Adding Machine Intelligence to Hybrid Memory Management

Thaleia Dimitra Doudali (IMDEA Software Institute)*

@ HMEM 2022

* Work done while at Georgia Tech (PhD).

The Era of Heterogeneous Hardware

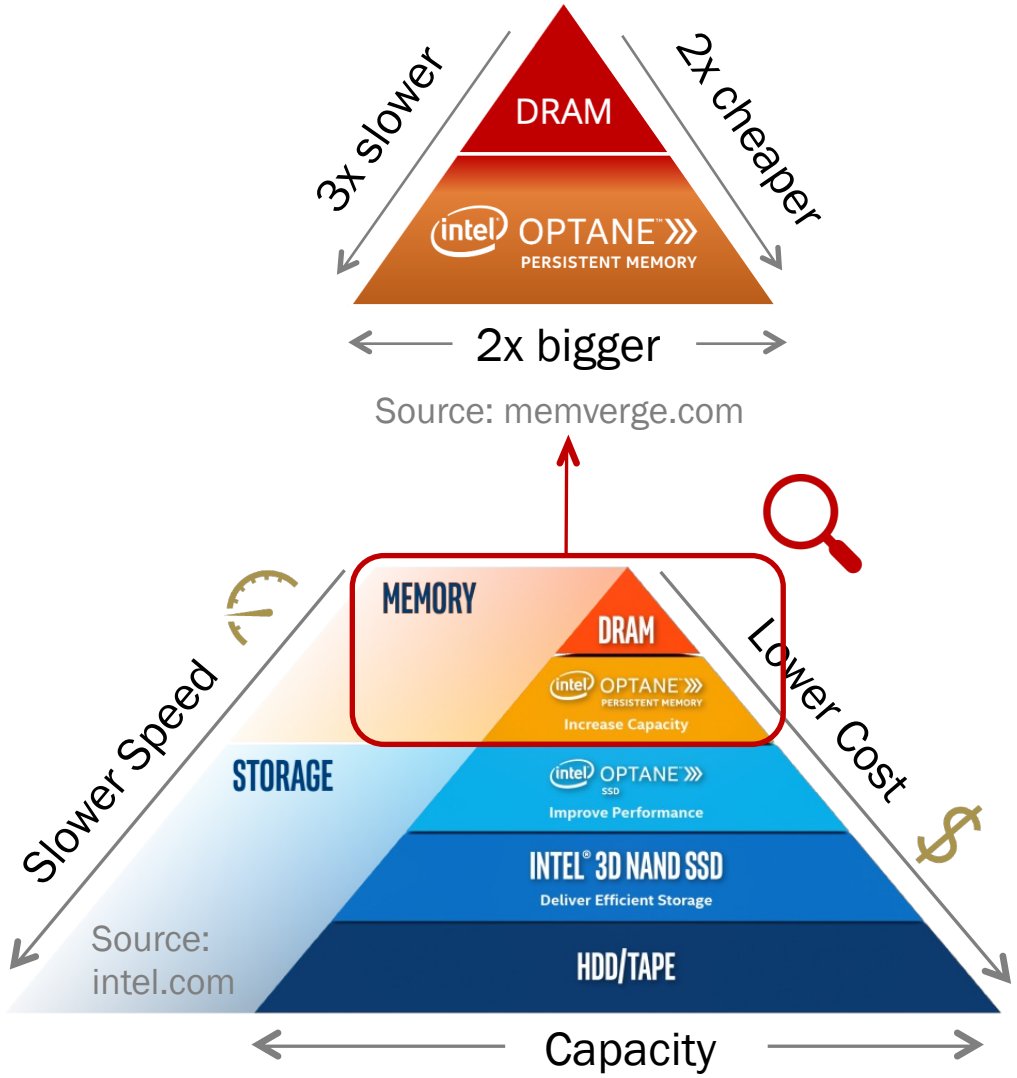

Characteristic	Technology	Hardware Vendors
Low Latency	MRAM	
High Bandwidth	HBM	 
Persistence	PMEM	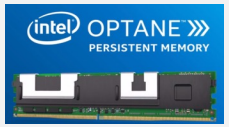

Examples of other heterogeneous memory technologies.

Hybrid Memory Management is Complex

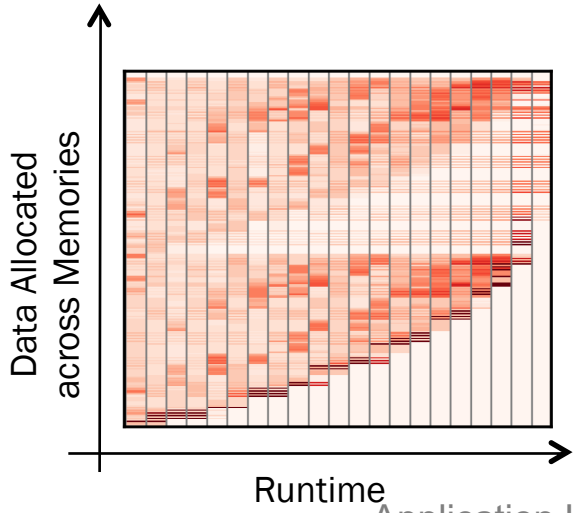

Dynamic Data Movements!

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

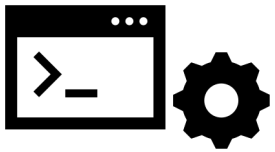

Resource Management Software

System Level

Hardware Level

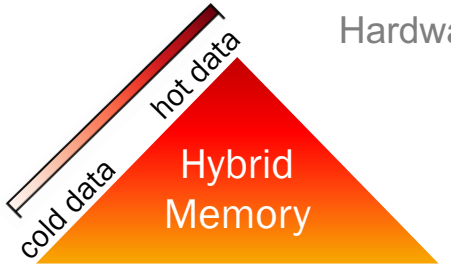

Why do we need more intelligent systems?

Application data sizes

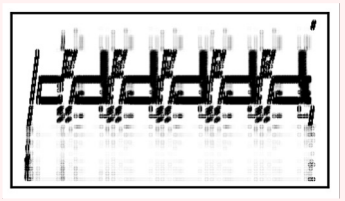

Complex data access patterns

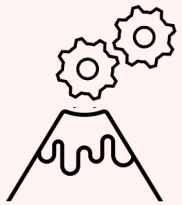

Exploded system parameter space

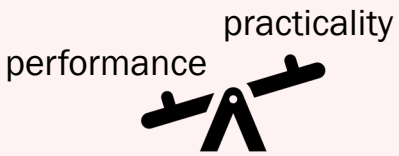

Hard to balance

Machine Intelligent Hybrid Memory Management

The Vision.

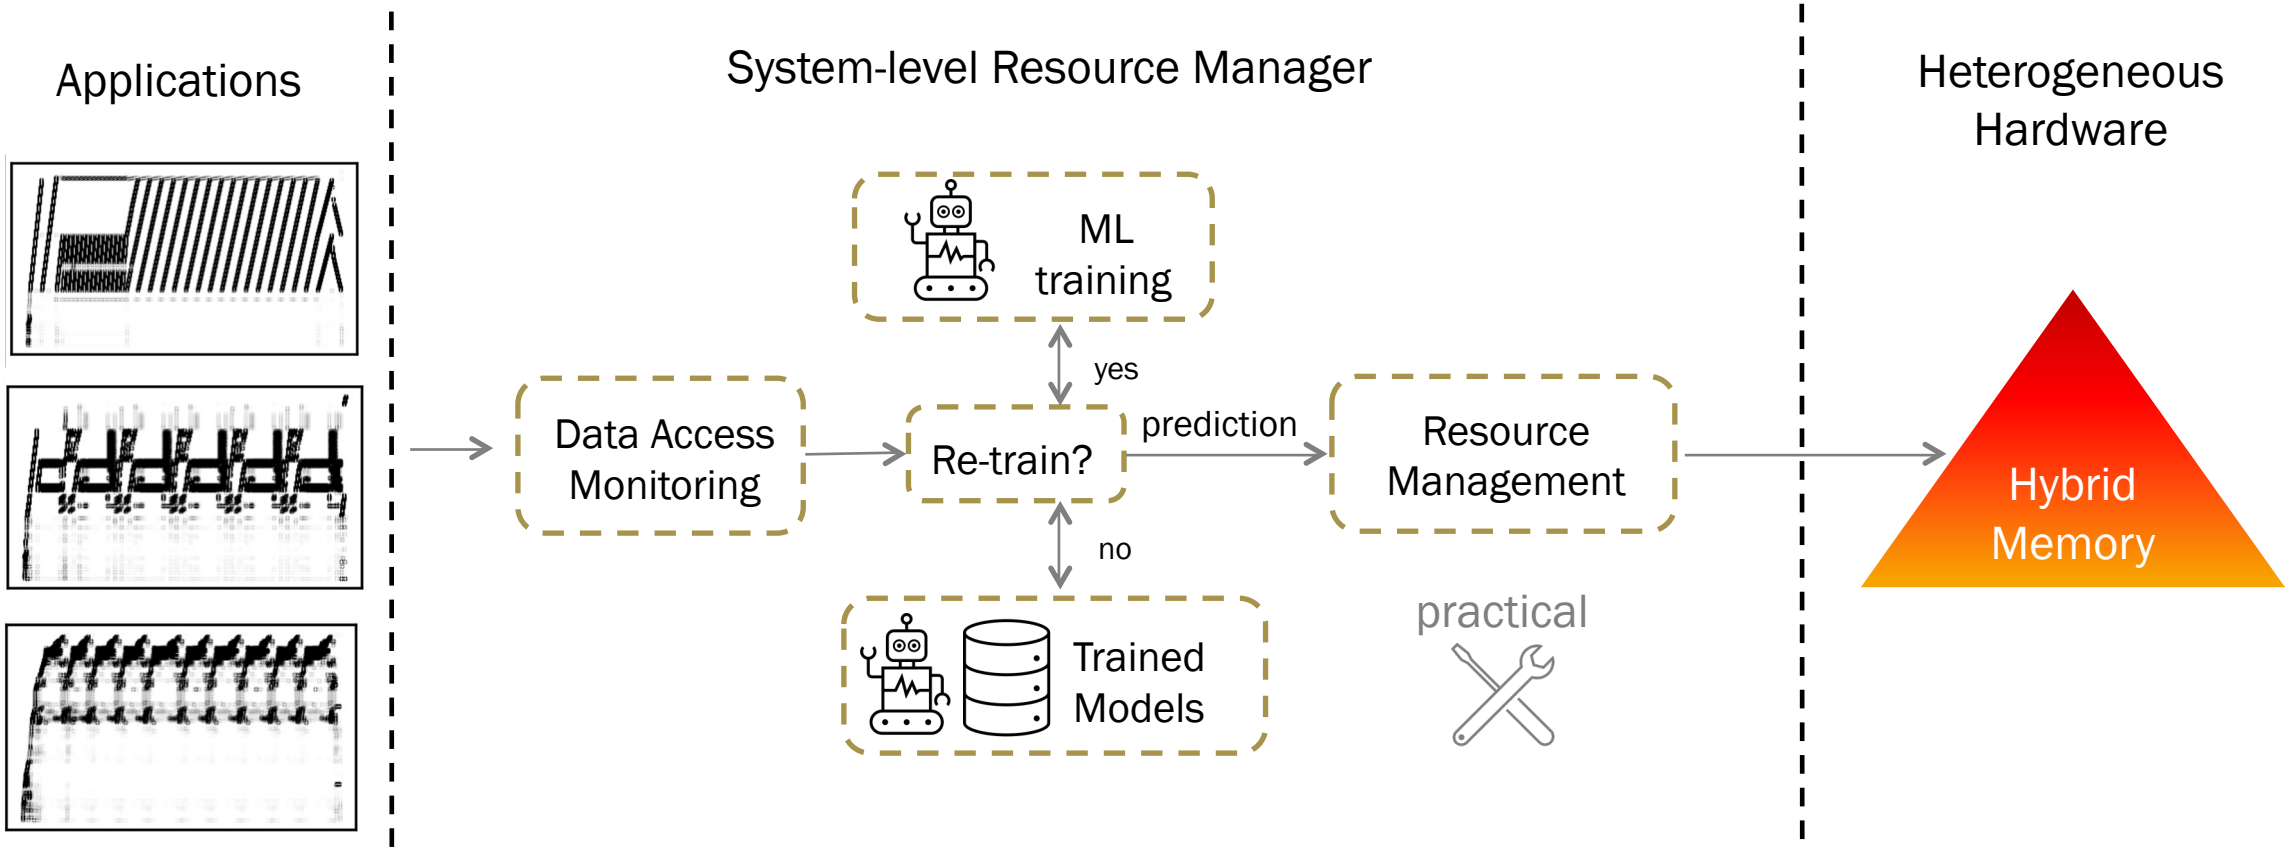

Machine Intelligent Hybrid Memory Management

Laying the grounds for the practical integration of ML.

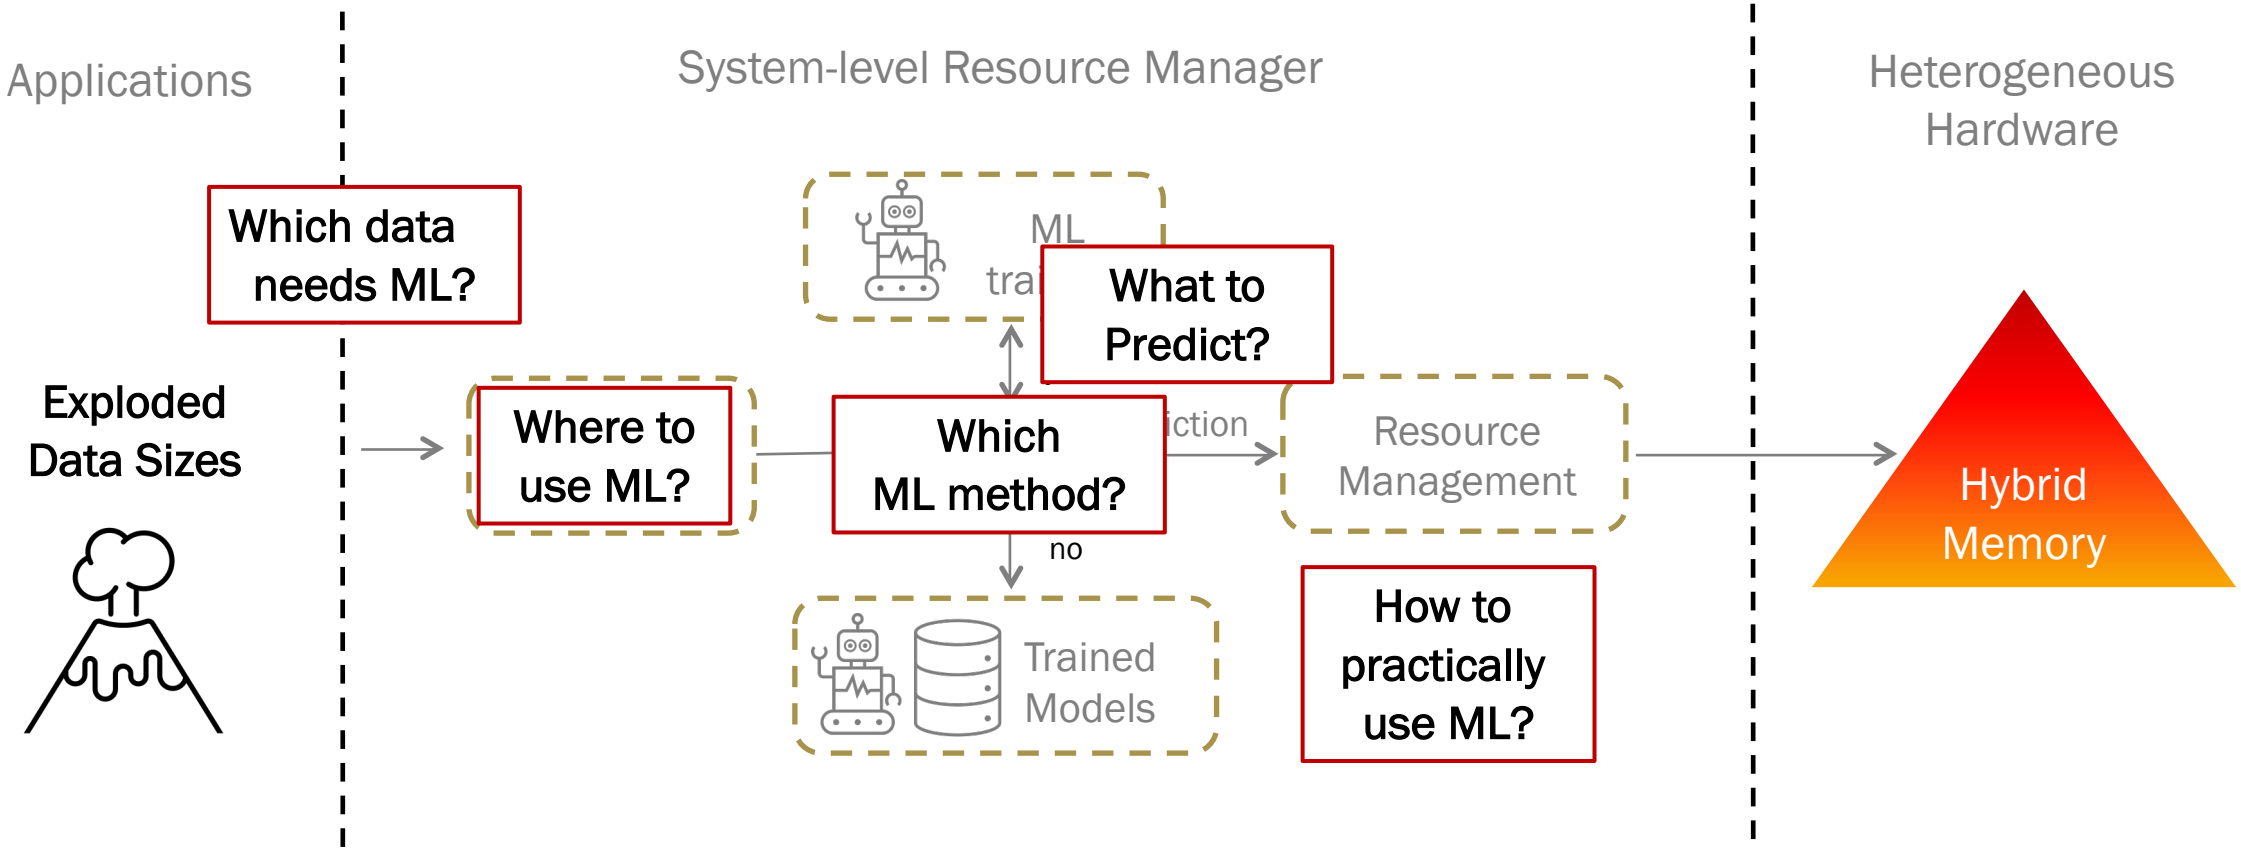

System Design of *Kleio*

Kleio is a hybrid memory page scheduler with machine intelligence. [Best Paper Award Finalist at HPDC 2019.]

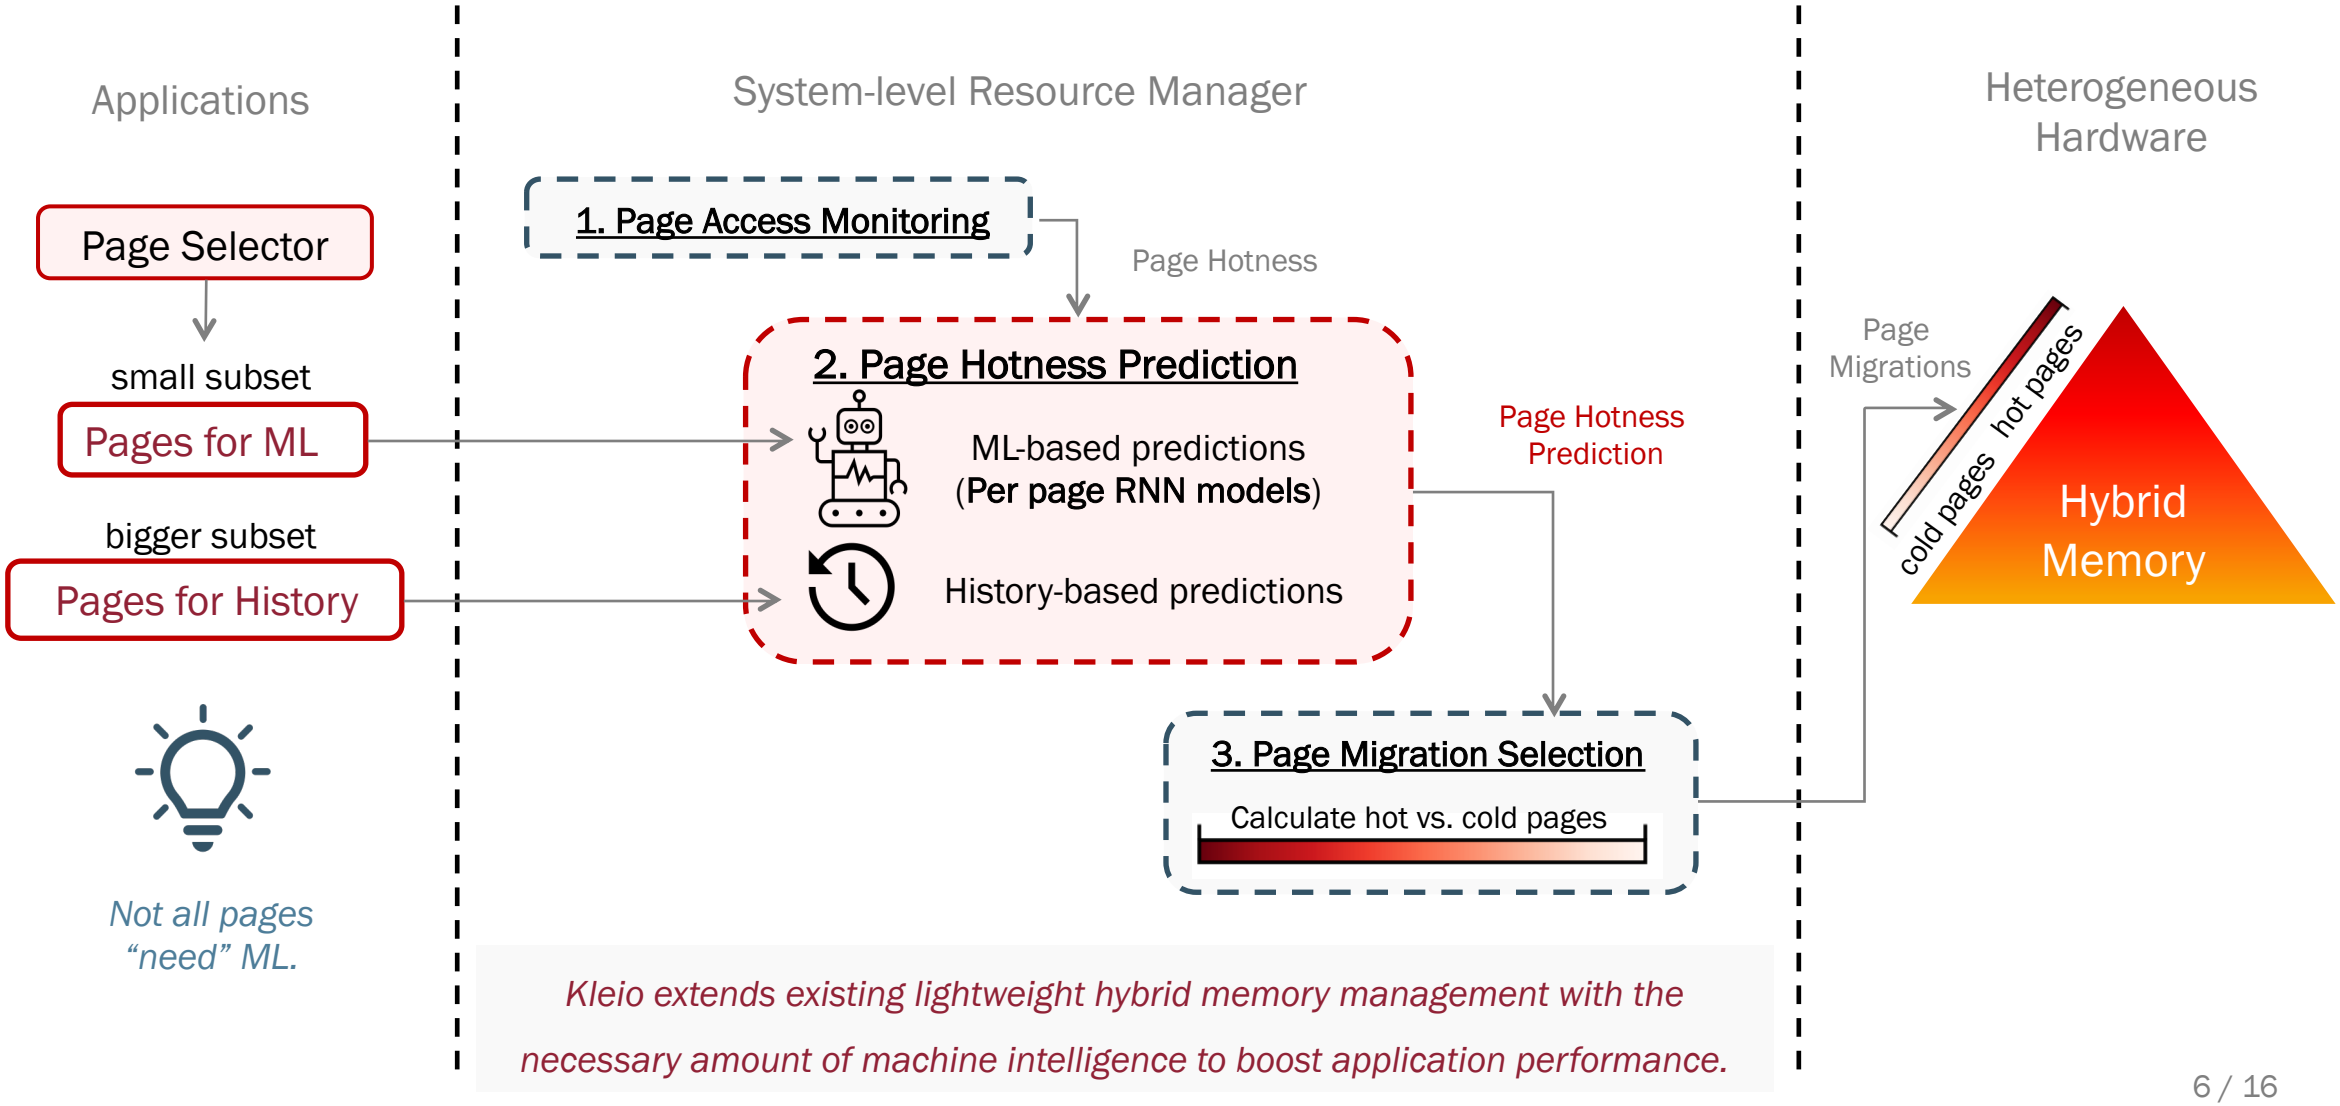

The Key(s) to a Practical and Efficient ML-based System Design

Apply ML **when** and **where** necessary.

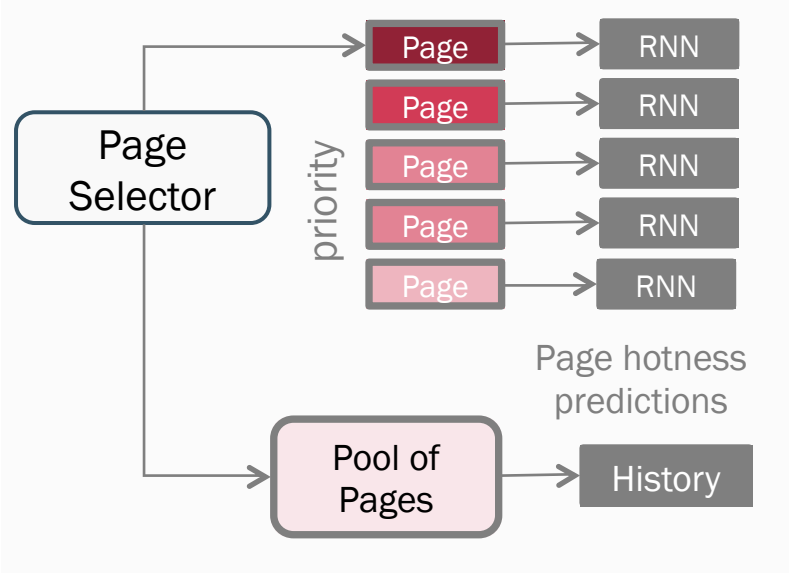

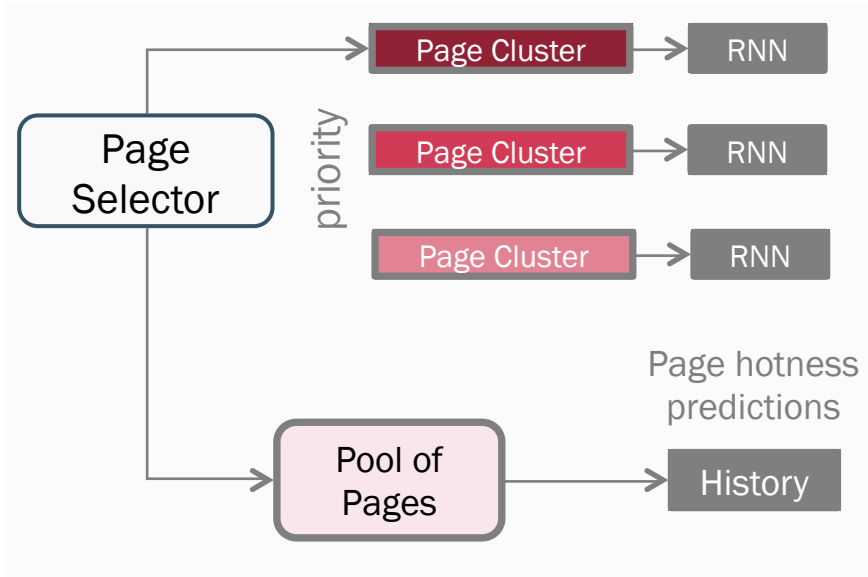

Apply ML on a small page subset.

↳ Foundations for practical use of ML.

Carefully select pages for ML.

↳ Application performance boost.

Small can still mean thousands of pages, because of the massive memory footprints of modern workloads.

Can we reduce the number of pages via clustering?

Insights from the System Design of **Coeus**

Coeus: Clustering (A)like Patterns for Practical Machine Intelligent Hybrid Memory Management . [CCGrid 2022]

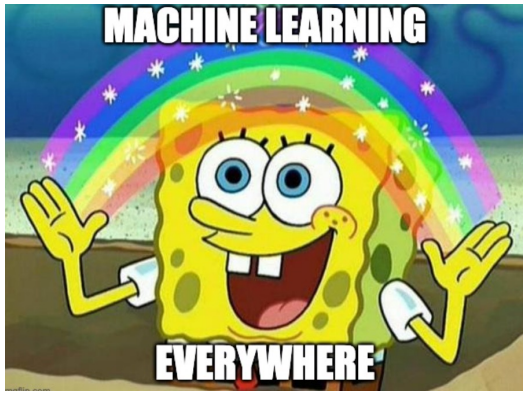

Clustering? Let's use ML!

For example, K-means.
- How many clusters?
- Clustered input to ML?

Not trivial to configure.

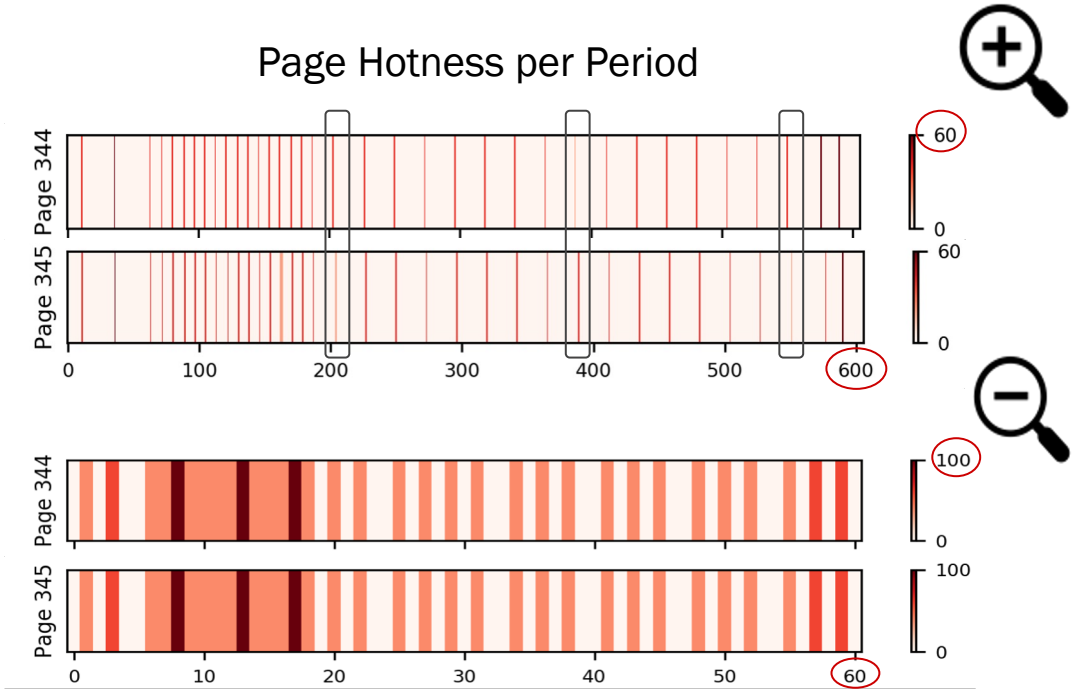

Let's use our human intelligence..

.. Kleio learns the patterns of page hotness across time periods.

Group pages with *identical* patterns
under a single ML model.

Key Idea

Which is the Right Granularity?

Longer periods result in more pages having identical patterns of page access hotness across time.

Longer periods may result in insufficient data movements and impact application performance.

Periods that align with the page reuse distance, maximize performance.

***Page Reuse Distance** = The time gap between two accesses to the same page.

Insight from ***Cori**: Dancing to the Right Beat of Periodic Data Movements over Hybrid Memory Systems.* [IPDPS 2021]

Cori is a lightweight tuning solution for hybrid memory page schedulers, that we will use to determine the “right granularity”.

Scaling ML to More Pages and Improving Performance

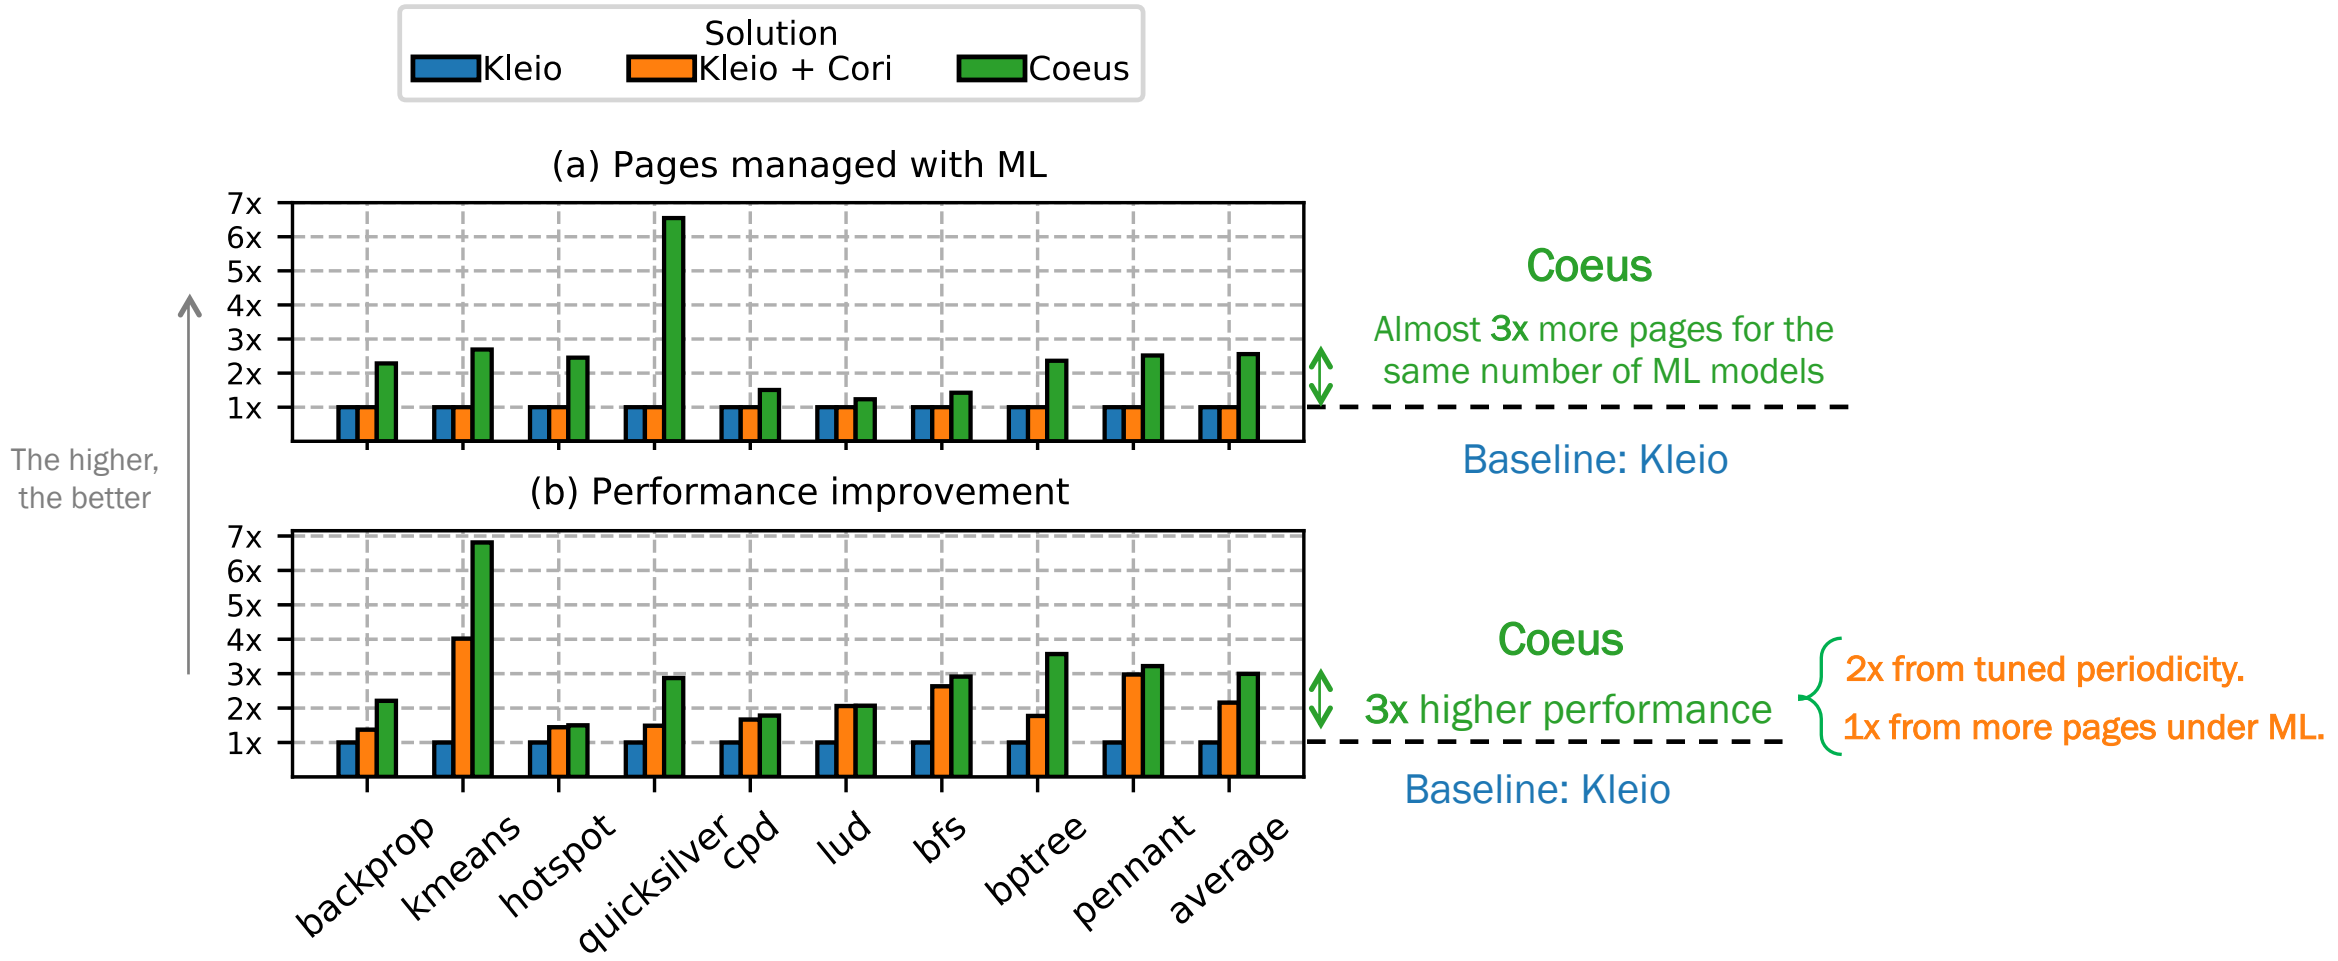

Reducing Runtime Overheads of ML-based Management

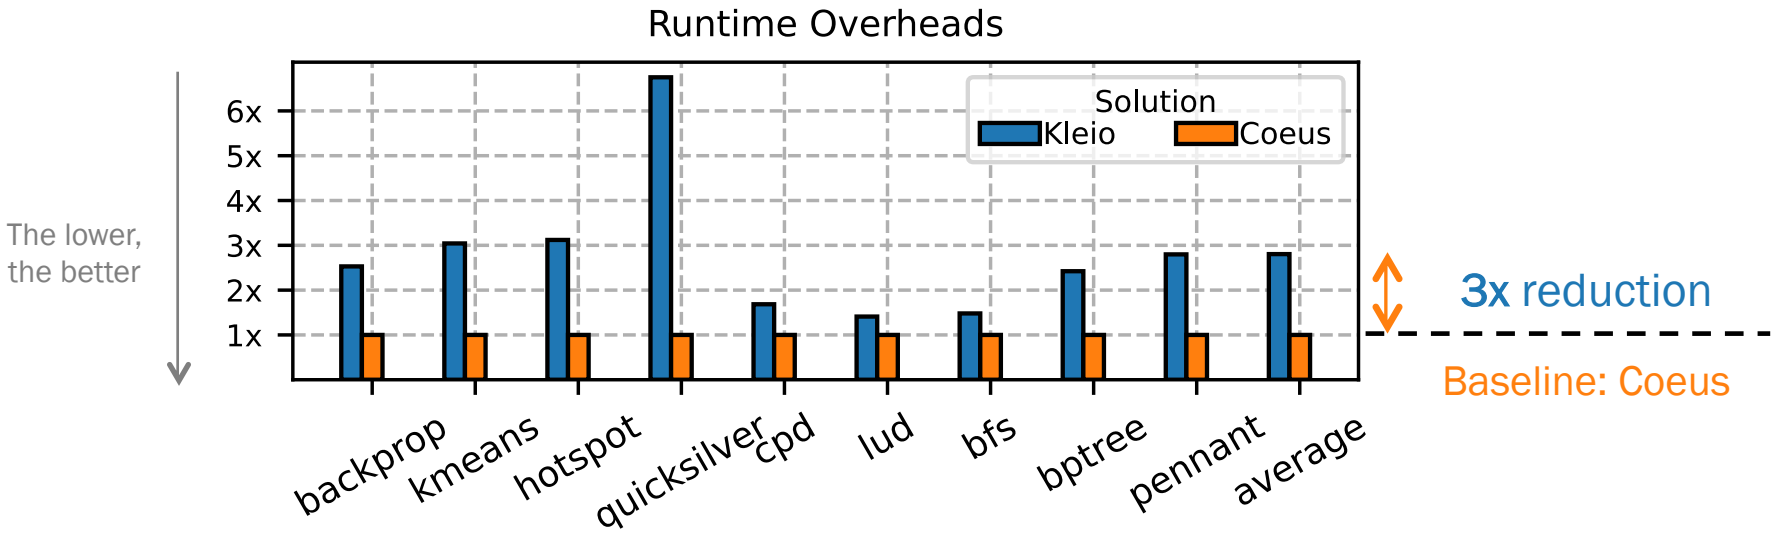

Coeus enables:

- ☑ Quick page clustering process.
- ☑ Fewer ML models deployed.

The Key(s) to a Practical and Efficient ML-based System Design

Apply ML **when** and **where** necessary.

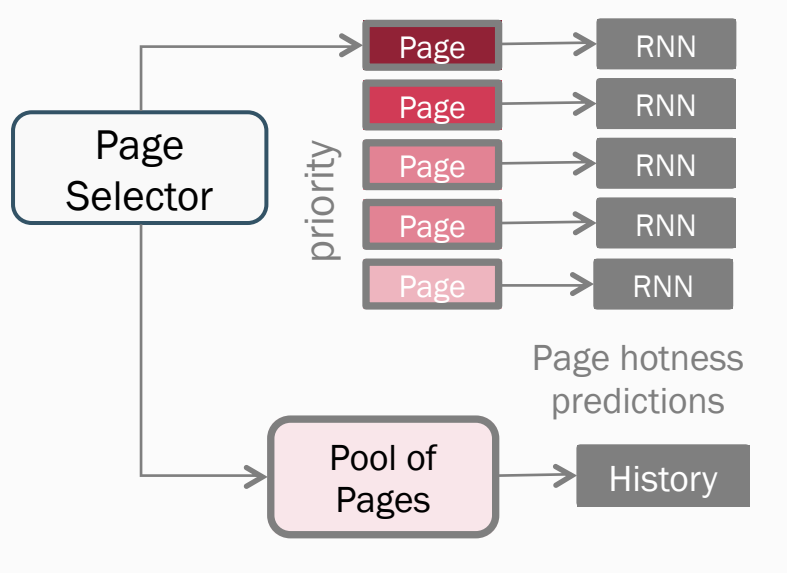

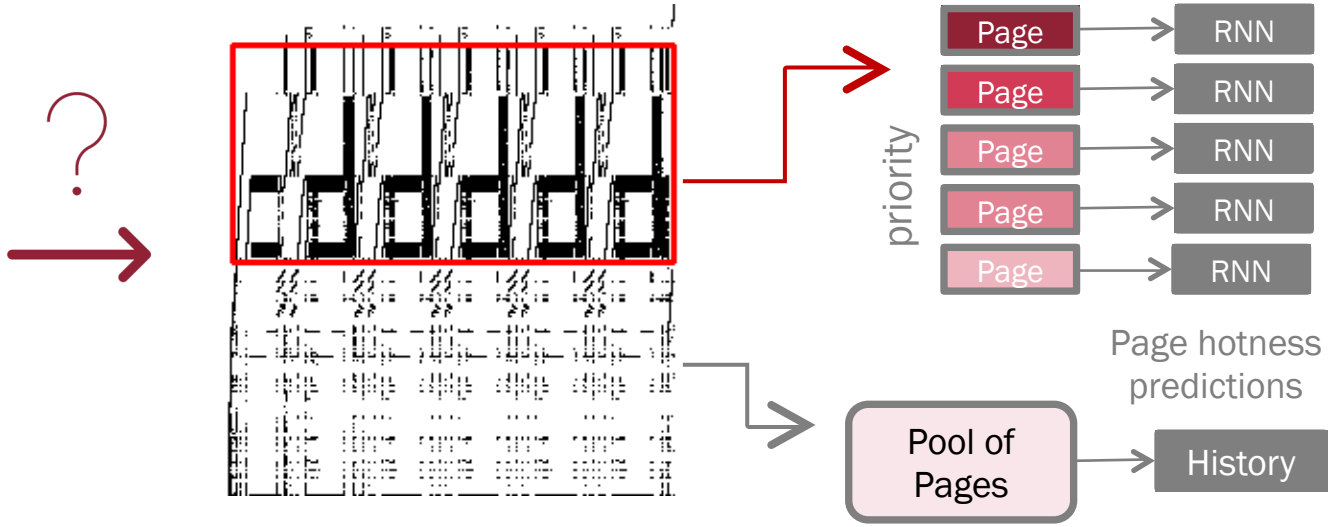

Apply ML on a small page subset.

↳ Foundations for practical use of ML.

Carefully select pages for ML.

↳ Application performance boost.

The page selection is not a lightweight process.
Performance modeling and estimations are used to maximize the effects of ML on application performance.

Can we accelerate the page selection process?

Insight from Visualizing Pages Selected for ML

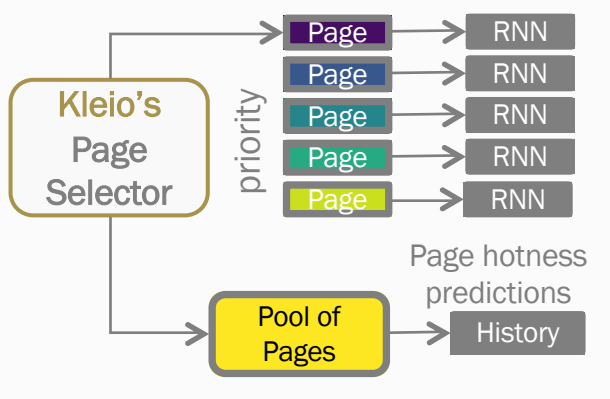

High Priority — Low Priority

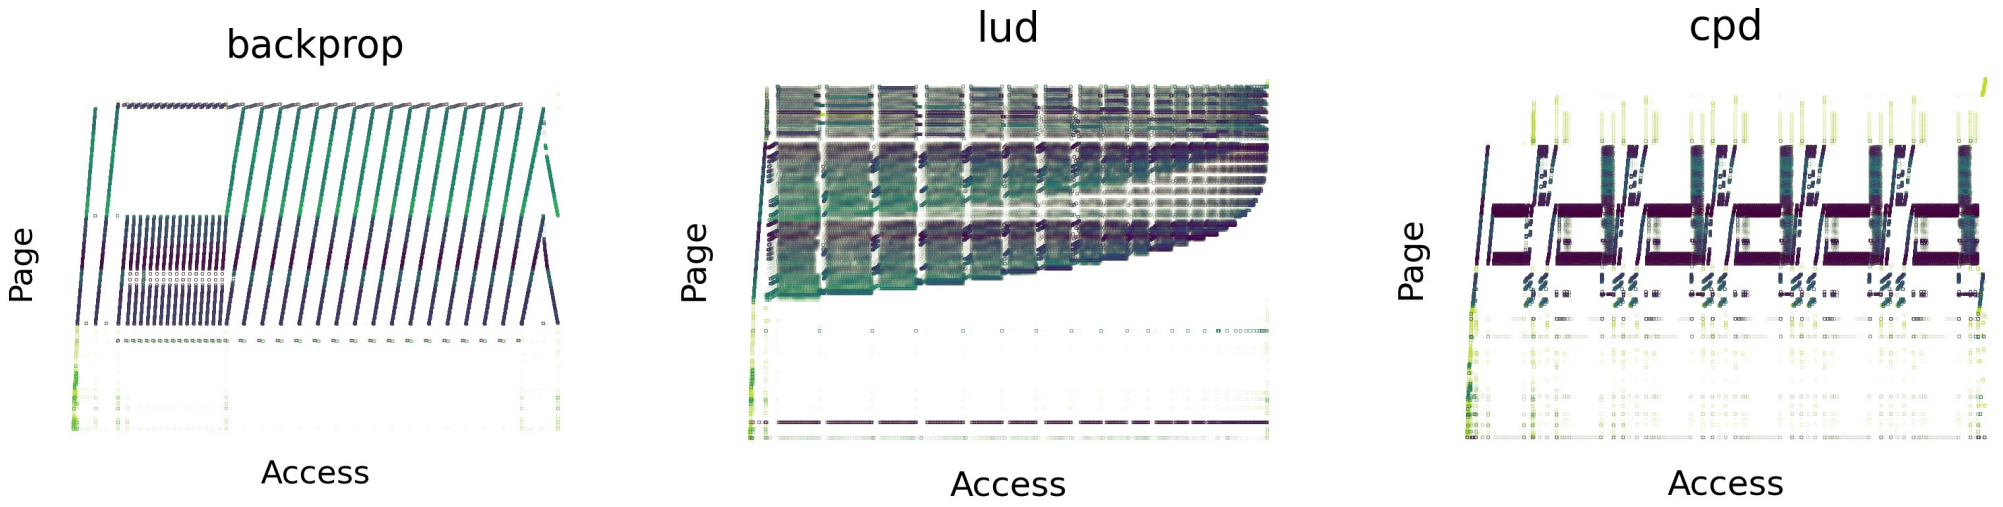

Neighboring pages that are part of distinct access patterns across *time* receive similar priority for ML.

Towards Image-based Page Selection

Research paper under submission.

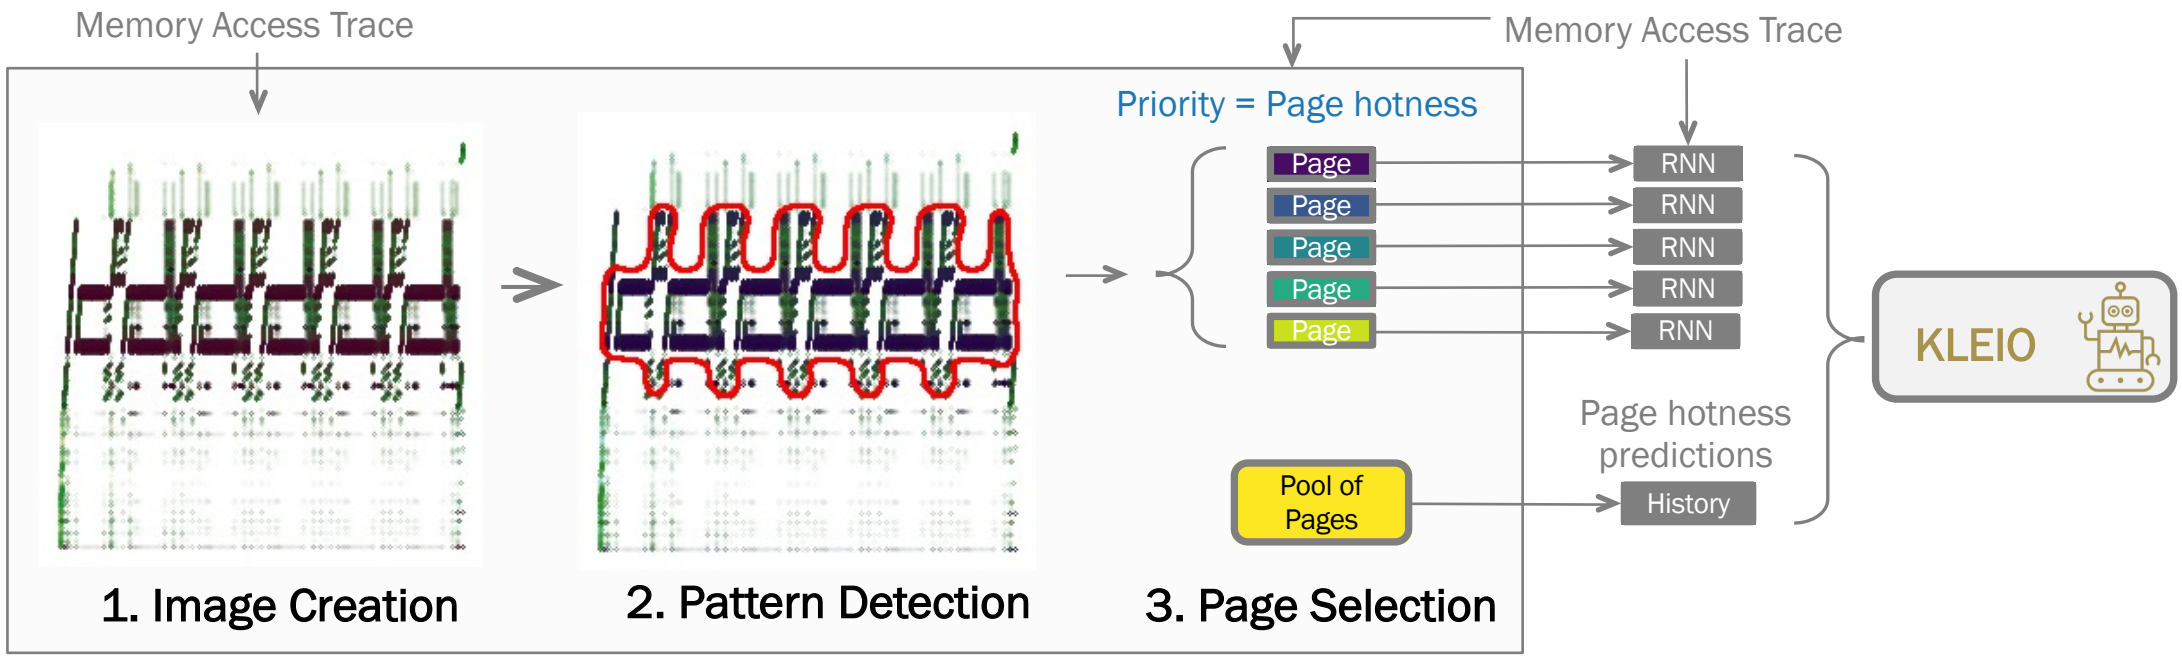

Our system reduces by **400x** the page selection times, from minutes down to seconds.

Computer Vision + Machine Learning for Systems

What can an image-based system pipeline look like?

E.g., learning memory access patterns.

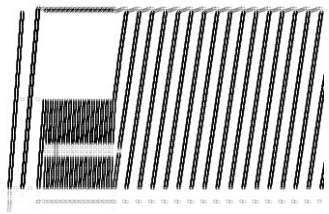

image + metadata

Pattern Recognition

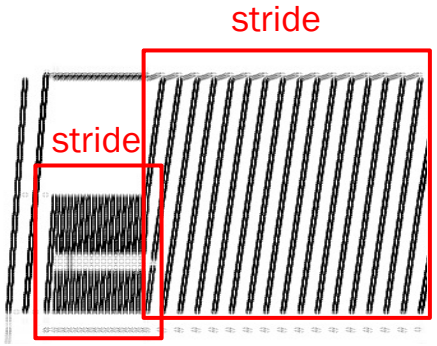

Class = “stride”

Choose based on pattern.

Pre-trained ML models

Pattern Prediction

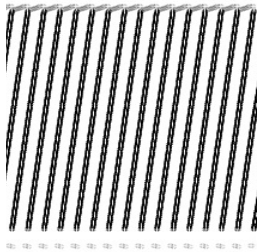

Prediction

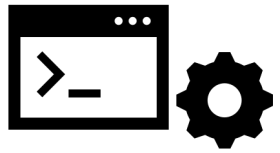

Memory Manager

System Level

Hardware Level

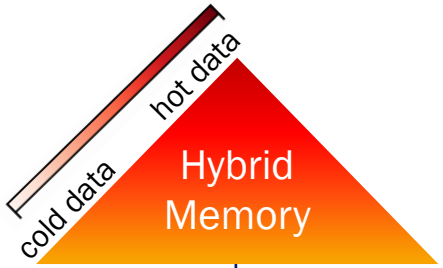

Not Accurate Prediction? Retrain.

Future Research Directions

My research lies at the intersection of Machine Learning and Systems.

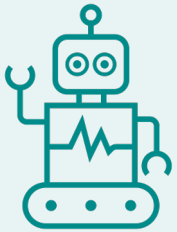

Machine Learning (**ML**)

ML ***for*** Systems

E.g., Online practical training, ML for different systems problems.

Systems ***for*** ML

E.g., Optimize memory management for RNNs / ML workloads.

Operating Systems (**OS**) Software

Computer Vision (**CV**)

ML + CV ***for*** Systems

E.g., Image-based pattern recognition and prediction of resource usage.