RecD: Deduplication for End-to-End Deep Learning Recommendation Model Training Infrastructure

Mark Zhao [1] **Dhruv Choudhary** [2] **Devashish Tyagi** [2] **Ajay Somani** [2] **Max Kaplan** [2] **Sung-Han Lin** [2]
Sarunya Pumma [2] **Jongsoo Park** [2] **Aarti Basant** [2] **Niket Agarwal** [3] **Carole-Jean Wu** [2] **Christos Kozyrakis** [1]

Abstract

We present *RecD* (*Rec*ommendation *D*eduplication), a suite of end-to-end infrastructure optimizations across the Deep Learning Recommendation Model (DLRM) training pipeline. *RecD* addresses immense storage, preprocessing, and training overheads caused by feature duplication inherent in industry-scale DLRM training datasets. Feature duplication arises because DLRM datasets are generated from interactions. While each user *session* can generate multiple training samples, many features' values do not change across these samples. We demonstrate how *RecD* exploits this property, end-to-end, across a deployed training pipeline. *RecD* optimizes data generation pipelines to decrease dataset storage and preprocessing resource demands and to maximize duplication within a training batch. *RecD* introduces a new tensor format, *InverseKeyedJaggedTensors* (IKJTs), to deduplicate feature values in each batch. We show how DLRM model architectures can leverage IKJTs to drastically increase training throughput. *RecD* improves the training and preprocessing throughput and storage efficiency by up to $2.48\times$, $1.79\times$, and $3.71\times$, respectively, in an industry-scale DLRM training system.

1 Introduction

Machine learning (ML) infrastructure is one of the most dominant components of industry-scale datacenters. For example, ML consumes over 70% of FLOPs at Google (Patterson et al., 2022). To support the computational demands of ML, and especially training, companies such as Google (Lardinois, 2022), Meta (Mudigere et al., 2022; Meta, 2022), and AWS (AWS, 2022) are deploying massive clusters consisting of tens of thousands of accelerators.

Deep learning recommendation model (DLRM) training is a principal industrial use-case for these clusters. For example, DLRM training dominates ML capacity across Meta's fleet (Naumov et al., 2020). This demand is driven by the ubiquity of DLRMs across industry, as they underlie critical services from Google (Anil et al., 2022; Li et al., 2020; Zhao et al., 2019), Taobao (Ge et al., 2018), Meta (Meta, 2019; Hazelwood et al., 2018; Acun et al., 2021), and others. DLRM training clusters are fed by a data storage and ingestion (DSI) pipeline — systems that generate, store, and preprocess training data — which can demand more power consumption than what is required by the training accelerators (trainers) themselves (Zhao et al., 2022). To

train larger, more complex, and more accurate models, it is critical to improve the performance and efficiency of the *end-to-end DLRM training pipeline*, from DSI to trainers.

To this end, this paper presents a suite of optimizations, called *RecD*, spanning the DLRM training pipeline. *RecD* exploits the inherent *session-centric* nature of DLRM datasets. DLRM training samples are generated from user interactions which query industrial recommendation models. Each user's *session* typically requires numerous inferences, and thus produces many training samples (Wang et al., 2021). However, the *features* that largely compose each sample likely remain static throughout each session. For example, an e-commerce DLRM dataset may contain a user feature representing the sequence of the last N items added to a shopper's cart. The e-commerce site may serve recommendations throughout a user's shopping session, but if the shopper does not add a new item, each of the session's samples will contain *the same values* for that feature.

While prior work has mentioned feature duplication (Ge et al., 2018; Gai et al., 2017), none has characterized its prevalence in industry-scale datasets nor provided solutions that optimize for it across the training pipeline. Current pipelines spend considerable resources storing, preprocessing, and training over duplicate features. These overheads constrain industry-scale training infrastructures from supporting larger datasets, longer features, and more complex modeling techniques (e.g., attention) that yield more accurate models (Ardalani et al., 2022; de Souza Pereira Moreira

---

[1]Stanford University, Stanford, California, USA [2]Meta, Menlo Park, California, USA [3]NVIDIA, Santa Clara, California, USA. Correspondence to: Mark Zhao <myzhao@cs.stanford.edu>.
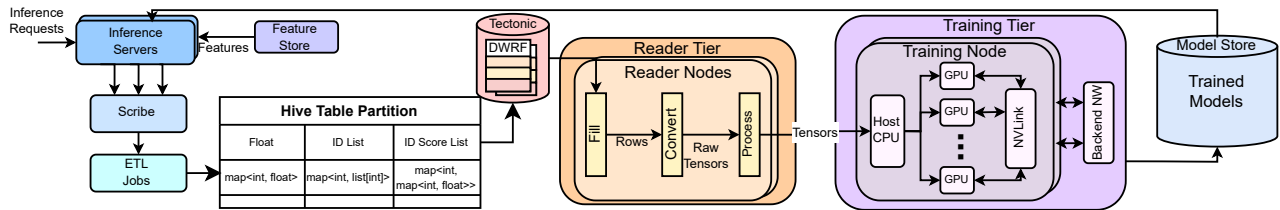
Figure 1. Industrial DLRM training pipeline composed of DSI and training services.

et al., 2021; Fang et al., 2020; Li et al., 2019).

We begin the paper with an in-depth characterization of how session-centricity generates significant feature value duplication in datasets used by an industrial DLRM training pipeline. Each session produces many samples (16.5 on average), and feature values are largely duplicated across the session's samples (81.6 − 89.4% on average). *RecD* addresses the significant storage, preprocessing, and training overheads caused by duplication throughout the DLRM training pipeline.

RecD begins at data generation by sharding raw inference logs by session ID to improve compression ratios in Scribe (Karpathiotakis et al., 2019), a distributed message passing system. These logs are ingested by ETL engines to produce training samples. *RecD* coalesces each session's samples within a training batch. Not only does this reduce dataset sizes due to native compression, it also allows *RecD* to convert each batch to a new tensor format during data reading, InverseKeyedJaggedTensors (IKJTs), that deduplicates feature values.

IKJTs require minimal resource overheads to generate and use for preprocessing and training. Meanwhile, they allow readers, which preprocess data, and trainers to operate on deduplicated tensors, significantly reducing resource demands across the training pipeline. We explore these benefits. We present how DLRM architectures can leverage IKJTs to reduce GPU compute, network, and memory resource requirements — improving training throughput and enabling more powerful modeling techniques. In summary:

- We provide a characterization using petabyte-scale industrial DLRM datasets showing how feature duplication is inherent in DLRM training pipelines. We discuss the opportunities and challenges of deduplication.

- We present necessary optimizations made in the data storage and ingestion pipeline to enable a novel tensor format, IKJTs, that deduplicates features in each training batch.

- We show how IKJTs improve DLRM training throughput and resource utilization by eliminating redundant compute, memory, and network usage during training.

- We evaluate on industrial DLRMs. *RecD* improves training and preprocessing throughput and storage efficiency

by up to 2.48×, 1.79×, and 3.71×, respectively.

2 BACKGROUND

Figure 1 shows an end-to-end industrial training pipeline (Zhao et al., 2022), with DSI and training services.

2.1 Data Storage and Ingestion

Data Generation. Training data is continuously generated from deployed recommendation services. User-facing services request batches of inferences throughout a user's *session*. For each batch of requests, features corresponding to the user and potentially recommended items are retrieved from a feature store and are used as input to the DLRM to generate relevant predictions. Since features continuously change, inference servers log features for each request to avoid data leakage (Kaufman et al., 2012). Given predictions, user-facing services generate relevant impressions of items and log events (i.e., impression outcomes). Logs are aggregated in Scribe, a global distributed messaging system (Karpathiotakis et al., 2019).

Streaming and batch processing engines, such as Spark (Zaharia et al., 2012), ingest data from Scribe. These engines join raw features and events to produce labeled samples. Training samples are subsequently landed into time partitioned (e.g., hourly) Hive tables (Thusoo et al., 2009). To maintain data freshness, new table partitions are constantly landed and old partitions are deleted.

Dataset Schema and Storage. Each training sample, corresponding to an impression and outcome, is stored as a structured row containing *features* and *labels*. Features represent almost all of the bytes within a sample. DLRMs require two types of features: dense and sparse. Dense features represent continuous values, such as time, and are stored as a map from feature key to a float value. Sparse features represent categorical values, such as item IDs, and are stored in map columns that map a feature key to its value, typically a variable-length list of item IDs. Compared to dense features, sparse features require significantly more storage, preprocessing, and training resources across the DLRM training pipeline (Zhao et al., 2022; Naumov et al., 2020; Sethi et al., 2022).

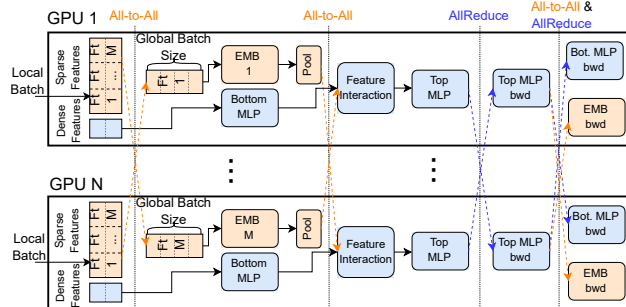Hive partitions are stored as columnar DWRF (Zhao et al.,

Figure 2. How DLRMs are synchronously trained using hybrid parallelism across multiple GPUs. Dashed lines show collective communication across model-parallel and data-parallel modules.

2022) files similar in format to ORC (ORC, 2022). Files are composed of regions, called *stripes*, that represent a small set of rows. Within a stripe, rows are stored as columnar streams. Feature columns are first flattened (i.e., each feature key becomes a separate column). Values for each flattened column (e.g., ID lists) are then encoded and compressed into streams. Files are stored in Tectonic, an exabyte-scale distributed filesystem (Pan et al., 2021).

Data Reading and Preprocessing. Each DLRM training job specifies its *dataset* (table partitions) and *preprocessing* needs via a DataLoader specification (PyTorch, 2023). A *reader tier*, composed of stateless *readers*, is launched for each job. Readers scan through the dataset partitions and preprocess rows into tensors. Each reader *fills* a batch of rows by reading files from Tectonic and *converts* the batch into raw tensors. The reader then *processes* the tensors according to the job's specifications by applying transformations such as normalization and hashing. Preprocessed tensors are sent to trainers. The number of readers for each job is scaled to meet trainers' ingestion bandwidth demands.

2.2 DLRM Training at Scale

Figure 2 shows how a typical DLRM (Naumov et al., 2019) is synchronously trained across multiple GPUs. DLRMs primarily consist of multilayer perceptrons (MLPs) and embedding tables (EMBs) composed into three main architectural components. EMBs ingest sparse feature lists and produce a dense activation vector for each list element. A pooling function (e.g., average, sum, or max) aggregates activations for each sparse feature. Meanwhile, a bottom MLP transforms dense features into a dense representation with the same dimensionality as embedding vectors. An interaction layer explicitly computes second-order interactions across dense and sparse features (e.g., via pairwise dot product). A top MLP and sigmoid processes the result to produce a probability output (e.g., click-through rate).

DLRMs are trained using hybrid parallelism across multi-

ple GPUs. MLPs are copied across GPUs in a distributed data parallel (DDP) fashion, while EMBs are sharded across GPUs via distributed model parallelism (DMP) due to their large size. During each training iteration, each GPU ingests a local batch from the reader tier. A sparse data distribution (SDD) step first aggregates the appropriate feature values, across all local batches, to the corresponding GPU using an all-to-all collective (NVIDIA, 2022) across all GPUs. After the EMB lookup and pooling, another all-to-all distributes embedding vectors back to their original GPUs, as feature interaction and the top MLP is data parallel. After calculating the loss, an all-reduce aggregates gradients to update MLPs. Similarly, an all-to-all synchronizes EMB parameter updates during the backward pass. Thus, the iteration time is determined by both the per-GPU compute and memory bandwidth resources (for MLPs, interactions, pooling, and EMB lookups), as well as the backend network bandwidth and latency (for collective communications).

Scaling Systems for DLRMs. Improving DLRM accuracy necessitates systems efficiency and throughput optimizations across the training pipeline. For example, (Ardalani et al., 2022) showed that data scaling significantly improves DLRM performance. Supporting growing dataset volumes requires not only more efficient storage, but also improved reader and trainer throughputs to complete training within a reasonable amount of time. Meanwhile, recent DLRM architectures focus on capturing users' long-term interests via a sequential history of interactions (Pi et al., 2019; Li et al., 2019; Chen et al., 2019). These architectures use long *sequence features* and attention mechanisms, such as transformers (Vaswani et al., 2017), to pool embeddings across many sequence features. They demand significant GPU compute, memory, and network resources. Thus, optimizing for DSI and training performance and efficiency is increasingly urgent as resource demands continue to grow.

3 UNDERSTANDING DATA REUSE

To understand the opportunity for *RecD*, we explore the prevalence of duplication within industry-scale datasets. We focus on sparse features because they *a)* are prone to duplication as we characterize next, and *b)* demand significantly more training pipeline resources than dense features (Section 2) and thus present a more attractive optimization target.

Duplication arises because sparse user features rarely change across impressions within a session[1]. For example, consider social media features f_{like} and f_{share}, which contain a user's last N posts they liked and shared, respectively. During a session, a user may view multiple posts (impressions). While each impression may generate a training sample, f_{like} and f_{share} will be exactly the same across

---

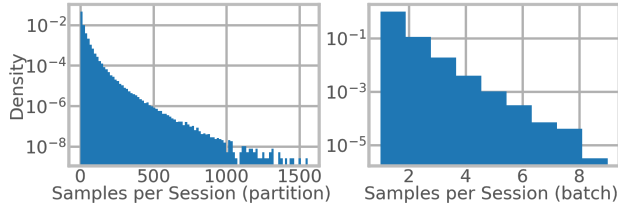[1]A session is a set of user impressions in a fixed time window.

Figure 3. Histogram of the number of samples per session within an hourly partition (left) and 4096 batch (right) from the partition.
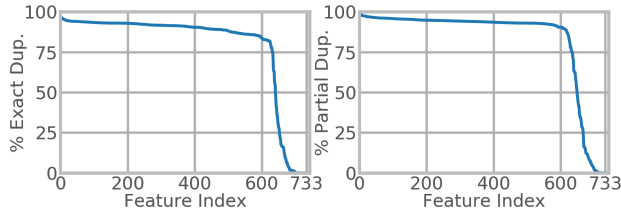

Figure 4. Percent of exact (left) and partial (right) duplicate values across sparse features within an hourly partition.

samples if the user did not like or share a post. Some features, such as f_{share}, may rarely change, and even if they do, their lists will be shifted with most elements being the same. In summary, duplication depends on *the number of samples per session*, and is a per-feature property depending on *how often each feature's value changes across samples*.

Each session generates multiple training samples, but they are distributed across training batches. Figure 3 shows a histogram of the number of training samples generated by each session in a representative, O(100PB)-scale industrial DLRM dataset. We show graphs for an hourly table partition and a batch of 4096 samples from the partition.

We observe that, on average, each session generates 16.5 samples within an hourly partition with a significant tail of over 1000 samples per session. While this shows that there is significant opportunity for deduplication, the training pipeline operates on a relatively small subset of training samples at a time (e.g., file splits and training batches). Deduplicating a session's samples requires us to co-locate the samples closely together within the table.

However, the data generation infrastructure typically orders samples based on inference time. The large volume of inference requests across services naturally interleaves samples across sessions. Figure 3 demonstrates how within each batch of 4096 samples, this interleaving results in only 1.15 samples per session on average. Optimizations must be co-designed alongside data generation infrastructure to fully coalesce and deduplicate samples within each session.

Within a session's samples, there is a large amount of exact and partial feature duplication. It is also important to validate that feature values (i.e., lists) themselves are largely duplicated across these samples. We first quantify

the amount of exact duplicate feature values across samples. Specifically, for given feature x, we analyze the percent of samples within the hourly partition that contained exactly the same list as another sample from the same session within the partition. For example, if feature x is never updated across 16.5 samples per session, we would expect a maximum of $15.5/16.5 = 93.9\%$ of exact duplicates in the partition. Figure 4 shows this result across 733 sparse features in the hourly partition. We observe that on average across all features, 80.0% of feature values are an exact duplicate. This validates our assumption that many DLRM features are not updated across a session's samples.

Specifically, DLRM sparse features largely reflect either user or item traits. User sparse features (e.g., last N liked item IDs) are largely duplicated across a user's samples. Item features (e.g., the item ID that is evaluated for recommendation) are less duplicated since many different items are ranked in a given session. Figure 4 shows this distinction – user features comprise the vast majority of dataset volume and accordingly represent the large subset of features with high duplication. Meanwhile, item features exhibit less duplication, representing the subset of features right of the knee. We expect increased reliance on user features, and thus higher feature duplication, since recommender systems are increasingly focusing on larger user interaction history features compared to item metadata (Section 2.2).

For highly-duplicated user features, even if its values change across samples, we expect the majority of its list IDs to remain the same. We thus repeated the analysis on an individual list ID basis. For example, suppose feature x contained 100 IDs across 2 training samples. x may be updated by appending a new ID and shifting its list by one, resulting in $99/200 = 49.5\%$ partial duplication. Figure 4 shows how on average across all features, 83.9% of feature values within each feature list are duplicated. Many non-exact duplicate samples within the session contain partial duplicates.

Finally, it is important to note that *not all feature lists have the same length*. To understand how many bytes are duplicated, we weigh each feature in our prior analysis by its respective average length. We find that 81.6% and 89.4% of all IDs in feature values (i.e., bytes) are exact and partial duplicates, respectively, suggesting that longer features have slightly more exact and partial duplicates.

Summary. The vast majority of feature values are duplicated within the industry-scale DLRM dataset. While significant deduplication opportunities exist, they require each session's samples to be co-located within training batches. Thus, trainer-only solutions are insufficient — optimizations must be co-designed across the end-to-end training pipeline.

Table 1. Overview of *RecD* optimizations made throughout the industry-scale training pipeline.

Optimization	Target System	Benefit
O1: Log Sharding (§4.1)	Scribe	Improves black-box compression ratios to reduce Scribe network RX/TX and storage demands.
O2: Cluster by Session (§4.1)	ETL	Session sample co-location enables readers/trainers to exploit duplicate features. Improves file compression ratios, reducing storage and read IOPS demands.
O3: Inverse KJTs (§4.2)	Readers	New tensor encoding allows downstream preprocessing/training operations to use deduplicated features, enabling significant resource savings.
O4: Deduplicated Preproc. (§4.3)	Readers	IKJT preprocessing modules reduce preprocessing compute demands. Deduplicated outputs require less NW bandwidth between readers and trainers.
O5: Deduplicated EMB (§5)	Trainers	Reduced per-iteration trainer compute/memory/NW demands by deduplicating EMB features, lookups, and activations.
O6: JaggedIndex-Select (§5)	Trainers	Reduced memory copy overheads by enabling index select without first converting jagged tensors to a dense representation.
O7: Deduplicated Compute (§5)	Trainers	Reduced compute for sparse feature modules (especially attention pooling) by allowing them to operate on deduplicated tensors.

4 *RecD* IN DATA STORAGE AND INGESTION

RecD implements these optimizations, summarized in Table 1, throughout the industrial pipeline shown in Figure 1.

4.1 Data Generation and Storage

Log Sharding. Scribe is a message passing service which logically aggregates and buffers raw logs from each inference server. To load balance, Scribe consistently hashes the message and routes each to a *shard* on a physical storage node, which buffers and compresses messages in memory and on disk. Unfortunately, the default hashing configuration distributes logs for each session randomly across shards. *RecD* configures Scribe to instead *use session IDs as the shard key*, improving the "compressibility" of data within each shard. Thus, we can both reduce the number of Scribe storage nodes and the amount of network bandwidth needed for downstream ETL jobs to ingest logs.

Clustering by Session. While improved sharding increases the locality of a session's logs, it does not guarantee that a session's training samples are adjacent within the dataset. This grouping is needed for downstream systems to deduplicate features. Thus, *RecD* adds a data generation ETL job, which clusters partitions by session ID and sorts by log timestamp. As with Scribe, we also expect two direct benefits from ETL clustering. First, each file's stripes are compressed using black-box compression, e.g. zstd (Zstandard, 2022). Ensuring that each stripe contains multiple rows for a given session increases compression ratios and thus reduces dataset storage requirements. Secondly, smaller files also reduce compute and network resources needed to read samples during online preprocessing.

4.2 Tensor Encoding for Deduplication

Figure 5 shows how reader nodes generate preprocessed tensors for each training job. Each reader reads batches of
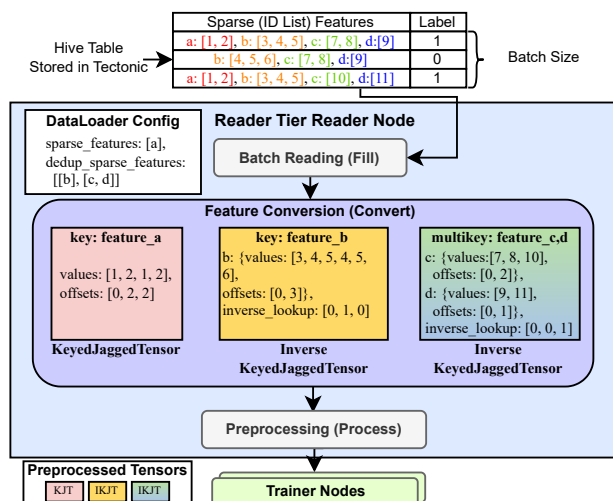

Figure 5. Each reader node extracts mini-batches from storage. Batches are converted to tensors, preprocessed, and sent to trainers.

samples, converts rows to tensors, and preprocesses tensors.

A *Feature Conversion* step copies data from raw batches of rows, read into memory as byte arrays, into structured tensors. The typical tensor format used for sparse features is a KeyedJaggedTensor (KJT) (PyTorch, 2022). A KJT maps a key (i.e., the feature key) to a JaggedTensor — a tensor with a jagged dimension (i.e., different length slices). For example, Figure 5 shows how a batch of 3 rows for feature a is transformed into a KJT with two slices representing the feature's $values$ and $offsets$ for the batch. The $offsets$ slice has an entry for each row, with $offsets[i]$ pointing to the starting index in the $values$ slice for row i. The length of the feature for row i is calculated from $offsets[i + 1] - offsets[i]$ (or $|values| - offsets[i]$ for the last row). In the example, feature a has $2\times$ duplication in the $values$ slice as rows 0 and 2 both contain $[1, 2]$.

InverseKeyedJaggedTensor. To deduplicate feature values, *RecD* introduces a new *inverse* KJT format. Figure 5 shows how ML engineers can specify a `dedup_sparse_features` field in the PyTorch DataLoader, which is a `List[List[featureKey]]`, containing lists of *feature groups* to deduplicate. *RecD* deduplicates each feature group to an InverseKeyedJaggedTensor (IKJT) during feature conversion. Of course, users can still generate KJTs for features not exhibiting high duplication.

RecD deduplicates features by detecting and avoiding duplicate copies during feature conversion. An IKJT instead adds an additional *inverse_lookup* slice, where *inverse_lookup*[i] points to the respective entry in the deduplicated *offsets* slice for row i in the batch. *offsets* encodes the *values* slice as before. In our example, feature b contains duplicate values for rows 0 and 2. Thus *inverse_lookup*[0] == *inverse_lookup*[2], with both pointing to *offset*[0] which encodes the duplicate values [3, 4, 5]. IKJTs avoid storing a second copy of [3, 4, 5] in the *values* slice. Since exact matches are the vast majority of duplication (Section 3), we focus on deduplicating exact matches and discuss supporting partial matches in Section 7.

Grouped IKJTs. Users can deduplicate multiple features within a single IKJT. Grouped IKJTs are designed for features updated synchronously across samples and thus share *inverse_lookup* values. For example, an e-commerce model may use two features which track the item ID and seller ID for items added to a user's cart. Since both features track the same item sequence, they are both updated at the same time (i.e., when a new item is added). As we explore in Section 5, grouped IKJTs are designed to enable additional optimizations during each training iteration.

In Figure 5, features c and d are deduplicated as a group. For both features, rows 0 and 1 are duplicates. Thus, *RecD* uses a common *inverse_lookup* to reference the *offsets* slice for both features, even if their respective offsets or values slices are different. For example, *inverse_lookup*[0] will map to [7, 8] for feature c and [9] for feature d. In the event that grouped feature values are not synchronously updated across samples, we will not deduplicate the corresponding unsynchronized rows to ensure that the *inverse_lookup* invariant is maintained.

Using IKJTs. Not all features may be worth deduplicating. To understand the value of deduplication, we use the following analytical model for a feature f. S is the average number of samples per session. B is the batch size. $d(f)$ is the probability that the f's value will remain the same across adjacent rows. $l(f)$ is the average length of f.

$$DedupeLen(f) = l(f) * B * (1 - (S - 1) * S^{-1} * d(f))$$
$$DedupeFactor(f) = l(f) * B / DedupLen(f)$$

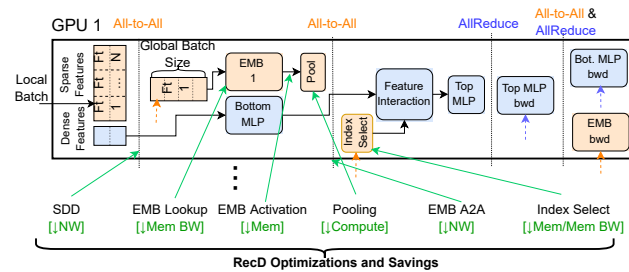Specifically, $DedupeLen(f)$ expresses the size of the

Figure 6. Overview of a training iteration, with green arrows showing where *RecD* optimizations save trainer resources. Dashed arrows show all-to-all and all-reduce communication.

values slice after deduplicating f for each training batch. The deduplication factor, $DedupeFactor(f)$, is calculated as the ratio of the original *values* slice length to $DedupeLen(f)$. For example, suppose $B = S = 3$, $l(b) = 3$, and $d(b) = 0.5$ for feature b in Figure 5. Deduplicating b results in $DedupeLen(b) = 6$ and $DedupeFactor(b) = 1.5$. The total amount of feature values deduplicated increases with higher S, $l(f)$, and $d(f)$, which aligns (i.e., increases) with data scaling trends (Section 2.2).

While *inverse_lookup* and *offsets* requires more elements than *offsets* alone (up to B), the overhead is negligible as for most features $l(f) * B >> B$. Furthermore, as we discuss in Section 5, because only *values* and *offsets* tensors are communicated across GPUs, IKJTs strictly decrease over-the-network tensor sizes during training. Finally, $DedupeFactor(f)$ provides an initial guidance on the impact of deduplicating f. We typically deduplicate features with $DedupeFactor(f) > 1.5$. However, the actual performance benefit depends on how well readers and trainers can use IKJTs, as we explore next and discuss in Section 7.

4.3 Preprocessing over IKJTs.

After feature conversion, each reader node preprocesses tensors using a set of user-provided TorchScript modules. If a user deduplicates a feature, we automatically add a wrapper that transparently supports preprocessing over IKJTs. Since the original function used KJTs, the wrapper simply provides the *offsets* and *values* slices from the IKJT held in memory, saving significant compute resources by avoiding preprocessing duplicate values. Deduplicated preprocessing functions also output IKJTs. This reduces network bandwidth requirements between reader and trainer nodes and allows trainers to further leverage IKJTs.

5 *RecD* IN TRAINING

Building on our newly proposed IKJT tensor format, we design a series of *RecD* PyTorch modules as direct replacements for DLRM embedding and pooling operations. The

IKJT format generated by readers enables a host of optimizations at the trainer, summarized in Table 1. As shown in Figure 6, these modules operate on *deduplicated* (i.e., IKJT) tensors during the forward pass, reducing resources spent operating over duplicate sparse feature values.

Sparse Data Distribution. After receiving a batch of samples, each GPU executes a sparse data distribution (SDD) step. Using an all-to-all collective, SDD coalesces a global batch only containing the respective features corresponding to each GPU's model-parallel EMBs. Previously, KJTs required sending significant amounts of duplicate feature values over the network. With *RecD*, deduplicated IKJT *value* and *offset* slices are sent instead (*inverse_lookup* slices are kept local). *RecD* thus reduces the amount of bytes distributed during SDD by $DedupeFactor(f)$ for each feature f. Since SDD runs before any embedding lookups, reducing the amount of data over the network in each iteration directly improves training throughput.

EMB Lookups. After SDD, each trainer needs to translate every feature value in the KJT into an embedding by performing a lookup in each EMB. By using IKJTs, the length of the *values* slice is reduced by $DedupeFactor(f)$, reducing the overall number of EMB lookups we need to perform in each iteration and thus required memory bandwidth.

EMB Inputs and Activations. Each GPU also needs to allocate significant dynamic memory to store the feature inputs and EMB activations of each sparse value. This is especially true of long length sequence models. For example, a single feature f with $l(f) = 1000$, $B = 4096$, and an EMB dimension of 128 would require $4096 * 1000 * 128 * 4B \approx 2GB$ of GPU memory to store activations. By performing lookups using IKJT values, we directly reduce the amount of dynamic GPU memory required by $DedupeFactor(f)$.

Deduplicated Pooling. DLRMs use a set of pooling modules (e.g., sum, avg.) that operate on the EMB activations prior to feature interaction. Recent trends have motivated more complex pooling modules, such as transformers and other attention mechanisms (Pi et al., 2019; Li et al., 2019), which operate over multiple long-length sequence features. These modules require significant GPU resources.

To reduce the computational and memory overhead for these sequential pooling modules, *RecD* allows users to run compute modules with IKJTs as inputs. Specifically, by ensuring that the *inverse_lookup* slice is shared across all features within an IKJT, we can deduplicate compute by simply operating on the deduplicated *values* and *offsets*. For example, assume a module element-wise sums values for each row across features c and d in the example in Figure 5. Using KJTs, the GPU computes $[7 + 8 + 9, 7 + 8 + 9, 10 + 11] = [24, 24, 21]$. With IKJTs, we instead compute $[7 + 8 + 9, 10 + 11] = [24, 21]$ and

simply use the shared *inverse_lookup* to expand the output to $[24, 24, 21]$. By applying this technique to expensive attention pooling modules, we reduce the compute demand by $DedupeFactor(f)$ for each sequence feature f.

Deduplicated EMB. Since the output of pooling layers are still in the IKJT format, we can get more network savings during the all-to-all that broadcasts pooled embeddings back to each GPU for feature interaction.

Jagged Index Select. Before feature interaction, IKJTs must be converted back to a KJT to be interacted with other non-deduplicated features. We use `torch.index_select` to perform this conversion. Prior to *RecD*, `index_select` could only operate on dense, not jagged tensors. We needed to first convert jagged tensors into dense tensors (e.g., via padding), incurring large memory overheads. We implemented a jagged `index_select` to operate over jagged tensors, eliminating this overhead.

Summary. As summarized in Figure 6, IKJTs enable a host of GPU network, memory, memory bandwidth, and compute optimizations during training. These optimizations improve training throughput and reduce GPU resource demands, allowing us to train more complex models at a faster rate.

6 EVALUATION

6.1 End-to-end Performance Optimizations

RecD improves the performance and efficiency of the entire training pipeline, including storage, readers, and trainers. To study each component, we used three representative industrial DLRMs, RM_1, RM_2, and RM_3, designed around the core DLRM architecture (Naumov et al., 2019). RM_1, RM_2, and RM_3 contain $O(10^9)$, $O(100^9)$, and $O(100^9)$ parameters with $O(10GB)$, $O(100GB)$, and $O(100GB)$ of embedding tables, respectively. Embedding dimensions range from 64-1024 across each RM.

We evaluate on a trainer tier consisting of ZionEX training nodes (Mudigere et al., 2022). Each ZionEX node contains 8 NVIDIA A100 GPUs with a total of 320 GB HBM and 12.4 TB/s of memory bandwidth. Intra-node communication across GPUs occurs via NVLink. Each GPU is equipped with a 200 Gbps RoCE NIC for inter-node communication over a dedicated RoCE backend network. Input data is supplied by 4 host CPU sockets, each with a 100 Gbps NIC that ingests data from a tier of DPP (Zhao et al., 2022) readers. Each reader is a general-purpose, x86 CPU server with 18 cores, a 12.5 Gbps NIC, and 64 GB of memory. Readers read data from each RM's respective $O(100PB)$ industrial dataset stored within the Tectonic file system (Pan et al., 2021).

For each RM, we used the default baseline configuration, with RM_1, RM_2, and RM_3 using a batch size of 2048,
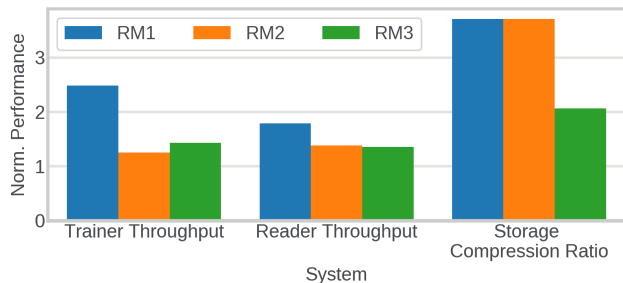
Figure 7. Trainer, reader, and storage performance using *RecD* across three representative RMs, normalized to their baselines.

Figure 8. Breakdown of trainer iteration latency across RMs, normalized to each RM's baseline latency using the same batch size.

2048, and 1152, and 48, 48, and 64 GPUs, respectively. We then enabled the full suite of *RecD* optimizations for each RM, which allowed us to increase the batch size for RM_1 and RM_3 to 6144 and 2048, respectively. For RM_2, we could not substantially increase batch size beyond 2048. For RM_1, we deduplicated 16 sequence features in 5 groups. RM_2 and RM_3 deduplicated 6 and 11 sequence features, respectively, in one group. Each RM also deduplicated an additional ≈ 100 features that were element-wise (e.g., sum, max) pooled. $DedupeFactor$ was $\approx 4 - 15$ for deduplicated features. We used a clustered table for *RecD* models containing the same data as the baseline table. We kept all other hyper parameters the same and scaled the number of readers to provide sufficient throughput to avoid data stalls in all configurations.

Figure 7 shows how trainer throughput, reader throughput, and storage compression ratio improved with respect to the baseline for each RM. Trainer throughput is the total samples per second processed by all trainers. Since we scale the number of readers based on trainer throughput, we report the samples per second processed on average by each reader. Finally, we report the compression ratio of the clustered table's Tectonic files relative to the baseline table. *RecD* improved trainer throughput by $2.48\times$, $1.25\times$, and $1.43\times$, significantly decreasing training job latencies. Similarly, each reader processed samples $1.79\times$, $1.38\times$, and $1.36\times$ faster, reducing the number of readers needed for each training job by the same amount. We explore in Section 6.2 and 6.3 why RM_1's increased use of sequence features allowed *RecD* to further increase trainer and reader throughput, respectively, compared to RM_2 and RM_3. Clustered tables improved the compression ratio by $3.71\times$ (RM_1 and RM_2 used the same table) and $2.06\times$, directly improving storage efficiency by reducing the number of storage nodes needed to store and serve each RM's dataset. RM_1 and RM_2's table exhibited higher samples per session than RM_3's table, leading to a larger increase in compression after co-locating each session's samples within a file stripe. Finally, we increased the compression ratio at Scribe from $1.50\times$ to $2.25\times$ by sharding logs by session ID.
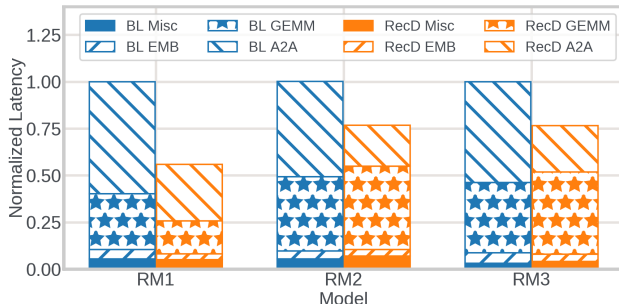
6.2 Why does *RecD* improve trainer throughput?

Iteration Breakdown. To understand where *RecD* improves training throughput, we ran a training job for each RM *using the same batch size as the baseline*. Figure 8 shows a breakdown of the iteration latency of the *RecD* training job normalized to the baseline iteration latency, averaged across all GPUs. Specifically, we show *exposed* latency (i.e., non-overlapping compute/communication), broken down by GPU time spent on EMB lookups, compute (GEMM), all-to-all communication (A2A), and other miscellaneous operations (e.g., all-reduce and reduce-scatter).

First, we observe that *RecD* halves exposed A2A communication across all RMs. A2A is a significant component of each training iteration. *RecD* significantly improves training throughput by reducing the amount of over-the-network bytes via IKJTs. Thus one reason for RM_1's larger training throughput gains is because it exposes more communication by using more sequence features that *RecD* optimizes for.

The second reason is because RM_1 uses expensive transformers to pool EMB activations for several user sequence features; *RecD* deduplicated the compute for these transformers by grouping each transformer's features together using IKJTs. This is evidenced by an additional reduction (12% of iteration latency) in the amount of time spent in GEMMs for RM_1. Meanwhile, RM_2 and RM_3 saw slight increases in exposed $GEMM$ time. This is because less of it was hidden as *RecD* reduced A2A latencies, as well as a slight increase due to the additional `index_select`. We also observe a small improvement across RMs due to faster EMB lookups ($1 - 2\%$ of iteration latency) by eliminating redundant lookups. While this did not significantly improve trainer throughput given the same batch size, using fewer EMB activations allowed us to increase batch size to improve trainer throughput, as we study next.

Finally, Figure 8 shows why translating $DedupeFactor$ to throughput gains is challenging. RM_1 and RM_2 used the same table and features with similar $DedupeFactor$s.
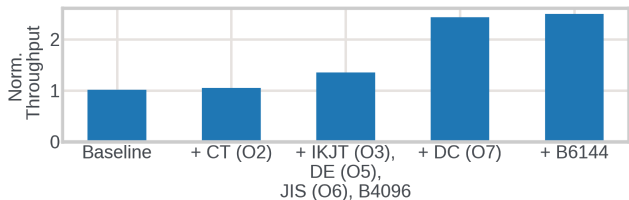
Figure 9. Ablation study showing normalized trainer throughput for RM_1. O=Table 1 Optimization (O1 & O4 not applicable). CT=Clustered Table, DE=Dedupe. EMB, JIS=Jagged Index Select, DC=Dedupe. Compute, B=Batch Size.

Table 2. Breakdown of trainer throughput (QPS) and efficiency for RM_1 with *RecD*, enabling larger batches and complex models.

Config.	Norm. QPS	Max Mem. Util.	Avg. Mem. Util.	Norm. Comp. Efficiency (flop/s/GPU)
Baseline	1.00	99.90	72.83	1.00
RecD	1.89	27.76	22.20	1.73
RecD + EMB D256	1.55	40.87	31.17	1.92
RecD + B6144	2.26	91.78	51.55	2.12

However, *RecD* reduced RM_1's iteration time by 44% compared to 23% for RM_2 due to the differences in model architectures and exposed compute/communication cycles as discussed above. We discuss observations on how ML engineers choose which features to deduplicate in Section 7.

Ablation Study. To understand how specific optimizations contribute to training throughput, we performed an ablation study using RM_1 as shown in Figure 9. First, simply using clustered tables provides no training throughput benefit. While clustering is necessary for *RecD*, it is not sufficient alone since KJTs still contain duplicate feature values. By using IKJTs (and jagged index select) to deduplicate EMB lookups and activations, we could increase batch size to 4096 and realized a $1.34\times$ gain to training throughput. We then used multiple IKJT groups, which further allowed us to deduplicate the compute required by expensive transformers, which led to a $2.42\times$ increase in throughput. Finally, this further allowed us to increase batch size to 6144, which resulted in a final $2.48\times$ increase in training throughput.

Trainer Resource Utilization. Because *RecD* reduces GPU resource requirements, we can also tune model hyper parameters in order to further improve model throughput and accuracy. To illustrate this, Table 2 shows the throughput, memory utilization, and GPU compute efficiency for RM_1 as we enabled *RecD*. Using a baseline batch size of 2048 required the entirety of GPU memory. *RecD* reduced the maximum and average memory utilization from 99.9%

Table 3. Reader ingest & egress bytes for a fixed # of samples.

Experiment	Read Bytes (GB)	Send Bytes (GB)
Baseline	538	837
with Cluster	179	837
with IKJT	179	713

to 27.76% and 72.83% to 22.2%, respectively. This allowed us memory headroom to devote to EMBs or larger batches to improve model accuracy or training throughput, respectively. For example, we were able to increase EMB dimensions from 128 to 256 or batch size from 2048 to 6144. Furthermore, Table 2 shows how *RecD* improves the utilization of GPU compute by increasing realized GPU FLOPS to $2.12\times$ the baseline. GPU streaming multiprocessors can achieve higher utilization because they spend less time waiting for exposed A2A communication due to smaller IKJTs.

Single-node Training. While *RecD* greatly increased distributed training throughput, we also evaluated *RecD*'s benefit for single-node training. To do so, we downsized RM_1 to fit within a single ZionEX training node and launched a training job with and without *RecD*. We observed a $2.18\times$ throughput increase in the single-node training setup by using *RecD*. *RecD* still benefits single-node training because it targets GPU memory, network, and compute resources (Section 5). While single-node training reduces the amount of exposed communication due to high-bandwidth NVLink interconnects, *RecD* still improves compute and memory efficiency, leading to improved training throughput. Since storage and readers are disaggregated, *RecD*'s benefits are the same for single-node training as shown in Figure 7.

Impacts to Accuracy. *RecD* itself largely does not affect model accuracy. Specifically, IKJTs encode the exact same logical data as KJTs and thus trainers can train on the exact same batches. The only *RecD* optimization that affects model accuracy is clustering tables by session ID. In fact, clustering leads to significant improvements in accuracy. This is because without clustering, duplicate examples from a session are distributed across batches. The model observes the same sparse feature values for a user across multiple iterations, leading to multiple sparse updates. This causes models to overfit for these features, negatively impacting generalization, especially for less popular (tail) values. By grouping similar samples within the same batch, the model sees each user's data only once, reducing the chance of overfitting less popular sparse feature values.

6.3 Why does *RecD* improve reader throughput?

Figure 7 also showed how *RecD* also improved the throughput of each reader, allowing us to provision fewer readers to feed trainers. To understand why, Figure 10 shows a breakdown of reader CPU compute time spent on *filling*,
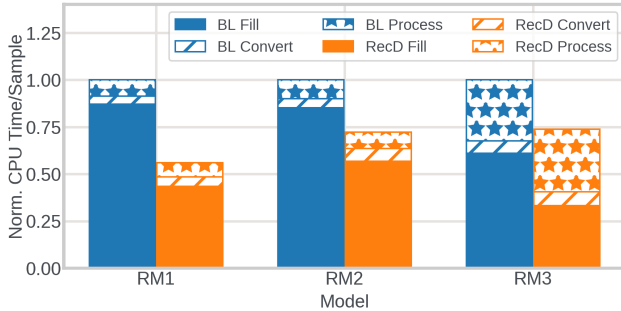
Figure 10. CPU time breakdown for a reader to preprocess a sample across RMs, normalized to each model's baseline CPU time.

converting and *processing* each sample, normalized to the baseline. For each RM, reader time is largely spent on fills: fetching data from Tectonic and decrypting, decompressing (zstd), and decoding bytes to form rows. As shown in Table 3, by clustering tables, *RecD* allows each reader to read significantly fewer bytes per sample. Readers need to spend on both reading and extracting data, reducing the CPU time spent on fills by 50%, 33%, and 46% for RM_1, RM_2, and RM_3, respectively. Furthermore, *RecD* can also reduce the CPU time required for processing, since preprocess operations take deduplicated IKJTs as input. RM_1 and RM_2 required 13% and 11% less time for processing, while RM_3 was effectively the same (3% increase).

RecD requires additional compute at readers to detect duplicate values (via hashing) during feature conversion. Fortunately, Figure 10 shows that this overhead is largely negligible. While the feature conversion time increased by 21%, 37%, and 11% for RM_1, RM_2, and RM_3, respectively, feature conversion requires a small amount of overall compute. Thus, the overall overhead of conversion is negligible (1%) and is easily offset by fill and process benefits.

6.4 Summary

Table 4 summarizes a breakdown of the impacts of each optimization presented in Table 1 for the end-to-end training pipeline performance of RM_1, as reported in our evaluation results. *RecD* presents a suite of optimizations where not only does each optimization yield direct benefits at its respective pipeline system (e.g., storage, readers, or trainers), but it also enables further downstream optimizations. O1 (sharding) and O2 (clustering) directly improve storage efficiency and reader throughput, but they also increase the opportunity for O3 (IKJTs) to deduplicate features. While O3 introduces slight reader overheads, these are nullified by O4, and both O3 and O4 enable trainer-side optimizations. These trainer-side optimizations (O5-O7) ultimately lead to a $2.48\times$ training throughput, as reported in Section 6.2. Collectively, *RecD* optimizations benefit the end-to-end DLRM

Table 4. Summary of impacts of each optimization from Table 1 on end-to-end training pipeline performance for RM_1.

Opt.	Effect on Performance
O1	**Storage**: Improves compression by 1.50x.
O2	**Storage**: (with O1) Improves compression by 3.71x. **Reader**: Reduces CPU fill time by 50% (improves reader throughput by 1.78x).
O3	Enables O4-O6. **Reader**: Increases CPU convert time by 21% (reduces reader throughput by 0.01x).
O4	Enables O5-O6. **Reader**: Reduces CPU process time by 13% (improves reader throughput by 0.01x).
O5	**Trainer**: (with O6 and B4096) Improves training throughput by 1.34x.
O6	**Trainer**: (with O5 and B4096) Improves training throughput by 1.34x.
O7	**Trainer**: (with O5 , O6, and B6144) Improves training throughput by 2.48x.

training pipeline, including storage, readers, and trainers.

7 DISCUSSION

Deciding Which Features to Deduplicate. ML engineers typically apply heuristics to decide which features to deduplicate. IKJTs introduce no trainer overheads aside from an additional index_select used to convert IKJTs to KJTs. Thus, the benefit of deduplicating a feature f must at least offset this overhead. While the specific "worth it" $DedupeFactor(f)$ threshold varies from model-to-model (due to model architecture differences illustrated in Section 6.2), ML engineers will typically start by deduplicating features with $DedupeFactor(f) > 1.5$, and apply standard hyper parameter tuning techniques based on observed trainer throughput to finalize the deduplicated feature set.

Boosting Dedupe Factors. $DedupeFactor(f)$ increases as a function of S, the average number of samples per session. While Section 3 showed that $S = 16.5$ for the characterized industry-scale dataset, we are exploring methods to increase S to yield benefits across the end-to-end training pipeline. For example, the current data generation pipeline *downsamples* (i.e., discards) training samples to keep datasets at a manageable size. However, downsampling is applied on a per-sample basis. By downsampling at a *per-session* basis, we can further increase S, increasing $DedupeFactor$ without affecting model accuracy.

Alternative Solutions and Generality. We considered alternative designs to exploit the session-centric characteristic of DLRM datasets. One promising avenue was to explicitly deduplicate samples within the table schema itself. Specifically, each user session requests inferences (one for each impression) in batches, and each batch uses the same features guaranteeing exact matches. Instead of generating a table row for each impression, we considered generating

one row (with one list for each feature) for each batch of impressions and encoding individual impressions as separate labels within a list. This would directly deduplicate feature values by generating fewer rows within the dataset.

We decided against this approach for several reasons. First, there would still be many duplicates as feature values are largely static even *across* inference batches. Secondly, we use a common dataset schema across DLRMs to ensure interoperability and developer velocity across models and datasets — introducing a new schema would require significant engineering and adoption effort across multiple services. *RecD* is transparent to the training infrastructure because it does not require schema changes, and it enables even more deduplication across request batches.

Furthermore, *RecD* optimizations are easily generalized to different environments, supporting myriad table schemas and model architectures. Enabling *RecD* reader and trainer optimizations only requires a *feature converter* module to convert arbitrary table schemas into the IKJT encoding, and an `index_select` call to convert IKJTs back to KJTs when necessary. Because IKJTs directly build on standardized jagged tensors (ragged tensors in TensorFlow (Tensor-Flow, 2022)), preprocessing functions and model architectures can operate on IKJTs as KJTs with minimal changes.

Supporting Partial IKJTs. Supporting exact matches captures the vast majority of duplication in industry-scale datasets — 81.6% of an estimated 93.9% maximum (Section 3). Even so, IKJTs are also easily extended to support partial deduplication, capturing an additional 7.8% of values, by leveraging the fact that partial matches are shifts. Partial IKJTs remove the $offsets$ slice, and instead encode each row's $[offset, length]$ in the $inverse\_lookup$ slice. In the example in Figure 5, feature b can by partially deduplicated via a partial IKJT consisting of $values = [3, 4, 5, 6]$ and $inverse\_lookup = [[0, 3], [1, 3], [0, 3]]$.

8 RELATED WORK

Duplication in DLRM Datasets. Gai et al. notes how DLRM datasets at Alibaba exhibits feature duplication (Gai et al., 2017). To exploit this, the authors mention a "common feature trick" that routes samples from similar users to the same worker in a parameter server training setup. The authors speed up training throughput by caching and reusing the parameter update for "common" features across each worker's samples. Follow-up work by Ge et al. cites using the "common feature trick" during training (Ge et al., 2018). Unfortunately, the authors provide scant details on how "common" features are generated, stored, or encoded. They also do not elucidate how model architectures can exploit duplicate features, nor how the "common feature trick" can extend beyond parameter servers to synchronous

training used in scale-out GPU training clusters (Mudigere et al., 2022). We provide an in-depth characterization of feature value duplication in industrial DLRM datasets. *RecD* deduplicates features across the end-to-end training pipeline by coalescing duplicate features in storage, compactly encoding them into IKJTs, and intelligently training on IKJTs.

Deduplication in ML. Data deduplication has been studied in ML training outside of DLRMs. Lee et al. studied how deduplication in a text corpus improved model accuracy for language tasks (Lee et al., 2022). Allamanis studied how duplication in code datasets degraded model performance for ML models for source code (Allamanis, 2019). To the best of our knowledge, our work is the first to study the systems implications of duplication in ML training datasets.

Database Systems. Data deduplication is a well-studied area in databases. IKJTs use a similar encoding mechanism to dictionary encoding commonly used in file formats such as Parquet (Apache, 2022). To coalesce duplicates within an IKJT, we rely on CLUSTER BY clauses supported by myriad database execution engines, such as Spark (Zaharia et al., 2012). *RecD* applies these concepts to enable and encode deduplicated tensors for ML training jobs.

Systems Optimizations for DLRM Training. Zhao et al. presented various optimizations to improve DSI efficiency for DLRM training at Meta (Zhao et al., 2022). RecShard (Sethi et al., 2022) and Adnan et al. (Adnan et al., 2021) leveraged skewed feature popularities to shard EMBs across GPUs, improving training throughput. Similarly, Fleche (Xie et al., 2022) is an embedding cache that caches EMBs on GPU HBM while relying on CPU DRAM for holding entire EMBs, targeting only single-GPU training. To avoid cache write conflicts, Fleche recognizes that many sparse feature IDs are duplicated within a batch and performs only a single cache lookup for each unique ID, similar to *RecD*'s ability to deduplicate EMB lookups. TT-Rec (Yin et al., 2021) demonstrated compression techniques for EMBs. *RecD* provides orthogonal optimizations to improve storage, reading, and training performance by deduplicating features across the DLRM training pipeline.

9 CONCLUSION

This paper presented *RecD*, a suite of optimizations for industry-scale, end-to-end DLRM training pipelines. We provide an in-depth characterization of how DLRM datasets exhibit inherent feature duplication. *RecD* coalesces duplicate features within a training batch, efficiently encodes them using IKJTs, and optimizes DLRM model architectures to train on deduplicated tensors. As a result, *RecD* improves training and preprocessing throughput and storage efficiency by up to $2.48\times$, $1.79\times$, and $3.71\times$, respectively.

REFERENCES

Acun, B., Murphy, M., Wang, X., Nie, J., Wu, C., and Hazelwood, K. Understanding training efficiency of deep learning recommendation models at scale. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 802–814, 2021.

Adnan, M., Maboud, Y. E., Mahajan, D., and Nair, P. J. Accelerating recommendation system training by leveraging popular choices. *Proc. VLDB Endow.*, 15(1):127–140, sep 2021. ISSN 2150-8097. doi: 10.14778/3485450. 3485462. URL https://doi.org/10.14778/3485450.3485462.

Allamanis, M. The adverse effects of code duplication in machine learning models of code. In *Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, Onward! 2019, pp. 143–153, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450369954. doi: 10.1145/ 3359591.3359735. URL https://doi.org/10.1145/3359591.3359735.

Anil, R., Gadanho, S., Huang, D., Jacob, N., Li, Z., Lin, D., Phillips, T., Pop, C., Regan, K., Shamir, G. I., Shivanna, R., and Yan, Q. On the factory floor: Ml engineering for industrial-scale ads recommendation models, 2022. URL https://arxiv.org/abs/2209.05310.

Apache. Apache parquet file encodings, 2022. URL https://parquet.apache.org/docs/file-format/data-pages/encodings/.

Ardalani, N., Wu, C.-J., Chen, Z., Bhushanam, B., and Aziz, A. Understanding scaling laws for recommendation models, 2022. URL https://arxiv.org/abs/2208.08489.

AWS. Aws trainium, 2022. URL https://aws.amazon.com/machine-learning/trainium/.

Chen, Q., Zhao, H., Li, W., Huang, P., and Ou, W. Behavior sequence transformer for e-commerce recommendation in alibaba. In *Proceedings of the 1st International Workshop on Deep Learning Practice for High-Dimensional Sparse Data*, DLP-KDD '19, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367837. doi: 10.1145/3326937.3341261. URL https://doi.org/10.1145/3326937.3341261.

de Souza Pereira Moreira, G., Rabhi, S., Lee, J. M., Ak, R., and Oldridge, E. Transformers4rec: Bridging the gap between nlp and sequential / session-based recommendation. In *Proceedings of the 15th ACM Conference on Recommender Systems*, RecSys '21, pp. 143–153, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384582. doi: 10. 1145/3460231.3474255. URL https://doi.org/10.1145/3460231.3474255.

Fang, H., Zhang, D., Shu, Y., and Guo, G. Deep learning for sequential recommendation: Algorithms, influential factors, and evaluations. *ACM Trans. Inf. Syst.*, 39(1), nov 2020. ISSN 1046-8188. doi: 10.1145/3426723. URL https://doi.org/10.1145/3426723.

Gai, K., Zhu, X., Li, H., Liu, K., and Wang, Z. Learning piece-wise linear models from large scale data for ad click prediction, 2017. URL https://arxiv.org/abs/1704.05194.

Ge, T., Zhao, L., Zhou, G., Chen, K., Liu, S., Yi, H., Hu, Z., Liu, B., Sun, P., Liu, H., Yi, P., Huang, S., Zhang, Z., Zhu, X., Zhang, Y., and Gai, K. Image matters: Visually modeling user behaviors using advanced model server. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, CIKM '18, pp. 2087–2095, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450360142. doi: 10.1145/3269206.3272007. URL https://doi.org/10.1145/3269206.3272007.

Hazelwood, K., Bird, S., Brooks, D., Chintala, S., Diril, U., Dzhulgakov, D., Fawzy, M., Jia, B., Jia, Y., Kalro, A., Law, J., Lee, K., Lu, J., Noordhuis, P., Smelyanskiy, M., Xiong, L., and Wang, X. Applied machine learning at facebook: A datacenter infrastructure perspective. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 620–629, 2018. doi: 10.1109/HPCA.2018.00059.

Karpathiotakis, M., Wernli, D., and Stojanovic, M. Scribe: Transporting petabytes per hour via a distributed, buffered queueing system, Oct 2019. URL https://engineering.fb.com/2019/10/07/data-infrastructure/scribe/.

Kaufman, S., Rosset, S., Perlich, C., and Stitelman, O. Leakage in data mining: Formulation, detection, and avoidance. *ACM Trans. Knowl. Discov. Data*, 6(4), dec 2012. ISSN 1556-4681. doi: 10.1145/ 2382577.2382579. URL https://doi.org/10.1145/2382577.2382579.

Lardinois, F. Google launches a 9 exaflop cluster of cloud TPU v4 pods into public preview. https://techcrunch.com/2022/05/11/google-launches-a-9-exaflop-cluster-of-cloud-tpu-v4-pods-into-public-preview/, May 2022.

Lee, K., Ippolito, D., Nystrom, A., Zhang, C., Eck, D., Callison-Burch, C., and Carlini, N. Deduplicating training data makes language models better. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 2022.

Li, C., Liu, Z., Wu, M., Xu, Y., Zhao, H., Huang, P., Kang, G., Chen, Q., Li, W., and Lee, D. L. Multi-interest network with dynamic routing for recommendation at tmall. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, CIKM '19, pp. 2615–2623, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450369763. doi: 10.1145/3357384.3357814. URL https://doi.org/10.1145/3357384.3357814.

Li, R., Qin, Z., Wang, X., Chen, S. J., and Metzler, D. *Stabilizing Neural Search Ranking Models*, pp. 2725–2732. Association for Computing Machinery, New York, NY, USA, 2020. ISBN 9781450370233. URL https://doi.org/10.1145/3366423.3380030.

Meta. Powered by ai: Instagram's explore recommender system, 2019. URL https://ai.facebook.com/blog/powered-by-ai-instagrams-explore-recommender-system/.

Meta. Introducing the ai research supercluster, 2022. URL https://ai.facebook.com/blog/ai-rsc/.

Mudigere, D., Hao, Y., Huang, J., Jia, Z., Tulloch, A., Sridharan, S., Liu, X., Ozdal, M., Nie, J., Park, J., Luo, L., Yang, J. A., Gao, L., Ivchenko, D., Basant, A., Hu, Y., Yang, J., Ardestani, E. K., Wang, X., Komuravelli, R., Chu, C.-H., Yilmaz, S., Li, H., Qian, J., Feng, Z., Ma, Y., Yang, J., Wen, E., Li, H., Yang, L., Sun, C., Zhao, W., Melts, D., Dhulipala, K., Kishore, K., Graf, T., Eisenman, A., Matam, K. K., Gangidi, A., Chen, G. J., Krishnan, M., Nayak, A., Nair, K., Muthiah, B., khorashadi, M., Bhattacharya, P., Lapukhov, P., Naumov, M., Mathews, A., Qiao, L., Smelyanskiy, M., Jia, B., and Rao, V. Software-hardware co-design for fast and scalable training of deep learning recommendation models. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*, ISCA '22, pp. 993–1011, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450386104. doi: 10.1145/3470496.3533727. URL https://doi.org/10.1145/3470496.3533727.

Naumov, M., Mudigere, D., Shi, H.-J. M., Huang, J., Sundaraman, N., Park, J., Wang, X., Gupta, U., Wu, C.-J., Azzolini, A. G., Dzhulgakov, D., Mallevich, A., Cherniavskii, I., Lu, Y., Krishnamoorthi, R., Yu, A., Kondratenko, V., Pereira, S., Chen, X., Chen, W., Rao, V., Jia, B., Xiong, L., and Smelyanskiy, M. Deep learning recommendation model for personalization and recommendation systems, 2019. URL https://arxiv.org/abs/1906.00091.

Naumov, M., Kim, J., Mudigere, D., Sridharan, S., Wang, X., Zhao, W., Yilmaz, S., Kim, C., Yuen, H., Ozdal, M., Nair, K., Gao, I., Su, B.-Y., Yang, J., and Smelyanskiy, M. Deep learning training in facebook data centers: Design of scale-up and scale-out systems, 2020. URL https://arxiv.org/abs/2003.09518.

NVIDIA. Nvidia nccl, 2022. URL https://developer.nvidia.com/nccl.

ORC, A. Apache orc: High-performance columnar storage for hadoop, 2022. URL https://orc.apache.org/.

Pan, S., Stavrinos, T., Zhang, Y., Sikaria, A., Zakharov, P., Sharma, A., P, S. S., Shuey, M., Wareing, R., Gangapuram, M., Cao, G., Preseau, C., Singh, P., Patiejunas, K., Tipton, J., Katz-Bassett, E., and Lloyd, W. Facebook's tectonic filesystem: Efficiency from exascale. In *19th USENIX Conference on File and Storage Technologies (FAST 21)*, pp. 217–231. USENIX Association, February 2021. ISBN 978-1-939133-20-5. URL https://www.usenix.org/conference/fast21/presentation/pan.

Patterson, D., Gonzalez, J., Hölzle, U., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D. R., Texier, M., and Dean, J. The carbon footprint of machine learning training will plateau, then shrink. *Computer*, 55(7):18–28, 2022. doi: 10.1109/MC.2022.3148714.

Pi, Q., Bian, W., Zhou, G., Zhu, X., and Gai, K. Practice on long sequential user behavior modeling for click-through rate prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pp. 2671–2679, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362016. doi: 10.1145/3292500.3330666. URL https://doi.org/10.1145/3292500.3330666.

PyTorch. Torchrec.sparse, 2022. URL https://pytorch.org/torchrec/torchrec.sparse.html.

PyTorch. torch.utils.data, 2023. URL https://pytorch.org/docs/stable/data.html.

Sethi, G., Acun, B., Agarwal, N., Kozyrakis, C., Trippel, C., and Wu, C.-J. Recshard: Statistical feature-based memory optimization for industry-scale neural recommendation. In *Proceedings of the 27th ACM Inter-*

national Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2022, pp. 344–358, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392051. doi: 10.1145/3503222.3507777. URL https://doi.org/10.1145/3503222.3507777.

TensorFlow. Ragged tensors, 2022. URL https://www.tensorflow.org/guide/ragged_tensor.

Thusoo, A., Sarma, J. S., Jain, N., Shao, Z., Chakka, P., Anthony, S., Liu, H., Wyckoff, P., and Murthy, R. Hive: A warehousing solution over a map-reduce framework. *Proc. VLDB Endow.*, 2(2):1626–1629, aug 2009. ISSN 2150-8097. doi: 10.14778/1687553.1687609. URL https://doi.org/10.14778/1687553.1687609.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

Wang, S., Cao, L., Wang, Y., Sheng, Q. Z., Orgun, M. A., and Lian, D. A survey on session-based recommender systems. *ACM Comput. Surv.*, 54(7), jul 2021. ISSN 0360-0300. doi: 10.1145/3465401. URL https://doi.org/10.1145/3465401.

Xie, M., Lu, Y., Lin, J., Wang, Q., Gao, J., Ren, K., and Shu, J. Fleche: An efficient gpu embedding cache for personalized recommendations. In *Proceedings of the Seventeenth European Conference on Computer Systems*, EuroSys '22, pp. 402–416, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391627. doi: 10.1145/3492321.3519554. URL https://doi.org/10.1145/3492321.3519554.

Yin, C., Acun, B., Wu, C.-J., and Liu, X. Tt-rec: Tensor train compression for deep learning recommendation models. *Proceedings of Machine Learning and Systems*, 3:448–462, 2021.

Zaharia, M., Chowdhury, M., Das, T., Dave, A., Ma, J., McCauly, M., Franklin, M. J., Shenker, S., and Stoica, I. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, pp. 15–28, San Jose, CA, April 2012. USENIX Association. ISBN 978-931971-92-8. URL https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/zaharia.

Zhao, M., Agarwal, N., Basant, A., Gedik, B., Pan, S., Ozdal, M., Komuravelli, R., Pan, J., Bao, T., Lu, H., Narayanan, S., Langman, J., Wilfong, K., Rastogi, H., Wu, C.-J., Kozyrakis, C., and Pol, P. Understanding data storage and ingestion for large-scale deep recommendation model training: Industrial product. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*, ISCA '22, pp. 1042–1057, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450386104. doi: 10.1145/3470496.3533044. URL https://doi.org/10.1145/3470496.3533044.

Zhao, Z., Hong, L., Wei, L., Chen, J., Nath, A., Andrews, S., Kumthekar, A., Sathiamoorthy, M., Yi, X., and Chi, E. Recommending what video to watch next: A multitask ranking system. In *Proceedings of the 13th ACM Conference on Recommender Systems*, RecSys '19, pp. 43–51, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362436. doi: 10.1145/3298689.3346997. URL https://doi.org/10.1145/3298689.3346997.

Zstandard. Zstandard, 2022. URL www.zstd.net.