–10014–3757

Figure 2 depicts the preprocessing time of 25 randomly selected samples for two workloads: *Image Segmentation* (Figure 2a) using the KiTS19 dataset and the 3D-UNet model, and *Object Detection* (Figure 2b) using the COCO dataset with the R-CNN model. While the average preprocessing time is approximately 500ms for image segmentation and 35ms for object detection, individual samples show wide variability, ranging from 10ms to 2.5s in image segmentation and from 10ms to 200ms in object detection, despite undergoing identical transformation pipelines. Table 2 presents these statistics in more detail for the two workloads across the full datasets. The long-tailed latencies, observed in Figure 2 and Table 2 (P90 column), result from a combination of input heterogeneity (e.g., varying resolutions, image sparsity, compression formats), transformation logic, and randomized data augmentations triggered by only a subset of samples.

Despite this variability, existing data loaders treat all samples equally during batch creation. For instance, in the Image Segmentation pipeline in Table 1, **RandomCrop** is the first and the slowest transformation (338ms on average). However, because transformations are applied sequentially, **RandomCrop** will dictate the total time for the entire pipeline, cause head-of-line blocking (§3.3), and exacerbate the variability.




Figure 3. CPU and GPU usage of the Object Detection workload when using two heuristics: (a) image size and (b) transformation reordering.

Takeaway 1. *Data loaders should take into account the variability of the preprocessing time across samples for an efficient training pipeline.*

3.2 Predicting Sample Processing Time

We now explore the effectiveness of applying heuristics to predict sample processing cost, based on *image size* and *transformation reordering*. Intuitively, image size could be used as a proxy for data preprocessing time. We show that even though in some cases this heuristic is accurate, it does not generalize across workloads. The transformation reordering heuristic is introduced by Pecan [18]. For these experiments, we extended PyTorch DataLoader with a custom load balancer that implements both heuristics. Figure 3 depicts the CPU and GPU usage over time of each setup.

Image size. In the image segmentation workload, we observe a strong correlation between input sample size and preprocessing time. For instance, a 299 MB sample takes 1.4s to preprocess, while a 63 MB sample requires only 0.3s. This correlation stems from the significant variation in size of 3D medical images, as well as transformations such as RandomCrop and RandomBrightness whose costs scale with the input dimensions. However, this correlation does not generalize. In the object detection workload, for instance, a 408 KB image may be preprocessed in just 13ms, while a 220 KB image may take 155ms. As shown in Figure 3a, applying the image size heuristic in this context fails to accurately distinguish between fast and slow samples, leading to increased fluctuations in GPU usage.

Transformation reordering. We evaluated a transformation reordering heuristic inspired by Pecan. Under the object detection workload, the heuristic modifies the position of the Resize transformation based on its effect on sample size: if Resize increases the data size, it is applied at the end in the pipeline; if it reduces the data size, it is placed at the beginning. The rationale is to minimize the volume of data processed by subsequent transformations. However, as shown in Figure 3b, this reordering had limited impact: the GPU utilization is only $\approx 3\%$ better compared to PyTorch DataLoader (Figure 8). This is because Pecan is designed for disaggregated environments where the primary goal is to




Figure 4. Impact of prefetch parameter on training time. Increasing the number of batches pre-fetched does not improve the training in both (a) Pytorch and (b) DALI.

reduce training cost across distributed compute nodes. The same strategy, however, is not guaranteed to succeed in a single-server, multi-GPU setting, on workloads with high variability in preprocessing times, as the reordering of transformations does not address the batch creation blocking.

Takeaway 2. *Even though simple heuristics may work for specific situations, they do not generalize across workloads. More generic mechanisms are needed to distinguish between fast and slow sample processing.*

3.3 Head-of-Line Blocking

We now analyze head-of-line blocking during data preprocessing. We conducted experiments using PyTorch DataLoader, configured with 12 parallel worker threads, while training the image segmentation workload.

Figure 1b shows that both CPU and GPU resources operate in an inverse relation, where periods of CPU activity align with GPU idleness. Indeed, the GPU frequently remains underutilized, with an average usage of 57.4%. This inefficiency stems from the sequential nature of the PyTorch DataLoader preprocessing pipeline, as depicted in Figure 1a. Specifically, each worker fetches data samples from persistent storage and applies the corresponding transformations. The order in which samples are processed and grouped in batches is predetermined. However, as observed in §3.1, while most samples are processed quickly, a single slow sample can delay the entire batch from being moved to the GPU. Because batch construction is synchronous, training cannot proceed until all samples in the batch are processed, causing the GPU to be idle for a significant portion of the training time.

Takeaway 3. *Sequential execution of preprocessing and training is inefficient and results in high GPU idleness and, therefore, worse end-to-end training times.*

3.4 Tuning Prefetching Parameters

To improve data loading performance, data loaders often implement prefetching strategies to load data samples in advance. In the PyTorch DataLoader, the prefetch_factor parameter determines how many batches each worker loads in advance from persistent storage. In DALI, the prefetch

`_queue_depth` controls the number of batches buffered between pipeline stages. To understand the impact of such mechanisms under heavy preprocessing workloads, we varied these parameters across the image segmentation, speech, and object detection workloads.

Figure 4 depicts the overall training time under different configurations. While both mechanisms aim to overlap data loading with training, results show that increasing their values yields limited benefits because neither system reduces the per-sample transformation cost. In the PyTorch DataLoader, increasing the prefetch factor offers minimal (even almost nil) improvement and may cause out-of-memory (OOM) errors due to excessive buffering. In addition, in DALI, where preprocessing is offloaded to the GPU, higher prefetch queue depth increases GPU memory usage and preprocessing load, which can interfere with training computations and ultimately prolong overall training time.

Takeaway 4. Existing mechanisms, such as prefetching, fail to improve GPU usage and end-to-end training time under sample preprocessing time variability.

3.5 Moving Preprocessing to GPU

To improve the data pipeline throughput, DALI offloads the data preprocessing phase to the GPU. This approach helps maintain high GPU utilization across workloads (as observed in Figure 8), contrasting with the underutilization often observed in PyTorch DataLoader. However, sharing the GPU between preprocessing and training leads to resource contention, which can interfere with the overall training performance, especially in compute-intensive workloads. While dedicating a separate GPU for preprocessing is a possible workaround, it reduces the number of GPUs available for training, an unfavorable compromise given that training is typically more compute-demanding.

Takeaway 5. If possible, it is better to avoid data preprocessing on the GPU, to be able to allocate the more scarce GPU cycles to training.

4 Design and Implementation

MINATOLoader is a general-purpose data loader that proactively prepares training data to enable efficient batch construction, without requiring prior knowledge of the dataset or preprocessing pipeline. It is designed for local training scenarios (*i.e.*, single node, multiple GPUs) and is provided as a drop-in replacement of PyTorch DataLoader, following the same API. MINATOLoader also supports distributed training (§6). Next, we present an overview of the MINATOLoader architecture (§4.1), detail the key techniques that underpin its design (§4.2–§4.3), and discuss implementation details (§4.4).




Figure 5. MINATOLoader high-level design. It continuously enqueues preprocessed samples into a specified queue based on a load balancer decision. Concurrently, the GPU dequeues preprocessed data for training. We show MINATOLoader for one GPU, but it generalizes to multi-GPU settings.

4.1 MINATOLoader in a Nutshell

Figure 5 depicts the high-level architecture of MINATOLoader, which is built around three key stages designed to prevent pipeline stalls and maximize GPU utilization: data loading, preprocessing, and batch construction. For clarity, we show how MINATOLoader operates for one GPU, but it generalizes to multi-GPU settings, as shown in our experiments (§5). First, the data loading stage is identical to PyTorch DataLoader, following the same API. Second, at the core of the system is a *dynamic, sample-aware load balancer* (§4.2) designed for mitigating head-of-line blocking caused by slow preprocessing samples (§3) by classifying samples as fast or slow on-the-fly during training. To achieve this, it maintains four queues: a *fast queue* for quick preprocessing samples, a *slow queue* for samples exceeding a predefined preprocessing timeout, a *batch queue* that assembles training batches using both fast and slow queues, and a *temp queue* used to temporarily store long samples that need to be preprocessed off the critical path. MINATOLoader maintains one *batch queue* per GPU, and separate *fast*, *slow*, and *temp* queues for each CPU worker. These queues enable decoupling data preparation from batch construction, allowing MINATOLoader to prioritize readily available samples and delay slower ones without stalling the pipeline. Further, MINATOLoader includes a *worker scheduler* (§4.3) that dynamically tunes the number of CPU preprocessing workers to match the throughput demands of the training workers on the GPUs (omitted in Figure 5 for clarity).

Like PyTorch, MINATOLoader requests samples in random order, but it builds batches from whichever samples finish preprocessing first, rather than imposing a strict order. Additionally, slow samples are never preempted: they continue preprocessing in background and are added to future batches as soon as they are ready. Importantly, MINATOLoader does not defer these samples to the very end, ensuring that its randomized batches remain statistically similar to PyTorch DataLoader, as shown in §5.6.

Operation flow. Figure 5 depicts the end-to-end workflow of MINATOLoader. The process begins with data samples being fetched from persistent storage by preprocessing workers

(1). Once loaded, each sample undergoes the model’s preprocessing pipeline, where a sequence of transformations is applied in parallel by CPU worker threads (2). As discussed in §3.1, preprocessing times can vary significantly across samples. To overcome this, MINATOLoader’s *load balancer* applies a per-sample timeout to classify samples as fast or slow (3). At the same time, MINATOLoader maintains lightweight profiling to adjust the threshold used to determine whether the samples are fast or slow (4, explained in §4.2). Samples that complete preprocessing within the timeout period are inserted into the *fast queue*, while those that exceed it are moved to the *temp queue*, where they continue applying the transformations (4). In the latter, as soon as slow samples finish preprocessing (occurring in parallel with the rest of the pipeline), they are transferred to the *slow queue*, as they are ready to be included in the batch construction (5). After that, a dedicated *batch queue* (one per GPU) constructs training batches by fetching samples from both fast and slow queues (6). Contrary to prior work, this design ensures slow samples do not block batch creation (*i.e.*, head-of-line blocking). Concurrently, multiple GPU worker threads continuously dequeue ready-to-train batches from the *batch queue* (6). This process overlaps with data loading and preprocessing to ensure GPUs are constantly fed with data for training, minimizing idle periods. Finally, in the background, MINATOLoader’s *worker scheduler* continuously monitors queue occupancy and CPU utilization, and dynamically adjusts the number of CPU preprocessing workers and GPU workers to sustain high training performance.

4.2 Avoiding Head-of-Line Blocking

MINATOLoader integrates the takeaways from §3 by designing a load balancer to classify samples. Initially, MINATOLoader optimistically assumes that all samples are fast. If a sample takes longer than a predefined timeout, it is flagged as slow. At that point, its preprocessing is paused, and the sample is migrated to the *temp queue*, where its processing continues in the background.

Figure 6 illustrates an example of this process, where a slow sample (*e.g.*, sample 2) is migrated during execution to the *slow queue*, and later included in a subsequent training batch. We show that even though the ordering of the samples in each batch can slightly change when compared to the default PyTorch DataLoader (Figure 1a), the accuracy of the model training is not impacted. Intuitively, the training accuracy follows a similar trend when using MINATOLoader because, in any case, samples need to be shuffled during training. This decoupled design enables the GPU to remain consistently busy by eliminating the dependency between the slowest sample and the training pipeline.

The slow sample timeout is determined through a simple offline workload profiling, used to make an educated guess regarding the cutoff threshold between fast and slow samples. However, if the dataset sample used for the offline

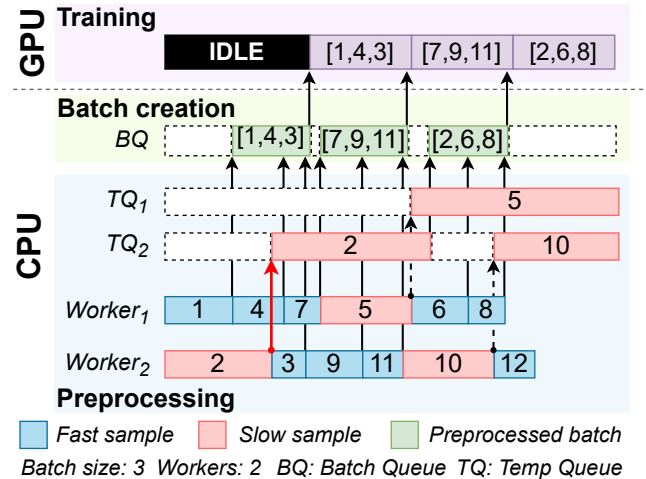


Figure 6. MINATOLoader preprocessing pipeline. Slow sample #2 does not delay the batch creation and gets migrated to the *temp queue* (red arrow). The fast and slow queues in the batch construction are omitted for brevity.

profiling is not representative of the entire workload, or if the workload drifts with time, MINATOLoader is able to automatically adjust the threshold. Before training begins, MINATOLoader performs a warm-up by running the model for a configurable period (*e.g.*, 10 minutes in our configuration). During this phase, it collects profiling statistics for each sample, including sample size, preprocessing time per transformation, total preprocessing time, and the number of transformations applied. At the end of the warm-up, the load balancer analyzes the recorded data and computes the 75th percentile of the total preprocessing times. This percentile is used as the default timeout value t_{out} by the load balancer, moving only the 25% slowest samples to *temp queue*. Unlike the median, which would split the dataset evenly, the 75th percentile strikes a better balance across workloads, focusing on true outliers and ensuring that *slow queue* remains smaller compared to *fast queue*. Although MINATOLoader uses the 75th percentile by default, the threshold t_{out} is adjustable. Indeed, if MINATOLoader detects that too many samples are incorrectly classified as slow (*e.g.*, due to a skewed distribution), it can automatically fall back to the 90th percentile. Moreover, MINATOLoader continues running the profiler in the background, in parallel with the preprocessing, adjusting the threshold as needed. This adaptive mechanism allows MINATOLoader to dynamically tune its timeout value for robust performance.

Algorithm 1 illustrates the core logic of MINATOLoader’s load balancer. Given a sample s and a list of transformations T , MINATOLoader applies the transformations sequentially while monitoring the elapsed time (7–10). If all transformations are completed within the timeout budget t_{out} , the sample is enqueued into the *fast_queue* (12). If the timeout t_{out} is exceeded, the system interrupts preprocessing, records

Algorithm 1: Load Balancer in MINATOLoader

Input: Transforms $T = [t_0, \dots, t_{n-1}]$, sample s , timeout t_{out}

```

1 Queues
2   fast_queue // Samples completed within timeout
3   temp_queue // Partially processed samples +
   transformation index
4   slow_queue // Completed timed-out samples
5   batch_queue // Batches ready for training
6 start ← current time ;
7 partial ← s // Initialize sample
8 for  $i = 0$  to  $n - 1$  do
9   partial ←  $t_i(\textit{partial})$  ;
10  if elapsed time since start >  $t_{out}$  then
11  | temp_queue.put(partial, i) // Timeout
12 fast_queue.put(partial) // Finished within timeout
13 Background Worker: Resume Timed-out Samples
14 while true do
15 | (partial, i) = temp_queue.pop() ;
16 | for  $j = i$  to  $n - 1$  do
17 | | partial ←  $t_j(\textit{partial})$ 
18 | | slow_queue.put(partial) // Now preprocessed
19 Batch Construction Thread
20 while true do
21 | batch ← [] ;
22 | while  $\text{len}(\textit{batch}) < \textit{batch\_size}$  do
23 | | if fast_queue not empty then
24 | | | sample ← fast_queue.pop() ;
25 | | | else if slow_queue not empty then
26 | | | | sample ← slow_queue.pop() ;
27 | | | else
28 | | | | sleep( $t$ ); // Not enough samples
29 | | | | append sample to batch ;
30 | | batch_queue.put(batch)
31 MINATOLoader__next__() Call
32 while true do
33 | if batch_queue not empty then
34 | | batch ← batch_queue.pop() ;
35 | | return batch // Send to GPU
36 | | else
37 | | | sleep( $t$ ); // Batch queue empty

```

the index of the transformation that was in progress, and enqueues it along with its index into the *temp_queue* (11). Since the last transformation was only partially applied, it must be re-executed to ensure that the sample is in a valid state to carry on the remaining transformations. This design avoids restarting the entire data preprocessing pipeline.

Afterwards, a background worker resumes the preprocessing of the slow samples from the recorded index and, once complete, places them into the *slow_queue* (14–18). Meanwhile, a batch construction thread continuously builds batches by pulling from both *fast_queue* and *slow_queue* (22–30). Finally, MINATOLoader continuously dequeues ready-to-train batches from the *batch_queue* (32–37). Empirically, we determine 10 ms as the optimal sleep time (28 and 37).

4.3 Keeping the GPUs Busy

In addition to mitigating head-of-line blocking, MINATOLoader also ensures that the GPUs stay busy by automatically determining the right number of CPU workers needed to keep up with each GPU. Each CPU worker is mapped to one CPU core. The optimal number of CPU workers can vary across workloads due to differences in data loading and preprocessing complexity. This is because preprocessing time is not uniform; some workloads involve more compute-intensive or variable transformations, which affect how many workers are needed to sustain throughput.

The CPU workers include the number of data loading workers, the slow-task workers, and the batch workers. The data loading CPU workers are responsible for loading data from disk, preprocessing fast samples, and moving slow samples to the *temp_queues*. Slow-task CPU workers run in the background, without blocking the pipeline. The batch CPU workers pull samples from the *slow* and *fast_queues* to construct batches ahead of time and ensure that the GPU workers always have prebuilt batches to ingest. In addition, MINATOLoader incorporates a prefetching mechanism that uses a CUDA stream to transfer the batch i from the *batch_queue* to the GPU memory before the GPU has finished executing the batch $i - 1$, which helps achieve continuous GPU training.

MINATOLoader starts by setting the number of CPU workers to 12 and *max_workers* to the total number of CPU cores. The initialization for the CPU workers can be tuned further depending on system characteristics such as dataset size and data preprocessing complexity. While the workload is running, MINATOLoader adjusts the number of CPU workers by monitoring average batch queue size and CPU utilization. The number of CPU workers is updated using Formula 1:

$$\text{workers} = \min(\text{max_workers}, \max(1, \text{workers}' + \Delta)) \quad (1)$$

$$\Delta = \alpha \cdot \left(1 - \frac{Q_{\text{size}}}{Q_{\text{max}}}\right) + \beta \cdot (C_{\text{usage}} - \theta_c) \quad (2)$$

where: Q_{size} is the moving average of the queue size, Q_{max} is the maximum queue capacity, C_{usage} is the average CPU utilization (normalized to $[0,1]$), θ_c is the CPU threshold (e.g., 0.7), and α and β are scaling factors controlling the sensitivity of queue and CPU adjustments, respectively.

Intuitively, Formula 1 increases the number of CPU workers when the queues are frequently empty ($Q_{\text{size}} \ll Q_{\text{max}}$) and/or CPU utilization is high, indicating a CPU bottleneck. Conversely, it reduces the number of workers when the

queues are full and CPU usage is low, which indicates over-provisioning. Formula 2 clips the final value of Δ to a small integer range (e.g., $[-2, +2]$) to avoid instability.

4.4 Implementation

MINATOLoader seamlessly integrates with PyTorch DataLoader interface, requiring no annotations or modifications to run, and preserves existing interfaces. It is implemented using 750 LoCs of Python: 350 LoCs for the data loading implementation and 300 LoCs for the data preprocessing.

Integration with PyTorch. Our implementation is based on PyTorch 2.7 and uses `torch.multiprocessing.Process` to launch processes, and `torch.multiprocessing.Queue` for the queues. Unlike threads, `torch.multiprocessing` provides isolated memory spaces. However, PyTorch provides support for sharing CPU tensors via shared memory.

`torch.multiprocessing` uses the `spawn` start method by default (which is safer for CUDA operations) and bypasses the Global Interpreter Lock (GIL). This allows full parallelism across cores, especially when executing Python-based data transformations. Another approach would be an implementation in C or C++ using Pybind11 to further accelerate data loading. However, it requires taking the GIL to access PyTorch preprocessing libraries, effectively introducing a synchronization bottleneck [11]. A final approach would be to reimplement all preprocessing operations in C/C++ and bypass PyTorch altogether; however, while likely improving performance, we argue that this approach is not practical for end users who typically call PyTorch libraries when defining their data preprocessing pipelines.

The `torch.multiprocessing.Queue` provides a process-safe, first-in-first-out (FIFO) communication channel between multiple producers and consumers. When multiple producers concurrently invoke `Queue.put()`, the operations are synchronized via internal locks, ensuring atomic insertion while preserving ordering. Similarly, multiple consumers calling `Queue.get()` block until items are available, with the operating system’s scheduler determining which consumer retrieves the next item.

5 Evaluation

Our evaluation answers the following questions:

- **Q1:** Can MINATOLoader improve training time and GPU utilization across different workloads (§5.2–§5.3)?
- **Q2:** How does MINATOLoader scale under varying hardware configurations (§5.4)?
- **Q3:** How does MINATOLoader behave under memory-constrained environments (§5.5)?
- **Q4:** Does MINATOLoader impact model accuracy (§5.6)?
- **Q5:** How does MINATOLoader behave under varying proportions of slow samples (§5.6)?

Table 3. Training configurations used for each workload.

Workload	#Epochs	#Iterations	Batch Size
Image segmentation	50	-	3
Object detection	-	1000	48
Speech recognition	-	1000	24

5.1 Experimental Environment

Hardware and OS configurations. We used the same hardware and OS configurations described in §3.

Data loaders. Experiments were executed using the PyTorch ML framework, configured under four distinct data loaders (§2.1) – PyTorch DataLoader [9], DALI [4], Pecan [18], and MINATOLoader. While Pecan was originally developed for TensorFlow, we reimplemented its `AutoOrder` policy (§2.1) in PyTorch to ensure a consistent and fair comparison across all systems. As this paper targets single-node multi-GPU scenarios, we have not used Pecan’s `AutoPlacement` policy.

Workloads and data loaders tuning. Experiments were conducted over four workloads from the MLPerf Training Benchmark [2], including image segmentation (3d-UNet), object detection (R-CNN), and speech recognition (RNN-T) (§2.2). The original pipeline of each workload is described in Table 1. Table 3 depicts the training configurations used for each model. We did our best to tune each system according to the specific characteristics of each workload’s pipeline, as described below (unless stated otherwise).

PyTorch DataLoader. We set the number of worker threads to 12 and used a `prefetch_factor` of 2, as increasing beyond these values had little effect on training time (§3.4).

DALI. We enabled DALI’s `exec_pipelined` and `exec_async` flags, which overlap computation stages through buffering and enable asynchronous execution between the DALI backend and the Python front end, respectively. Similarly to PyTorch DataLoader, we used the default `prefetch_queue_depth` of 2. The number of worker threads was set to the total number of CPU cores available on the machine. Furthermore, to account for GPU-accelerated preprocessing and ensure a fair comparison, we measured the average execution time of speech workload transformations (e.g., spectrogram, normalization) in both PyTorch DataLoader and DALI under identical conditions. We found that DALI was consistently 10× faster. Based on this, we scaled down the *LightStep* and *HeavyStep* transformations (Table 1) by a factor of 10 for DALI: *LightStep* was set to 0.05s, and *HeavyStep* to 0.3s (Speech-3s) and 1s (Speech-10s).

Pecan. In the speech recognition workload, Pecan’s `AutoOrder` policy moves the `Pad` transformation to the end of the pipeline, as it is an inflationary transformation. In the object detection workload, `Resize` is either inflationary or deflationary depending on the input size; thus, Pecan moves it to the end of the pipeline if it inflates the sample, or leaves it in




Figure 7. Throughput (MB/s) of PyTorch DataLoader, Pecan, DALI, and MINATOLoader on four ML workloads using A100 GPUs (*Config. A*). Solid vertical lines represent the end of the training for each data loader.

the original position otherwise. For the image segmentation workload, the AutoOrder algorithm is not applied because the transformations are already optimally ordered.

MINATOLoader. The initial worker configuration is of 1 GPU worker per GPU and 12 CPU workers per GPU worker. We use a prefetch factor of 2, and set the sample timeout threshold to the 75th percentile of all samples’ preprocessing time. All maximum queue sizes are set to 100.

Collected metrics. In all experiments, we measured the overall training time (in seconds) and model throughput (in MB/s), GPU and CPU utilization, memory usage, and disk read bandwidth. We represent the model throughput as the cumulative size in MB of data samples trained per second. We used the `nvidia-smi` tool from the NVIDIA Management Library [8] to monitor GPU utilization, and `dstat` [10] to collect CPU, memory, and disk usage.

5.2 End-to-end Training Performance

We start by evaluating the impact of MINATOLoader on end-to-end training performance. Figures 7 and 9 depict the throughput and training time of all systems. Due to space constraints, we omit the throughput plots for the *Config. B* testbed, composed of V100 GPUs. For the image segmentation workload, since the transformations are already ordered, Pecan’s results are identical to those of PyTorch DataLoader and therefore omitted.

Results show that MINATOLoader consistently achieves the highest throughput across all workloads. In the image segmentation workload (Figure 7a), MINATOLoader improves throughput by 2.5× over PyTorch DataLoader and 1.3× over DALI. For object detection (Figure 7b), MINATOLoader increases the throughput of PyTorch DataLoader and Pecan up to 2× and 1.6× of DALI. In the speech recognition workloads (Figures 7c and 7d), which comprise the most time-consuming transformations, MINATOLoader outperforms PyTorch DataLoader and Pecan by 3.5× to 5.5×, and DALI by 2× on average. These performance gains stem from its ability to eliminate head-of-line blocking in the pipeline by prioritizing fast samples and minimizing GPU idleness.

Moreover, as depicted in Figure 9, MINATOLoader achieves the shortest training times across all workloads. Over the

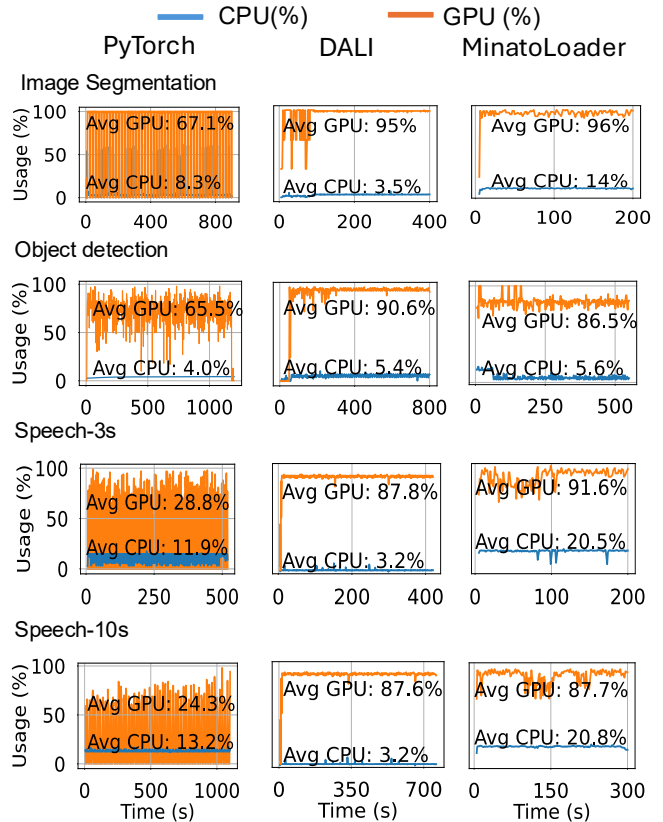


Figure 8. CPU and GPU usage for all systems across all workloads using 4xA100 GPUs.

Config. A testbed (4xA100 GPUs), MINATOLoader improves training time up to 7.5×, 4.9×, and 3×, when compared to PyTorch DataLoader, Pecan, and DALI, respectively. Over *Config. B* testbed (8xV100 GPUs), using an older GPU architecture, MINATOLoader still achieves substantial gains, improving training time up to 4.9× compared to PyTorch DataLoader and Pecan, and 2.6× over DALI.

5.3 Compute Resources Usage

We now compare the GPU and CPU utilization of MINATOLoader against the baselines. Due to space constraints, we report results for *Config. A* (A100 GPUs) in Figure 8; results for *Config. B* follow similar trends. Further, we omit Pecan

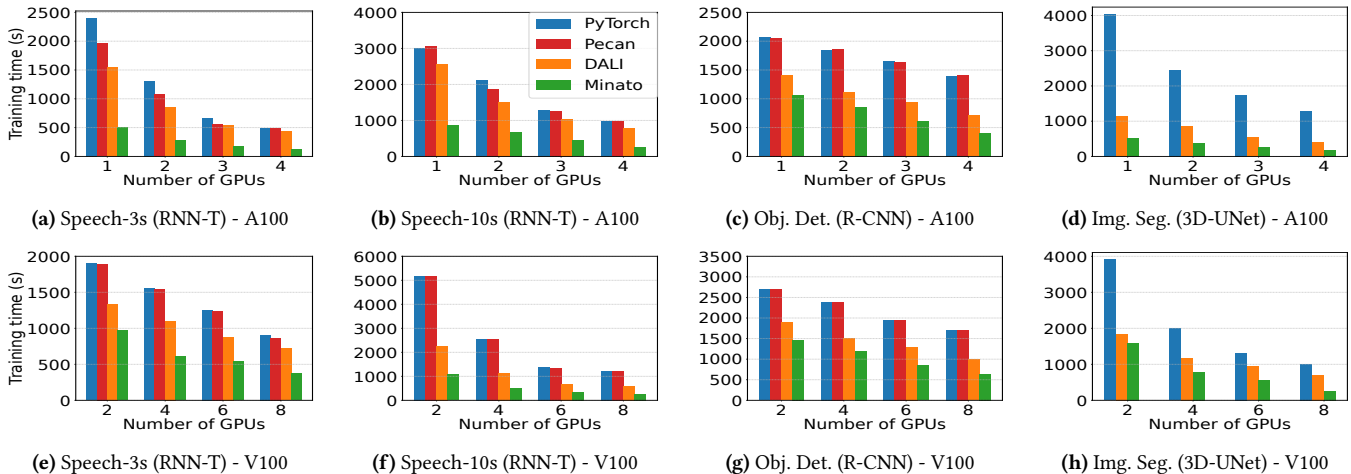


Figure 9. Training time (in seconds) of PyTorch DataLoader, Pecan, DALI, and MINATOLoader on four ML workloads executed over a varying number of A100 (top) and V100 GPUs (bottom).

from this analysis as its utilization closely mirrors that of PyTorch DataLoader, thus drawing similar conclusions.

Among the baselines, DALI achieves the highest GPU utilization by performing both preprocessing and training on the GPU, effectively maximizing usage. Similarly, MINATOLoader achieves a near-100% GPU usage, averaging 90.45% across all workloads. However, unlike DALI, MINATOLoader achieves this without offloading preprocessing to the GPU, whose usage exclusively reflects training activity. This highlights the effectiveness of MINATOLoader’s load balancer and worker scheduler components in mitigating head-of-line blocking and minimizing GPU idleness.

Moreover, we observe drops in DALI’s GPU usage during the first epoch of the image segmentation workload, also manifested in its throughput (Figure 7a). These drops stem from I/O contention on the shared filesystem in *Config. A* server. Specifically, as DALI aggressively loads and decodes data directly into GPU memory, it can temporarily saturate I/O bandwidth during the initial batch loading phase, creating read contention and pipeline stalls. Once data is cached, GPU usage stabilizes. As for PyTorch DataLoader, it experiences poor GPU utilization (averaging 46.4%) due to its synchronous behavior, causing head-of-line blocking. Finally, as expected, MINATOLoader exhibits slightly higher CPU utilization than PyTorch DataLoader (up to 20% utilization), due to its load balancer and adaptive worker scheduler.

5.4 Scalability

We now evaluate MINATOLoader’s scalability by varying the number of GPUs from 1× to 4×A100 (*Config. A*, Figures 9a–9d), and from 2× to 8×V100 (*Config. B*, Figures 9e–9h).

As expected, training time decreases with more GPUs in all data loaders. However, MINATOLoader consistently outperforms all systems across all workloads and testbeds. Moreover, some baselines fail to scale efficiently. For instance,

on 1×A100, DALI trains Speech-3s 25.6% faster than Pecan, but Pecan catches up on 3×A100. In contrast, MINATOLoader not only maintains its improvement as GPU count increases, but even achieves comparable or better training performance with 1×A100 than the baselines configured with 4×A100, effectively addressing the dominant bottleneck. For example, in the image segmentation workload, MINATOLoader on a single A100 reduces training time by up to 60% compared to PyTorch DataLoader and DALI using all GPUs.

5.5 Performance Under Memory Constraints

We now evaluate the effectiveness of MINATOLoader’s design on memory-constrained environments. We generated a 230GB dataset from KiTS19 by replicating it. We trained the image segmentation workload (3D-UNet) for 10 epochs, while limiting *Config. B*’s memory to 80GB using Linux cgroups (*i.e.*, approximately one-third of the dataset size), forcing all data loaders to access persistent storage. We measured the GPU and CPU usage and disk read bandwidth. Figure 10 presents the results. Again, Pecan is omitted from these results because it behaves similarly to PyTorch.

The PyTorch DataLoader (top row) experiences frequent memory pressure, resulting in continuous but volatile disk reads, reduced GPU usage (averaging around 57%), and prolonged training time (≈ 650 seconds). DALI (middle row) improves GPU usage (81.2% on average) but exhibits multiple drops below 50%, leading to a training time of around 500 seconds. To avoid substantial GPU usage drops, DALI performs data loading from persistent storage on the CPU. On the other hand, MINATOLoader (bottom row) maintains consistently high GPU utilization (82.1% on average) and completes training in 330 seconds. Disk I/O activity remains stable and high (maximizing the NVMe bandwidth) with MINATOLoader, only showing periodic drops at the end of every epoch due to model validation. This demonstrates

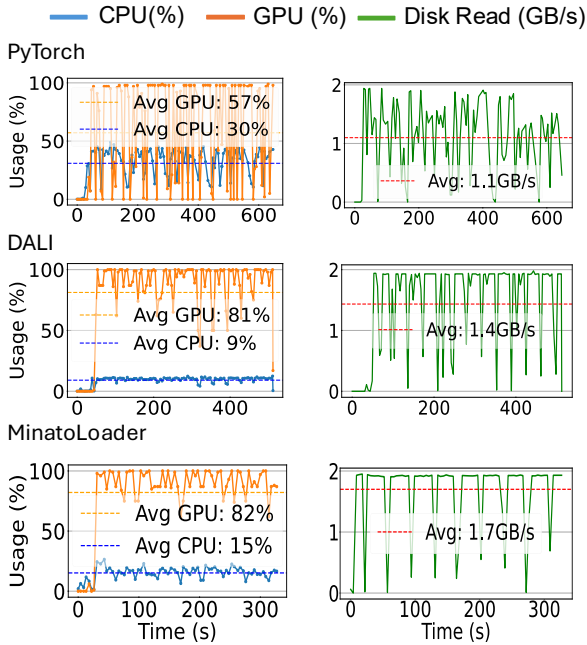


Figure 10. CPU and GPU usage (left) and disk read (right) of the image segmentation workload with PyTorch DataLoader ($\approx 650s$), DALI ($\approx 500s$), and MINATOLoader ($\approx 330s$) when training a 230GB dataset under an 80GB memory limit.

that MINATOLoader can handle memory-constrained scenarios thanks to its design that leverages multiple queues to pipeline data preprocessing and training.

5.6 Sensitivity Analysis

We now evaluate distinct features of MINATOLoader.

Accuracy Preserving with MINATOLoader. To demonstrate that MINATOLoader does not affect the model performance, we conducted experiments measuring the accuracy of the image segmentation (3D-UNet) and object detection (Mask R-CNN) workloads. Due to space limitations and because speech recognition was primarily used as a microbenchmark, we omit the training results. Achieving good accuracy generally requires training for a large number of iterations. For example, training the Mask R-CNN model requires between 90,000 and 180,000 iterations to reach convergence [30], which is ≈ 14 days of training on our testbed. Since the goal of MINATOLoader is not to improve the accuracy, it suffices to show here that it preserves the same model accuracy trend as other data loaders, while accelerating training. We trained the Mask R-CNN model for 45,000 iterations using PyTorch DataLoader and MINATOLoader. With MINATOLoader, the model reached the same 6% accuracy but 60% faster, completing in 5 hours and 12 minutes, compared to 13 hours and 55 minutes with PyTorch DataLoader. The results, shown in Figure 11a, confirm that MINATOLoader maintains a similar accuracy to PyTorch DataLoader.

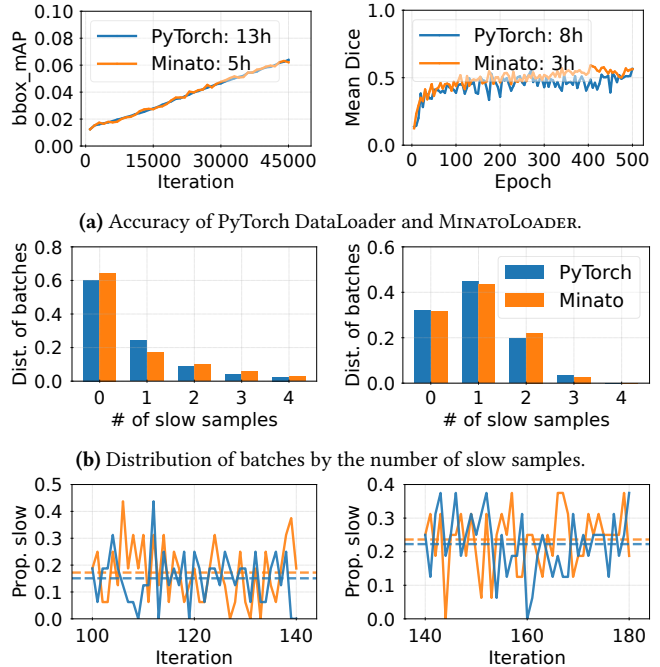


Figure 11. Sensitivity analysis of PyTorch DataLoader (in blue) and MINATOLoader (in orange) with object detection (left) and image segmentation (right) workloads.

For image segmentation, we trained the model until accuracy stabilized, which required approximately 500 epochs. This took 8 hours and 2 minutes with PyTorch DataLoader and only 3 hours and 52 minutes with MINATOLoader. As illustrated in Figure 11a, MINATOLoader preserves the full accuracy trend and reaches the same final accuracy of 58%.

Batch composition analysis. To complement the accuracy results, we conducted a batch composition study to validate that MINATOLoader’s strategy does not introduce bias. Specifically, we evaluated (i) the distribution of batches by the number of slow samples they contain (Figure 11b), *i.e.*, the proportion of batches with a given number of samples, and (ii) the proportion of slow samples over training iterations (Figure 11c). Experiments were conducted on the object detection (R-CNN) and image segmentation (3D-UNet) workloads with a batch size of 4.

Results show that MINATOLoader constructs mixed batches comparable to PyTorch DataLoader and does not introduce systematic bias. Figure 11b shows a similar distribution of batches across systems for both workloads, indicating that MINATOLoader preserves the natural ratio of slow samples. Figure 11c further shows that slow samples are incorporated into batches as soon as they are ready, rather than being deferred or preempted. For example, PyTorch DataLoader yields an average proportion of slow samples of 0.15 and 0.23

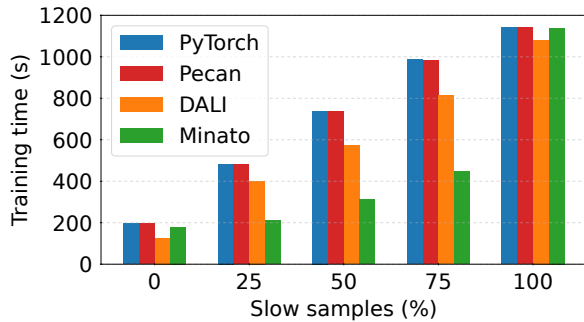


Figure 12. Training time of PyTorch DataLoader, Pecan, DALI, and MINATOLoader across different proportions of slow samples.

under the object detection and image segmentation workloads, respectively, while MINATOLoader presents 0.17 and 0.24. These results confirm that MINATOLoader preserves fairness in sample usage and maintains a batch distribution closely aligned with PyTorch DataLoader.

Cluster of slow samples. We now evaluate MINATOLoader under varying proportions of slow and fast samples. To this end, we modified the Speech-3s workload so that the *HeavyStep* transformation is applied to a configurable proportion of the dataset, rather than every five samples. Experiments were conducted with increasing proportion of slow samples from 0% up to 100%. Figure 12 depicts the training time of PyTorch DataLoader, DALI, Pecan, and MINATOLoader. As expected, in the edge cases (*i.e.*, 0% and 100%), MINATOLoader performs similarly to PyTorch DataLoader and Pecan, since all samples have uniform preprocessing cost. The benefits of MINATOLoader appear in the intermediate range (25%-75%), whose strategy leverages the variability across data samples, outperforming existing solutions up to 2.4 \times .

6 Discussion

Order-sensitive scenarios. Some training scenarios require strict sample ordering. In multimodal models (*e.g.*, image-text or audio-text), it is essential to maintain alignment between paired modalities. In the speech recognition workload (§2.2), MINATOLoader orders samples based on audio preprocessing time but always processes the audio-text pair together, preventing mismatches. This approach generalizes to other multimodal settings. Similarly, in curriculum learning [43], where the global sample order is semantically important (*e.g.*, easier examples preceding harder ones), MINATOLoader can be configured to disable reordering. In this mode, it behaves like PyTorch DataLoader, preserving the required sample order. While this disables MINATOLoader’s reordering advantage, it guarantees correctness in scenarios that depend on strict order.

Distributed training. While our evaluation focused on a single server, MINATOLoader’s design generalizes for distributed training with multiple nodes and GPUs. In such

scenarios, each server has a PyTorch instance accessing data from local or shared storage, with MINATOLoader integrated into each instance. Under data and model parallelism, MINATOLoader retains its preprocessing and batch construction benefits, and performs similarly to PyTorch’s DataLoader during the data loading phase, as shown in Figure 5.

Suitable workloads. MINATOLoader provides the largest benefits for workloads where preprocessing is heavy, such as image, speech, or video training pipelines. In contrast, text-based NLP workloads, including large language models, are typically not input-bound, as their preprocessing is lightweight (*e.g.*, tokenization) or performed offline. In these cases, MINATOLoader offers limited additional benefit and behaves similarly to the default PyTorch DataLoader. A key future challenge for efficient ML preprocessing lies in multimodal workloads [26], where different data modalities impose distinct preprocessing requirements, making this stage critical for training performance.

7 Related Work

Numerous systems aim to improve ML training efficiency by optimizing various stages of the data pipeline. Some focus on the input pipeline [18, 23, 35, 45–47], while others cache preprocessed data [17, 24, 34]. Some offload preprocessing to GPUs or FPGAs [4, 21, 39], and others use storage-centric solutions to accelerate data preprocessing [25, 48]. Finally, there is work on tools to characterize the performance of data preprocessing for ML [13, 14], as well as big data systems techniques that address data preprocessing [44].

Input pipeline optimizations. FastFlow [45] and tf.data [35] alleviate input data stalls by offloading preprocessing tasks to remote CPU workers in disaggregated environments. While these systems aim to improve training throughput, they treat the preprocessing step as a black box and do not analyze its internal overhead, unlike MINATOLoader. Pecan [18] is the most closely related to our work. Pecan introduces a transformation reordering policy and a policy to distribute preprocessing to worker threads in disaggregated scenarios. In contrast, MINATOLoader operates on a single machine with multiple GPUs and dynamically adapts to per-sample preprocessing variability.

Caching optimizations. Cachew [17] and CoordL [34] focus on caching preprocessed data during training to reduce input pipeline stalls. CoordL provides an analysis of the preprocessing pipeline and characterizes sources of data stalls. Unlike MINATOLoader, neither system adapts dynamically to per-sample preprocessing time variability during training.

Using accelerators for preprocessing. Meta’s Presto [25] proposes in-storage preprocessing for recommender systems. Other approaches, including DALI [4], FusionFlow [21], and TrainBox [39], offload preprocessing to specialized hardware such as GPUs or FPGAs. While this can accelerate transformations, it poses challenges when handling user-defined

functions, which are often difficult to port to these environments. In contrast, MINATOLoader performs preprocessing on the CPU and generalizes to a broad range of tasks.

Tools for bottleneck diagnostics in preprocessing. Lotus [13] and the MLPerf Storage tool [14] are dedicated profilers for identifying bottlenecks in data preprocessing pipelines and are orthogonal to this work.

Big data systems. MINATOLoader targets the preprocessing computation and batch construction phases, whereas techniques applied by traditional big data systems (e.g., data skew/straggler, dynamic repartitioning, workload estimation) target how data is organized, being complementary. Speculative execution is not applicable in ML preprocessing. Adaptive scheduling is the closest technique to MINATOLoader. However, while ML preprocessing has simpler dependencies between samples and workers, it has stricter latency requirements than adaptive scheduling. Prior work [35] also notes this difference. For example, Spark Streaming [12] recommends batch granularities of at least 50ms, while ML training requires step times below 1ms. Further, in production, Spark is used earlier in the pipeline for ETL tasks [46], with outputs written to storage and later consumed by PyTorch. This separation of concerns reflects current design choices: Spark addresses offline preprocessing, while PyTorch performs online transformations.

8 Conclusion

We presented MINATOLoader, a general-purpose data loader that accelerates machine learning training by addressing head-of-line blocking in the data preprocessing pipeline. Unlike existing solutions, MINATOLoader dynamically adapts to per-sample processing variability. Through a sample-aware load balancer and adaptive worker scheduling, MINATOLoader sustains high throughput without requiring any prior knowledge of the dataset or transformations. Our evaluation across diverse workloads and hardware platforms demonstrates that MINATOLoader consistently improves training time, achieves up to 7.5× speedup over PyTorch DataLoader and Pecan, all while preserving model accuracy.

Acknowledgments

We thank our shepherd, Muhammad Ali Gulzar, and the anonymous reviewers for their valuable feedback. This research was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) CREATE 584767-2024; a gift from MLCommons; by computational resources from FCT IP at Deucalion supercomputer, jointly funded by EuroHPC JU and Portugal; and by the European Regional Development Fund (ERDF) through the Innovation and Digital Transition Programme (COMPETE 2030) under Portugal 2030, within the scope of the project CDMS, reference 17409 (COMPETE2030-FEDER-01193000).

References

- [1] [n. d.]. KiTS19 Challenge Dataset. <https://kits19.grand-challenge.org/data/>. Accessed: [Jan 12, 2025].
- [2] [n. d.]. MLPerf Training Benchmark Suite V3.1 Results. <https://mlcommons.org/benchmarks/training/>. Accessed: [May 5, 2025].
- [3] [n. d.]. NumPy - The fundamental Package for scientific computing with Python. <https://numpy.org/>. Accessed: [May 05, 2025].
- [4] [n. d.]. NVIDIA Data Loading Library (DALI). <https://developer.nvidia.com/dali>. Accessed: [May 5, 2025].
- [5] [n. d.]. Pandas: Powerful Python Data Analysis Toolkit. <https://pypi.org/project/pandas/>. Accessed: [May 05, 2025].
- [6] [n. d.]. Scikit-learn - Machine Learning in Python. <https://scikit-learn.org/stable/>. Accessed: [May 05, 2025].
- [7] 2020. NVIDIA A100 Tensor Core GPU. <https://www.nvidia.com/en-us/data-center/a100/> Accessed: May 11, 2025.
- [8] 2020. NVIDIA Management Library (NVML). <https://developer.nvidia.com/management-library-nvml> Accessed: May 11, 2025.
- [9] 2025. Datasets & DataLoaders. https://pytorch.org/tutorials/beginner/basics/data_tutorial.html. Accessed: May 14, 2025.
- [10] 2025. dstat - Versatile Tool for Generating System Resource Metrics. <https://linux.die.net/man/1/dstat> Accessed: May 14, 2025.
- [11] 2025. Global Interpreter Lock. <https://pybind11.readthedocs.io/en/stable/advanced/misc.html>. Accessed: May 14, 2025.
- [12] 2025. Spark Streaming Programming Guide. <https://spark.apache.org/docs/latest/streaming-programming-guide.html>.
- [13] Rajveer Bachkaniwala, Harshith Lanka, Kexin Rong, and Ada Gavrilovska. 2024. Lotus: Characterization of Machine Learning Preprocessing Pipelines via Framework and Hardware Profiling. In *IEEE International Symposium on Workload Characterization*. IEEE, 30–43. <https://doi.org/10.1109/IISWC63097.2024.00013>
- [14] Oana Balmau. 2022. Characterizing I/O in Machine Learning with MLPerf Storage. *SIGMOD Rec.* 51, 3 (2022), 47–48. <https://doi.org/10.1145/3572751.3572765>
- [15] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 2016. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In *19th International Conference on Medical Image Computing and Computer-Assisted Intervention*, Vol. 9901. 424–432. https://doi.org/10.1007/978-3-319-46723-8_49
- [16] Yanjie Gao, Yichen He, Xinze Li, Bo Zhao, Haoxiang Lin, Yoyo Liang, Jing Zhong, Hongyu Zhang, Jingzhou Wang, Yonghua Zeng, Keli Gui, Jie Tong, and Mao Yang. 2024. An Empirical Study on Low GPU Utilization of Deep Learning Jobs. In *46th IEEE/ACM International Conference on Software Engineering*. ACM, 96:1–96:13. <https://doi.org/10.1145/3597503.3639232>
- [17] Dan Graur, Damien Aymon, Dan Kluser, Tanguy Albrici, Chandramohan A. Thekkath, and Ana Klimovic. 2022. Cachew: Machine Learning Input Data Processing as a Service. In *2022 USENIX Annual Technical Conference*. USENIX Association, 689–706. <https://www.usenix.org/conference/atc22/presentation/graur>
- [18] Dan Graur, Oto Mraz, Muyu Li, Mohammad Sepehr Pourghannad, Chandramohan A. Thekkath, and Ana Klimovic. 2024. Pecan: Cost-Efficient ML Data Preprocessing with Automatic Transformation Ordering and Hybrid Placement. In *2024 USENIX Annual Technical Conference*. USENIX Association, 649–665. <https://www.usenix.org/conference/atc24/presentation/graur>
- [19] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. 2022. Chasing Carbon: The Elusive Environmental Footprint of Computing. *IEEE Micro* 42, 4 (2022), 37–47. <https://doi.org/10.1109/MM.2022.3163226>
- [20] Nicholas Heller, Niranjan Sathianathan, Arveen Kalapara, Edward Walczak, Keenan Moore, Heather Kaluzniak, Joel Rosenberg, Paul Blake, Zachary Rengel, Makinna Oestreich, Joshua Dean, Michael Tradewell, Aneri Shah, Resha Tejpal, Zachary Edgerton, Matthew Peterson, Shaneabbas Raza, Subodh Regmi, Nikolaos Papanikolopoulos,

- and Christopher Weight. 2020. The KiTS19 Challenge Data: 300 Kidney Tumor Cases with Clinical Context, CT Semantic Segmentations, and Surgical Outcomes. arXiv:1904.00445 <https://arxiv.org/abs/1904.00445>
- [21] Taeyoon Kim, Chanho Park, Mansur Mukimbekov, Heelim Hong, Minseok Kim, Ze Jin, Changdae Kim, Ji-Yong Shin, and Myeongjae Jeon. 2023. FusionFlow: Accelerating Data Preparation for Machine Learning with Hybrid CPU-GPU Processing. *Proceedings of the VLDB Endowment* 17, 4 (2023), 863–876. <https://doi.org/10.14778/3636218.3636238>
- [22] Sotiris Kotsiantis, Dimitris Kanellopoulos, and P. Pintelas. 2006. Data Preprocessing for Supervised Learning. *International Journal of Computer Science* (2006).
- [23] 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, 11 (2021), 2627–2641. <https://doi.org/10.14778/3476249.3476308>
- [24] Gyewon Lee, Irene Lee, Hyeonmin Ha, Kyung-Geun Lee, Hwarim Hyun, Ahnjae Shin, and Byung-Gon Chun. 2021. Refurbish Your Training Data: Reusing Partially Augmented Samples for Faster Deep Neural Network Training. In *2021 USENIX Annual Technical Conference*. USENIX Association, 537–550. <https://www.usenix.org/conference/atc21/presentation/lee>
- [25] Yunjae Lee, Hyeeseong Kim, and Minsoo Rhu. 2024. PreSto: An In-Storage Data Preprocessing System for Training Recommendation Models. In *51st ACM/IEEE Annual International Symposium on Computer Architecture*. IEEE, 340–353. <https://doi.org/10.1109/ISCA59077.2024.00033>
- [26] Zijing Liang, Yanjie Xu, Yifan Hong, Penghui Shang, Qi Wang, Qiang Fu, and Ke Liu. 2024. A survey of multimodal large language models. In *Proceedings of the 3rd International Conference on Computer, Artificial Intelligence and Control Engineering*.
- [27] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. 8693 (2014), 740–755. https://doi.org/10.1007/978-3-319-10602-1_48
- [28] Stefan Maetschke, Ruwan Bandara Tennakoon, Christian Vecchiola, and Rahil Garnavi. 2017. nuts-flow/ml: data pre-processing for deep learning. (2017). arXiv:1708.06046 <http://arxiv.org/abs/1708.06046>
- [29] Takaki Makino, Hank Liao, Yannis M. Assael, Brendan Shillingford, Basilio Garcia, Otavio Braga, and Olivier Siohan. 2019. Recurrent Neural Network Transducer for Audio-Visual Speech Recognition. In *IEEE Automatic Speech Recognition and Understanding Workshop*. IEEE, 905–912. <https://doi.org/10.1109/ASRU46091.2019.9004036>
- [30] Francisco Massa and Ross Girshick. [n. d.]. maskrcnn-benchmark: Fast, modular reference implementation of Instance Segmentation and Object Detection algorithms in PyTorch. <https://github.com/facebookresearch/maskrcnn-benchmark>. Accessed: May 11, 2025.
- [31] Peter Mattson, Christine Cheng, Gregory Damos, Cody Coleman, Paulius Micikevicius, David Patterson, Hanlin Tang, Gu-Yeon Wei, Peter Bailis, Victor Bittorf, et al. 2020. MLPerf Training Benchmark. *Proceedings of Machine Learning and Systems* (2020).
- [32] Mark Mazumder, Colby Banbury, Xiaozhe Yao, Bojan Karlaš, William Gaviria Rojas, Sudnya Damos, Greg Damos, Lynn He, Alicia Parrish, Hannah Rose Kirk, et al. 2023. Dataperf: Benchmarks for data-centric AI development. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023*.
- [33] MLCommons. [n. d.]. MLPerf Benchmarking Suite - PyTorch implementation for image segmentation. https://github.com/mlcommons/training/tree/master/image_segmentation/pytorch. Accessed: [May 5, 2025].
- [34] Jayashree Mohan, Amar Phanishayee, Ashish Raniwala, and Vijay Chidambaram. 2021. Analyzing and Mitigating Data Stalls in DNN Training. *Proceedings of the VLDB Endowment* 14, 5 (2021), 771–784. <https://doi.org/10.14778/3446095.3446100>
- [35] Derek Gordon Murray, Jiri Simsa, Ana Klimovic, and Ihor Indyk. 2021. tf.data: A Machine Learning Data Processing Framework. *Proceedings of the VLDB Endowment* 14, 12 (2021), 2945–2958. <https://doi.org/10.14778/3476311.3476374>
- [36] Rahma Nouaji and Stella Bitchebe. 2025. *Rahm-no/MinatoLoader: MinatoLoader v1.0.1*. <https://doi.org/10.5281/zenodo.17201356>
- [37] Rahma Nouaji, Stella Bitchebe, and Oana Balmau. 2024. SpeedyLoader: Efficient Pipelining of Data Preprocessing and Machine Learning Training. In *Proceedings of the 4th Workshop on Machine Learning and Systems*. ACM, 65–72. <https://doi.org/10.1145/3642970.3655824>
- [38] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 5206–5210. <https://doi.org/10.1109/ICASSP.2015.7178964>
- [39] Pyeongsu Park, Heetaek Jeong, and Jangwoo Kim. 2020. TrainBox: An Extreme-Scale Neural Network Training Server Architecture by Systematically Balancing Operations. In *53rd Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE, 825–838. <https://doi.org/10.1109/MICRO50266.2020.00072>
- [40] G. Thippa Reddy, M. Praveen Kumar Reddy, Kuruva Lakshmana, Rajesh Kaluri, Dharmendra Singh Rajput, Gautam Srivastava, and Thar Baker. 2020. Analysis of Dimensionality Reduction Techniques on Big Data. *IEEE Access* 8 (2020), 54776–54788. <https://doi.org/10.1109/ACCESS.2020.2980942>
- [41] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen K. Paritosh, and Lora Aroyo. 2021. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. In *Conference on Human Factors in Computing Systems*. ACM, 39:1–39:15. <https://doi.org/10.1145/3411764.3445518>
- [42] Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data* 6 (2019), 60. <https://doi.org/10.1186/S40537-019-0197-0>
- [43] Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. 2022. Curriculum Learning: A Survey. *International Journal of Computer Vision* 130, 6 (2022), 1526–1565. <https://doi.org/10.1007/S11263-022-01611-X>
- [44] Apache Spark. 2019. Spark.
- [45] Taegeon Um, Byungsoo Oh, Byeongchan Seo, Minhyeok Kweun, Goeun Kim, and Woo-Yeon Lee. 2023. FastFlow: Accelerating Deep Learning Model Training with Smart Offloading of Input Data Pipeline. *Proceedings of the VLDB Endowment* 16, 5 (2023), 1086–1099. <https://doi.org/10.14778/3579075.3579083>
- [46] Mark Zhao, Niket Agarwal, Aarti Basant, Bugra Gedik, Satadru Pan, Mustafa Ozdal, Rakesh Komuravelli, Jerry Pan, Tianshu Bao, Haowei Lu, Sundaram Narayanan, Jack Langman, Kevin Wilfong, Harsha Rastogi, Carole-Jean Wu, Christos Kozyrakis, and Parik Pol. 2022. Understanding data storage and ingestion for large-scale deep recommendation model training: industrial product. In *The 49th Annual International Symposium on Computer Architecture*. ACM, 1042–1057. <https://doi.org/10.1145/3470496.3533044>
- [47] Mark Zhao, Dhruv Choudhary, Devashish Tyagi, Ajay Somani, Max Kaplan, Sung-Han Lin, Sarunya Pumma, Jongsoo Park, Aarti Basant, Niket Agarwal, Carole-Jean Wu, and Christos Kozyrakis. 2023. RecD: Deduplication for End-to-End Deep Learning Recommendation Model Training Infrastructure. arXiv:2211.05239 <https://arxiv.org/abs/2211.05239>
- [48] Mark Zhao, Satadru Pan, Niket Agarwal, Zhaoduo Wen, David Xu, Anand Natarajan, Pavan Kumar, Shiva Shankar P., Ritesh Tijoriwala, Karan Asher, Hao Wu, Aarti Basant, Daniel Ford, Delia David, Nezhil Yigitbasi, Pratap Singh, and Carole-Jean Wu. 2023. Tectonic-Shift: A Composite Storage Fabric for Large-Scale ML Training. In *2023 USENIX Annual Technical Conference*. USENIX Association, 433–449.

A Artifact Appendix

A.1 Abstract

This artifact provides a comparative evaluation of three systems - MINATOLoader, PyTorch DataLoader, and NVIDIA DALI - on the 3D-UNet workload.

A.2 Description & Requirements

A.2.1 How to access. The artifact is publicly available at: <https://github.com/Rahm-no/MinatoLoader> and <https://doi.org/10.5281/zenodo.17048007>. The repository contains all the necessary code, Dockerfile, scripts, and instructions to reproduce a minimal example from the paper.

A.2.2 Hardware dependencies. The experiments in this artifact were run using a single node machine with $8 \times$ NVIDIA Tesla V100-SXM2-32GB, 503 GB of RAM, and 446 GB SSD. If the reviewers want access to our setup, we will grant them.

A.2.3 Software dependencies. All software dependencies (CUDA 12.6, cuDNN 8.9.7, PyTorch 2.4.1, DALI, Python 3.8, system libraries) are packaged inside the Docker image. The Docker version used is 28.1.1.

A.2.4 Benchmarks. The artifact experiments use the 2019 Kidney Tumor Segmentation Challenge (KiTS19) dataset [1]. The workload is a 3D-UNet model variant derived from ML-Commons Training Image Segmentation [33]. Additional evaluations in the paper used COCO [27] (object detection) and LibriSpeech [38] (speech recognition).

A.3 Set-up

1. Clone the repository:

```
git clone https://github.com/Rahm-no/MinatoLoader.git
cd MinatoLoader
```

2. Build the Docker image (≈ 5 minutes):

```
docker build -t minato:latest .
```

3. Download the KiTS19 dataset (≈ 48 minutes, 27 GB):

```
cd raw-data-dir/kits19
pip3 install -r requirements.txt
python3 -m starter_code.get_imaging
```

4. Start the container:

```
./start_container.sh
```

5. Preprocess the dataset (≈ 13 minutes, 29 GB):

```
python3 preprocess_dataset.py \
  --data_dir /raw_data \
  --results_dir /data
```

A.4 Evaluation workflow

A.4.1 Major Claims.

- (C1): MINATOLoader reduces end-to-end training time compared to PyTorch DataLoader and NVIDIA DALI. This is validated by runtime measurements across the three systems (**§Experiment (E1)**).
- (C2): MINATOLoader improves GPU utilization compared to PyTorch DataLoader. This is demonstrated by utilization traces showing that PyTorch suffers from significant under-utilization, while MinatoLoader sustains consistently higher GPU usage (**§Experiment(E2)**).

A.4.2 Experiments. Both experiments are executed automatically when running the `run_all.sh` script, which sequentially launches each system and collects training throughput and training time, and GPU/CPU utilization statistics.

§Experiment (E1): [Training Time] [~ 10 minutes]: Measures and compares the time-to-train for PyTorch, NVIDIA DALI, and MinatoLoader on the 3D-UNet workload.

[Preparation] Build the Docker image using the provided Dockerfile and start the container with

```
./start_container.sh
```

[Execution] Run all systems sequentially (PyTorch, DALI, MinatoLoader) on 8 GPUs for 10 epochs using:

```
./scripts/run_all.sh NUM_GPUS
```

[Results] After execution, training times are automatically appended to `results/results_allsystems.csv`.

Expected runtimes on $8 \times$ V100 GPUs (10 epochs): PyTorch: ≈ 210 sec, DALI: ≈ 151 sec, and MINATOLoader: ≈ 81 sec.

These results are consistent with the paper claim (C1) that MINATOLoader provides a speedup of $2.6 \times$ compared to PyTorch and $1.9 \times$ compared to DALI. They are lower than the numbers reported in Section 5 because we used a minimal setup here (*i.e.*, 10 epochs instead of 50).

The CSV file allows direct comparison across systems. The expected outcome is that MinatoLoader achieves the lowest training time due to reduced GPU stalls, whereas PyTorch shows the worst performance due to poor utilization. You can choose to plot training time results using

```
python3 scripts/plot_figure.py
```

§Experiment (E2): [Resource Utilization Analysis] [~ 10 minutes]: Evaluates the average CPU and GPU utilization of PyTorch, NVIDIA DALI, and MinatoLoader during training of the 3D-UNet workload.

[Preparation & Execution] Same as E1. **[Results]** Once execution completes successfully, three CSV files (one per system) are generated in the results folder. Each file contains the average resource utilization over time for the corresponding system. You can plot the results to visualize differences in utilization between the three data loaders by running:

```
python3 scripts/plot_usage.py
```

The plots show that NVIDIA DALI consistently achieves high GPU utilization because preprocessing is also performed on the GPU. By contrast, PyTorch exhibits frequent GPU idle

periods accompanied by CPU peaks, indicating that while the CPU is busy with preprocessing, the GPU is left waiting. For MINATOLoader, the GPU remains consistently and highly utilized throughout training, demonstrating efficient overlap of preprocessing and computation.