Smart

Fast

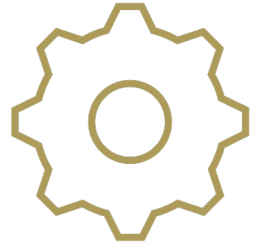

Systems

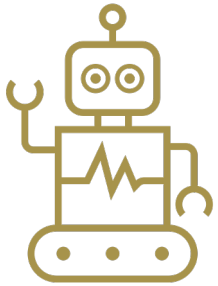

Machine Learning

Computer Vision

Building Smart and Fast Systems using Machine Learning and Computer Vision.

Thaleia Dimitra Doudali

Assistant Research Professor @IMDEA Software Institute

About Me

Born and raised in Greece.

2015

Undergrad in ECE at NTUA, Athens, Greece.

2021

PhD in CS at Georgia Tech, Atlanta, USA.

Advised by Ada Gavrilovska.

Start: October 2021

Assistant Professor at IMDEA, Madrid, Spain.

About My Research

I build software systems that manage heterogeneous hardware resources using machine learning and computer vision.

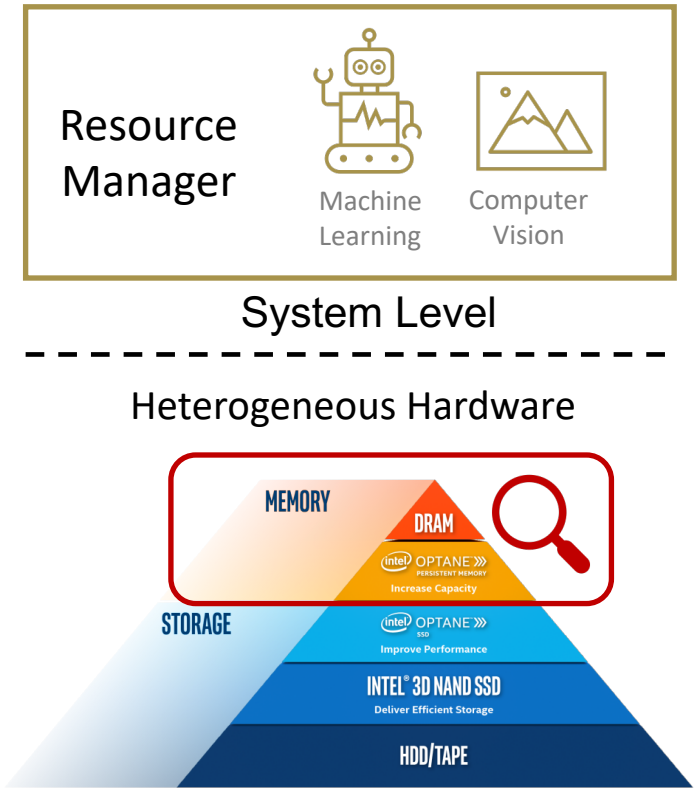

The Era of Data

"More than **65 ZB** of data will be created, captured, copied, and consumed in the world this year."

Source: International Data Corporation, March 2021.

Exploded Data Sizes

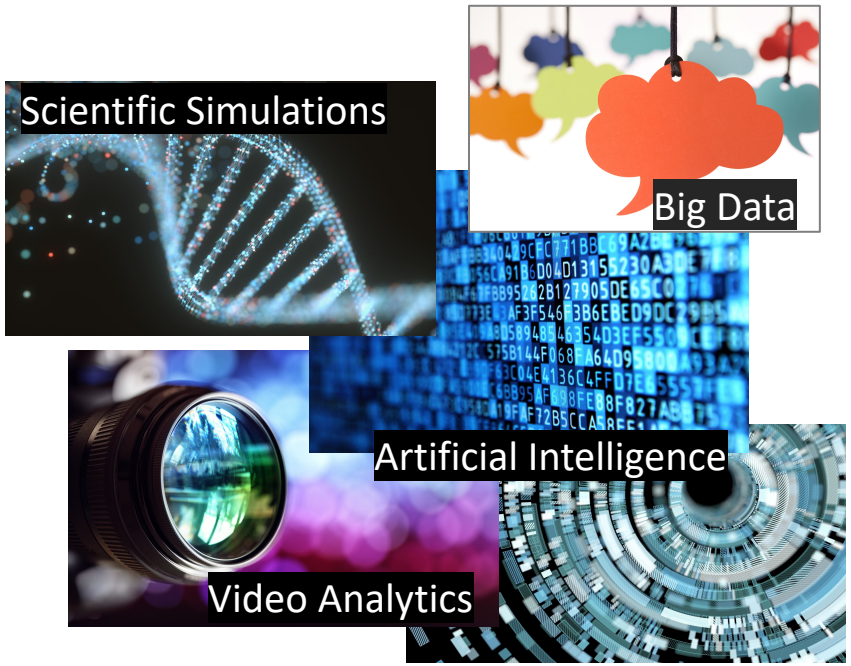

Data Analytics Pipeline

ZBs of data

Need for speed and massive storage capacities!

The Era of Heterogeneous Hardware

Compute Acceleration

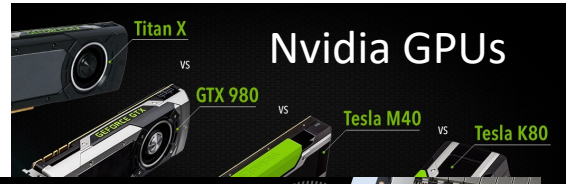

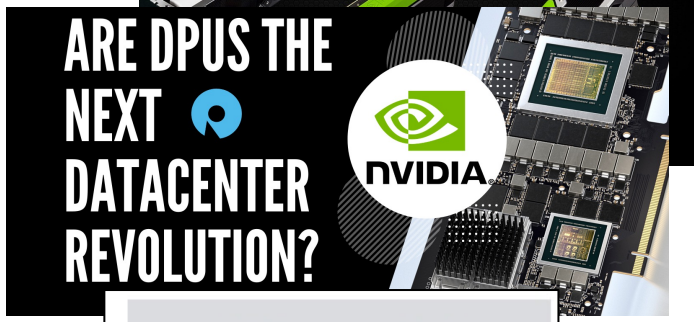

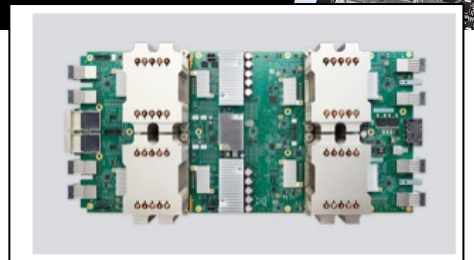

Cloud TPU v2

180 teraflops

64 GB High Bandwidth Memory (HBM)

Google

Data Storage Acceleration

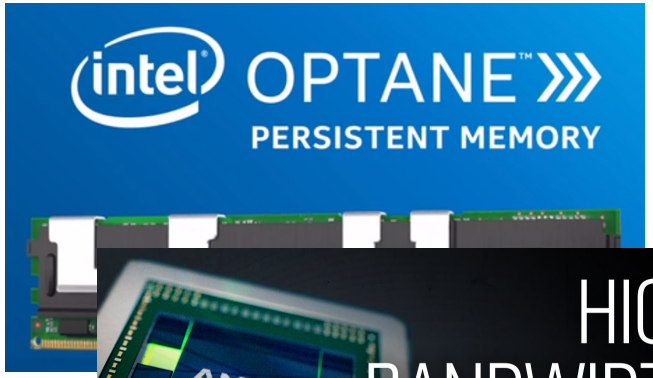

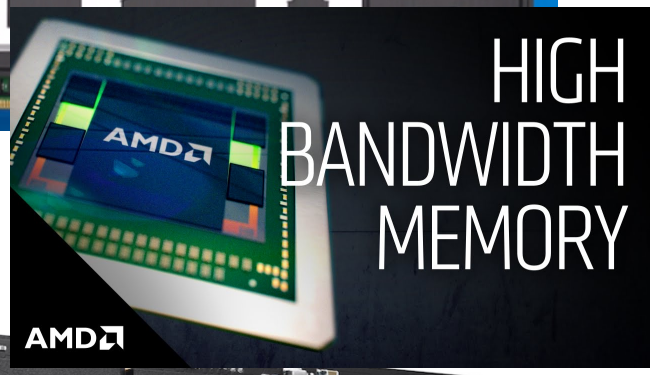

Network Acceleration

Mellanox Innova™-2 Flex
Open Programmable SmartNIC

Mellanox Innova™ 2 FLEX

Interconnection Standards

CXL Compute Express Link™

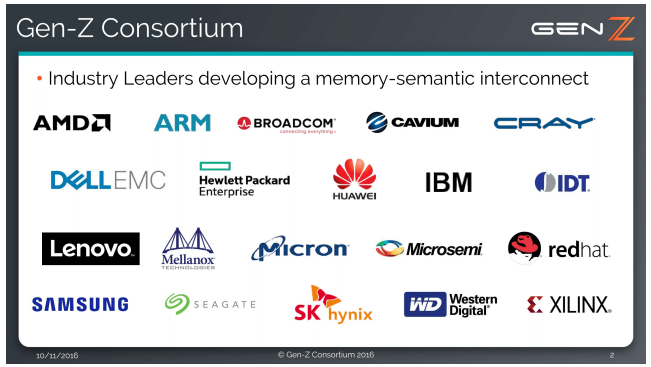

Heterogeneity Across Computing Platforms

Supercomputers

Exascale Era

HPCwire

Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

- Home
- Technologies
- Sectors

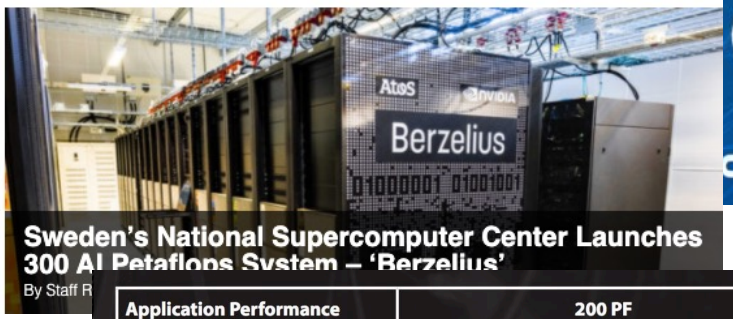

Sweden's National Supercomputer Center Launches 300 AI Petaflops System – 'Berzelius'

By Staff R

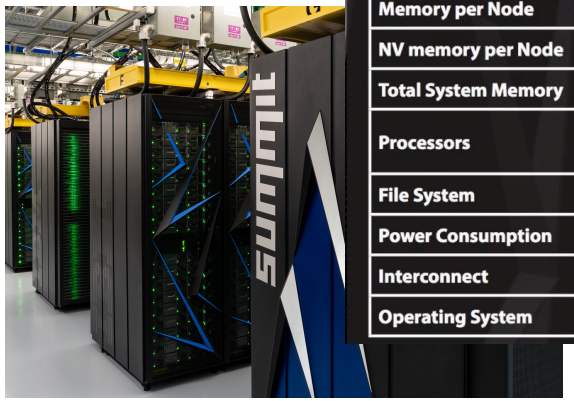

Application Performance	200 PF
Number of Nodes	4,608
Node performance	42 TF
Memory per Node	512 GB DDR4 + 96 GB HBM2
NV memory per Node	1600 GB
Total System Memory	>10 PB DDR4 + HBM2 + Non-volatile
Processors	2 IBM POWER9™ 9,216 CPUs 6 NVIDIA Volta™ 27,648 GPUs
File System	250 PB, 2.5 TB/s, GPFS™
Power Consumption	13 MW
Interconnect	Mellanox EDR 100G InfiniBand
Operating System	Red Hat Enterprise Linux (RHEL) version 7.4

Datacenters

Available first on Google Cloud: Intel Optane DC Persistent Memory

A2 VMs now GA—the largest GPU cloud instances with NVIDIA A100 GPUs

Personal Devices

Heterogeneity Trade-offs

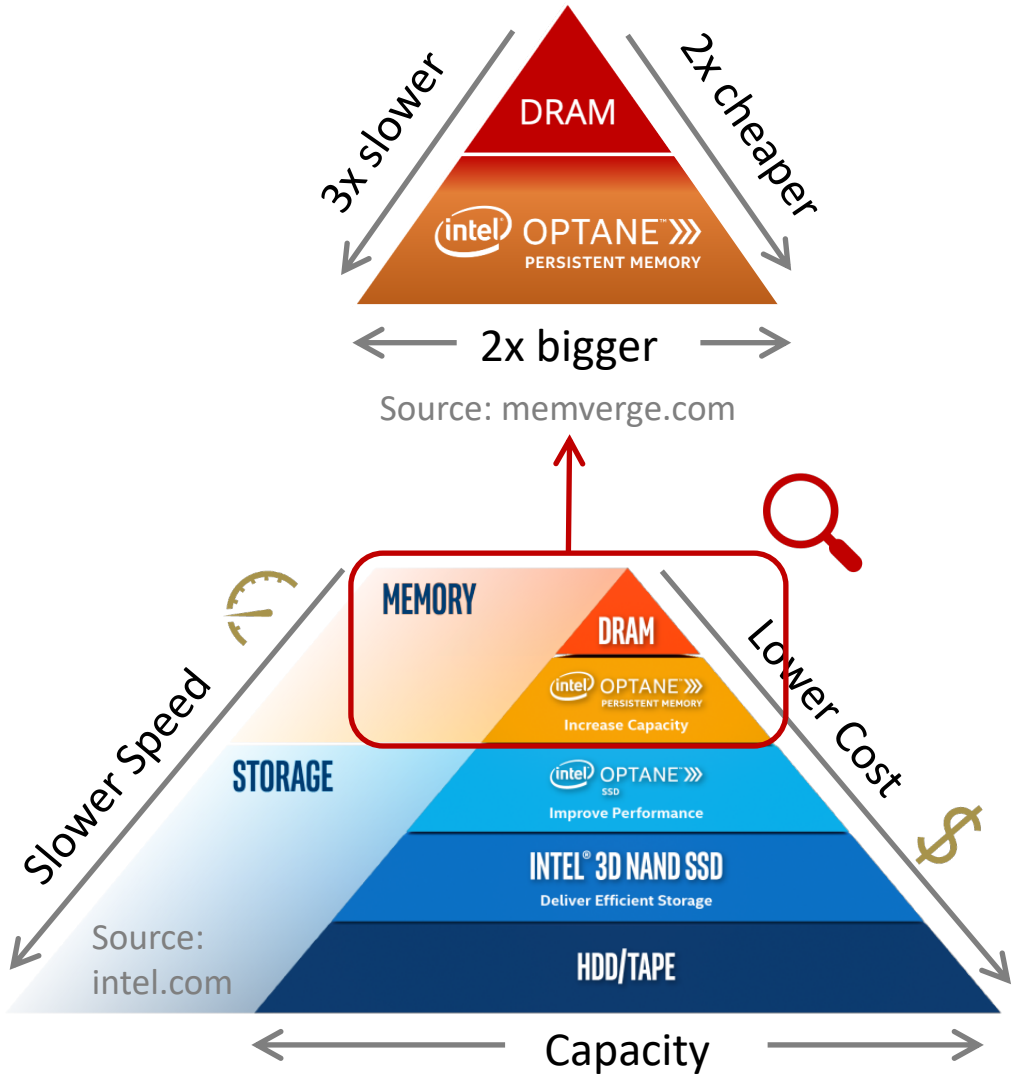

Characteristic	Technology	Hardware Vendors
Low Latency	MRAM	
High Bandwidth	HBM	 
Persistence	PMEM	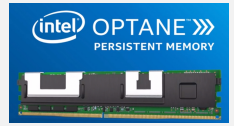

Examples of other heterogeneous memory technologies.

Building Software to Maximize the Hardware Efficiency

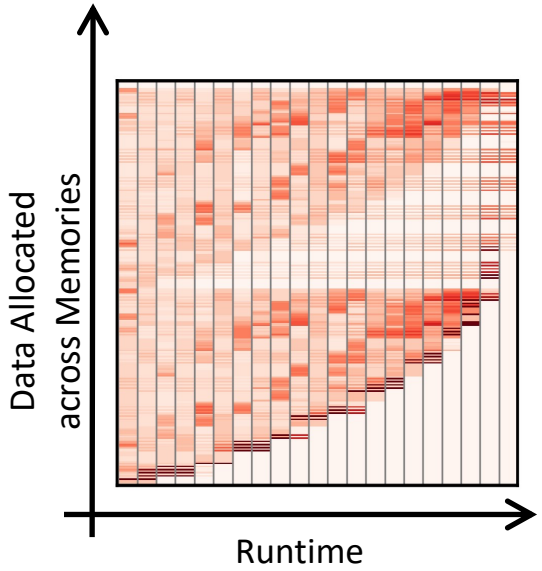

Dynamic Data Movements!

It is a **complex decision mix** to manage the data allocated across memories.

E.g., Which / How much / Where / When to move data?

Application Level

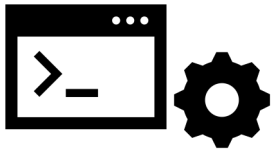

Resource Management Software

System Level

Hardware Level

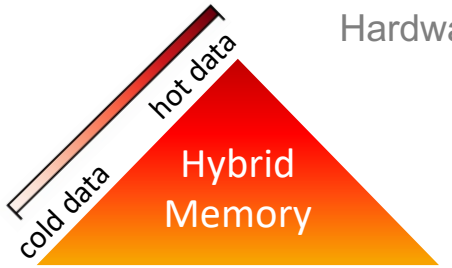

Why do we need smarter and faster systems?

Application data sizes

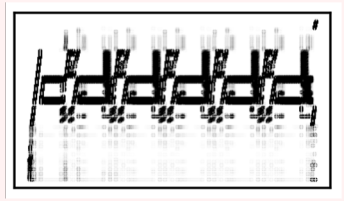

Complex data access patterns

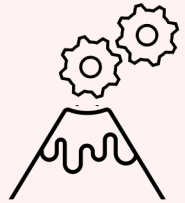

Exploded system parameter space

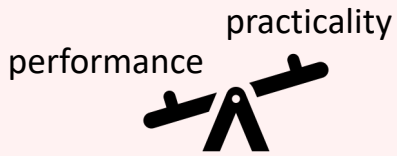

Hard to balance

Talk Outline

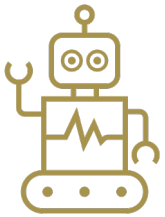

Building *Smart* Systems
Foundations for practical Machine Learning (ML)-based Management

Building *Fast* Systems
Reducing ML-based Management Overheads with Visualization

Open Research Questions

Talk Outline

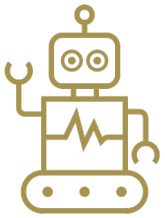

Building *Smart* Systems
Foundations for practical Machine Learning (ML)-based Management

Building *Fast* Systems
Reducing ML-based Management Overheads with Visualization

Open Research Questions

The Vision

ML-augmented heterogeneous resource manager.

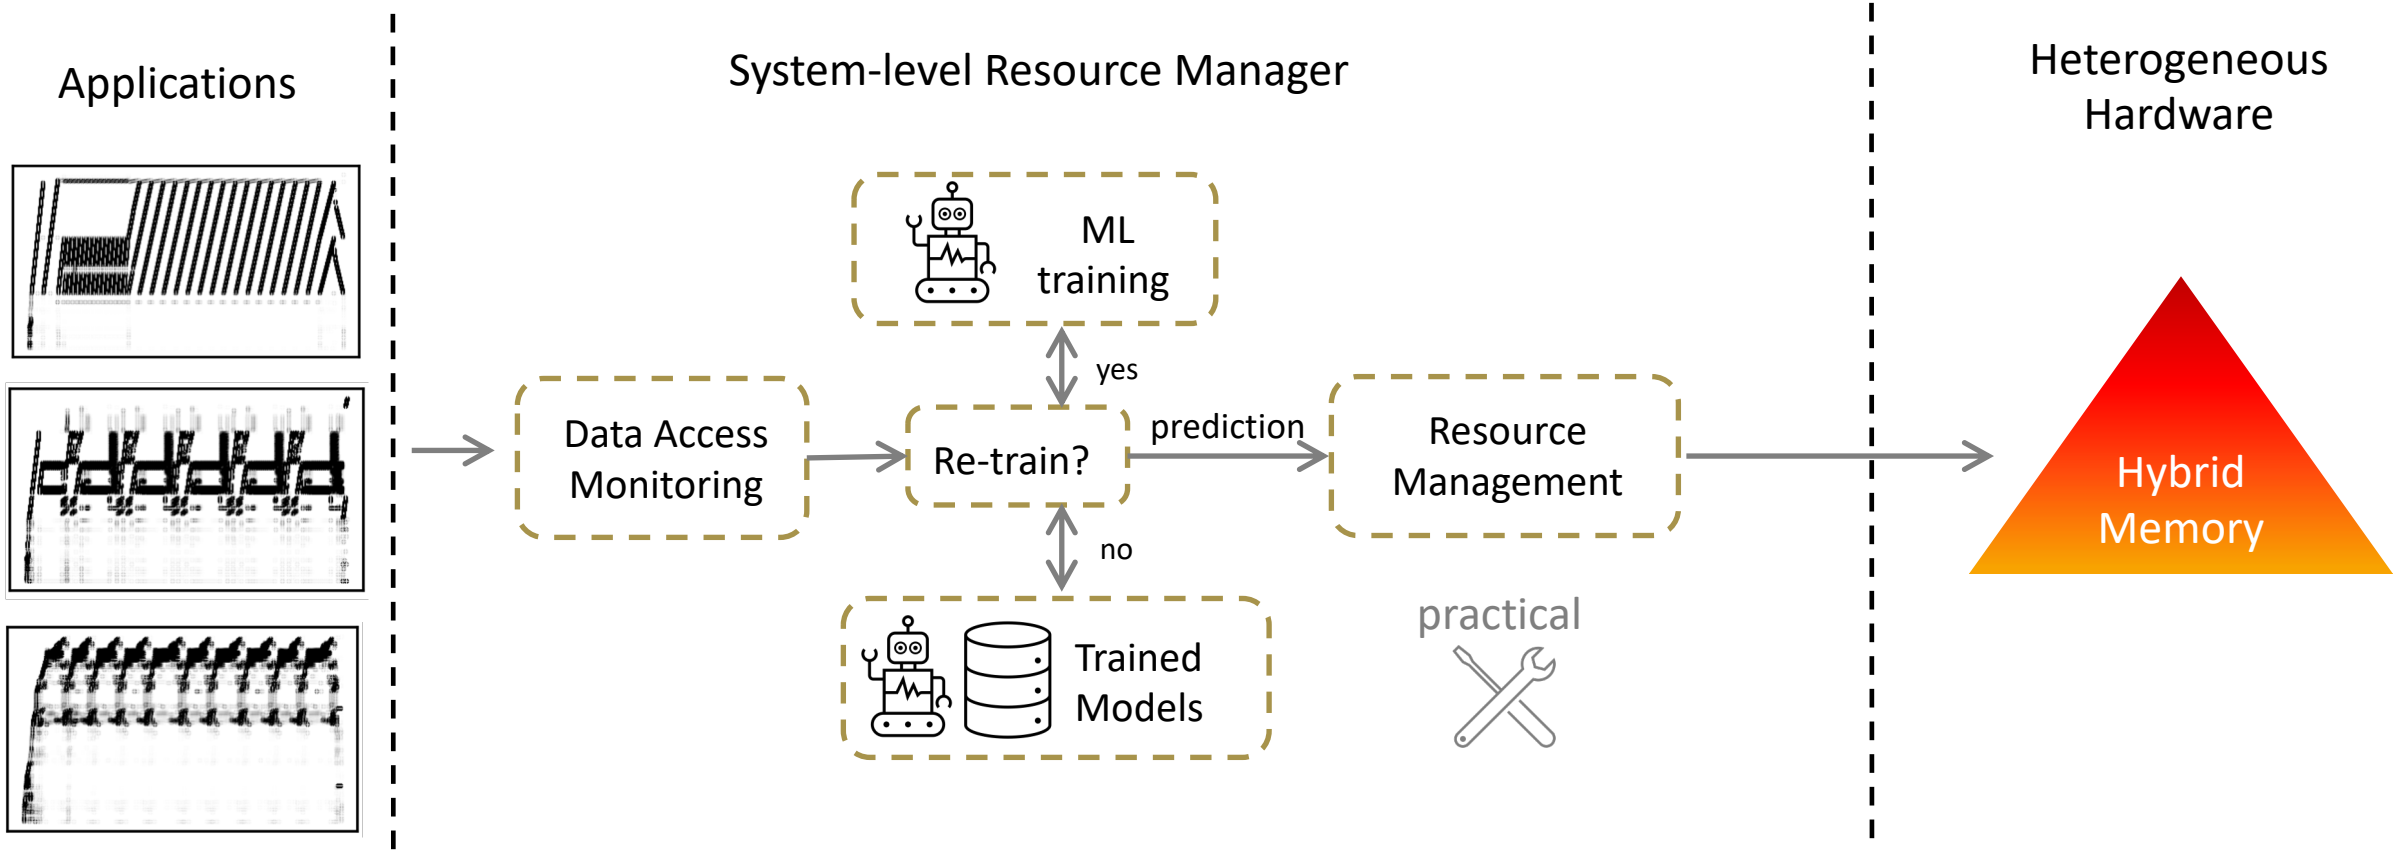

Contributions Towards the Vision

Laying the grounds for the *practical* integration of ML.

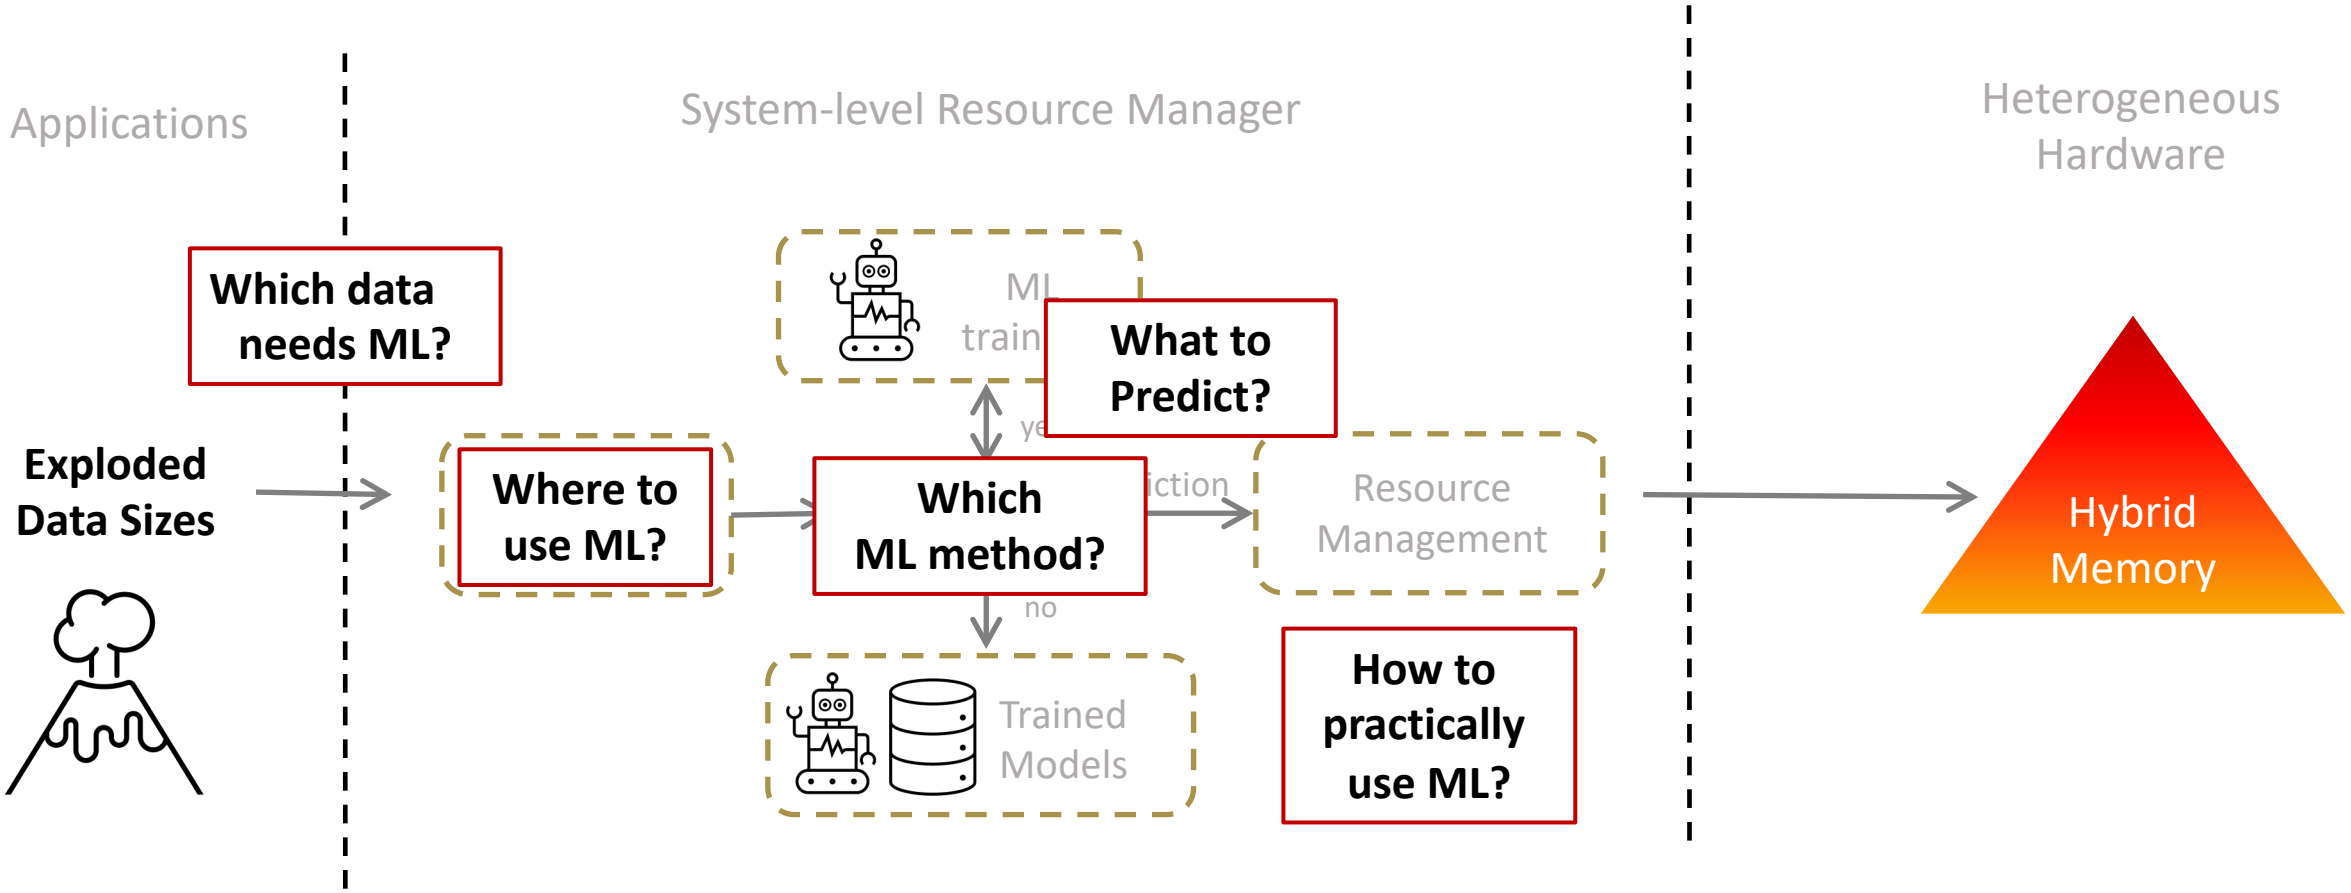

System design of **Kleio**

Kleio is a hybrid memory page scheduler with machine intelligence. [HPDC '19]

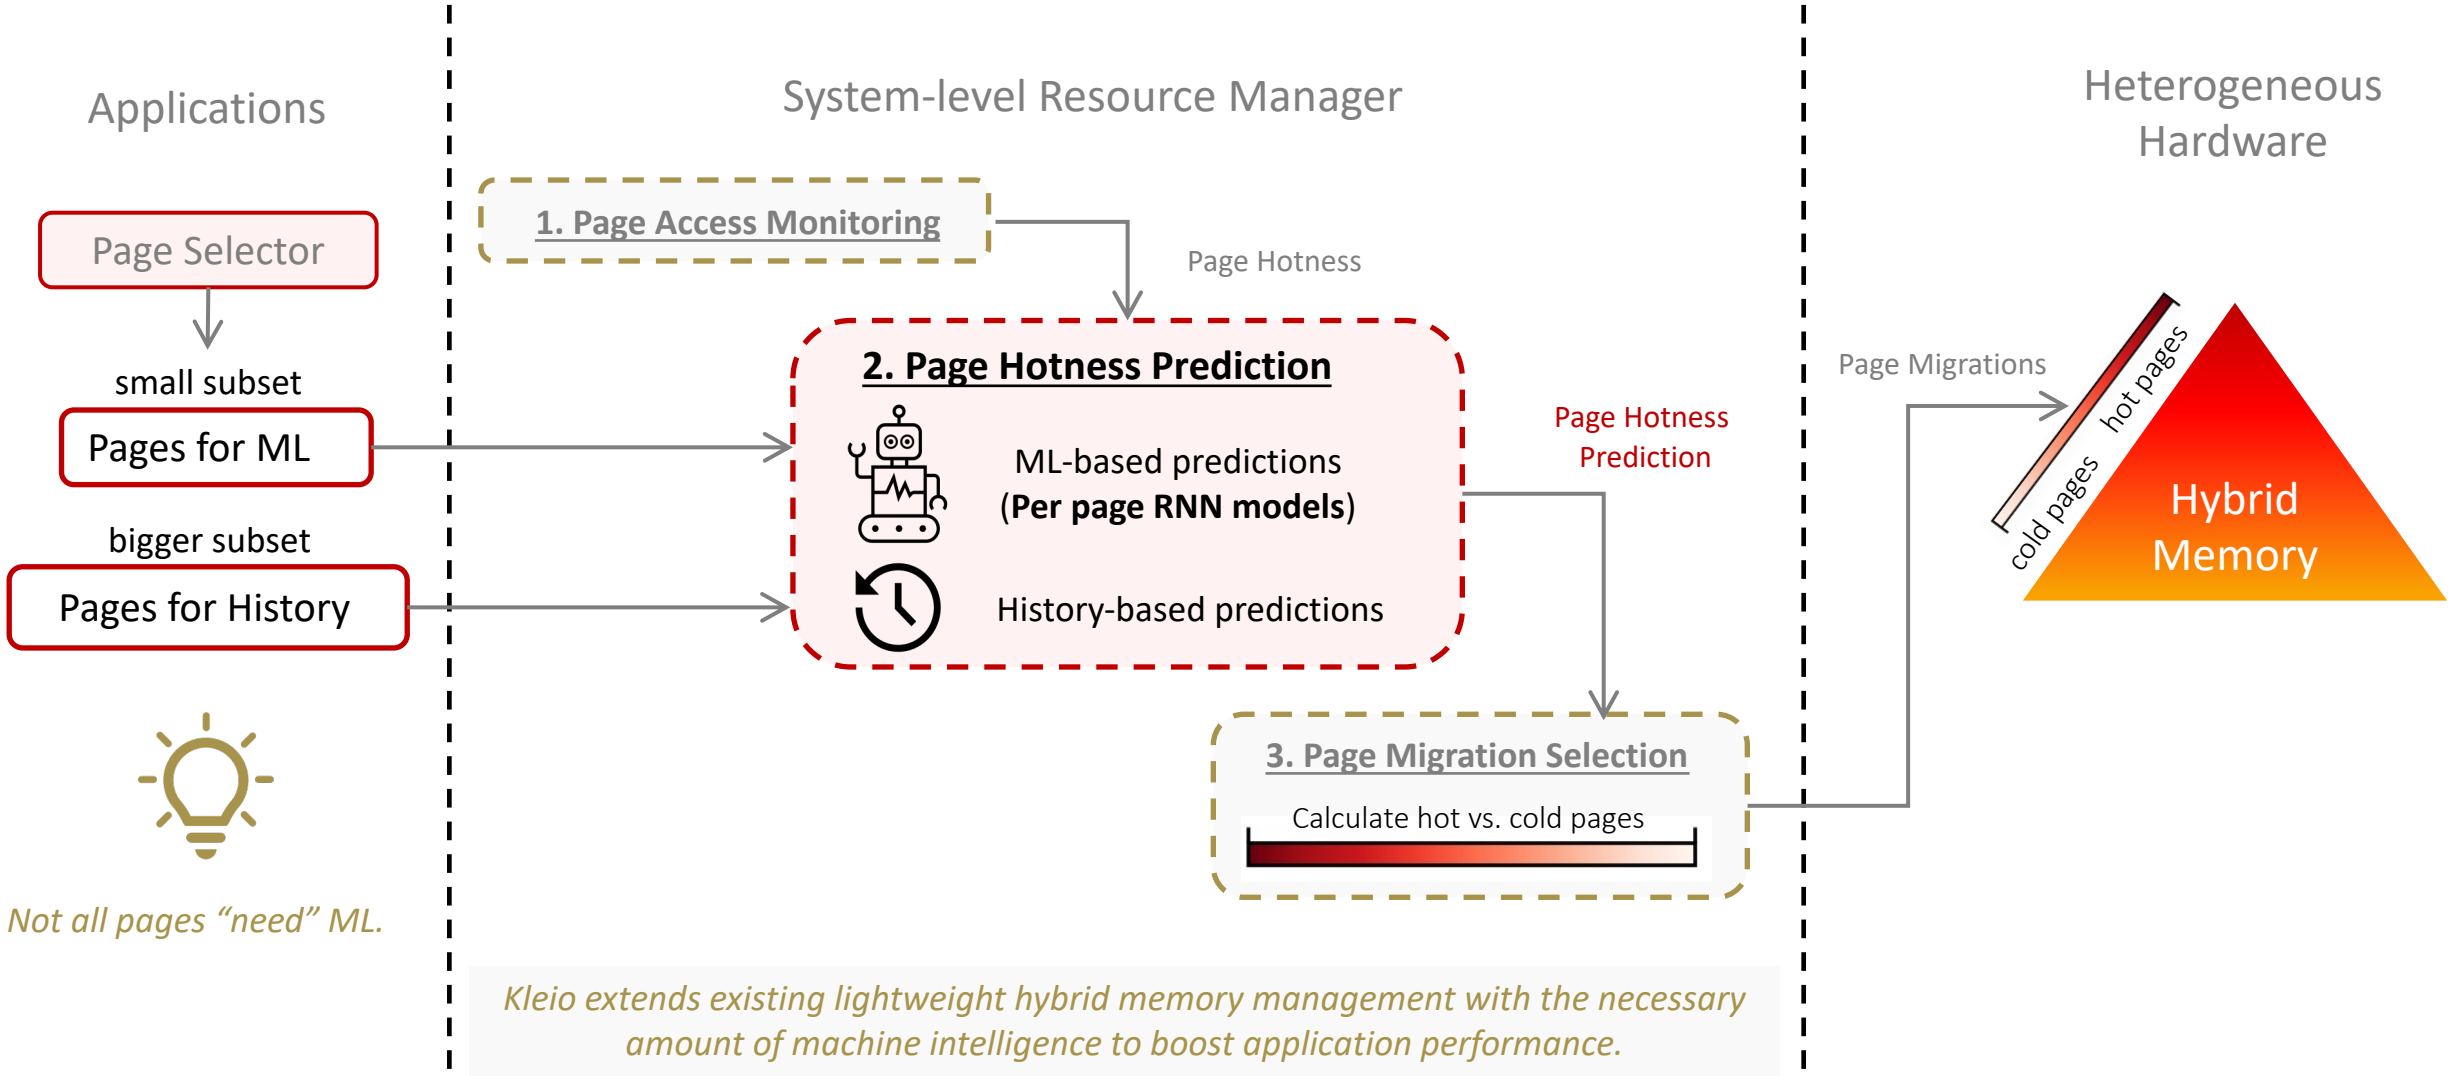

The Key to Success

Selecting a small page subset for ML-based management.

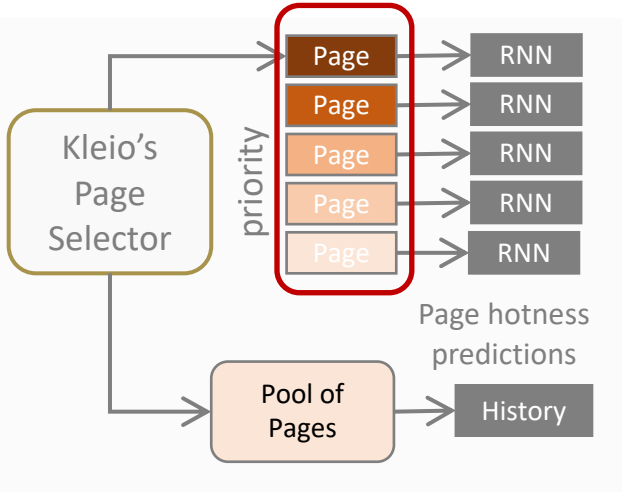

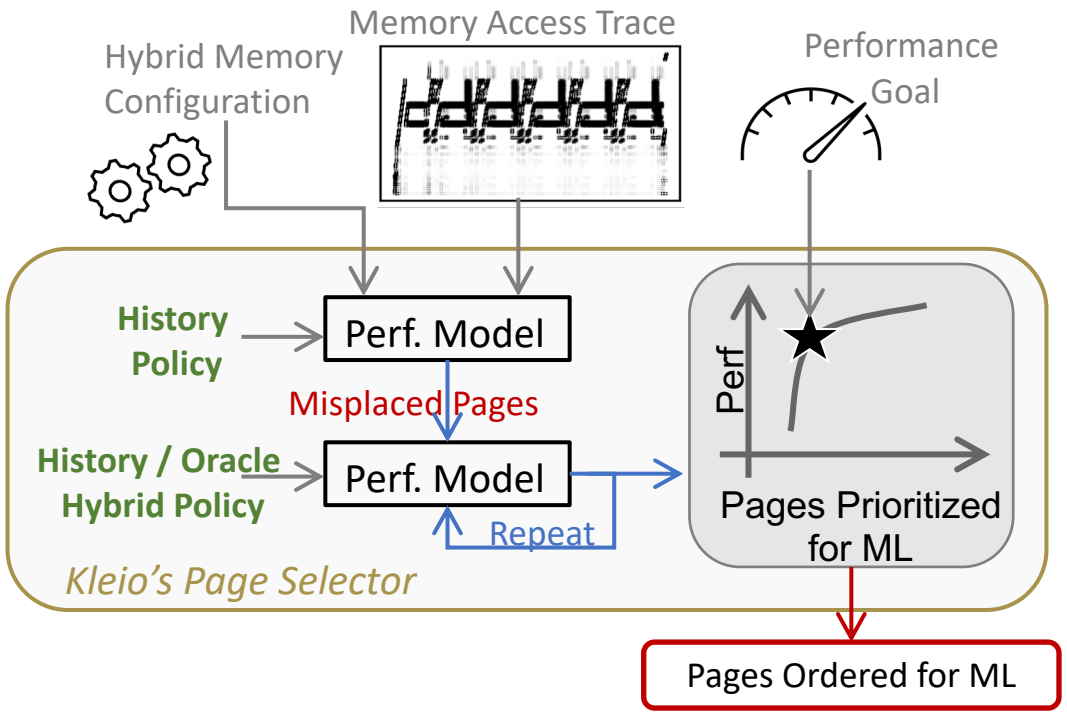

It is not a lightweight process, but necessary to maximize the effects of ML on application performance.

Kleio bridges the performance gap left by existing solutions by 80%, on average.

Talk Outline

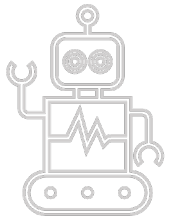

Building *Smart* Systems
Foundations for practical Machine Learning (ML)-based Management

Building *Fast* Systems
Reducing ML-based Management Overheads with Visualization

Open Research Questions

Reducing Operational Overheads of ML-based Management

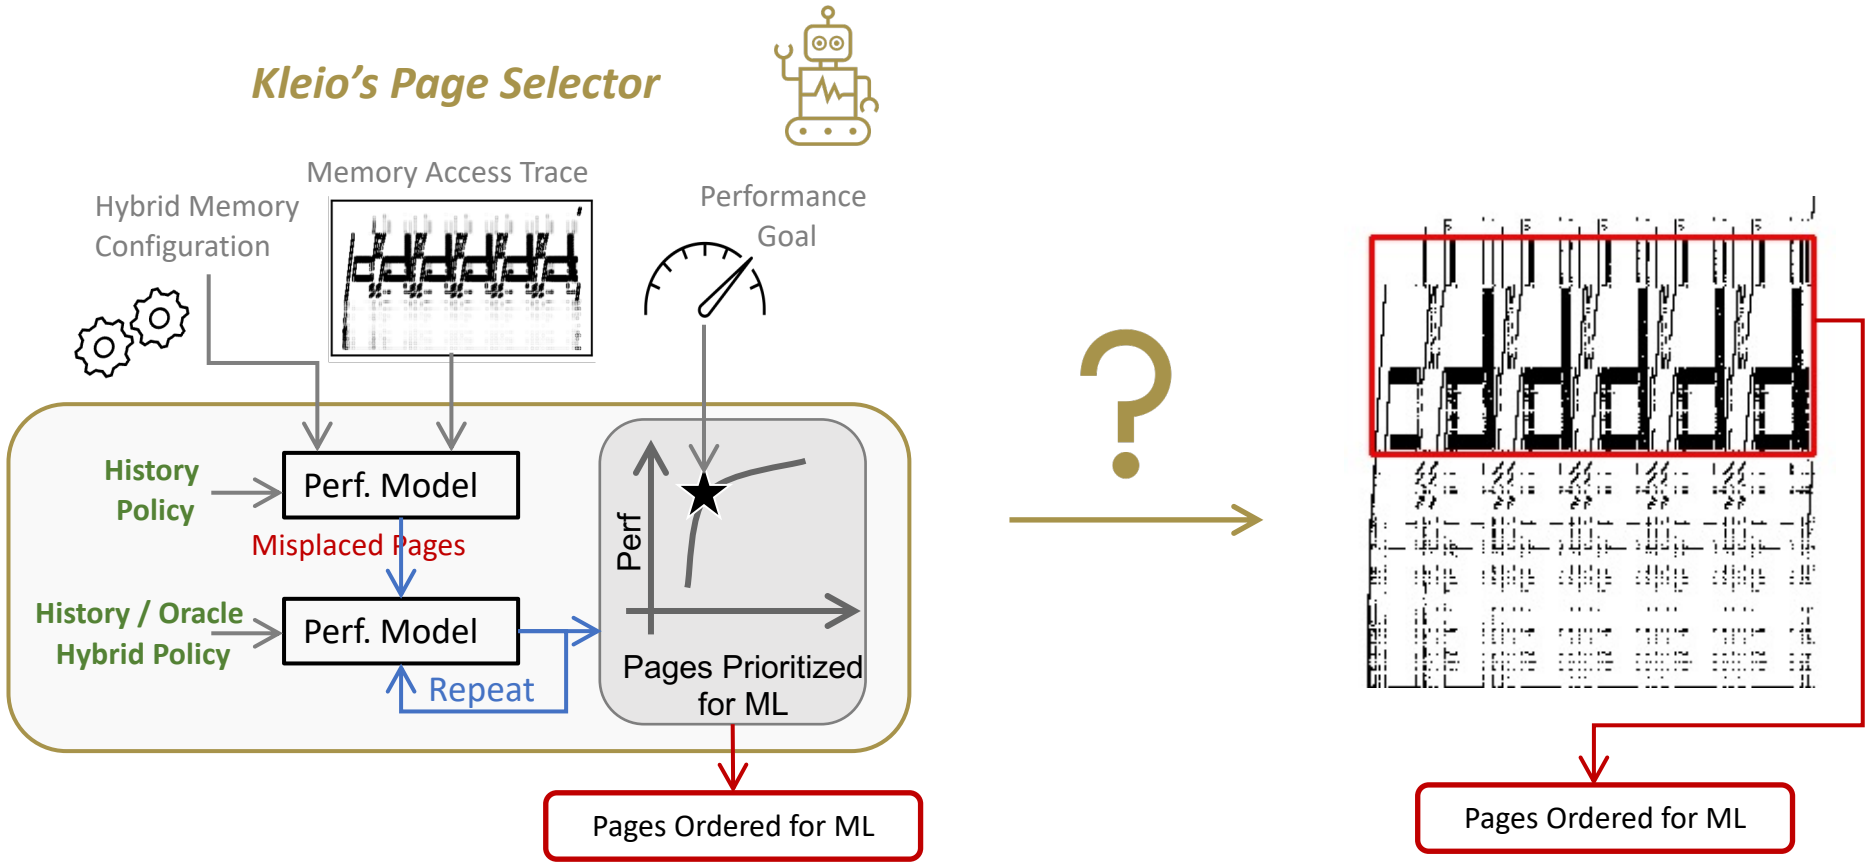

Can we accelerate the page selection process via image-based decisions?

Why images?

Personal Experience

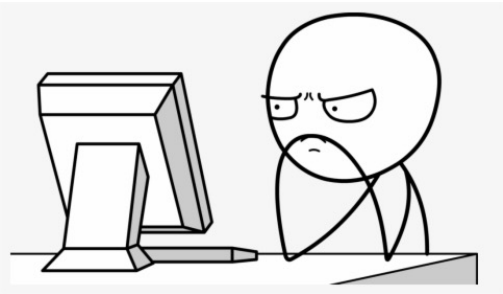

I came up with great observations and insights, by visualizing memory access patterns.

Feature Extraction

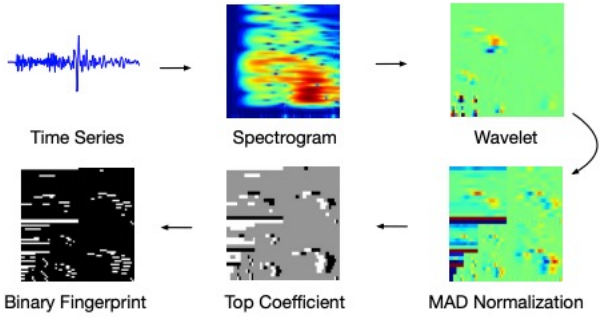

Figure 3: The fingerprinting algorithm encodes time-frequency features of the original time series into binary vectors.

Source: Kexin Rong et al. at VLDB '18.

Earthquake Detection: Extract Frequencies of Seismic Waves.

Image-based ML Classifiers

Figure 1: Typical workstation of a professional trader. Credit: Photoagriculture / Shutterstock.com.

Source: J.P. Morgan AI labs.

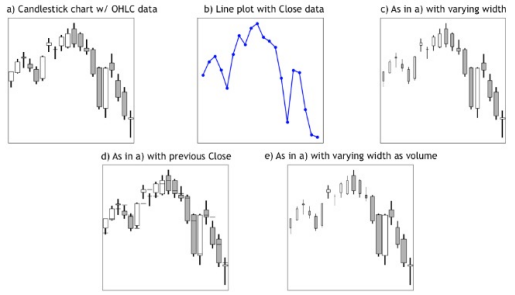

Figure 4: Various visual representations of the same time-series data.

Finance: Trading by learning time series data as images.

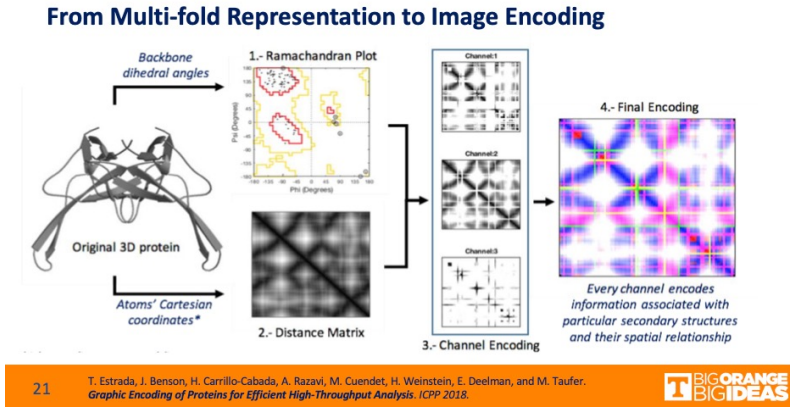

Source: Presentation from Michela Taufer.

Bioinformatics: Learn protein functions.

Visualizing Pages Selected for ML

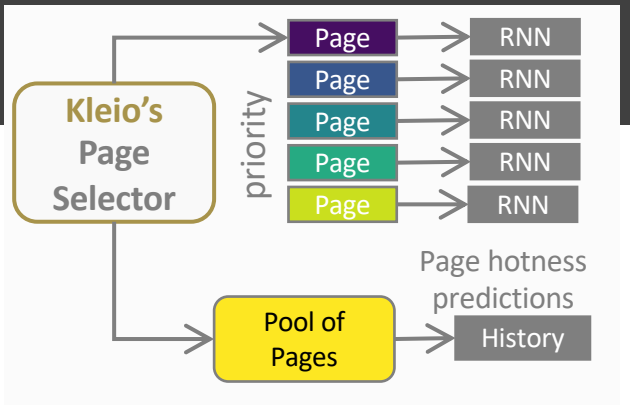

High Priority — Low Priority

backprop

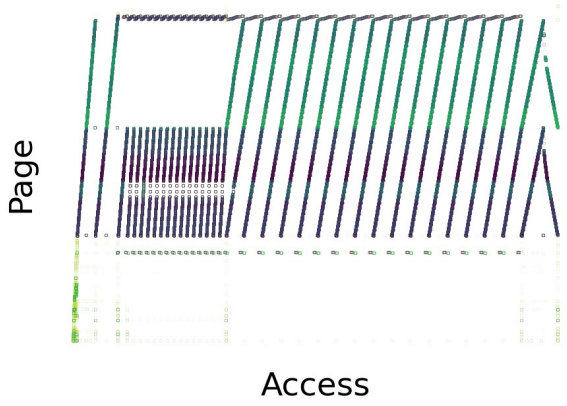

lud

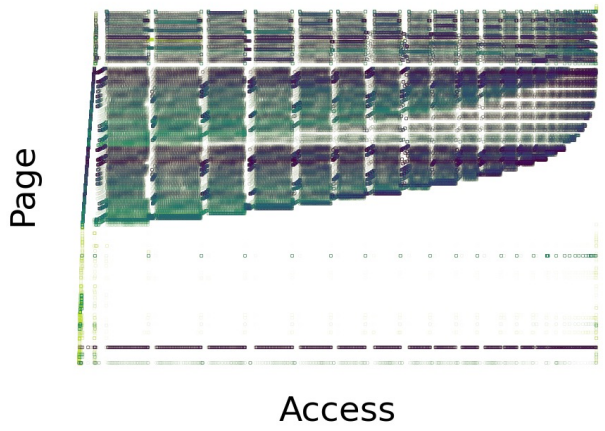

cpd

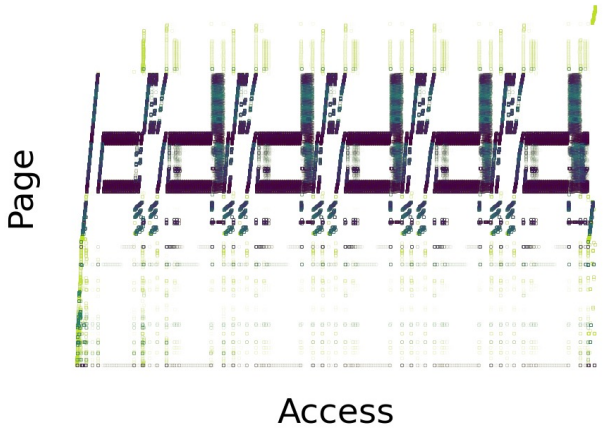

Neighboring pages that are part of distinct access patterns across *time* receive similar priority for ML.

Towards Image-based Page Selection

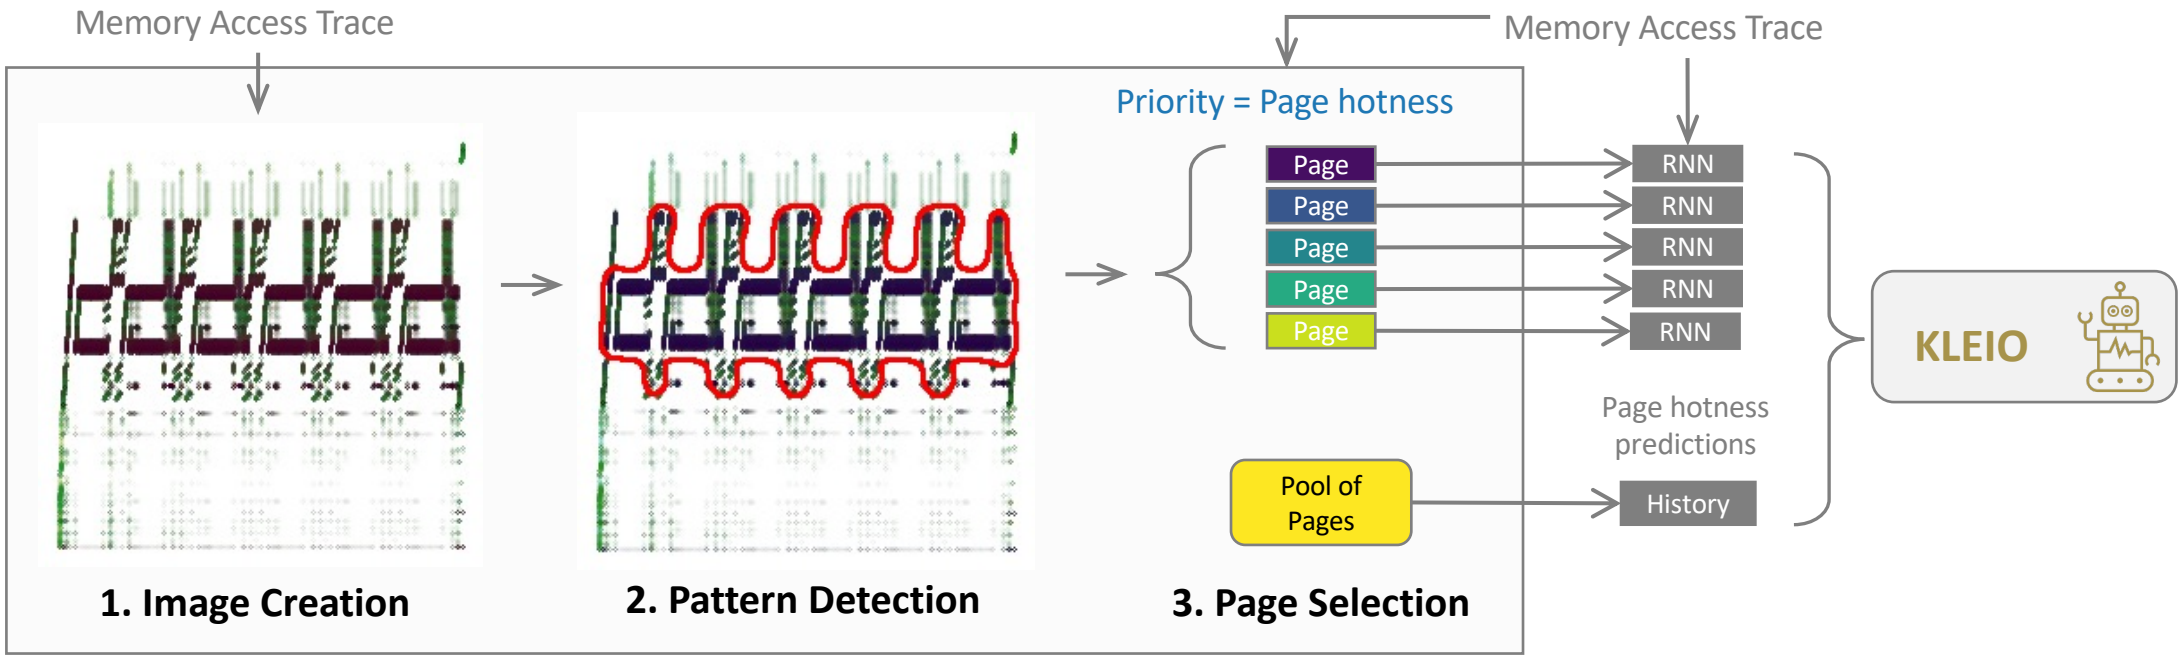

Page Selection Comparison

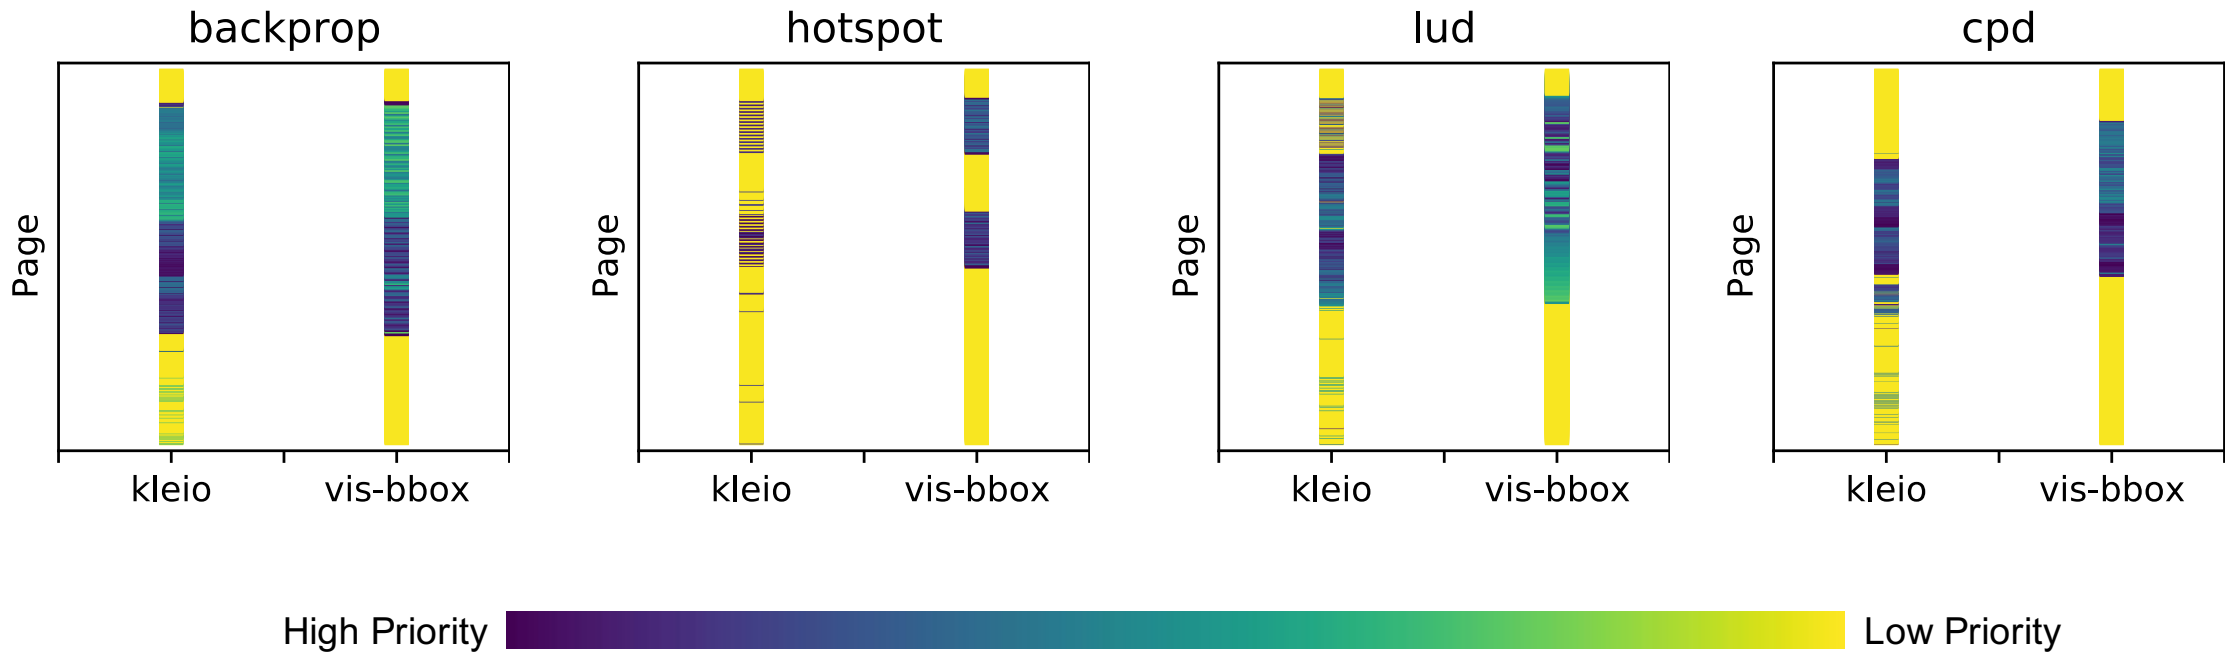

Similar page orderings between our initial approach (vis-bbox) and the performance-based selection of Kleio.

Performance Evaluation

1. Application Performance

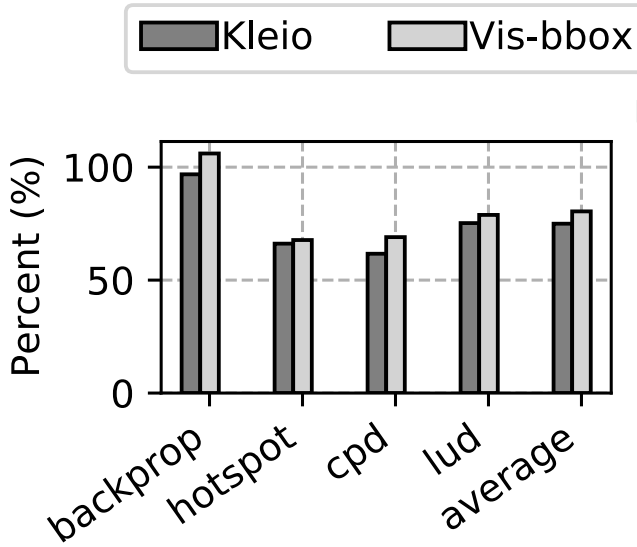

Performance slowdown from a DRAM-only platform.

Lower is better.

2. Page Selection Time

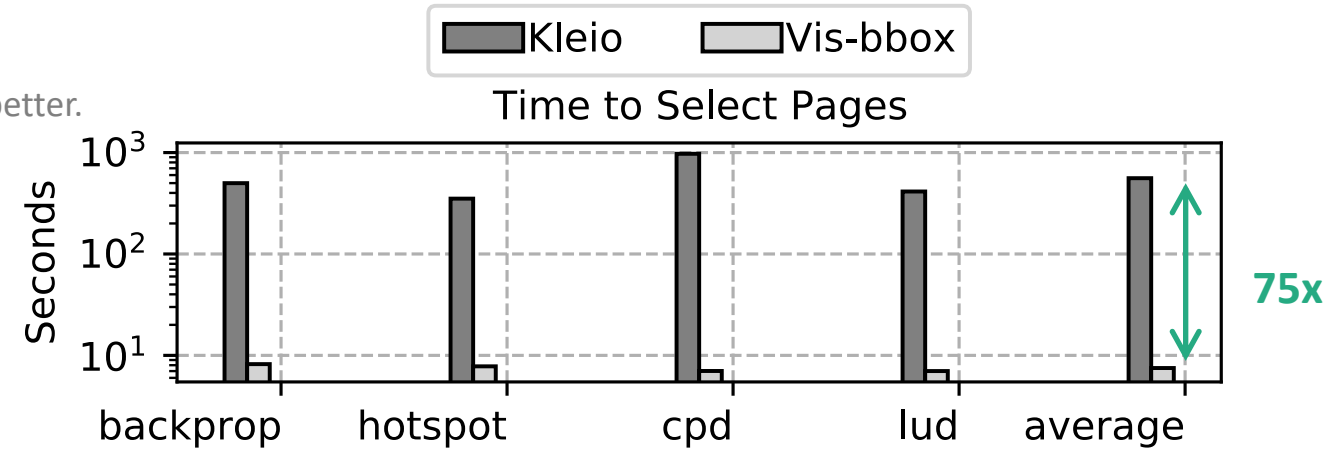

Comparable application performance levels.

Accelerates the time to select pages for ML-based management.

Talk Outline

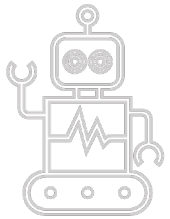

Building *Smart* Systems
Foundations for practical Machine Learning (ML)-based Management

Building *Fast* Systems
Reducing ML-based Management Overheads with Visualization

Open Research Questions

ML-augmented Heterogeneous Resource Manager

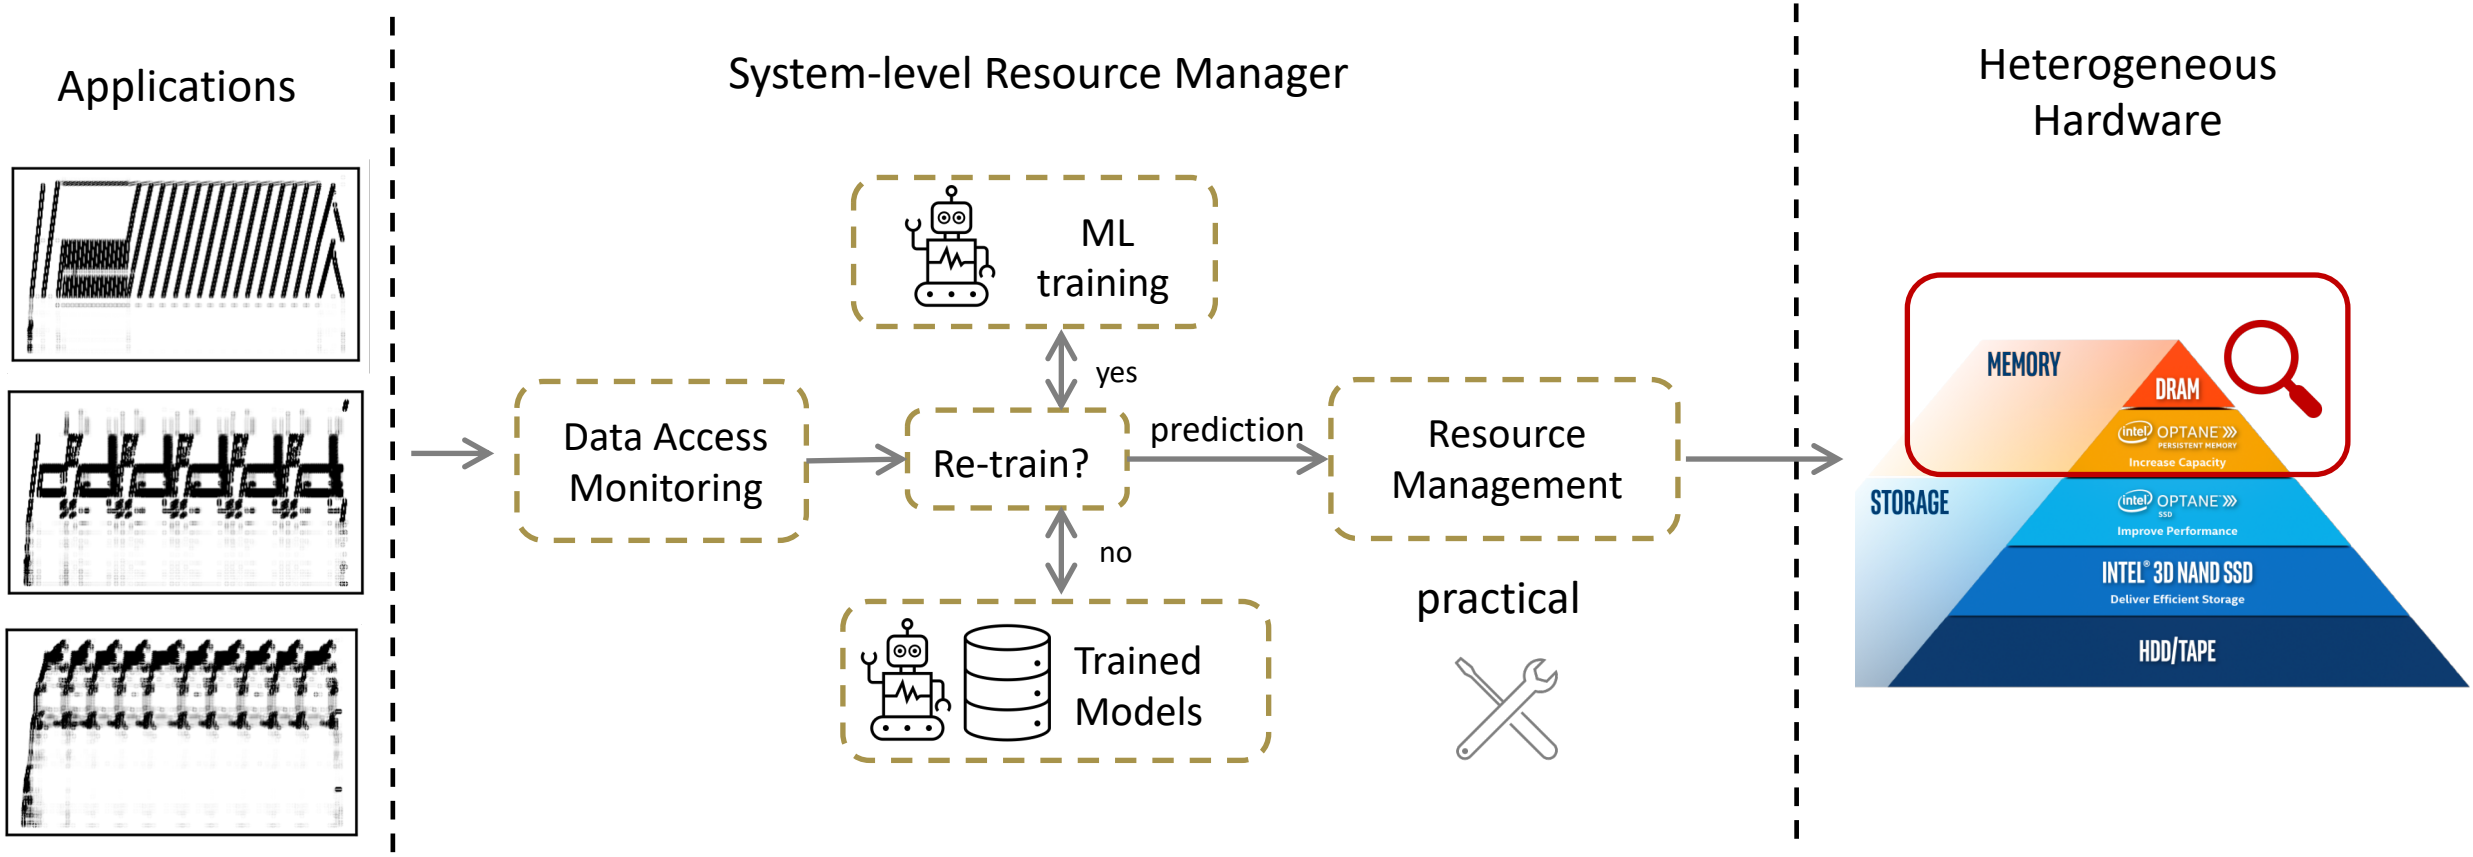

Remaining Challenges
Fully integrated adaptive resource manager.

Applications

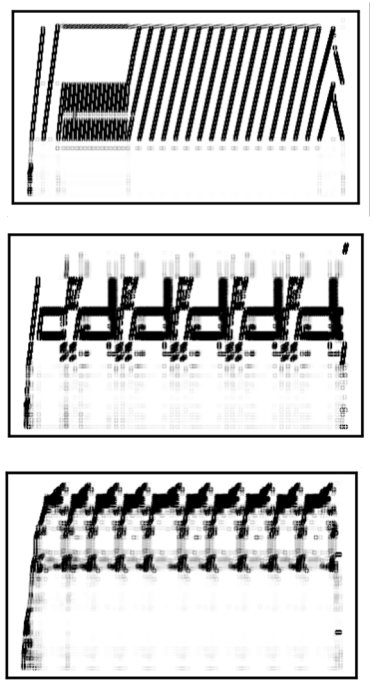

System-level Resource Manager

ML training

Pattern – Model Matching?

Data Access Monitoring

Re-train, use trained model, create new model?

Resource Management

no

Trained Models

practical

How to accelerate training?

Where to Store?

How to amortize the training cost?

Heterogeneous Hardware

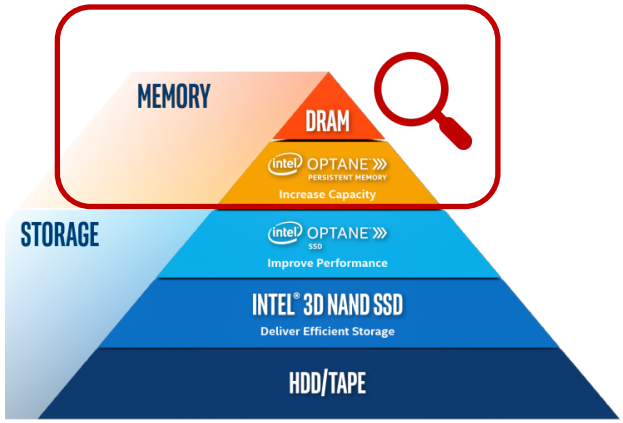

Computer Vision + Machine Learning
Pattern Recognition.

Applications

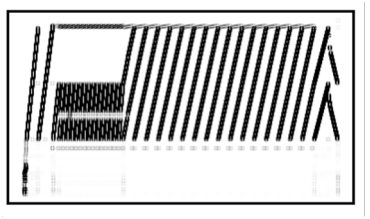

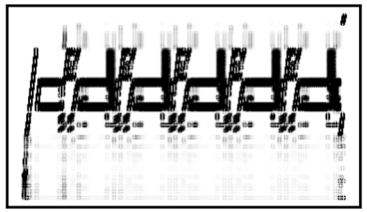

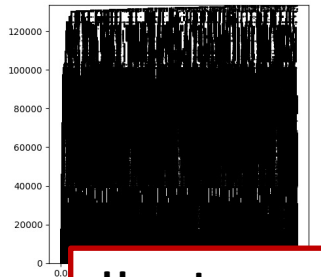

How to zoom-in?

Image Resolution?

System-level Resource Manager

Pattern Recognition

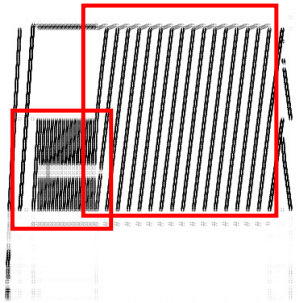

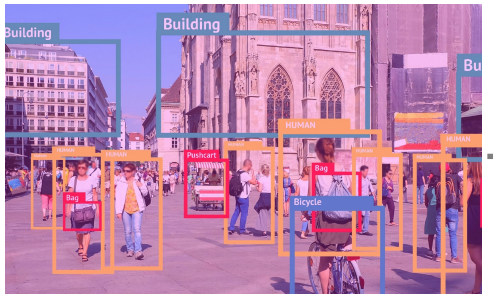

Pattern Class

Example Use Cases

- Select ML model
- Take ML related action
- Resource Management Decision

Build an “ImageNet” for memory access patterns.

Train classifiers to take actions upon recognizing a pattern.

What Classes to define?

Metadata?

Storage?

Label guidelines for open contributions?

Heterogeneous Hardware

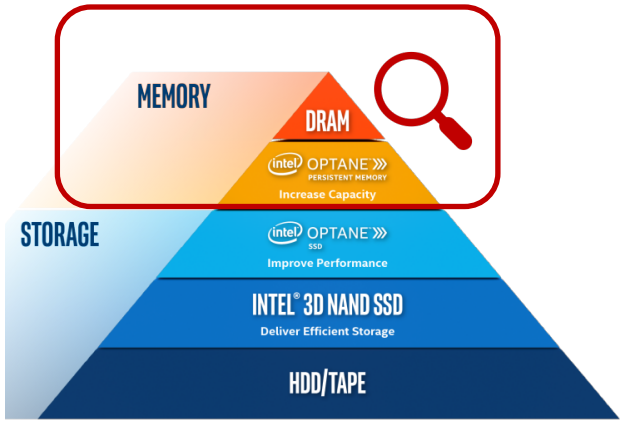

Computer Vision + Machine Learning
Pattern Prediction.

Applications

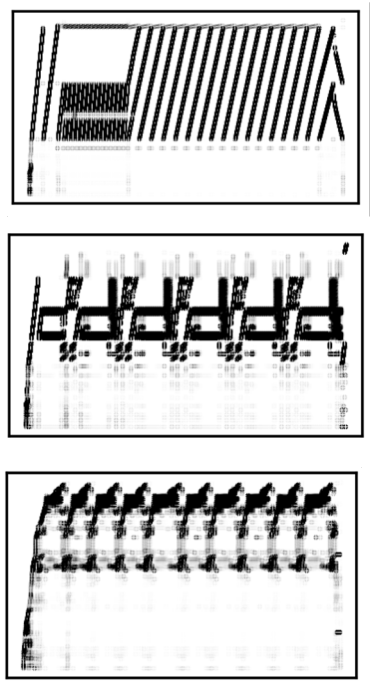

System-level Resource Manager

Pattern Prediction

ML models → Future memory access patterns → Resource Management Decision

Ground truth

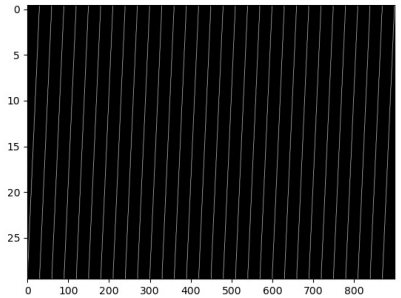

Prediction

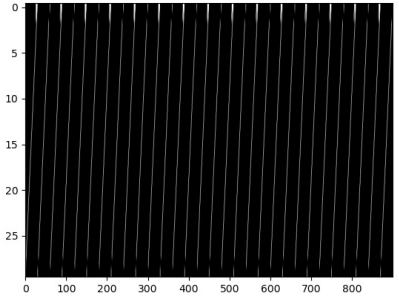

Prediction result of a Convolutional + Recurrent Neural Network.

Heterogeneous Hardware

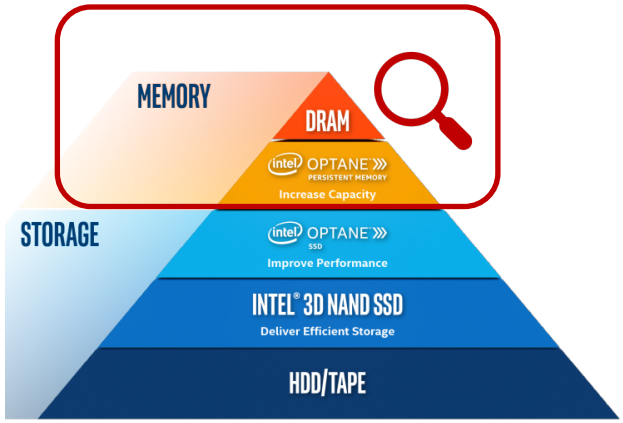

Compare with RNN-only?

1 model per app or per pattern?

Page-to-image mapping

"ML models over images have better predictive capabilities."

From OSDI '21 Keynote from J.P. Morgan AI labs.

Scan this to find more about my work.

Smart

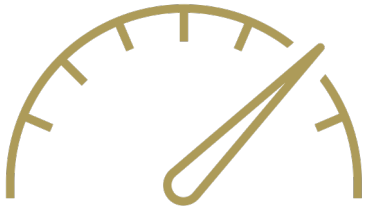

Fast

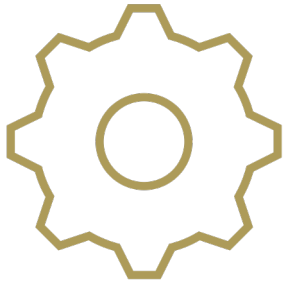

Systems

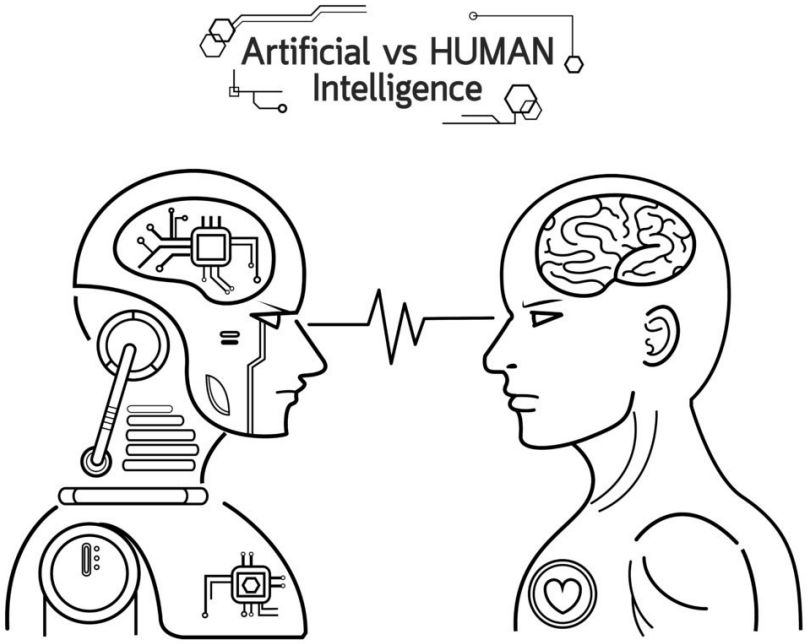

How can we use our human intelligence to build **practical** systems that leverage machine learning and computer vision?

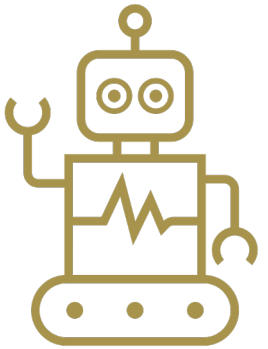

Machine Learning

Computer Vision