Characterizing I/O in Machine Learning Workloads

Oana Balmau

Resource-Aware ML Day @ ITU, February 13th, 2023

McGill

Humanity produces a lot of data

Data expected to
grow exponentially

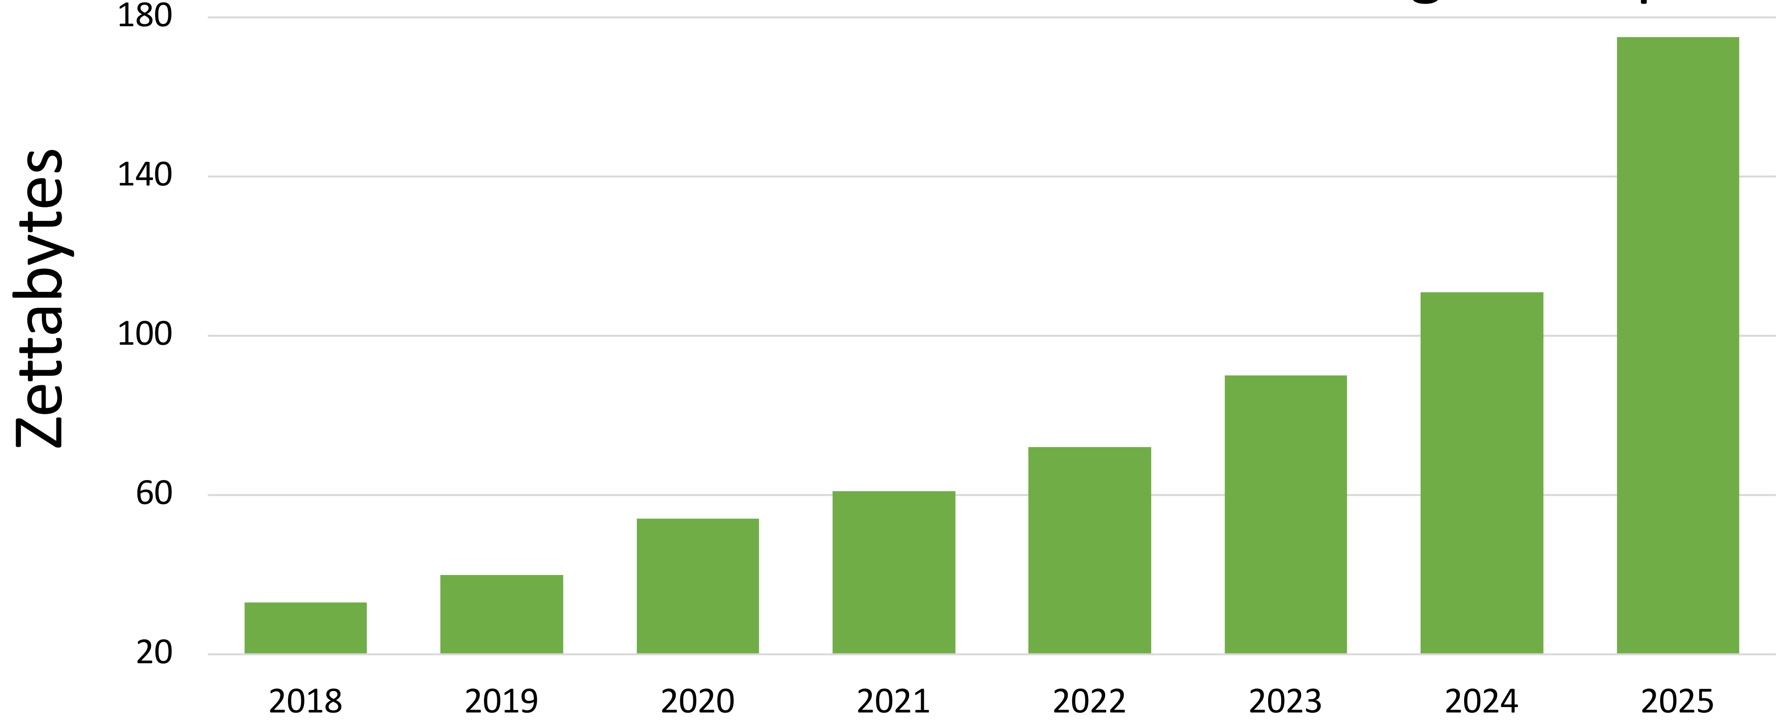

Source: IDC 2022

Humanity produces a lot of data

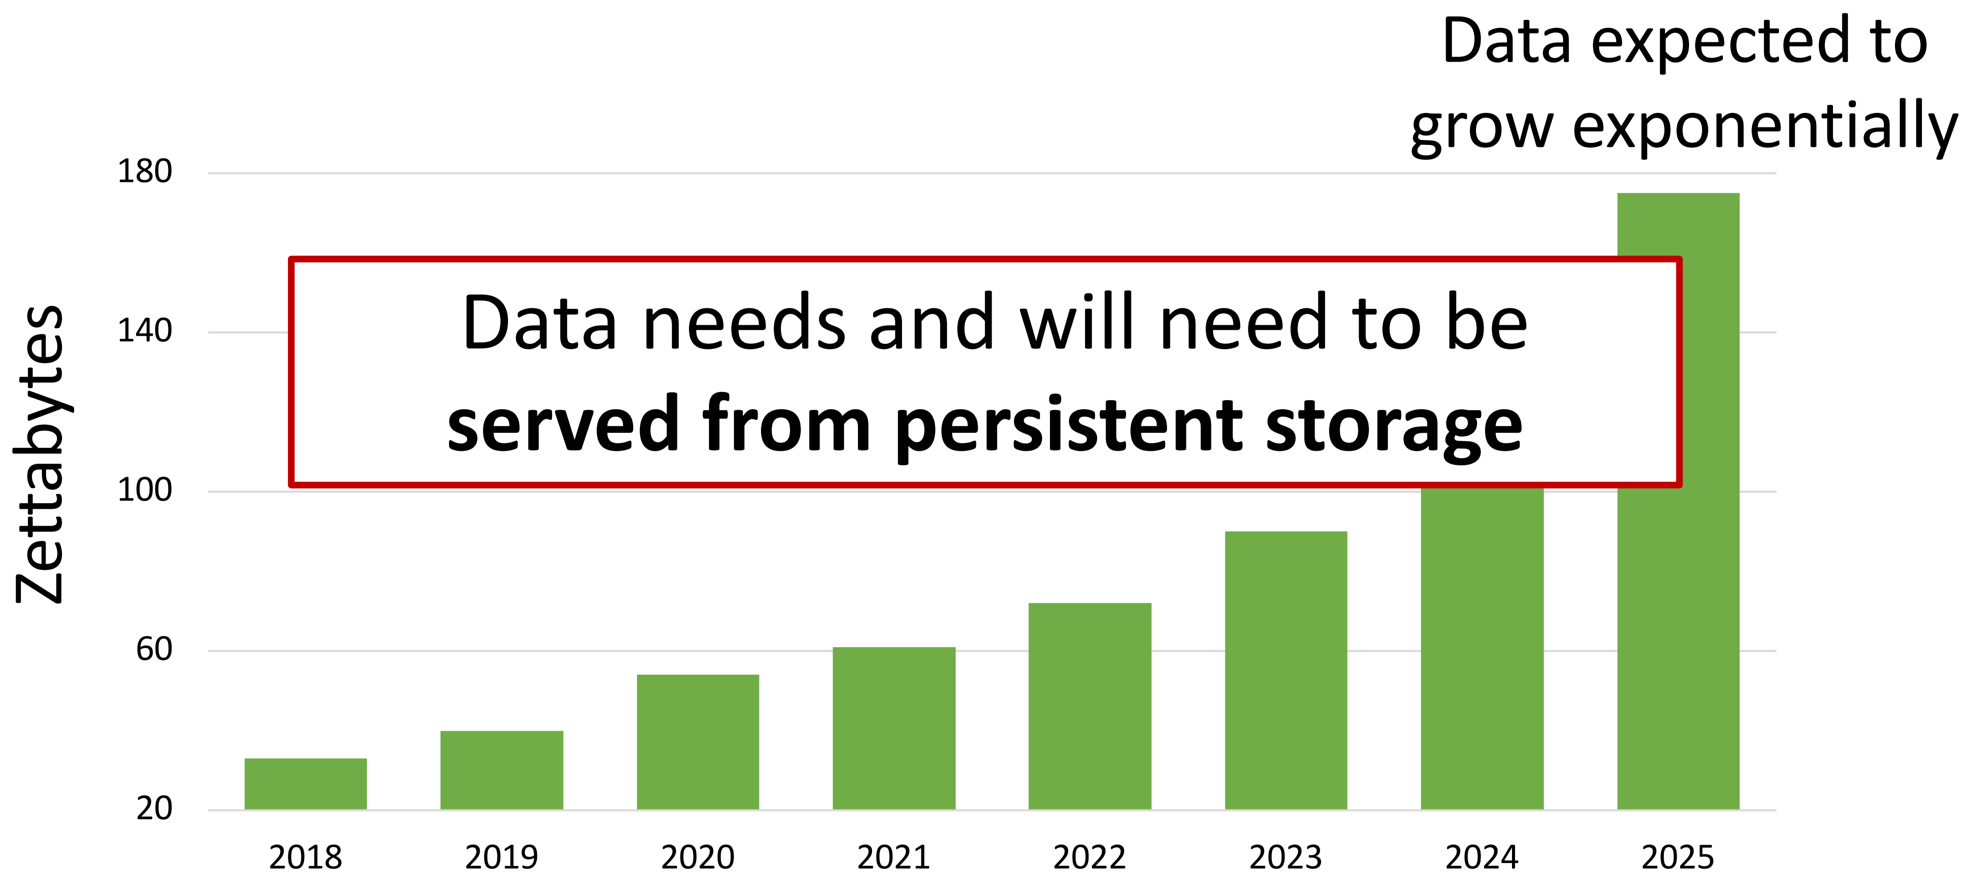

Source: IDC 2022

Data is the moving force of ML algorithms

… but in many projects the **storage decision is an afterthought**

Data is the moving force of ML algorithms

… but in many projects the **storage decision is an afterthought**

Why create an ML Storage benchmark?

Current ML/AI benchmarks

Many existing ML/AI benchmarks

DeepMind Lab

MLPerf

OpenAI

DLBT

PMLDB

DAWNBench

Current ML/AI benchmarks

- Focus on **end-to-end testing**

  → hard to isolate value of each component

- Insist on **training and inference** speed

  → tend to simplify storage

  → ignore pre-processing

- **Expensive accelerators** needed to run
- Require **extensive entry knowledge**

DeepMind Lab

MLPerf

OpenAI

DLBT

PMLDB

DAWNBench

Why create an ML Storage benchmark?

- **Understand storage bottlenecks in ML workloads and propose optimizations**

- **Help AI/ML researchers and practitioners make an informed *storage* decision**

MLPerf Storage working group
Who are we?

Mix of industry and academia

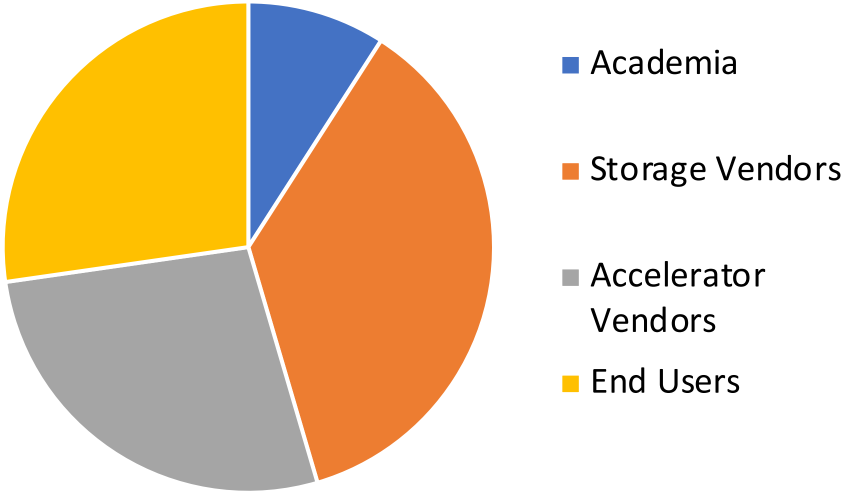

McGill
NVIDIA
intel
NUTANIX
WEKA
Red Hat
PANASAS
Micron
Argonne NATIONAL LABORATORY
tenstorrent

Benchmark Vision

Existing benchmarks	**Our work**
Focus on **end-to-end testing**	Focus on **storage impact in ML/AI**
Simplified storage setup	Realistic **storage & pre-processing** settings
Expensive accelerators needed to run	**No accelerator required** to run
Require **extensive entry knowledge**	**Minimal AI/ML knowledge** required

Stages of the ML Pipeline

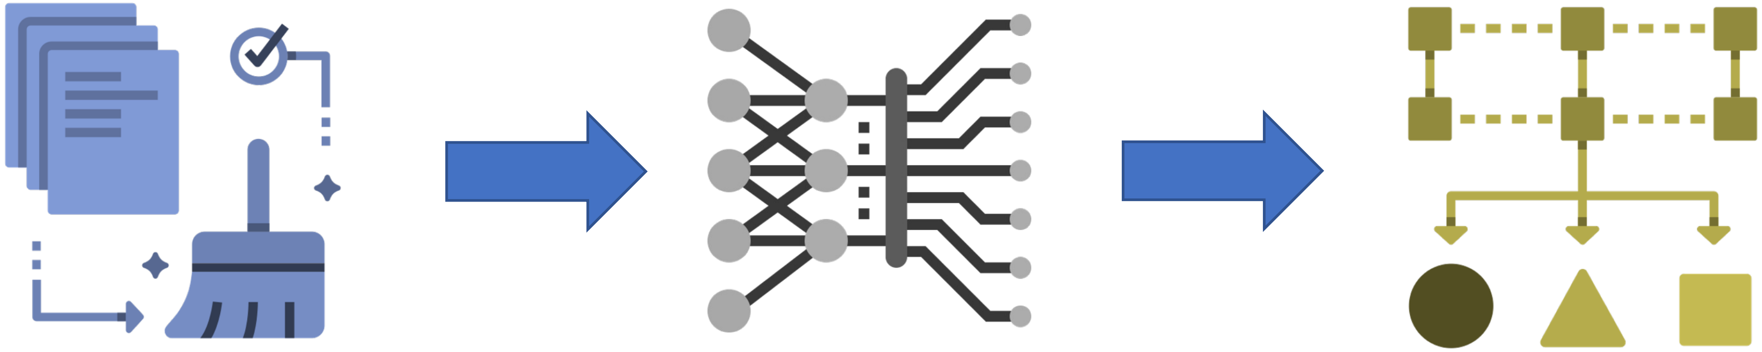

Data cleaning & pre-processing → **Training** → **Inference**

Stages of the ML Pipeline

I/O intensive [1,2] – Our focus

*As much as **50% of the Watts** can go into storage and data cleaning [2]*

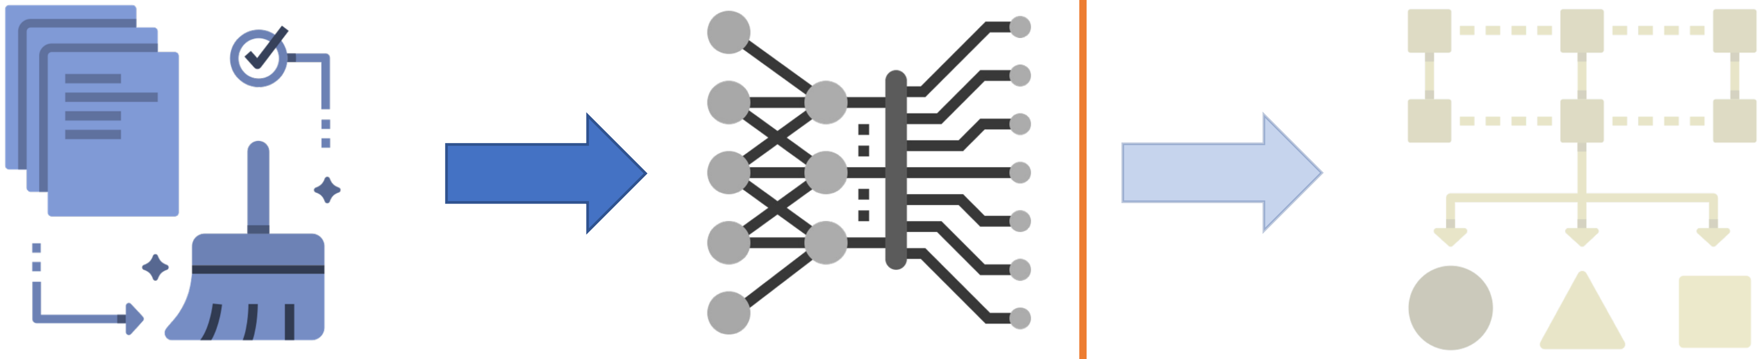

Data cleaning & pre-processing

Training

Inference

*[1] Murray et al. **tf.data: A Machine Learning Data Processing Framework**, VLDB 21.*
*[2] Zhao et a. **Understanding Data Storage and Ingestion for Large-Scale Deep Recommendation Model Training** ISCA 22.*

Data Pipeline in ML: Pre-processing

Storage resources

Disk

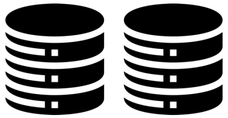

Memory

Compute resources

CPUs

Accelerators (GPU, TPU)

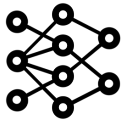

Train model

Data Pipeline in ML: Pre-processing

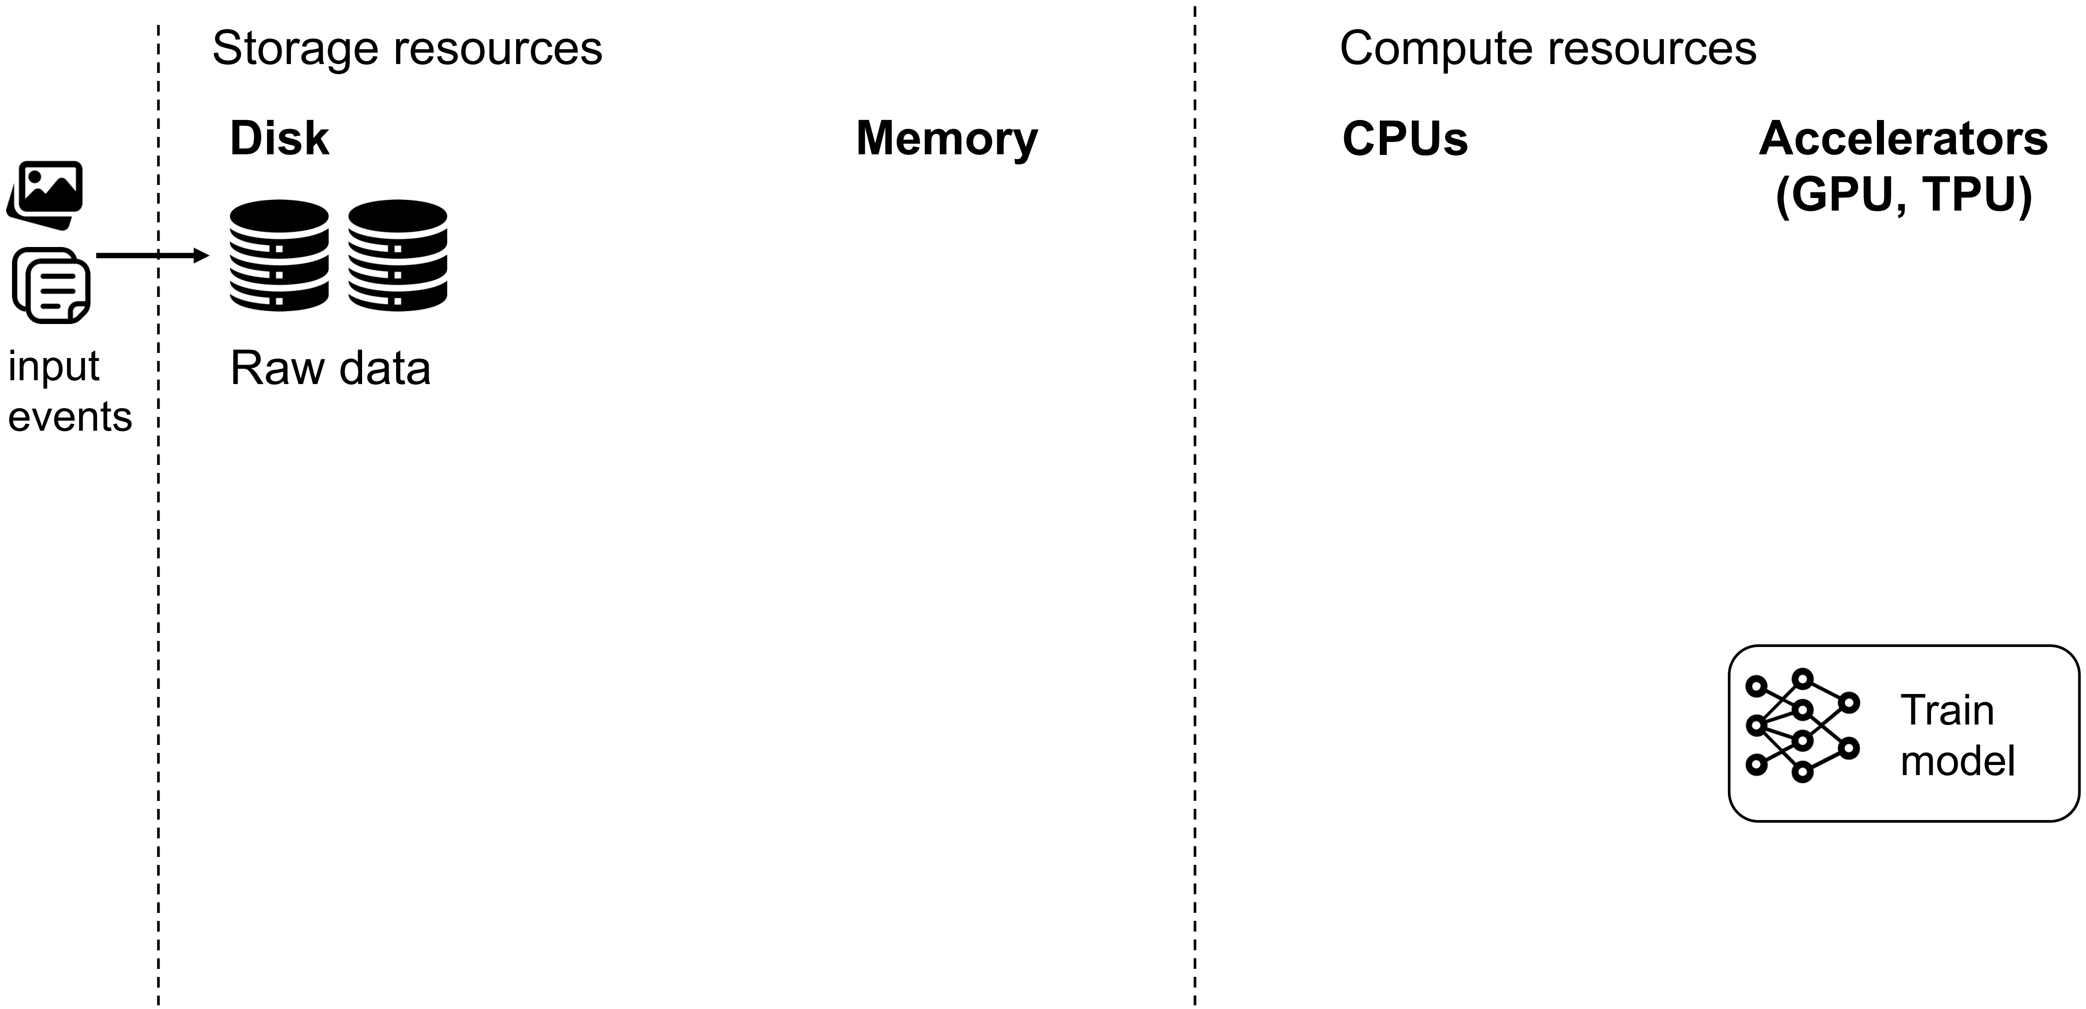

Data Pipeline in ML: Pre-processing

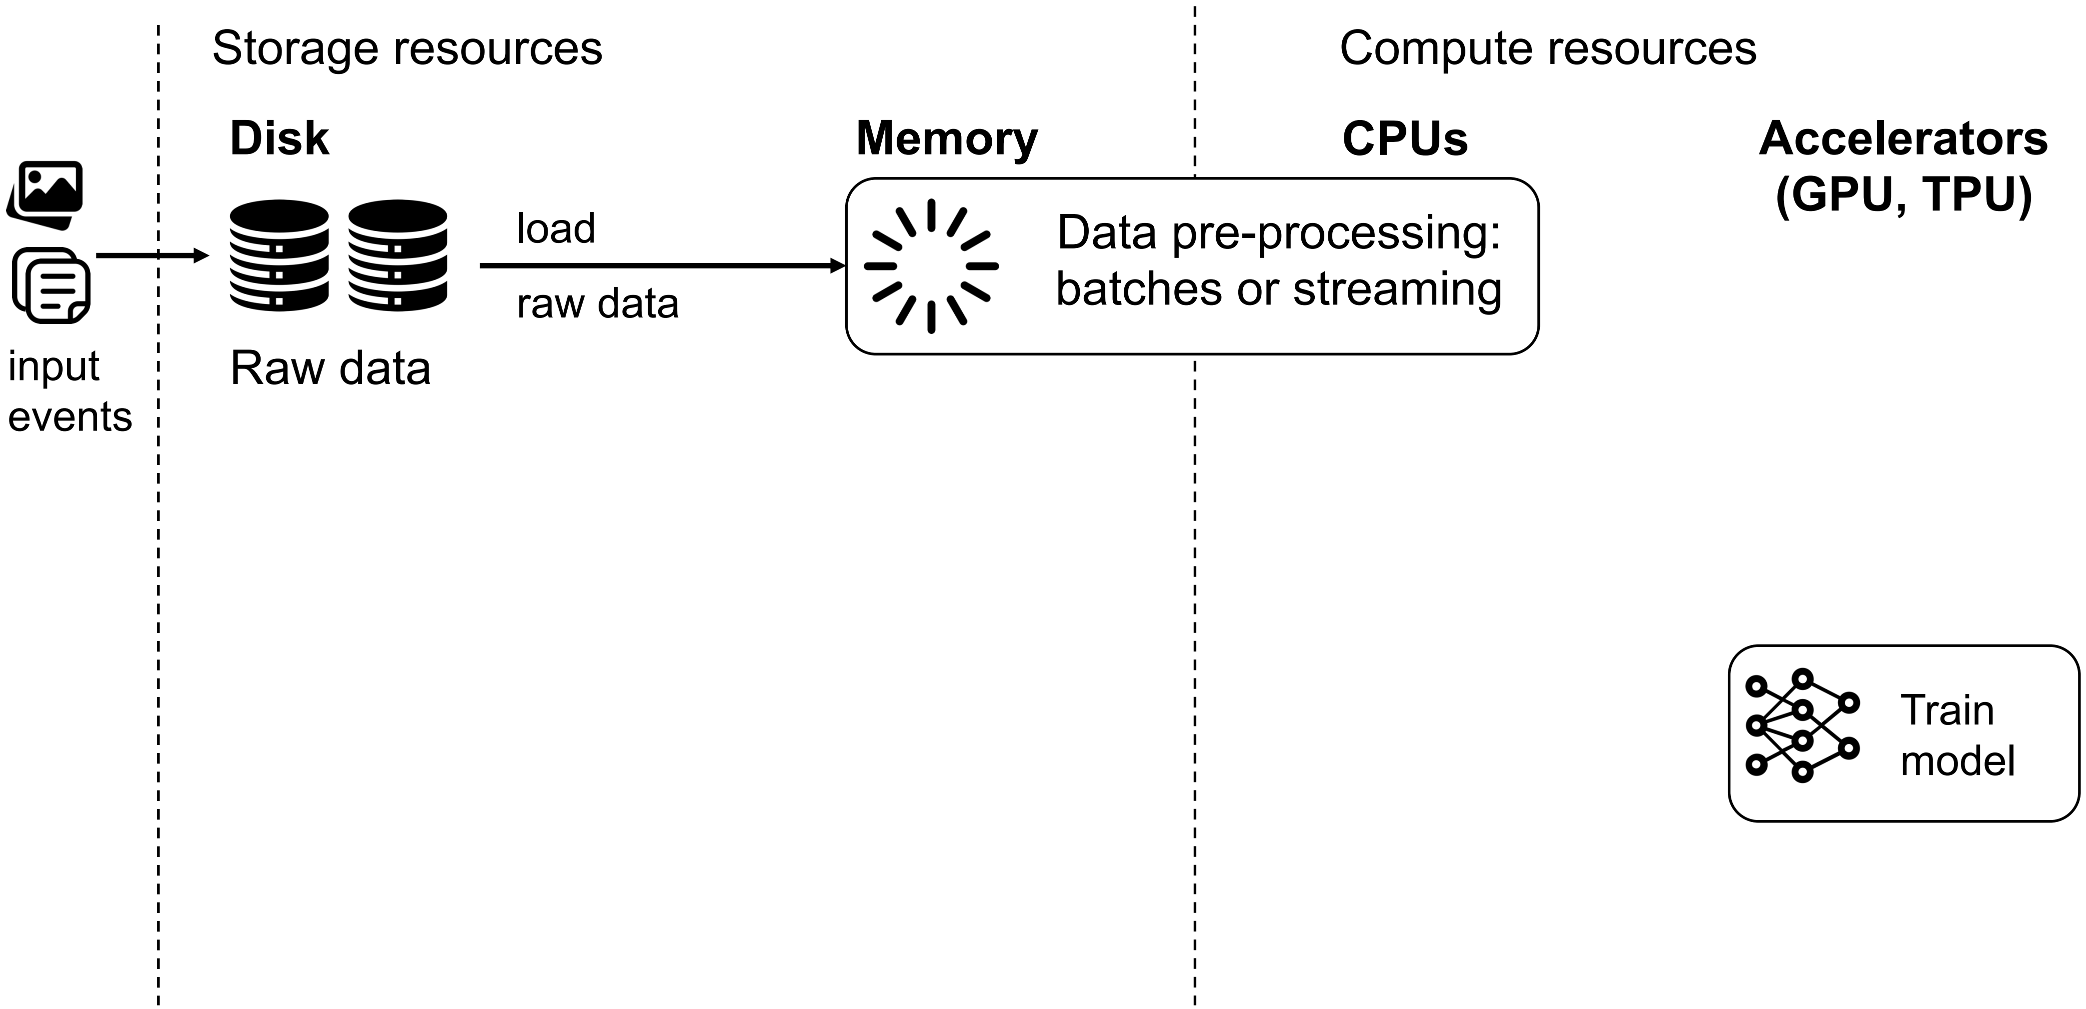

Data Pipeline in ML: Pre-processing

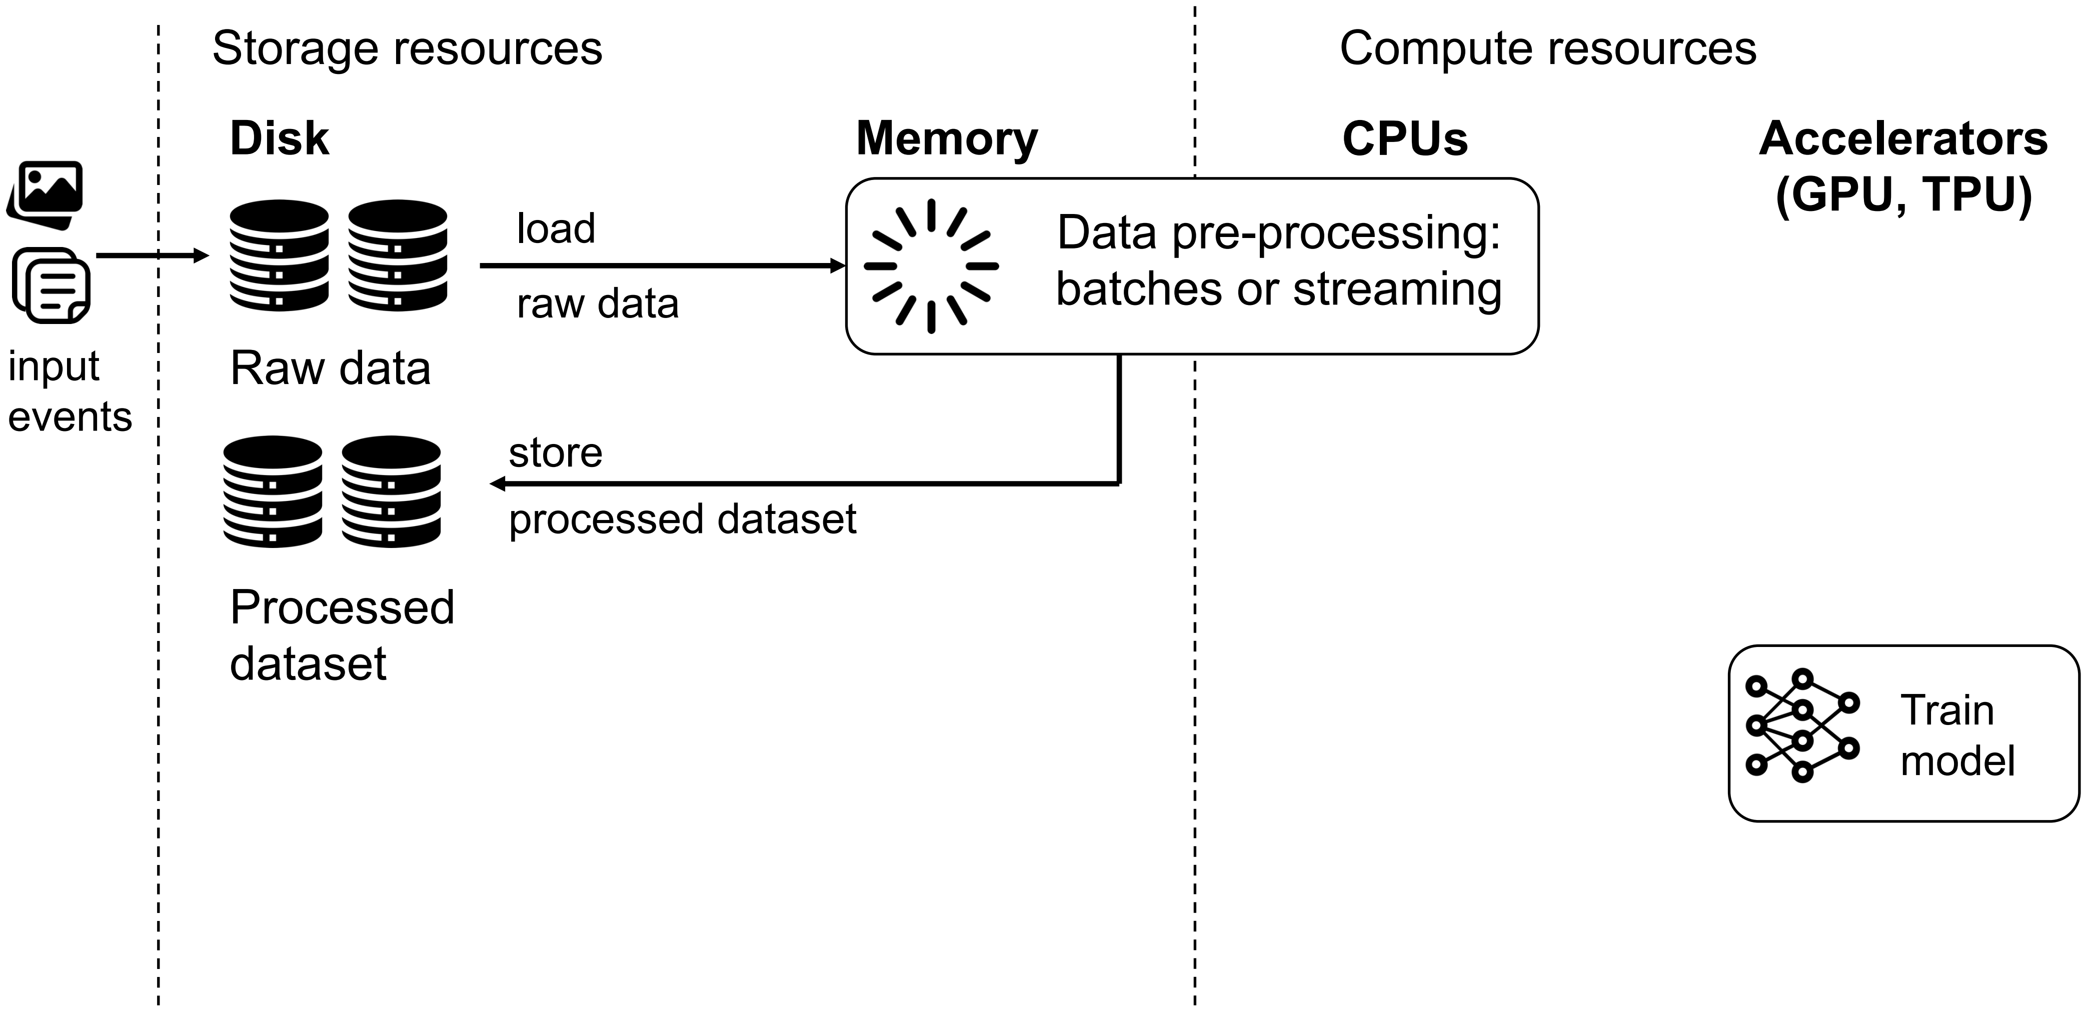

Data Pipeline in ML: Pre-processing

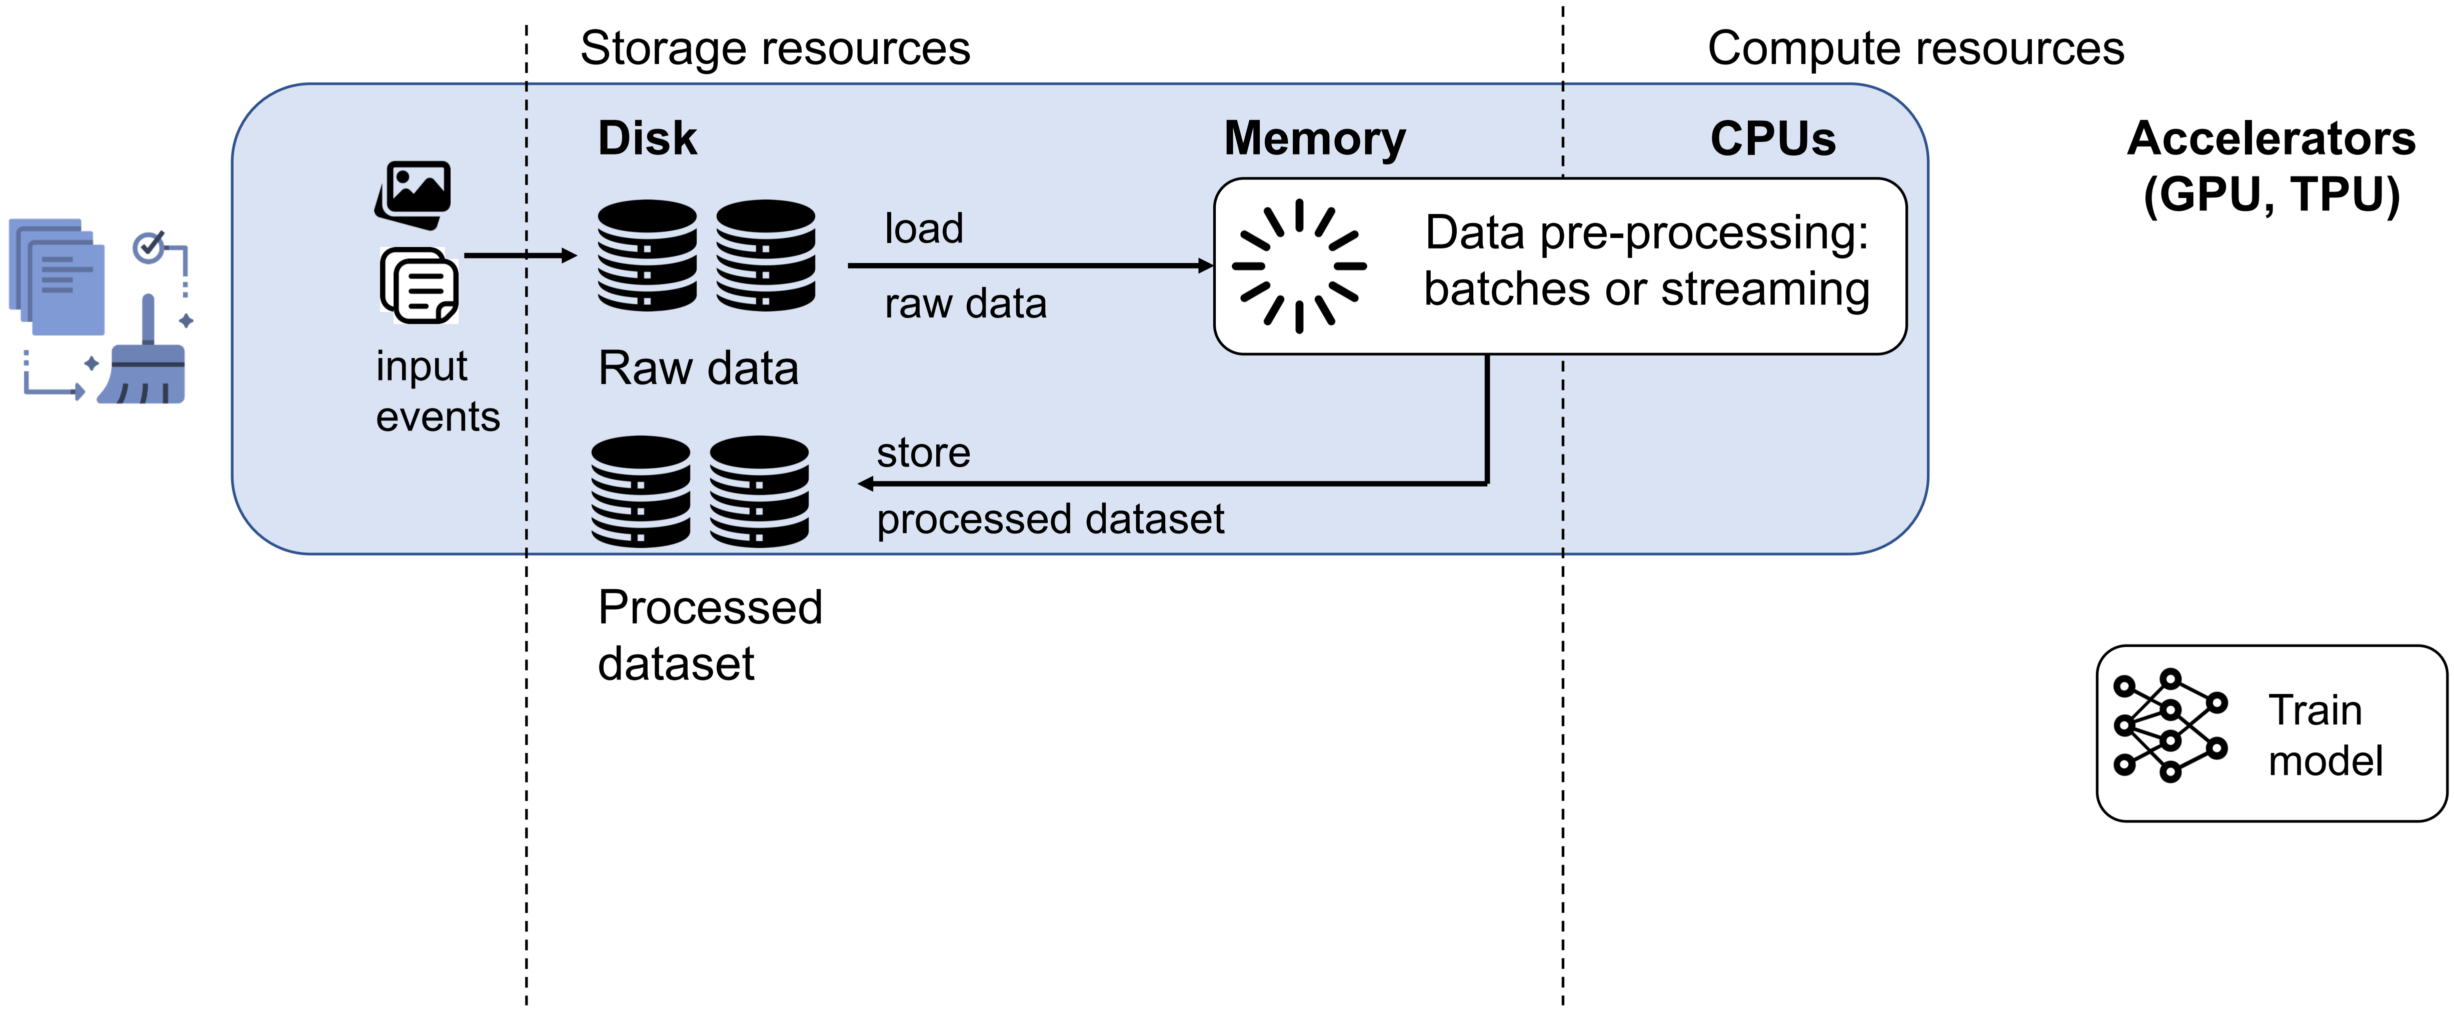

MLPerf Storage V0.5

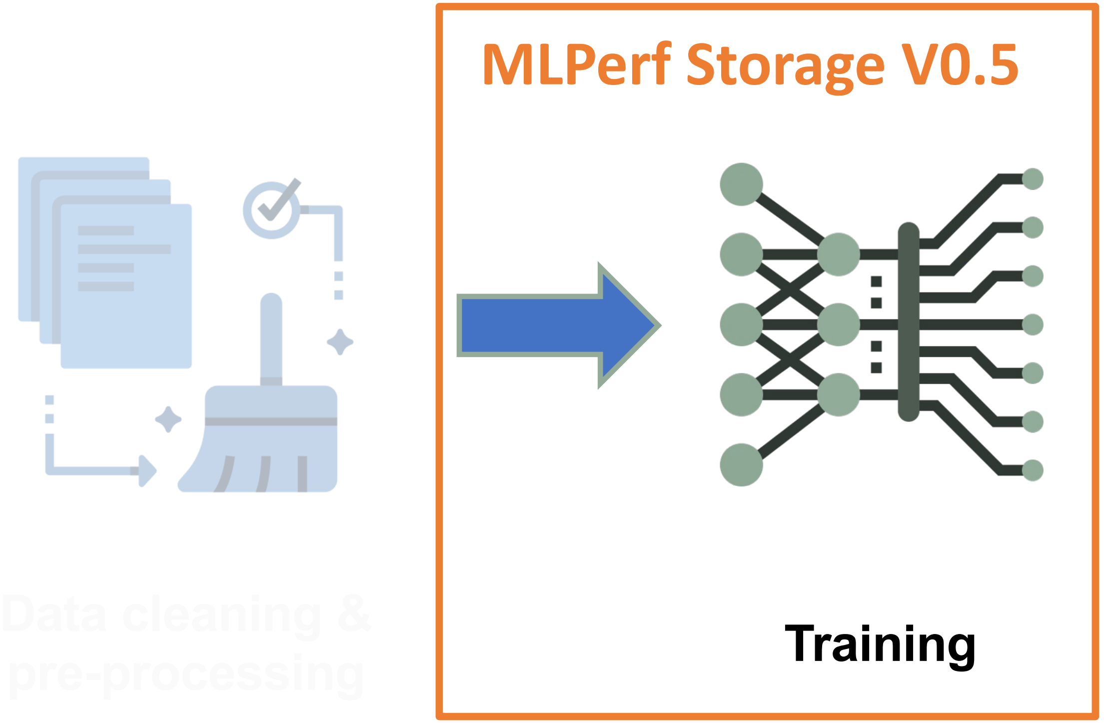

Focus on **storage impact in ML/AI**

Realistic **storage** settings in **training phase**

No accelerator required to run

Minimal AI/ML knowledge required

oana.balmau@cs.mcgill.ca

Data pipeline in ML: Training

Storage resources

Disk

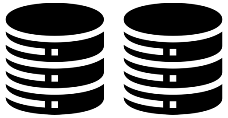

Cleaned
dataset

System Memory (DRAM)

Compute resources

CPUs

Accelerators (GPU, ASIC)

Data pipeline in ML: Training

Storage resources

Disk

Cleaned
dataset

TensorFlow
PYTORCH

load
data

System Memory (DRAM)

Cache
data

Compute resources

CPUs

Accelerators (GPU, ASIC)

Data pipeline in ML: Training

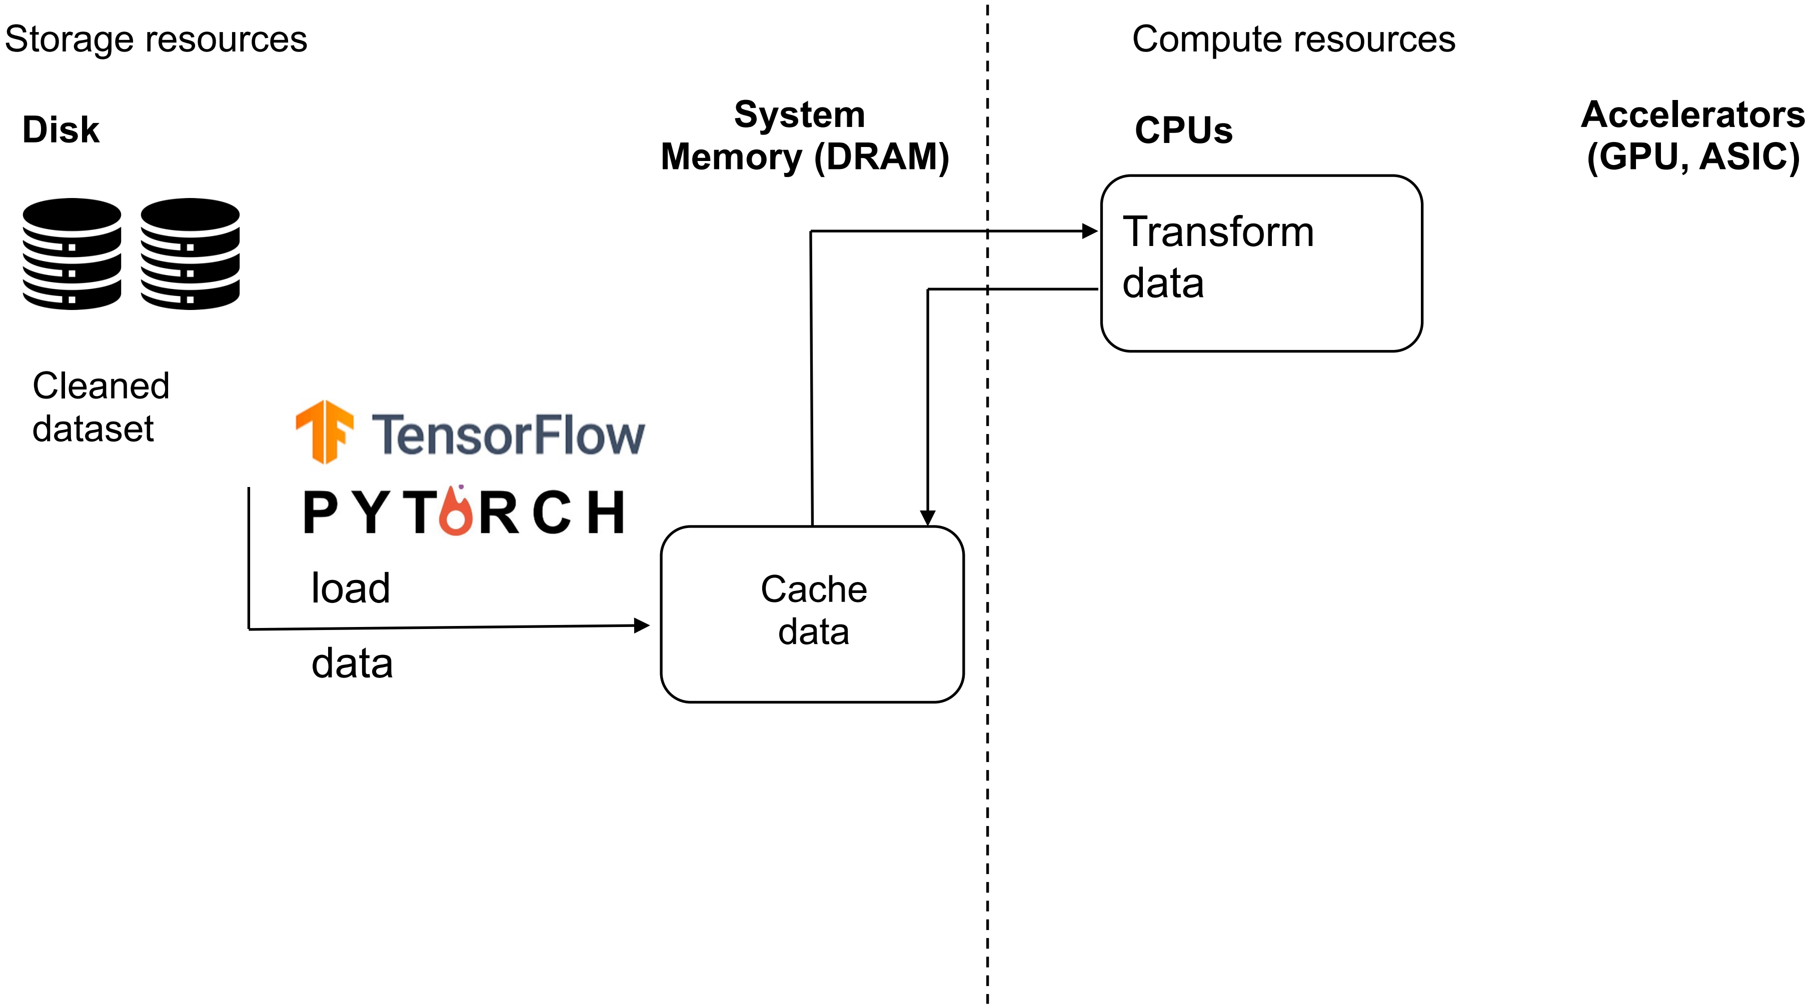

Data pipeline in ML: Training

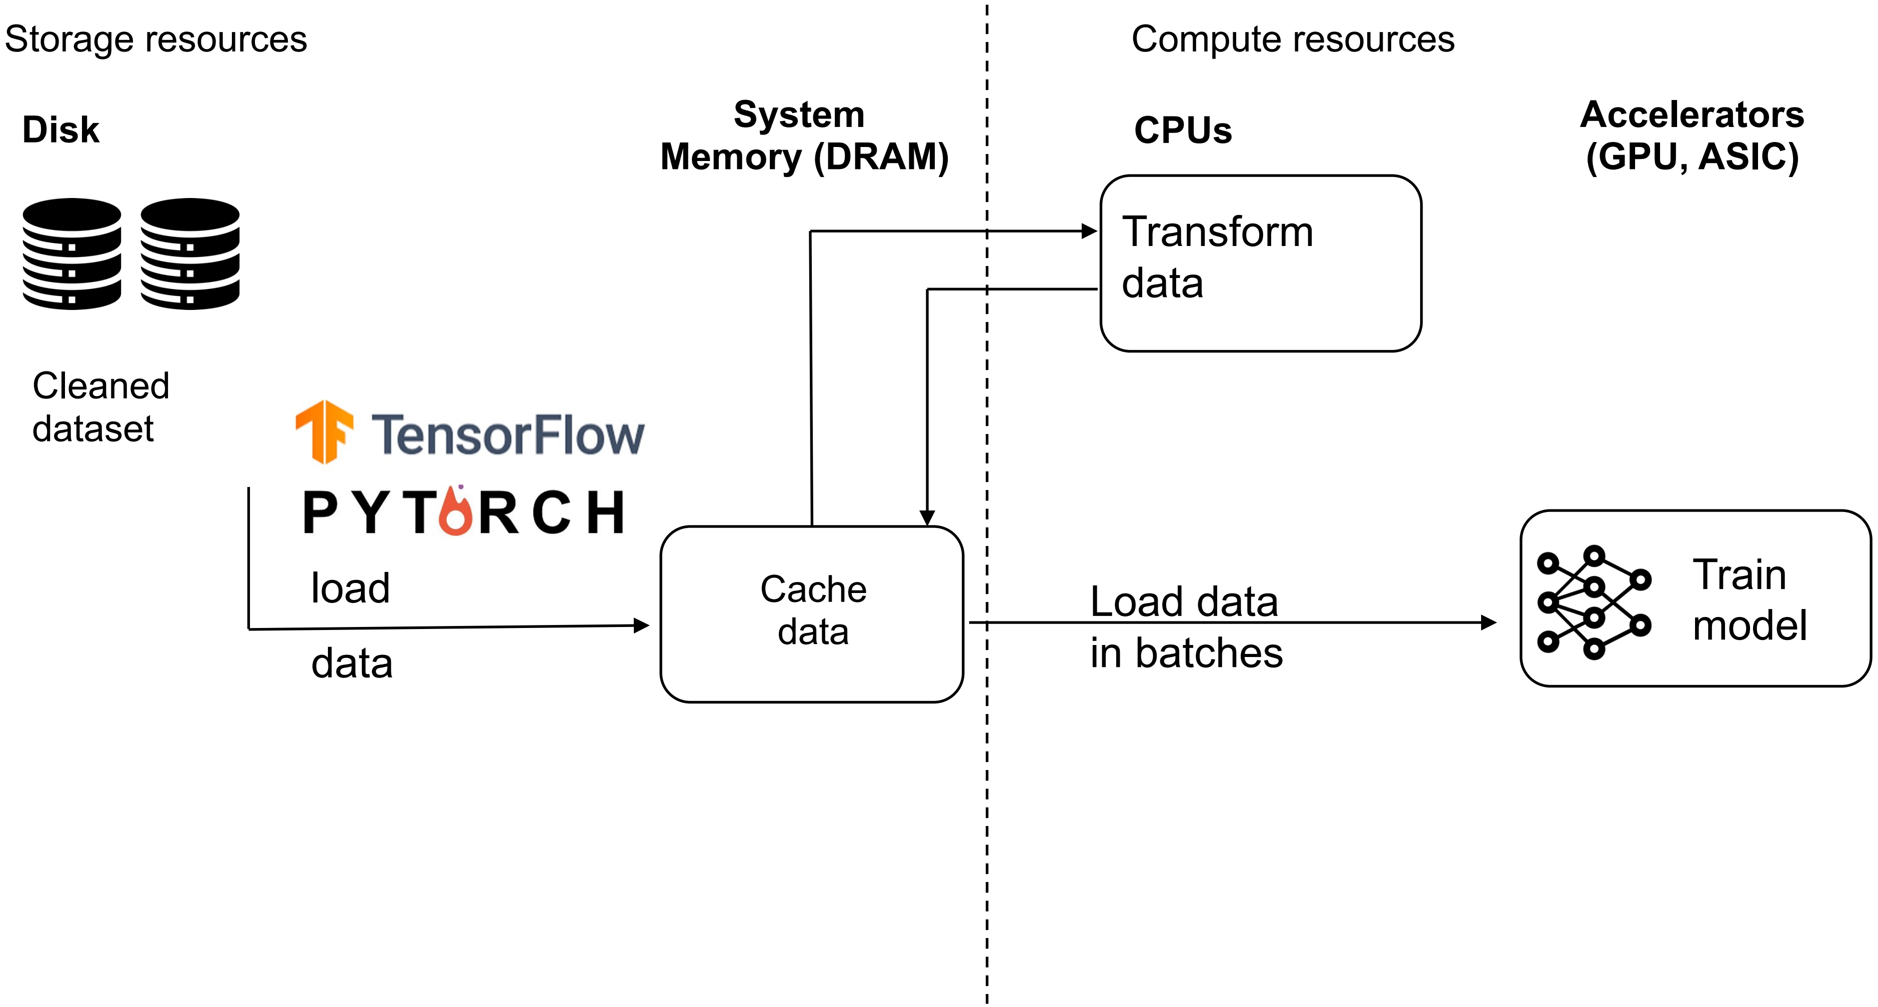

MLPerf Storage V0.5 – workloads

Workload	**Image segmentation**	**Natural language processing**	**Recommender Systems**
Model	Unet3D	BERT	DLRM
Seed data	KiTS19 Set of images	Wikipedia 2020 Text	Criteo Terabyte Click logs
Framework	Pytorch	Tensorflow	Pytorch
I/O behavior	Random access inside many small files	Sequential access of small subset of files, streamed.	Random access inside one large file

https://github.com/mlcommons/storage

Preview package

- **Single node**
- Many **simulated accelerators.**
- **Synthetic datasets** generated from real dataset seed.
- **Local storage**

oana.balmau@cs.mcgill.ca

Dataset Size to Memory Ratio is Important

Experiment setup
- DGX-1 server
  - 8 x V100 GPUs, 32GB GPU memory
  - 512GB DRAM
- Image segmentation workload:
  - Unet3D, Pytorch
  - MLPerf Training implementation
  - KiTS19 dataset

Dataset fits in system memory

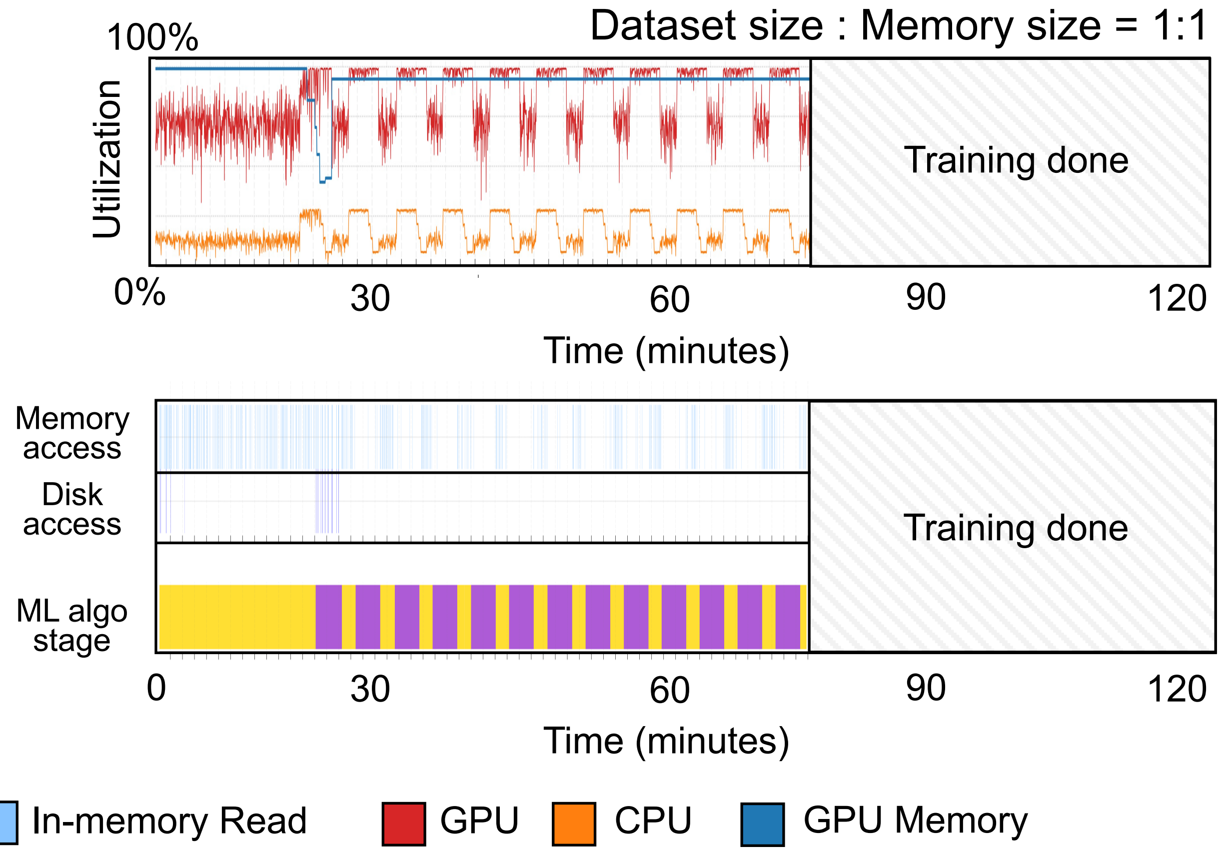

Dataset Size to Memory Ratio is Important

Experiment setup

- DGX-1 server
  - 8 x V100 GPUs, 32GB GPU memory
  - 512GB DRAM
- Image segmentation workload:
  - Unet3D, Pytorch
  - MLPerf Training implementation
  - KiTS19 dataset

High GPU utilization →

Little disk access →

Dataset fits in system memory

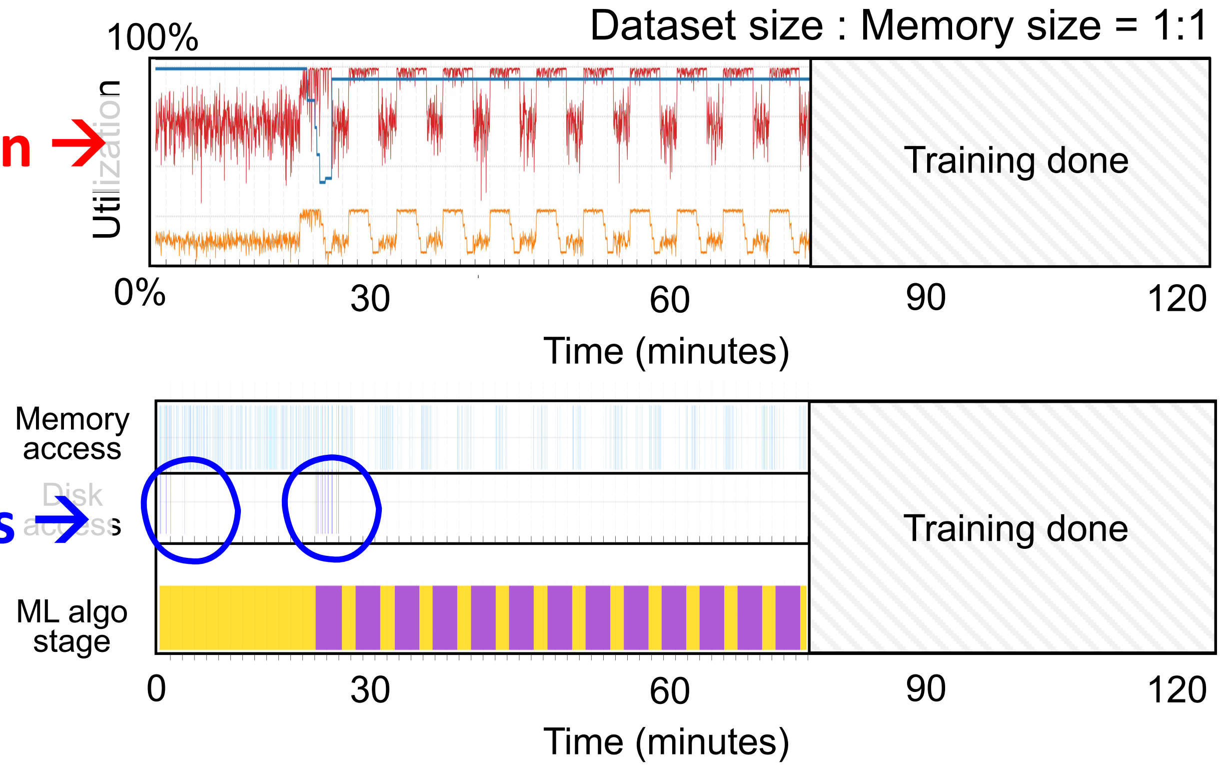

ML Training | ML Evaluation | Disk I/O Read | In-memory Read | GPU | CPU | GPU Memory

Dataset Size to Memory Ratio is Important

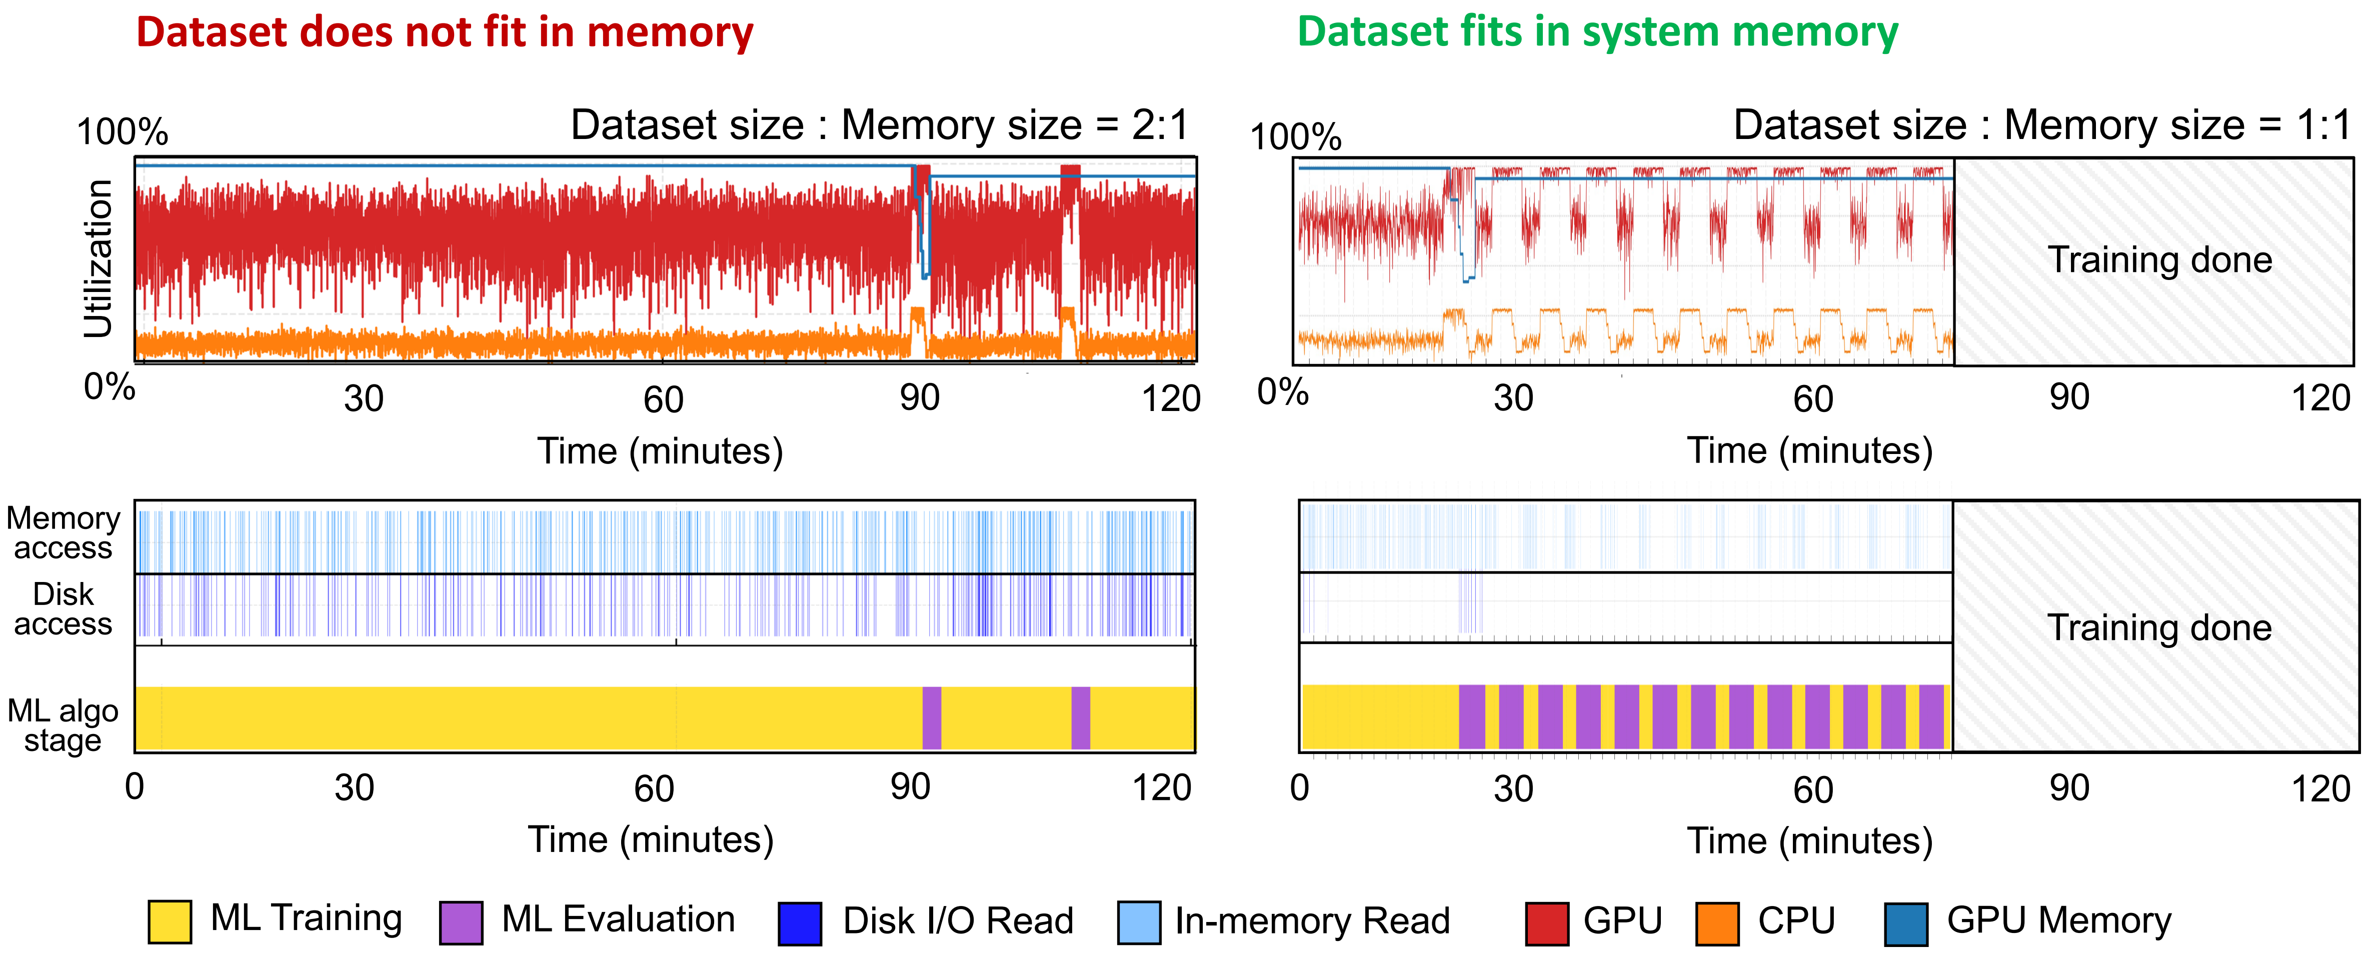

Dataset Size to Memory Ratio is Important

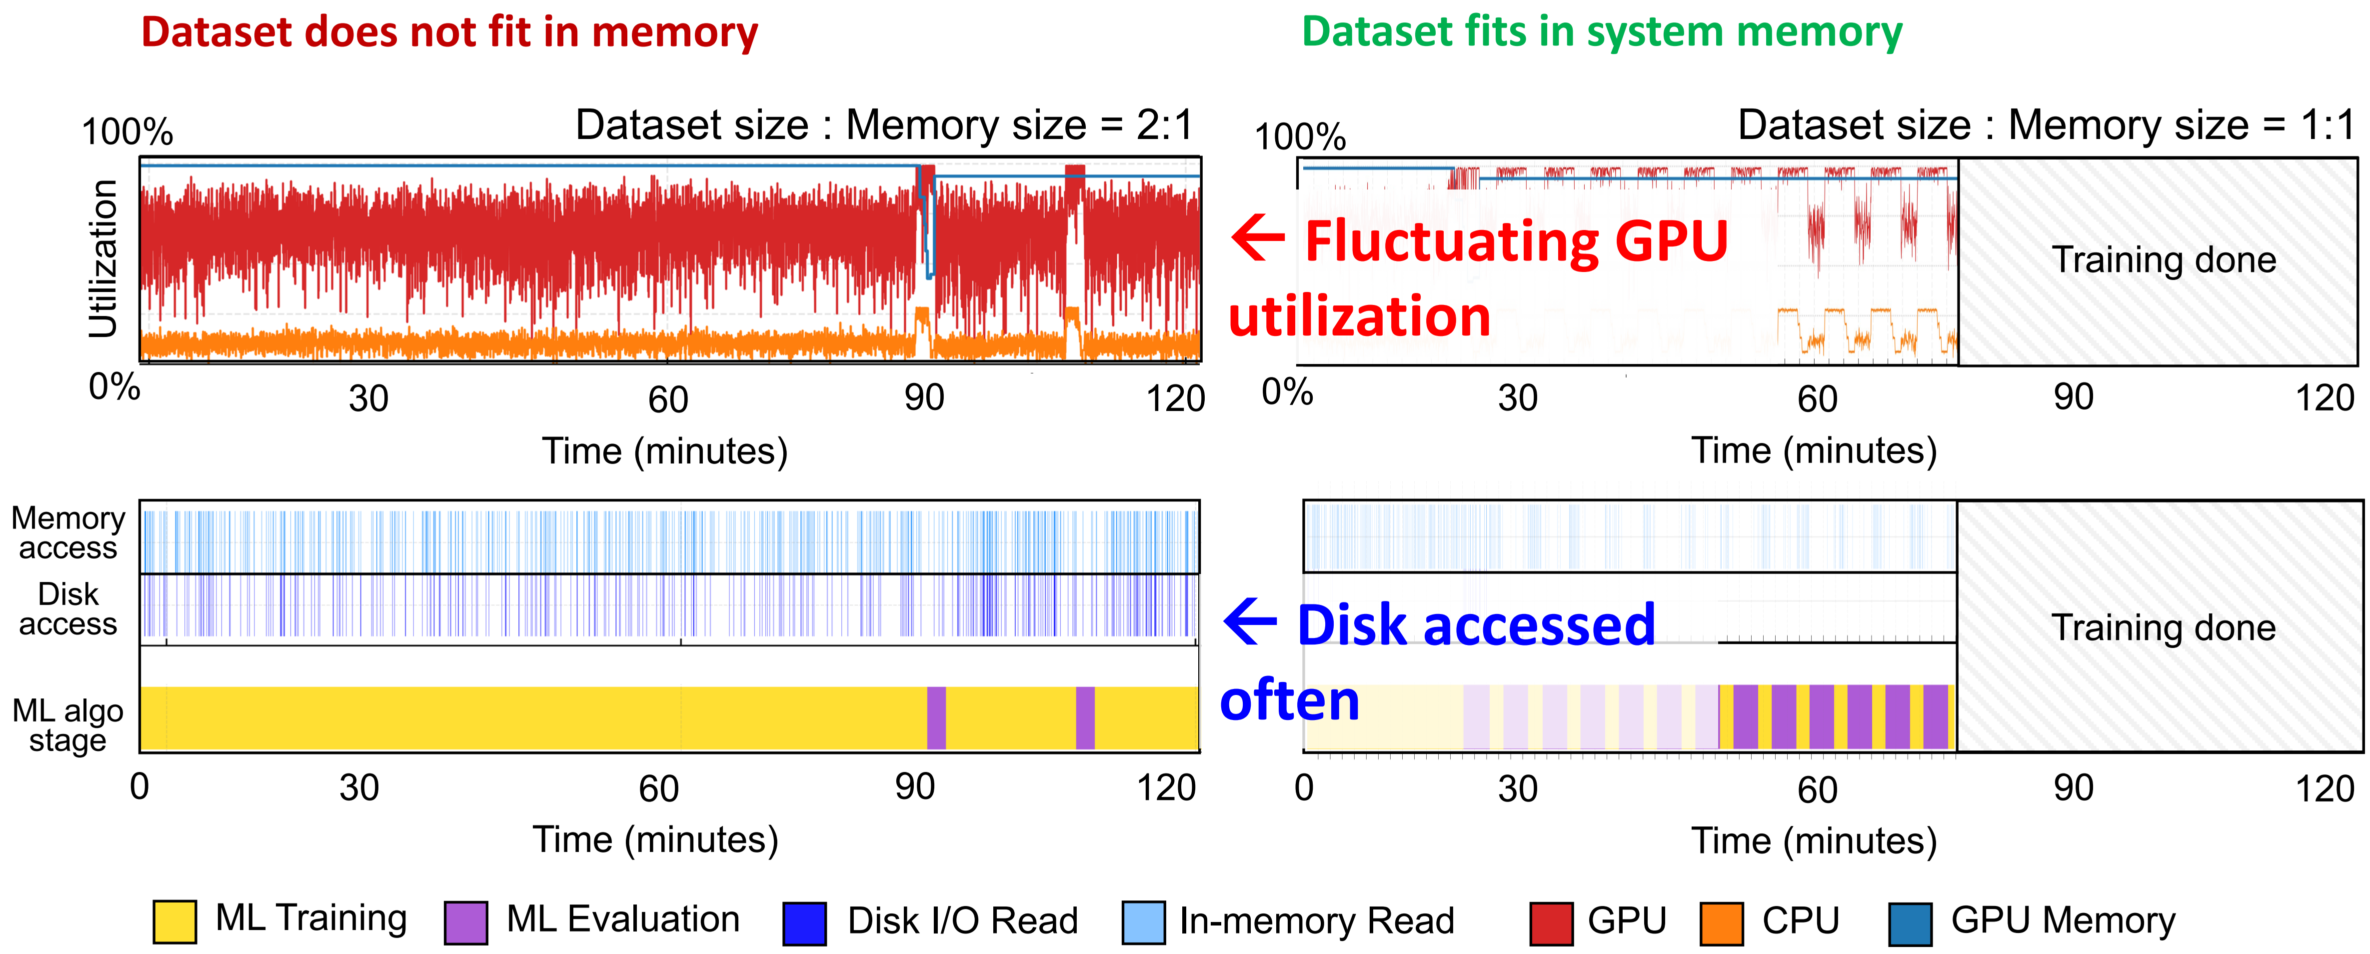

Data pipeline in MLPerf Storage benchmark

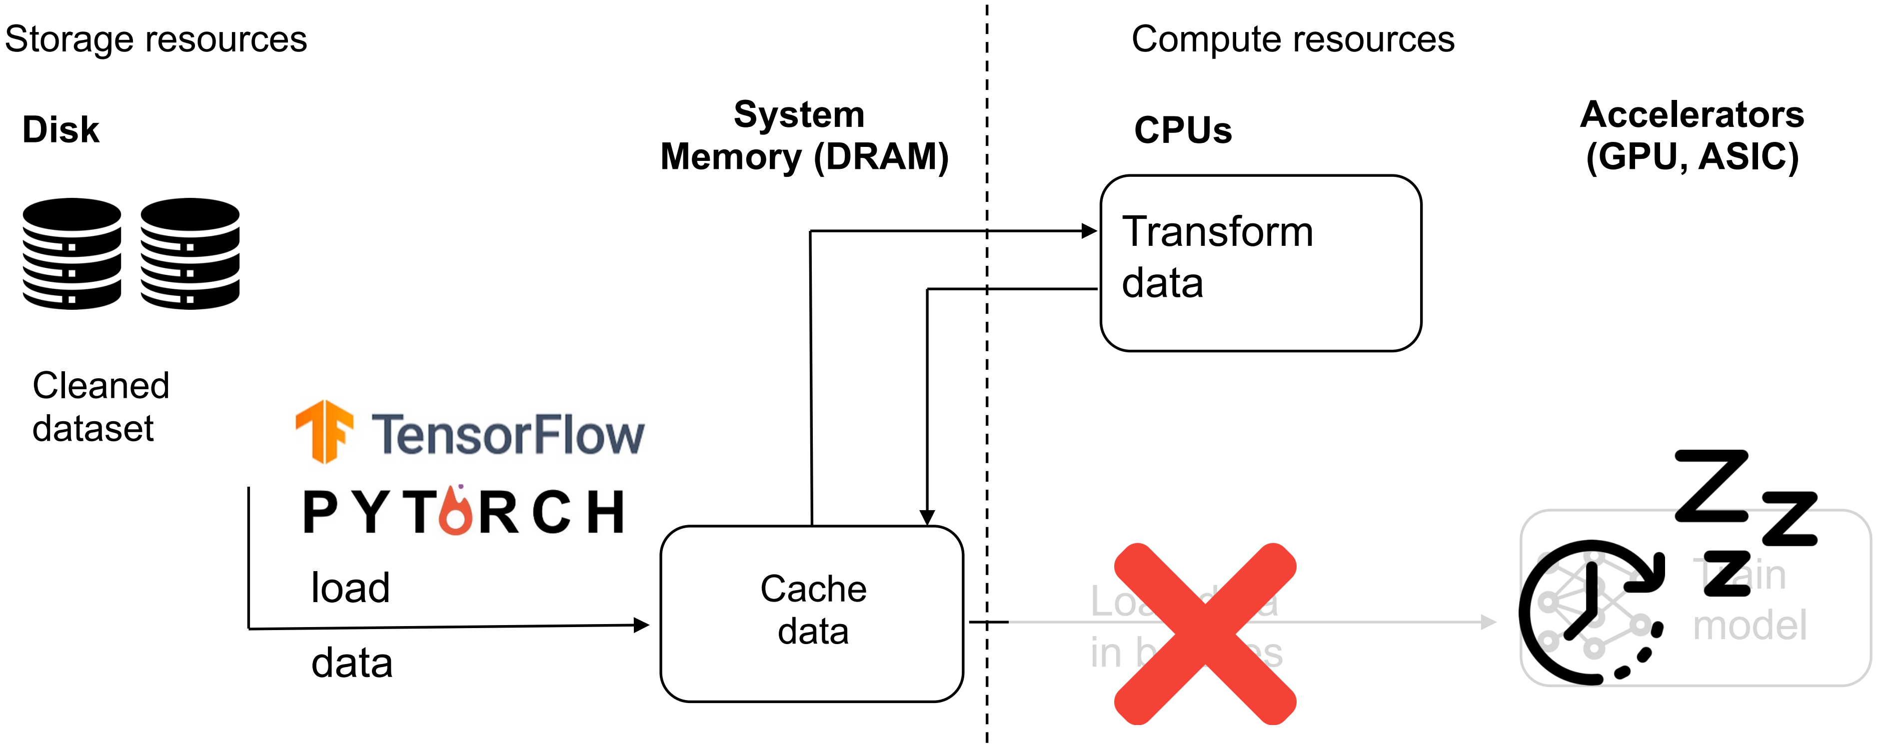

Benchmark is built as an extension of DLIO [1]

[1] H. Devarajan, H. Zheng, et al. DLIO: A Data-Centric Benchmark for Scientific Deep Learning Applications, CCGrid '21.

Data pipeline in MLPerf Storage benchmark

✓ Realistic storage settings: nothing changes in data pipeline, apart from training computation.

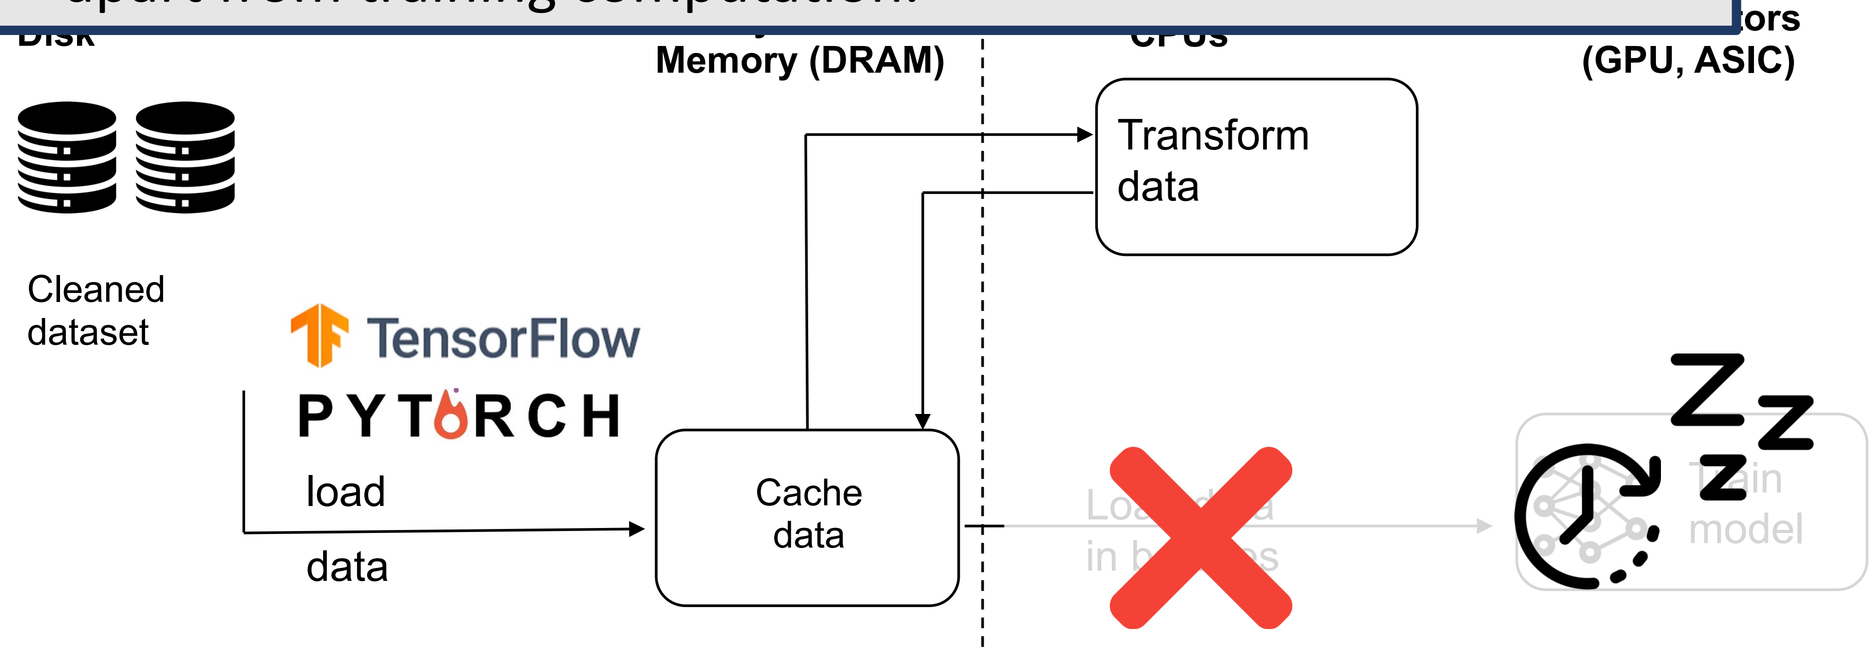

Benchmark is built as an extension of DLIO [1]

[1] H. Devarajan, H. Zheng, et al. DLIO: A Data-Centric Benchmark for Scientific Deep Learning Applications, CCGrid '21.

Data pipeline in MLPerf Storage benchmark

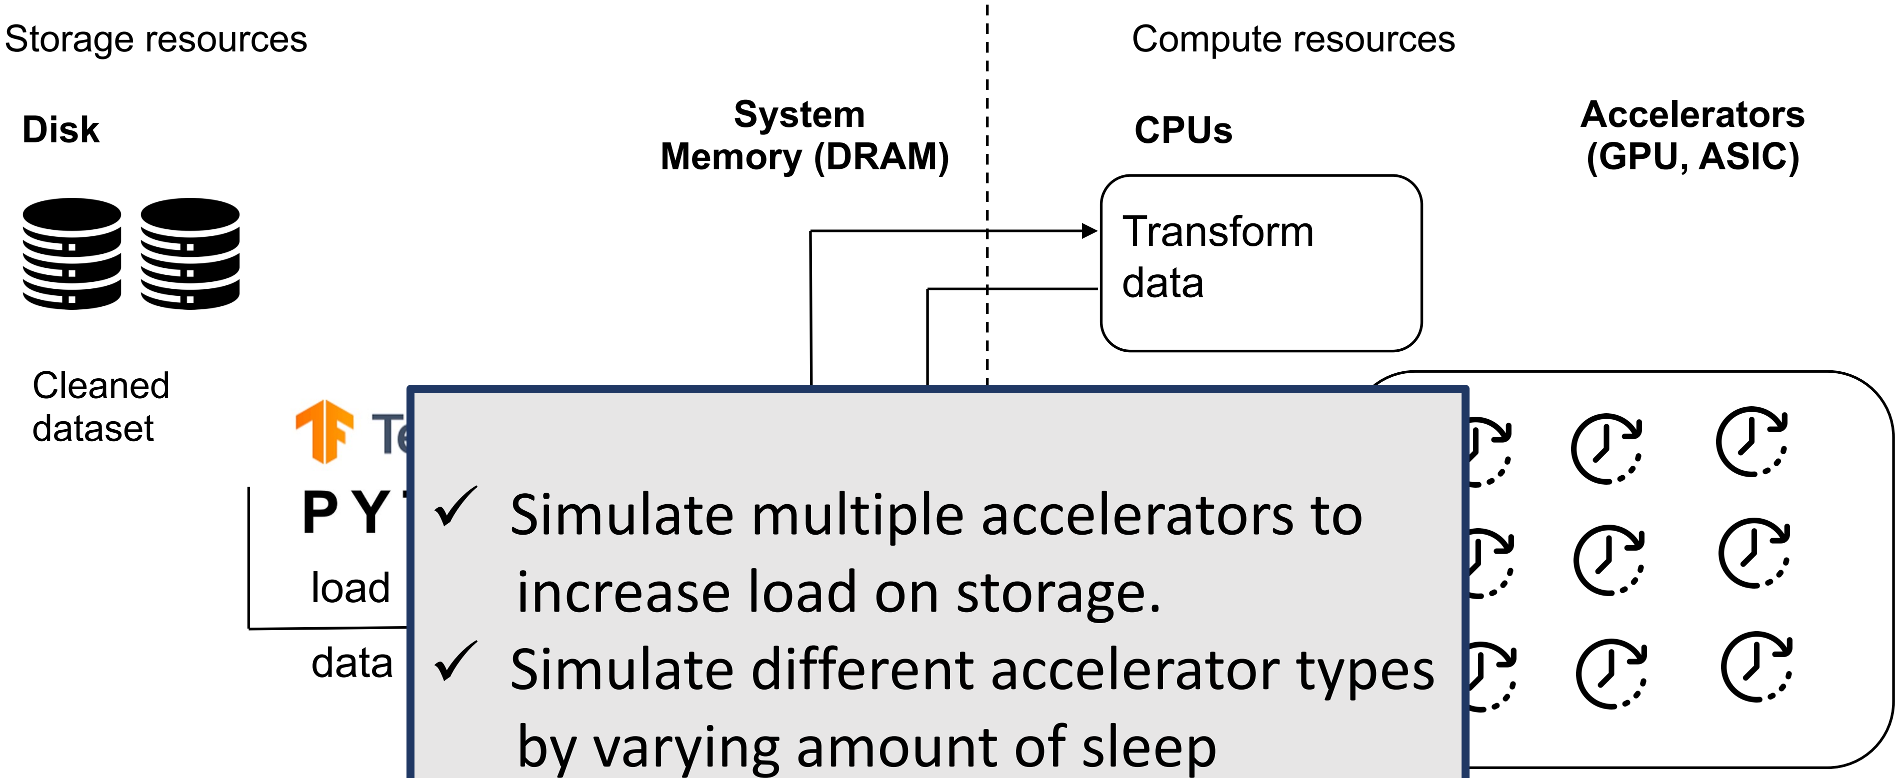

- ✓ Simulate multiple accelerators to increase load on storage.
- ✓ Simulate different accelerator types by varying amount of sleep

Benchmark is built as an extension of DLIO [1]

[1] H. Devarajan, H. Zheng, et al. DLIO: A Data-Centric Benchmark for Scientific Deep Learning Applications, CCGrid '21.

Simulating training time does not impact I/O patterns

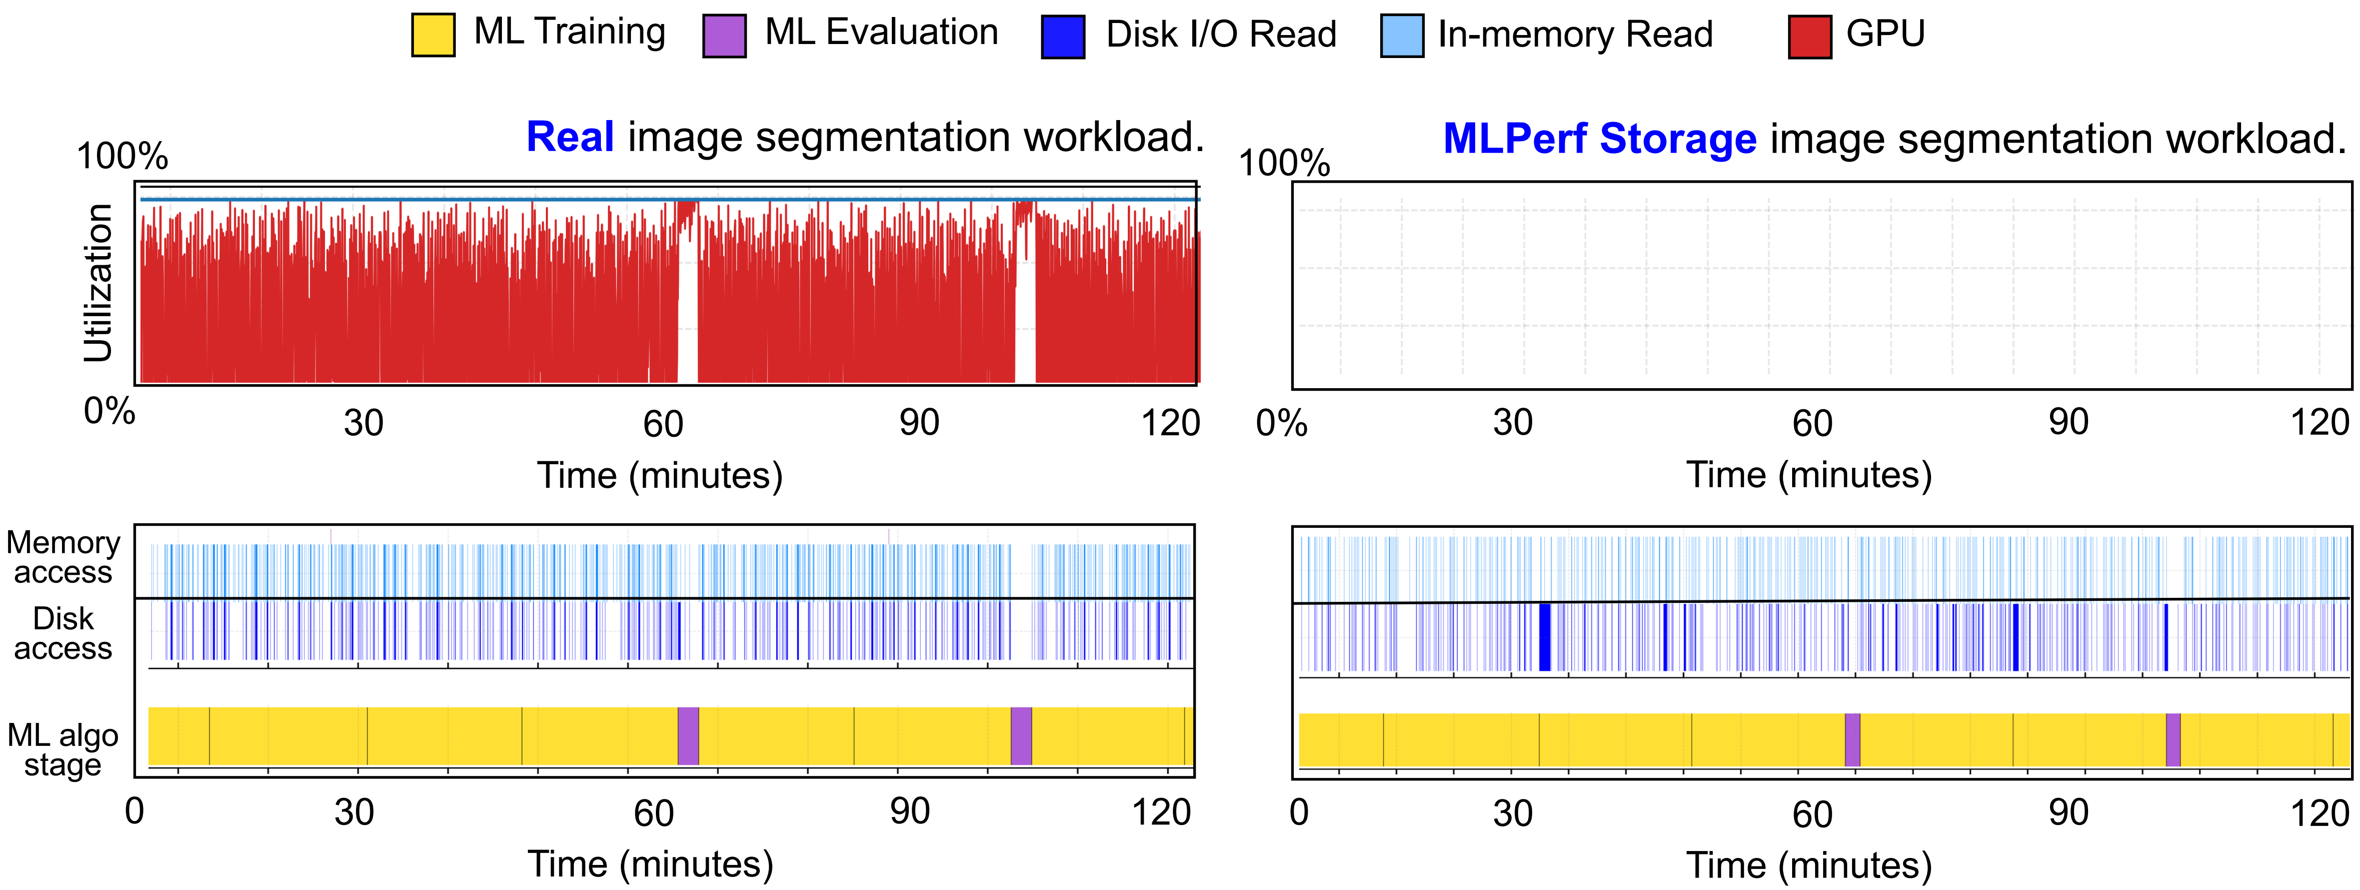

Experiment setup: DGX-1 with 8xV100 GPUs, 512GB DRAM. Dataset : KiTS19, Dataset size:Memory size ratio 2:1

Simulating training time does not impact I/O patterns

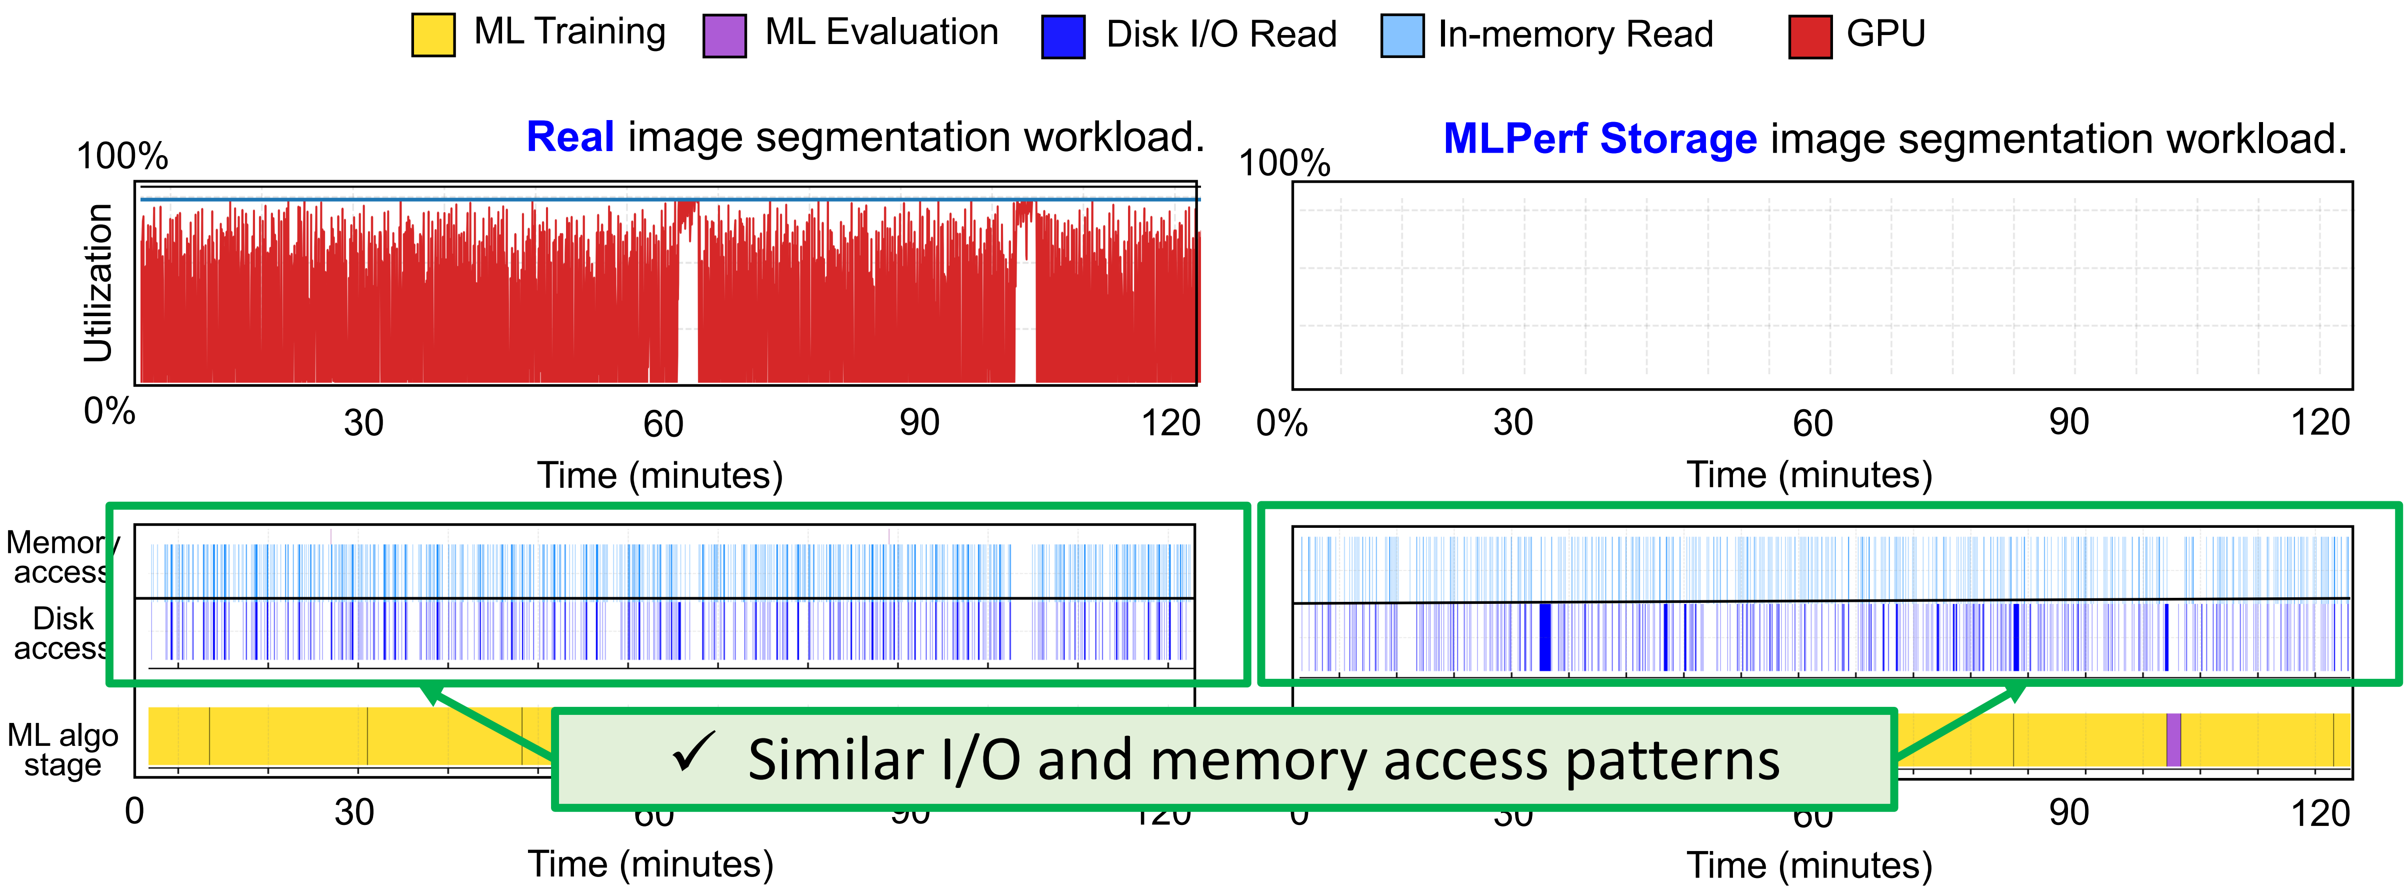

Experiment setup: DGX-1 with 8xV100 GPUs, 512GB DRAM. Dataset : KiTS19, Dataset size:Memory size ratio 2:1

Simulating training time does not impact I/O patterns

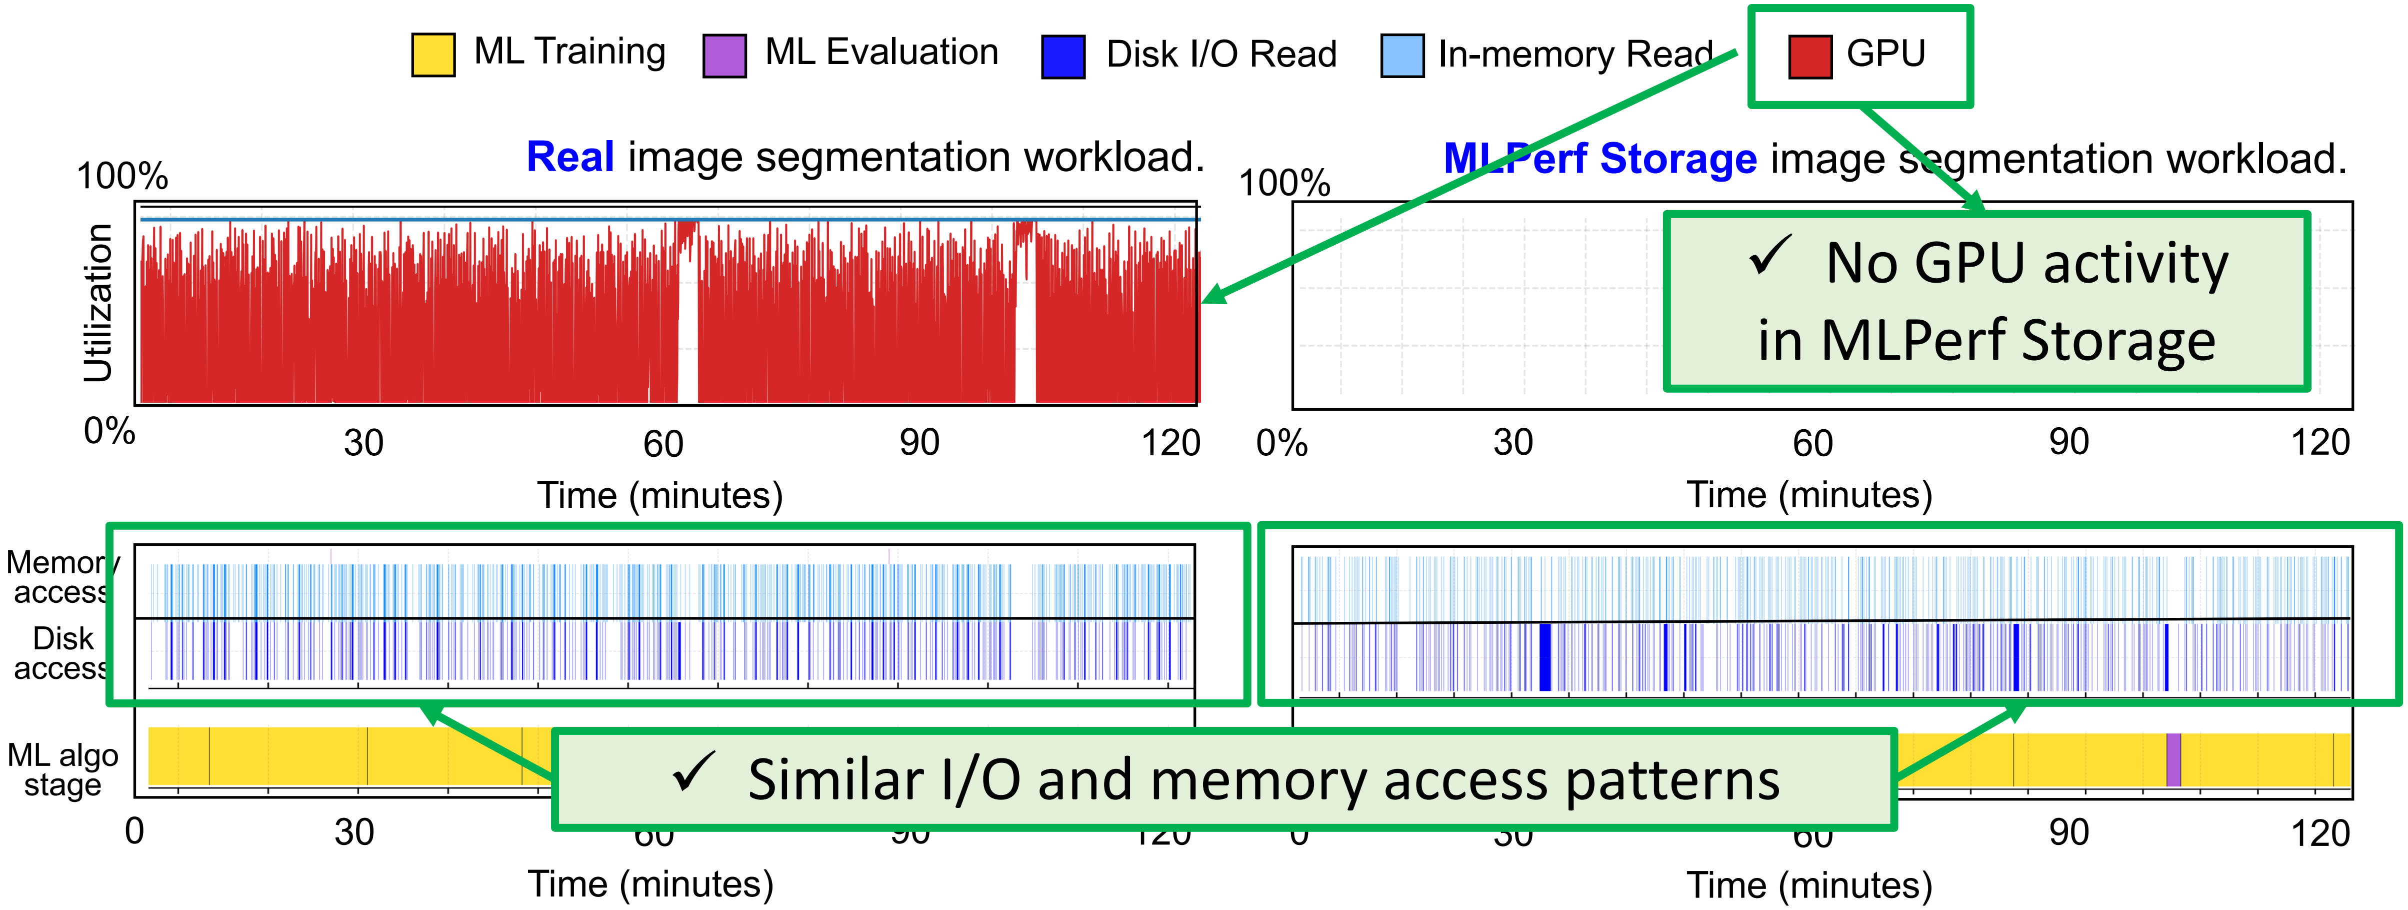

Experiment setup: DGX-1 with 8xV100 GPUs, 512GB DRAM. Dataset : KiTS19, Dataset size:Memory size ratio 2:1

Next Steps

Collect **processing times** for different accelerator types.

Open benchmark for submissions.

→ *https://github.com/mlcommons/storage/tree/v0.5-branch*

Parallelism

Trace and benchmark **ML pre-processing phase.**

Key Takeaways – MLPerf Storage

MLPerf Storage is a new benchmark

Realistic **storage** settings

No accelerators required to run

Follow MLPerf Storage repository for updates:

https://github.com/mlcommons/storage

Get involved
mlcommons.org/en/get-involved/

We appreciate your feedback

Share your thoughts
Email oana.balmau@cs.mcgill.ca

Thanks to all working group co-chairs!

Curtis Anderson
Panasas

Huihuo Zheng
Argonne National Labs

Johnu George,
Nutanix

McGill DISCS Lab

DISCS

discslab.cs.mcgill.ca
gitlab.cs.mcgill.ca/discs-lab

Postdoctoral Researcher

Dr. Stella Bitchebe

PhD Candidates:

Nelson Bore

Masters Students

Sebastian Rolon

Loïc Ho-Von

Aayush Kapur

Aidan Goldfarb

Rahma Nouaji

Undergraduate Students

Zachary Doucet

Christian Zhao

Zhongjie Wu

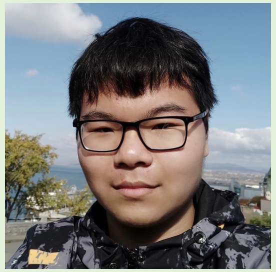

Jiaxuan Chen

Changjun Zhou

Olivier Michaud