

Paper Review- cedar: Optimized and Unified Machine Learning Input Data

Summary

This paper presents cedar, a unified and optimized framework for ML input data pipelines. Input pipelines read raw data, apply transformations like cropping or tokenizing, and quickly deliver mini batches to GPUs for training. A critical bottleneck in modern training is that pipelines often fail to keep up with GPUs, leading to stalls and wasted resources. Existing systems are fragmented, framework-specific, and lack systematic, context-aware optimization. cedar addresses this with a Python-native API where pipelines are built from composable operators across PyTorch, TensorFlow, JAX, and other libraries. Its key innovations include an extensible Optimizer that systematically searches, prunes, and applies multiple optimizations (reordering, caching, fusion, offloading, prefetching) using profiling and historical metadata, and semantic preserving reorderings guided by lightweight user hints. At runtime, a dynamic Scaler adjusts parallelism and right-sizes resources to meet demand efficiently. Experiments on CV, NLP, and speech pipelines show cedar improves throughput by 1.87x - 10.65x over state-of-the-art systems while reducing resource costs.

Evaluation

The paper tackles a highly significant problem: inefficient input pipelines that stall GPUs and waste resources, directly limiting ML scalability (example: Meta's DPP can require dozens of CPU servers to support a single GPU server). cedar's novelty lies in its principled, context-aware approach that unifies diverse optimizations instead of applying them in isolation. The Optimizer's cost-model-driven search and pruning resembles a database query optimizer and represents a real advance for input data systems. The use of semantic-preserving reorderings guided by user hints is particularly innovative, while its framework-agnostic design fills a major gap in current tools. The evaluation is detailed, spanning eight real-world pipelines across CV, NLP, and speech, and compares against a broad set of state-of-the-art systems (tf.data, tf.data service, Ray Data, FastFlow, Plumber, PyTorch DataLoader). Results show that cedar provides clear throughput gains and strong evidence of generality. Ablation, caching, and dynamic scaling studies further validate its contributions. The paper is also well-written and structured, with clear figures that enhance readability.

Main Takeaways

1. cedar shows that real speedups come from combining multiple optimizations together, not using them in isolation.
2. By supporting PyTorch, TensorFlow, JAX, and diverse libraries, cedar solves the fragmentation problem in today's ML pipelines.
3. Lightweight user hints let cedar respect randomness and semantics while still applying efficient optimizations like reordering and caching.

Strengths

1. Cedar's dynamic Scaler turns throughput gains into real resource efficiency by dynamically scaling resources, reducing wasted CPU and GPU cycles, making training faster and more cost-effective.
2. Systematic optimization via an extensible Optimizer that explores the combinatorial search space of reordering, fusion, caching, offloading, and prefetching.
3. Low optimization overheads with tracing and plan generation completed in under 6 seconds, making its techniques practical for long-running real world training jobs.

Weaknesses

1. While hints are powerful, requiring users to mark dependencies and randomness can be difficult for complex pipelines or less experienced users. Automated inference would reduce this workload.
2. The reordering cost model assumes linear scaling with input size, which may not hold for non-linear UDFs. Without customization, this could lead to suboptimal plans.
3. The metadata store is local to each job, preventing cross-job or cross-user learning. A shared or cloud-backed store could capture historical performance statistics to further improve plan selection and reduce profiling overhead.

Discussion Topics

1. Can cedar automatically infer dependencies and randomness (example: with AI or static analysis), reducing user effort?
2. How could cedar integrate GPU-based preprocessing (example: DALI, FusionFlow), and what changes would its cost model need to compare CPU-GPU tradeoffs?