AI-powered Databases:

From data deluge to rapid insights

Renata Borovica-Gajic
renata.borovica@unimelb.edu.au
https://renata.borovica-gajic.com/

EPFL, IC Colloquium

Work supported by: Australian Government Australian Research Council L'Oréal FWIS’23

THE UNIVERSITY OF MELBOURNE

Data proliferation

Global data generated annually*

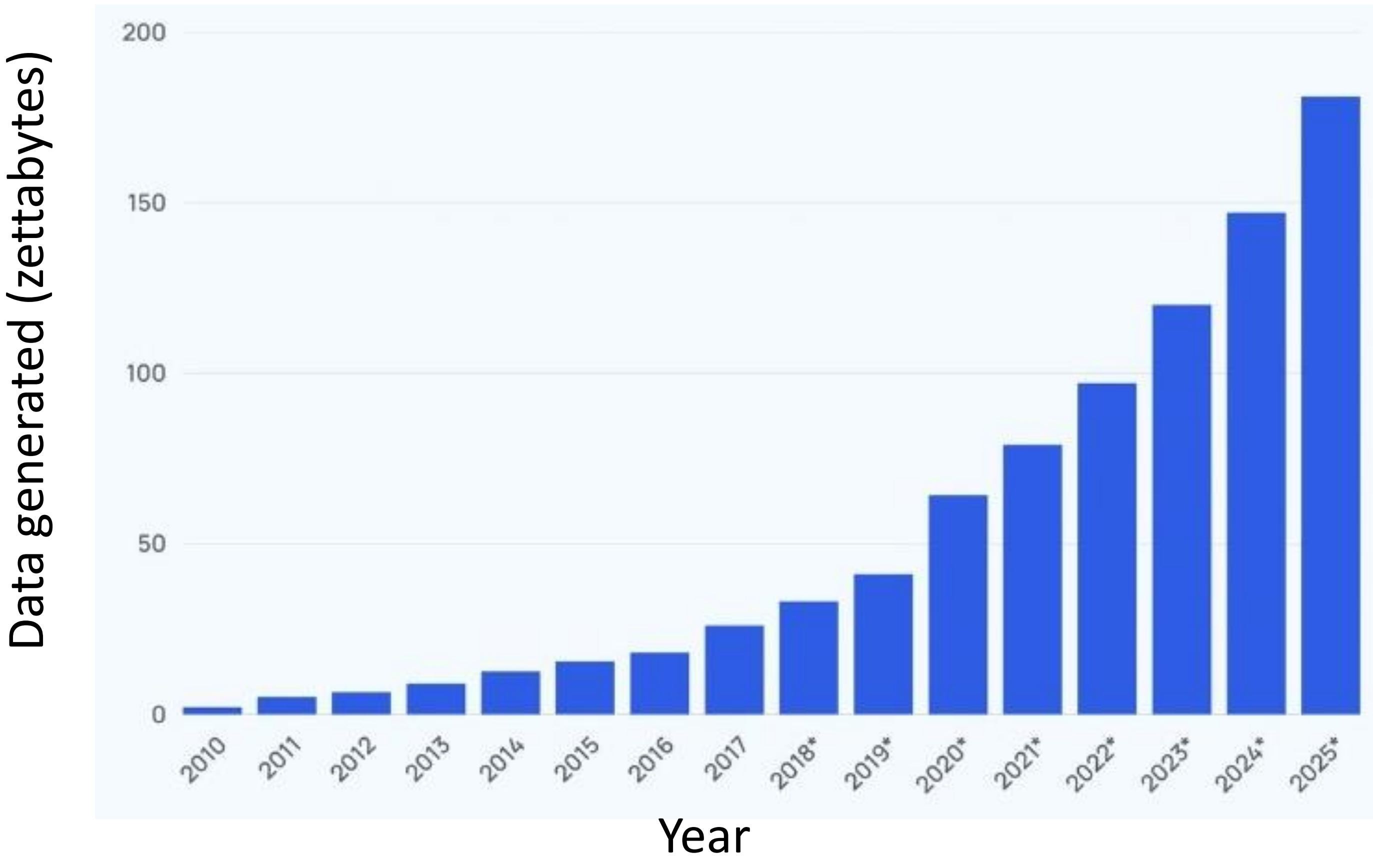

* Amount of data generating daily (Exploding Topics, 2024)

₸ "IDC FutureScape, 2024: Worldwide Future of Digital Infrastructure 2024 Predictions"

Data proliferation

Global data generated annually*

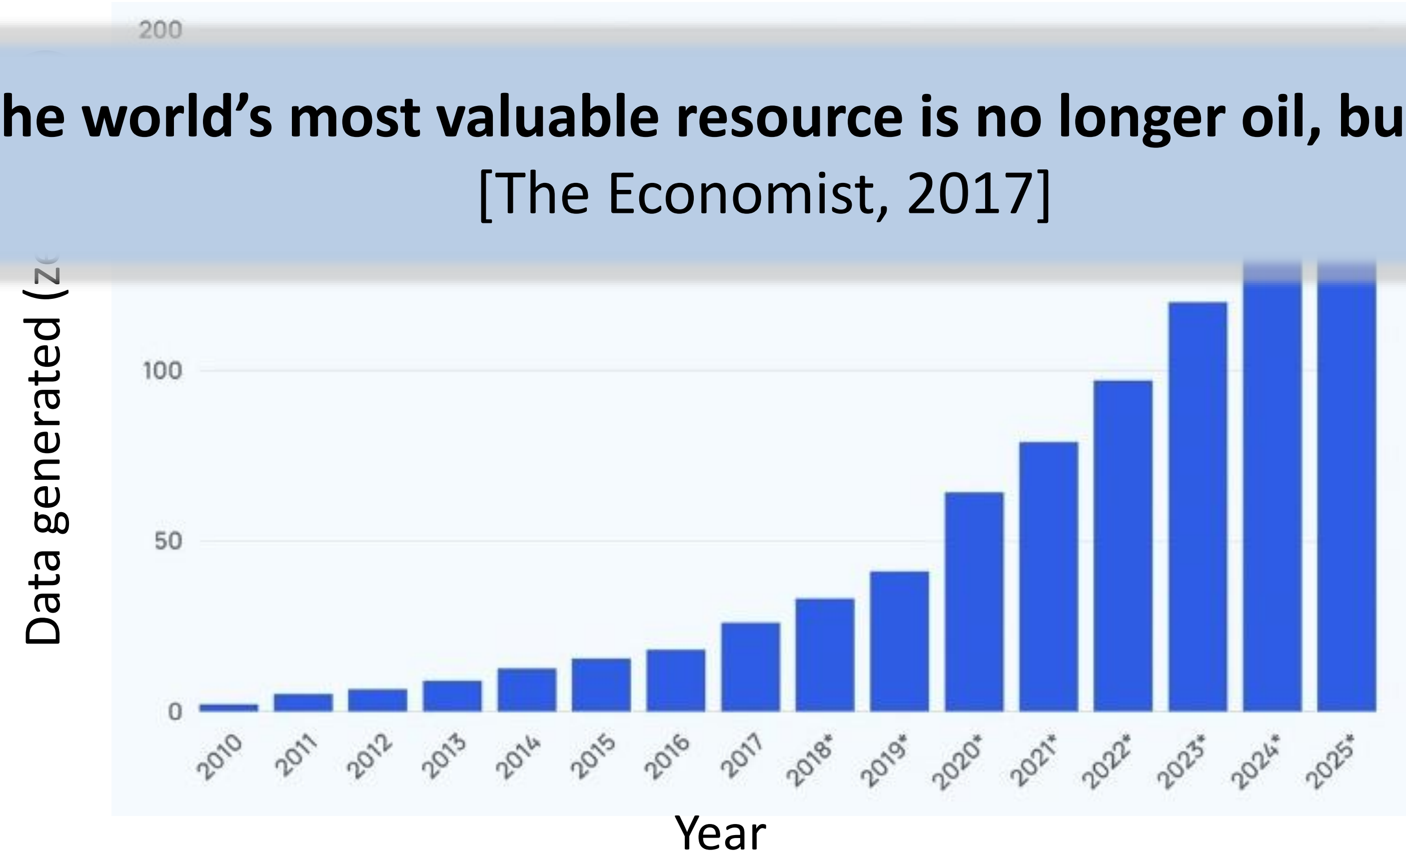

“The world’s most valuable resource is no longer oil, but data”
[The Economist, 2017]

* Amount of data generating daily (Exploding Topics, 2024)

ᚁ “IDC FutureScape, 2024: Worldwide Future of Digital Infrastructure 2024 Predictions”

Data proliferation

Global data generated annually*

“The world’s most valuable resource is no longer oil, but data”
[The Economist, 2017]

“IDC's 2024 predictions for the future of digital infrastructure point to greater emphasis on **fit-for-purpose platforms and services**… By 2025, 70% of companies will invest in alternative computing technologies **to drive business differentiation by compressing time to value of insights from complex data sets…** ”
Ŧ[IDC FutureScape, 2024]

* Amount of data generating daily (Exploding Topics, 2024)

Ŧ “IDC FutureScape, 2024: Worldwide Future of Digital Infrastructure 2024 Predictions”

Need for efficient data exploration

Mining data to uncover patterns, and gather insights

ͳ https://www.abc.net.au/news/science/2022-06-15/black-hole-fastest-growing-past-nine-billionyears/101149598

Need for efficient data exploration

Mining data to uncover patterns, and gather insights

“Recently, the brightest and fastest-growing supermassive black hole of the past 9 billion years was discovered. The researchers have mentioned that "people have been looking for these kinds of objects since the 1960s", and "**somehow, this one seemed to have escaped all our previous efforts to find it**"”

ͲT[ABC News, 2022]

Ͳ https://www.abc.net.au/news/science/2022-06-15/black-hole-fastest-growing-past-nine-billionyears/101149598

From data to insight with databases

typical workflow…

SQL Query 1

Select max(a7), avg(a12) from A join B where A.a1 > 7 and B.r = 45;

results

a7	**a12**
45	23.5

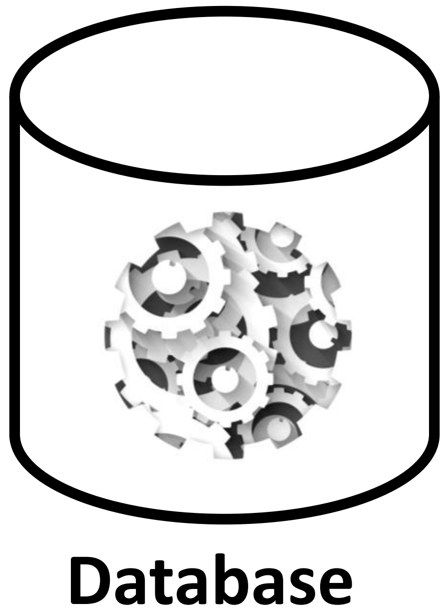

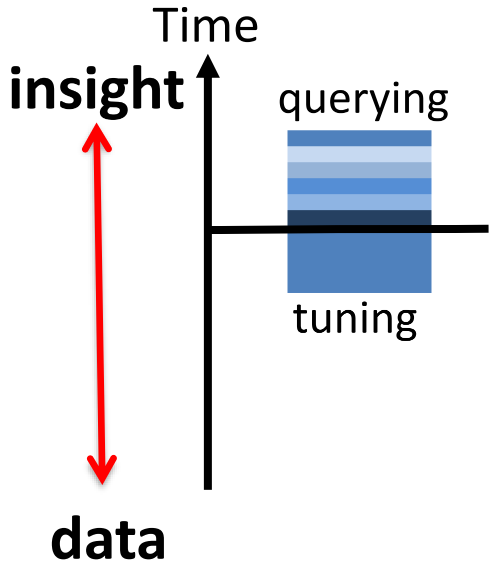

Database goal: minimize data to insight time

How do we minimize data to insight time?

Physical design tuning…

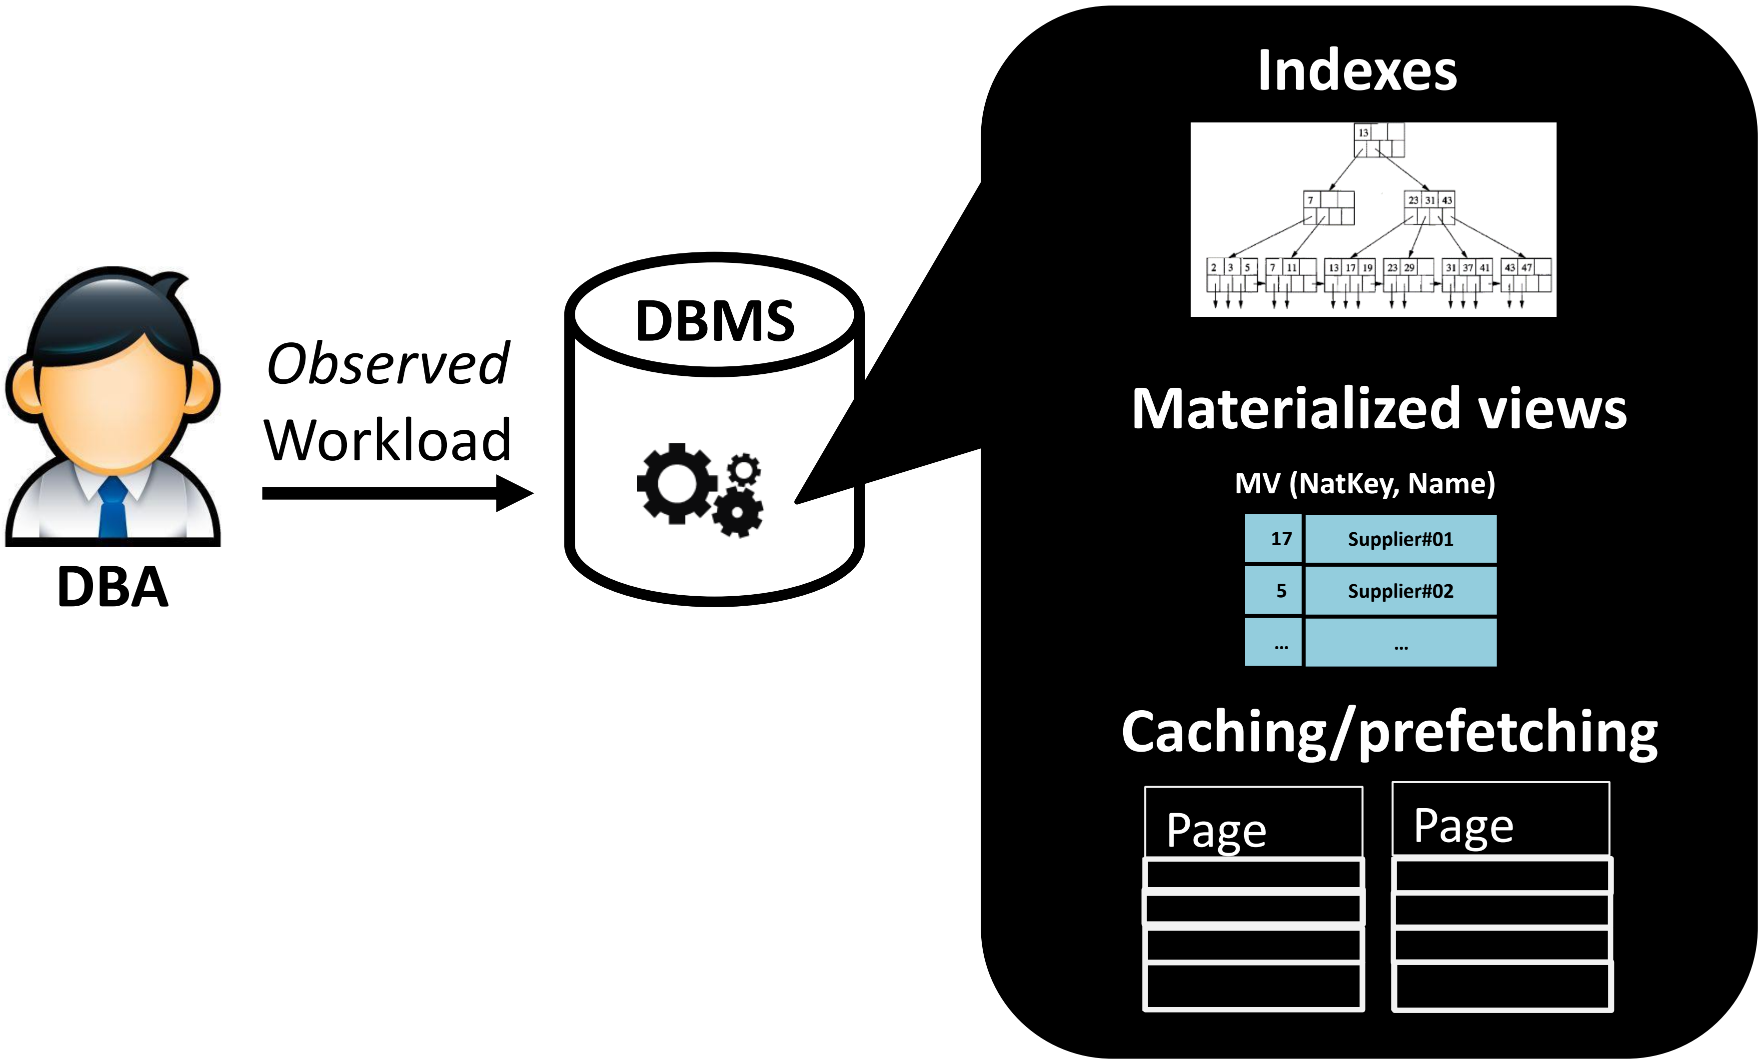

Data exploration properties

- Users are domain experts but **not DB(A) experts**
- **Ad hoc** queries in search of **unknown insights**
- Need for **interactivity** and **adaptivity**

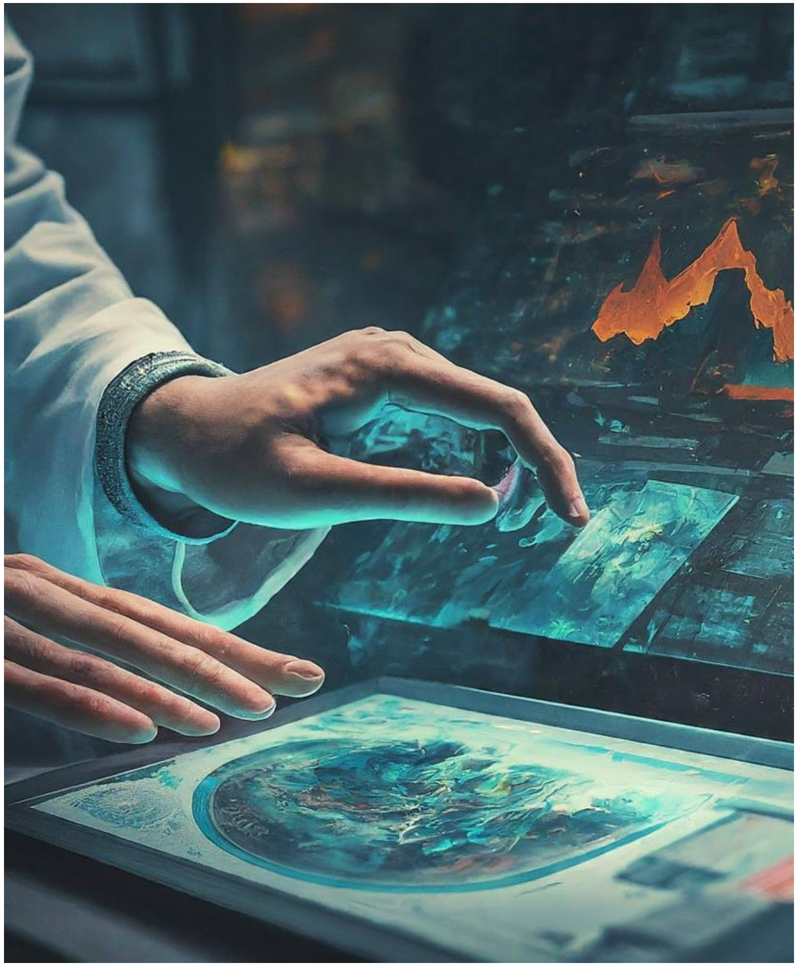

Credit: generated with Gemini

How do we minimize data to insight time in data exploration?

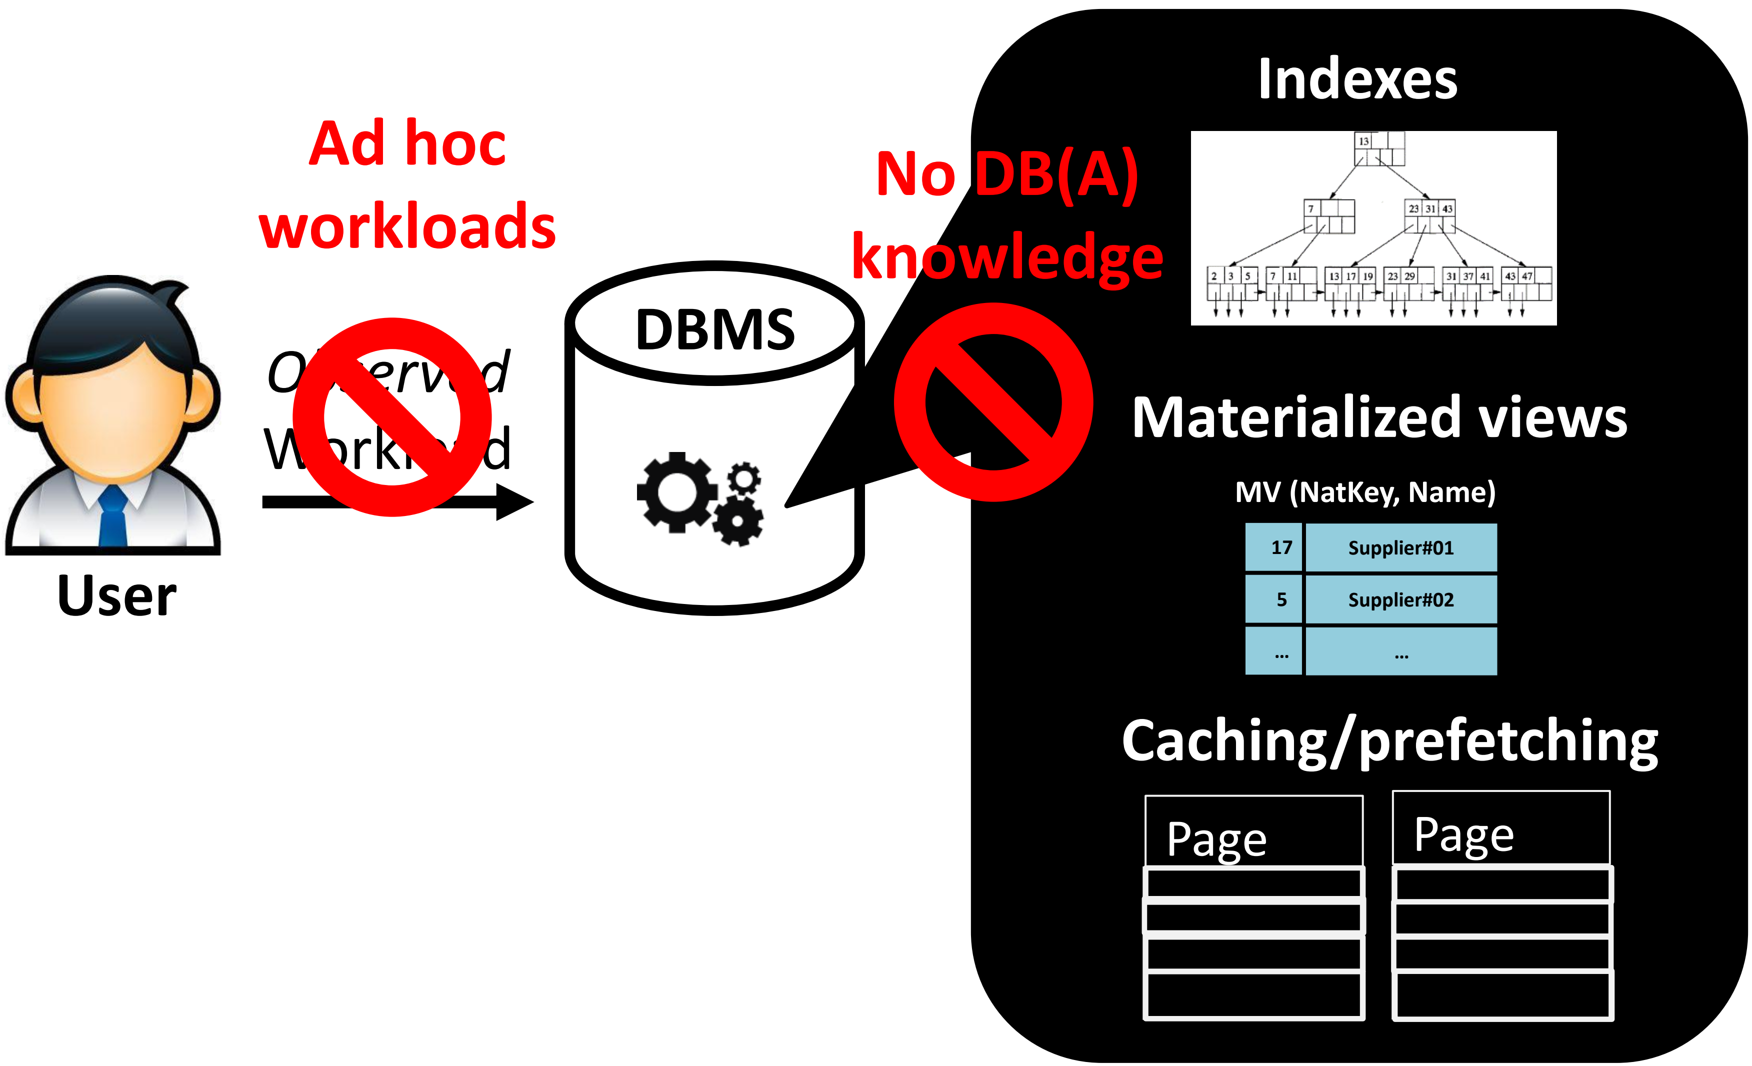

Research gap

Current databases **cannot offer support** for (omnipresent) **data exploration use cases** where users issue **unpredictable queries** in search of **unknown insights**.

Solution

Custom-tailored (AI-driven) databases can **automatically learn from user interactions** with the database and **optimize** its **performance**.

Outline

- **Select physical design structures**

  [ICDE'21, ICDM'21, VLDB'22, TKDE'23, ICDM'24]

- **Tune the layout of physical design structures**

  [SIGMOD'23, ICDE'24]

- **Prefetch data ahead of time**

  [VLDB'24]

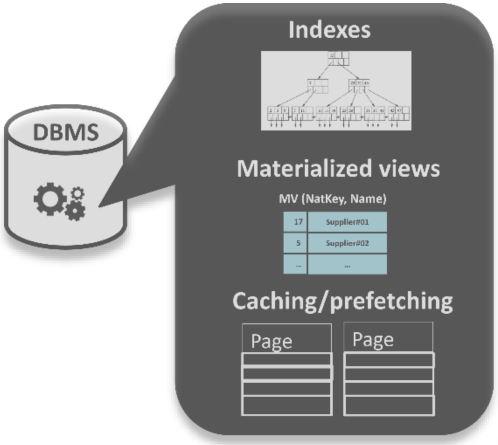

Outline

- **Select physical design structures**

  [ICDE'21, ICDM'21, VLDB'22, TKDE'23, ICDM'24]

- **Tune the layout of physical design structures**

  [SIGMOD'23, ICDE'24]

- **Prefetch data ahead of time**

  [VLDB'24]

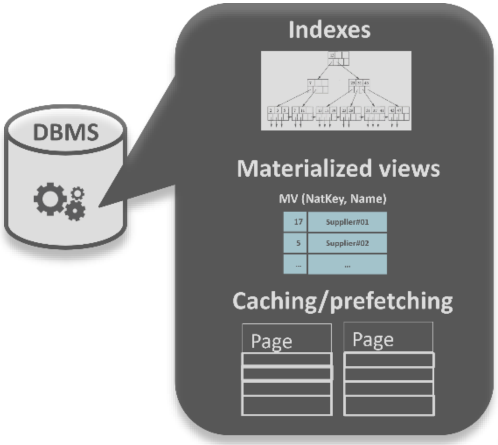

Physical design (PD) tuning is hard

[VLDBJ'18, ICDE'15, DBTest'12]

Setting: TPC-H, SF10, DBMS-X, Tuning tool 5GB space for indexes

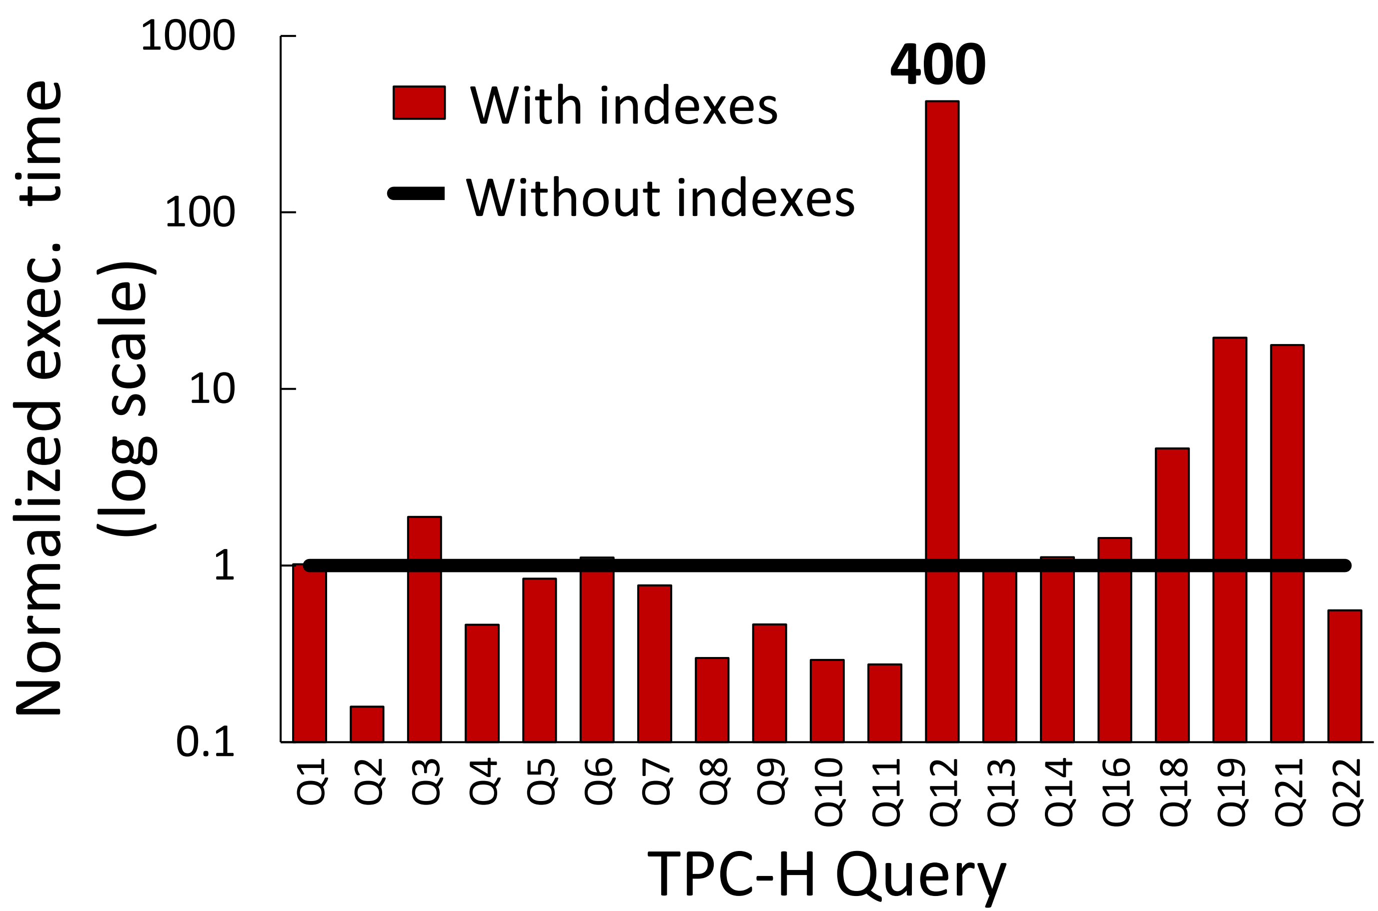

And results can be unpredictable

Physical design tuning under looking glass

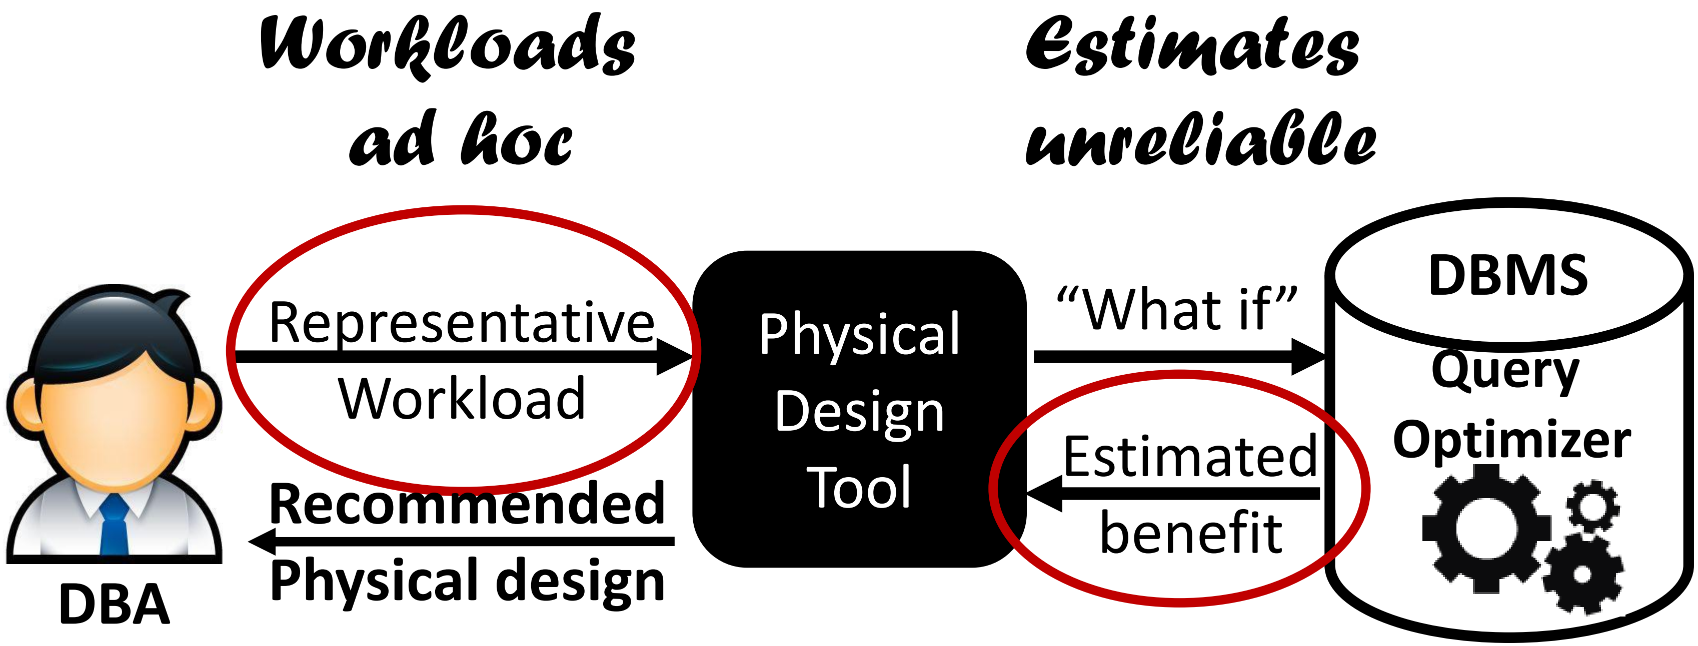

Broken pipeline....

Machine learning to the rescue

Embarking the (M) learning train…

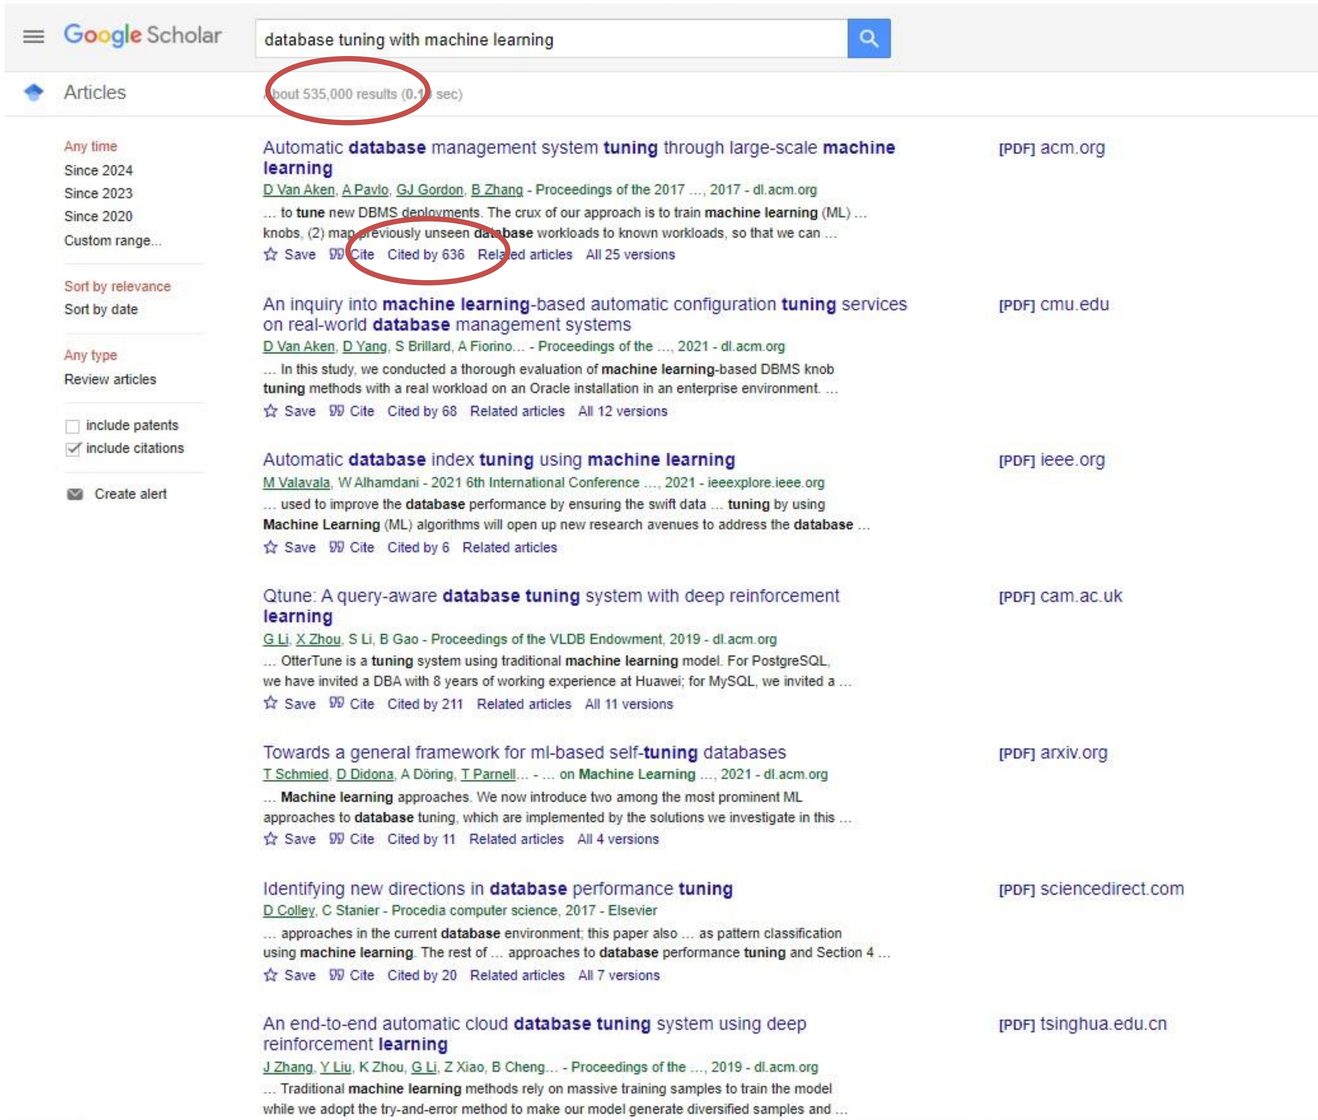

Google Scholar

database tuning with machine learning

Articles

About 535,000 results (0.10 sec)

Any time
Since 2024
Since 2023
Since 2020
Custom range...

Sort by relevance
Sort by date

Any type
Review articles

include patents
include citations

Create alert

Automatic **database** management system **tuning** through large-scale **machine learning** — [PDF] acm.org
D Van Aken, A Pavlo, GJ Gordon, B Zhang - Proceedings of the 2017 ..., 2017 - dl.acm.org
... to **tune** new DBMS deployments. The crux of our approach is to train **machine learning** (ML) ... knobs, (2) map previously unseen **database** workloads to known workloads, so that we can ...
Save Cite Cited by 636 Related articles All 25 versions

An inquiry into **machine learning**-based automatic configuration **tuning** services on real-world **database** management systems — [PDF] cmu.edu
D Van Aken, D Yang, S Brillard, A Fiorino... - Proceedings of the ..., 2021 - dl.acm.org
... In this study, we conducted a thorough evaluation of **machine learning**-based DBMS knob **tuning** methods with a real workload on an Oracle installation in an enterprise environment. ...
Save Cite Cited by 68 Related articles All 12 versions

Automatic **database** index **tuning** using **machine learning** — [PDF] ieee.org
M Valavala, W Alhamdani - 2021 6th International Conference ..., 2021 - ieeexplore.ieee.org
... used to improve the **database** performance by ensuring the swift data ... **tuning** by using **Machine Learning** (ML) algorithms will open up new research avenues to address the **database** ...
Save Cite Cited by 6 Related articles

Qtune: A query-aware **database tuning** system with deep reinforcement **learning** — [PDF] cam.ac.uk
G Li, X Zhou, S Li, B Gao - Proceedings of the VLDB Endowment, 2019 - dl.acm.org
... OtterTune is a **tuning** system using traditional **machine learning** model. For PostgreSQL, we have invited a DBA with 8 years of working experience at Huawei; for MySQL, we invited a ...
Save Cite Cited by 211 Related articles All 11 versions

Towards a general framework for ml-based self-**tuning** databases — [PDF] arxiv.org
T Schmied, D Didona, A Döring, T Parnell... - ... on **Machine Learning** ..., 2021 - dl.acm.org
... **Machine learning** approaches. We now introduce two among the most prominent ML approaches to **database** tuning, which are implemented by the solutions we investigate in this ...
Save Cite Cited by 11 Related articles All 4 versions

Identifying new directions in **database** performance **tuning** — [PDF] sciencedirect.com
D Colley, C Stanier - Procedia computer science, 2017 - Elsevier
... approaches in the current **database** environment; this paper also ... as pattern classification using **machine learning**. The rest of ... approaches to **database** performance **tuning** and Section 4 ...
Save Cite Cited by 20 Related articles All 7 versions

An end-to-end automatic cloud **database tuning** system using deep reinforcement **learning** — [PDF] tsinghua.edu.cn
J Zhang, Y Liu, K Zhou, G Li, Z Xiao, B Cheng... - Proceedings of the ..., 2019 - dl.acm.org
... Traditional **machine learning** methods rely on massive training samples to train the model while we adopt the try-and-error method to make our model generate diversified samples and ...

Multi-armed bandits (MAB) for PD tuning

- Pull an arm (slot machine) observe a reward (win/lose)
- Explore vs exploit
- Find a sequence of arms to maximize reward
- Many variants, but **C²UCB** most interesting

Optimism in the face of uncertainty

Index tuning with Multi-Armed Bandits MAB (C^2UCB)

[ICDE'21]

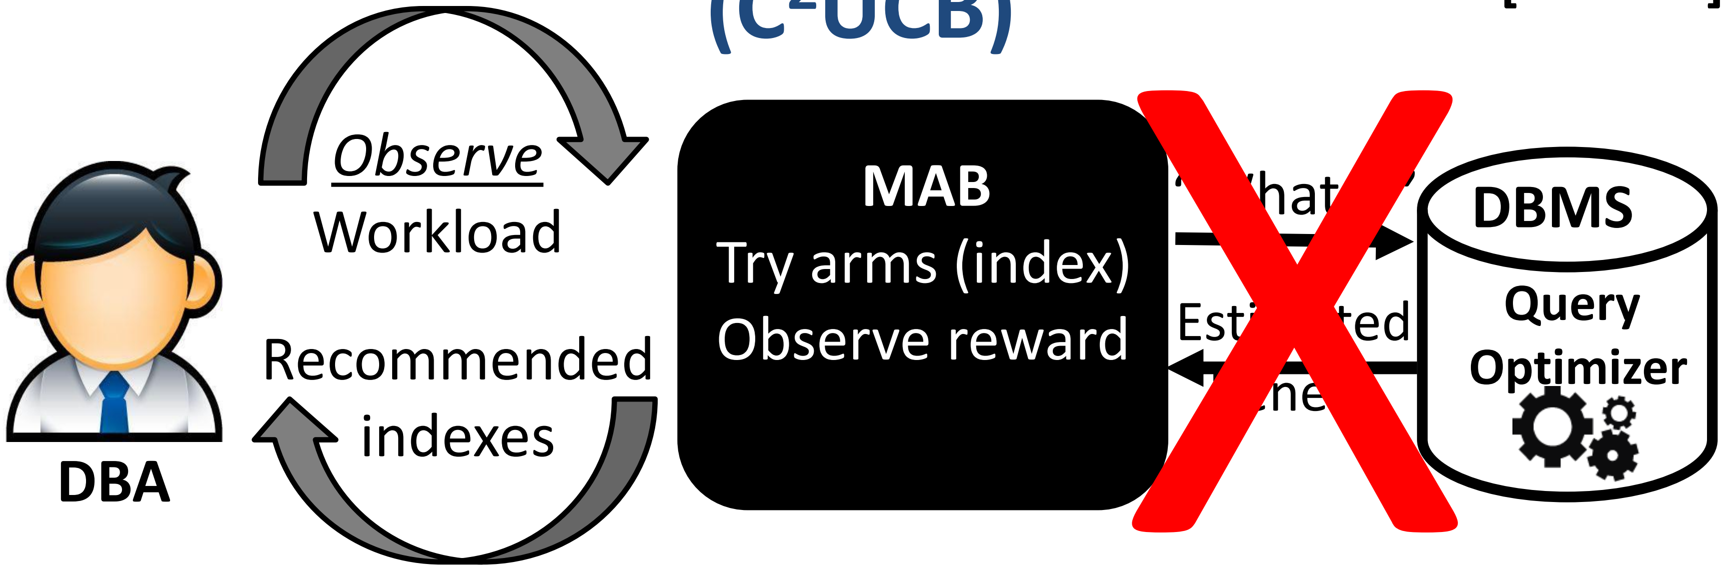

- **UCB** *guarantees* to converge to optimal policy (**effectiveness**)
- **C** *(contextual)* learns benefit of arms *without* pulling them (**efficiency**)
- **C** *(combinatorial)* pulls *a set* of arms per round given constraints (**efficiency**)

Safety guarantees with fast convergence

[ICDE'21] DBA bandits: Self-driving index tuning under ad-hoc, analytical workloads with safety guarantees. M. Perera, B. Oetomo, B. Rubinstein, and R. Borovica-Gajic.

MAB under looking glass... [ICDE'21]

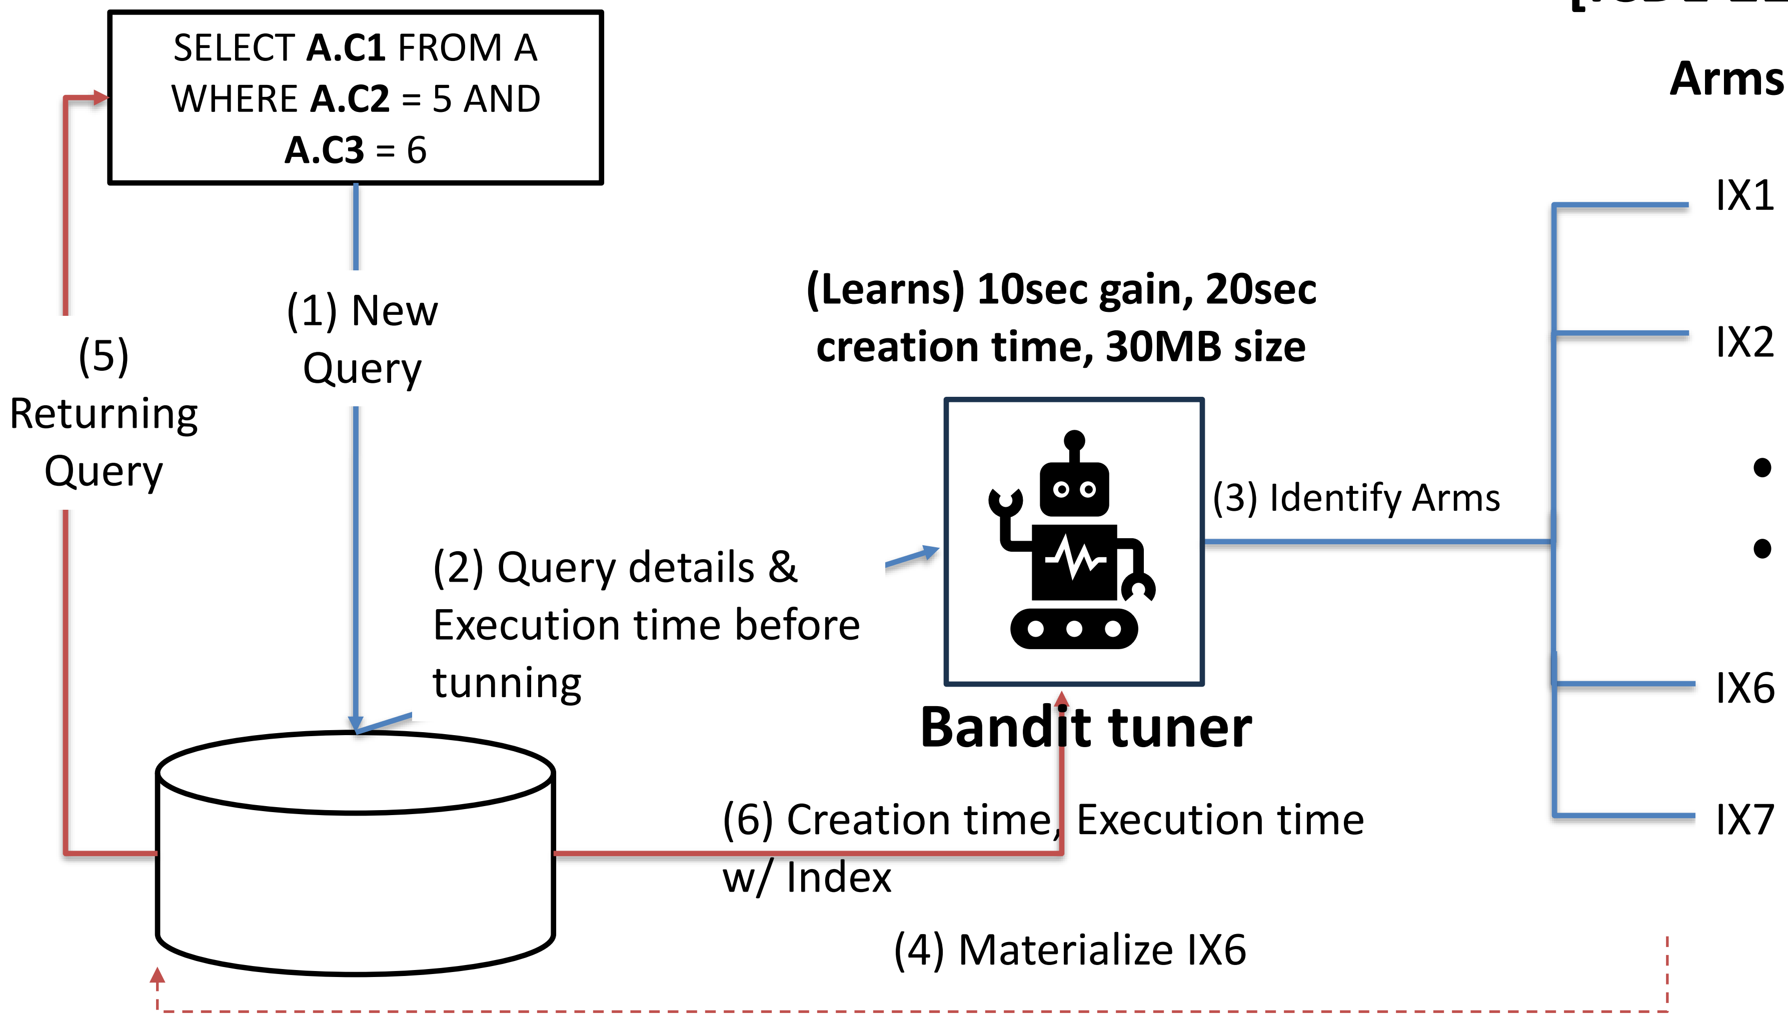

Automated tuning with provable guarantees

[ICDE'21] DBA bandits: Self-driving index tuning under ad-hoc, analytical workloads with safety guarantees. M. Perera, B. Oetomo, B. Rubinstein, and R. Borovica-Gajic.

MAB to the rescue

Setting: TPC-H, SF10, DBMS-X, Multi-armed bandits (MAB) for index tuning

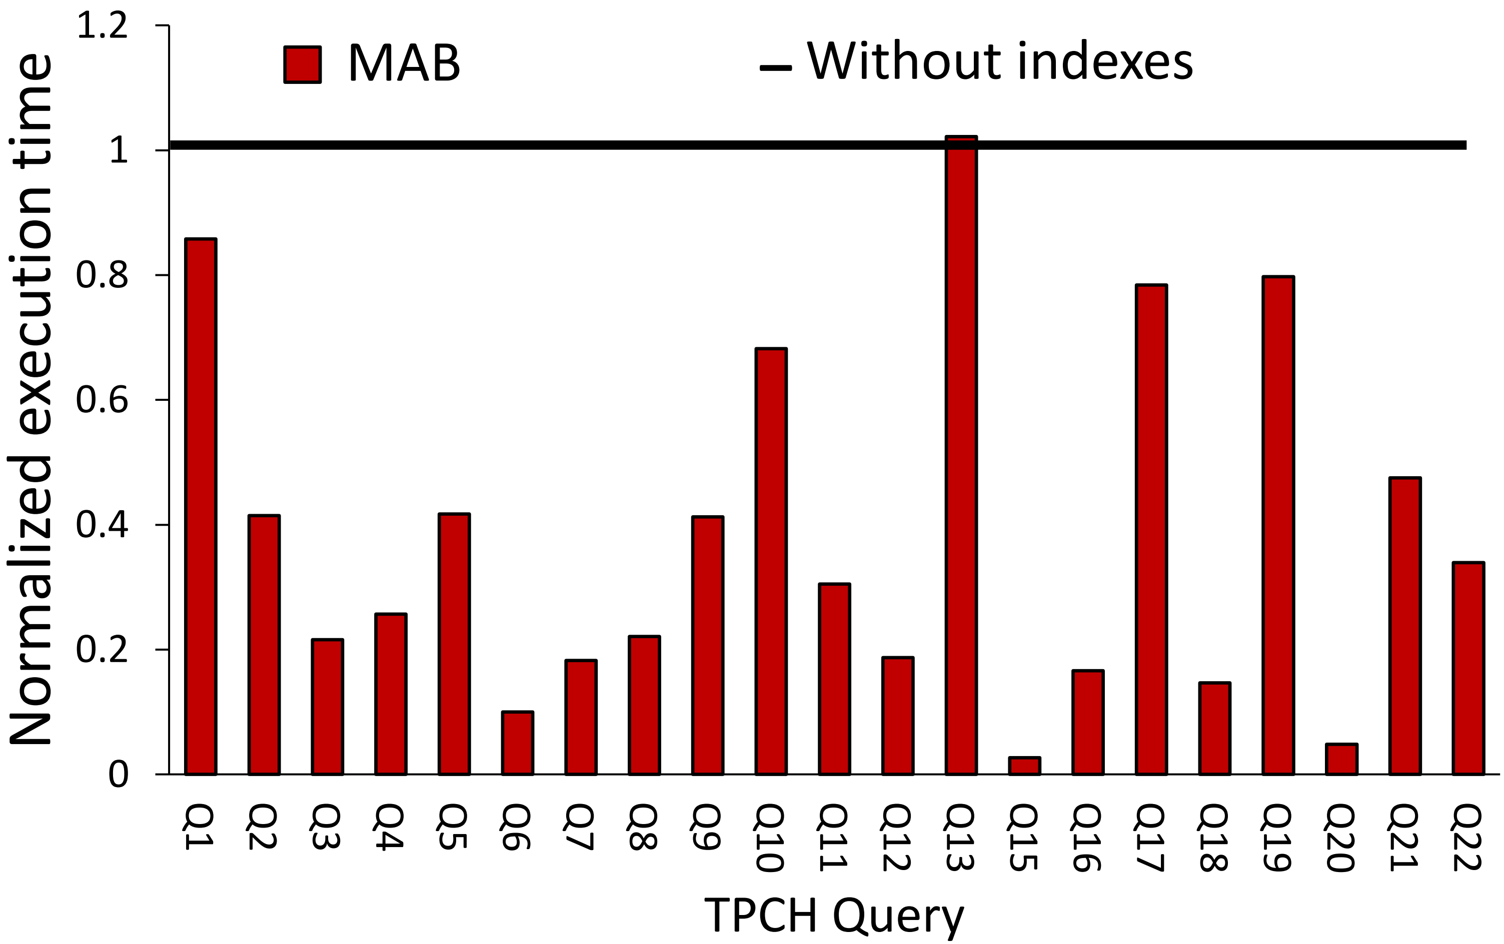

3x Speed up vs. previous 22x slowdown

MAB in action

[ICDE'21]

Setting: TPCH, TPCH skew, TPC DS, SSB (10GB); IMDb(6GB) datasets static (repetitive) vs random (ad hoc) queries, MAB vs PDTool, 25 rounds

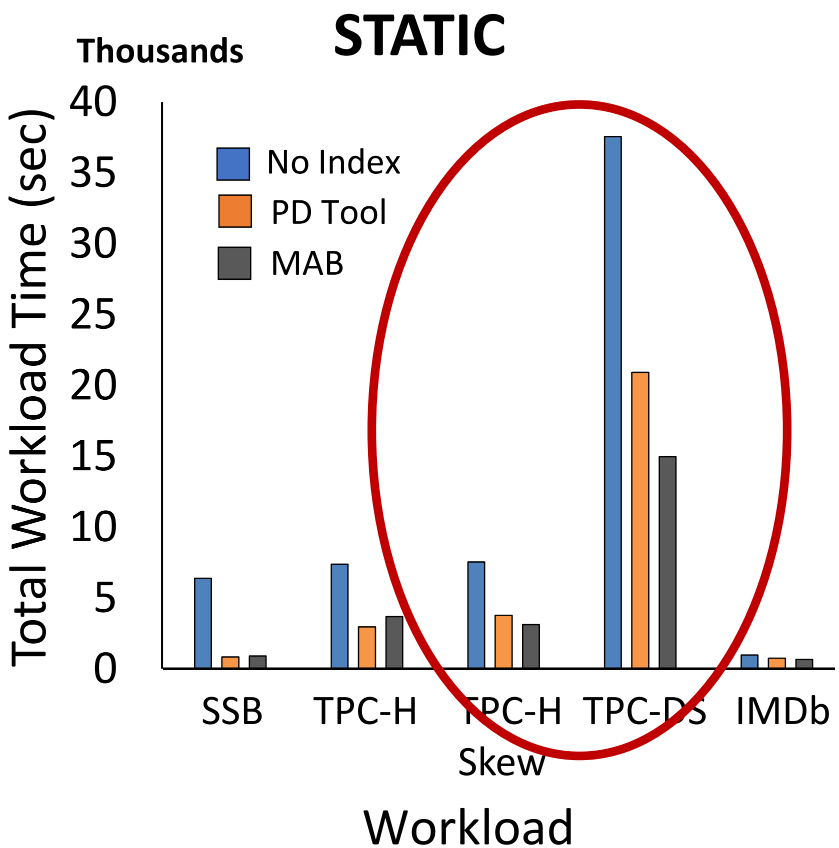

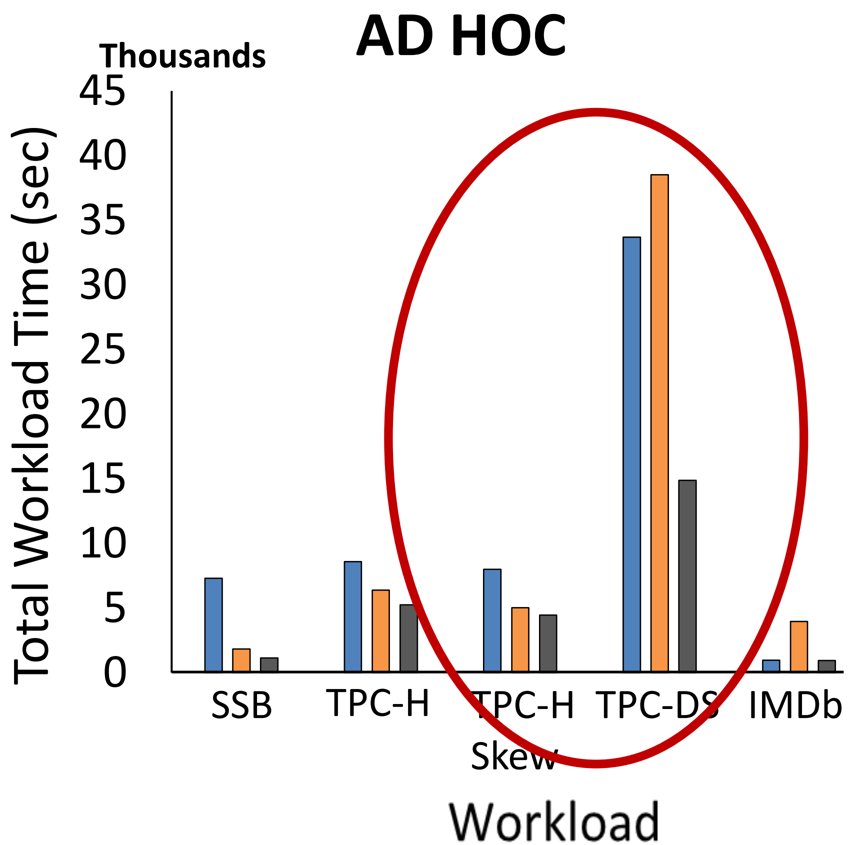

MAB robust against complex unpredictable workloads and skew

MAB for ~~Index~~ Tuning: An Example

Physical Design

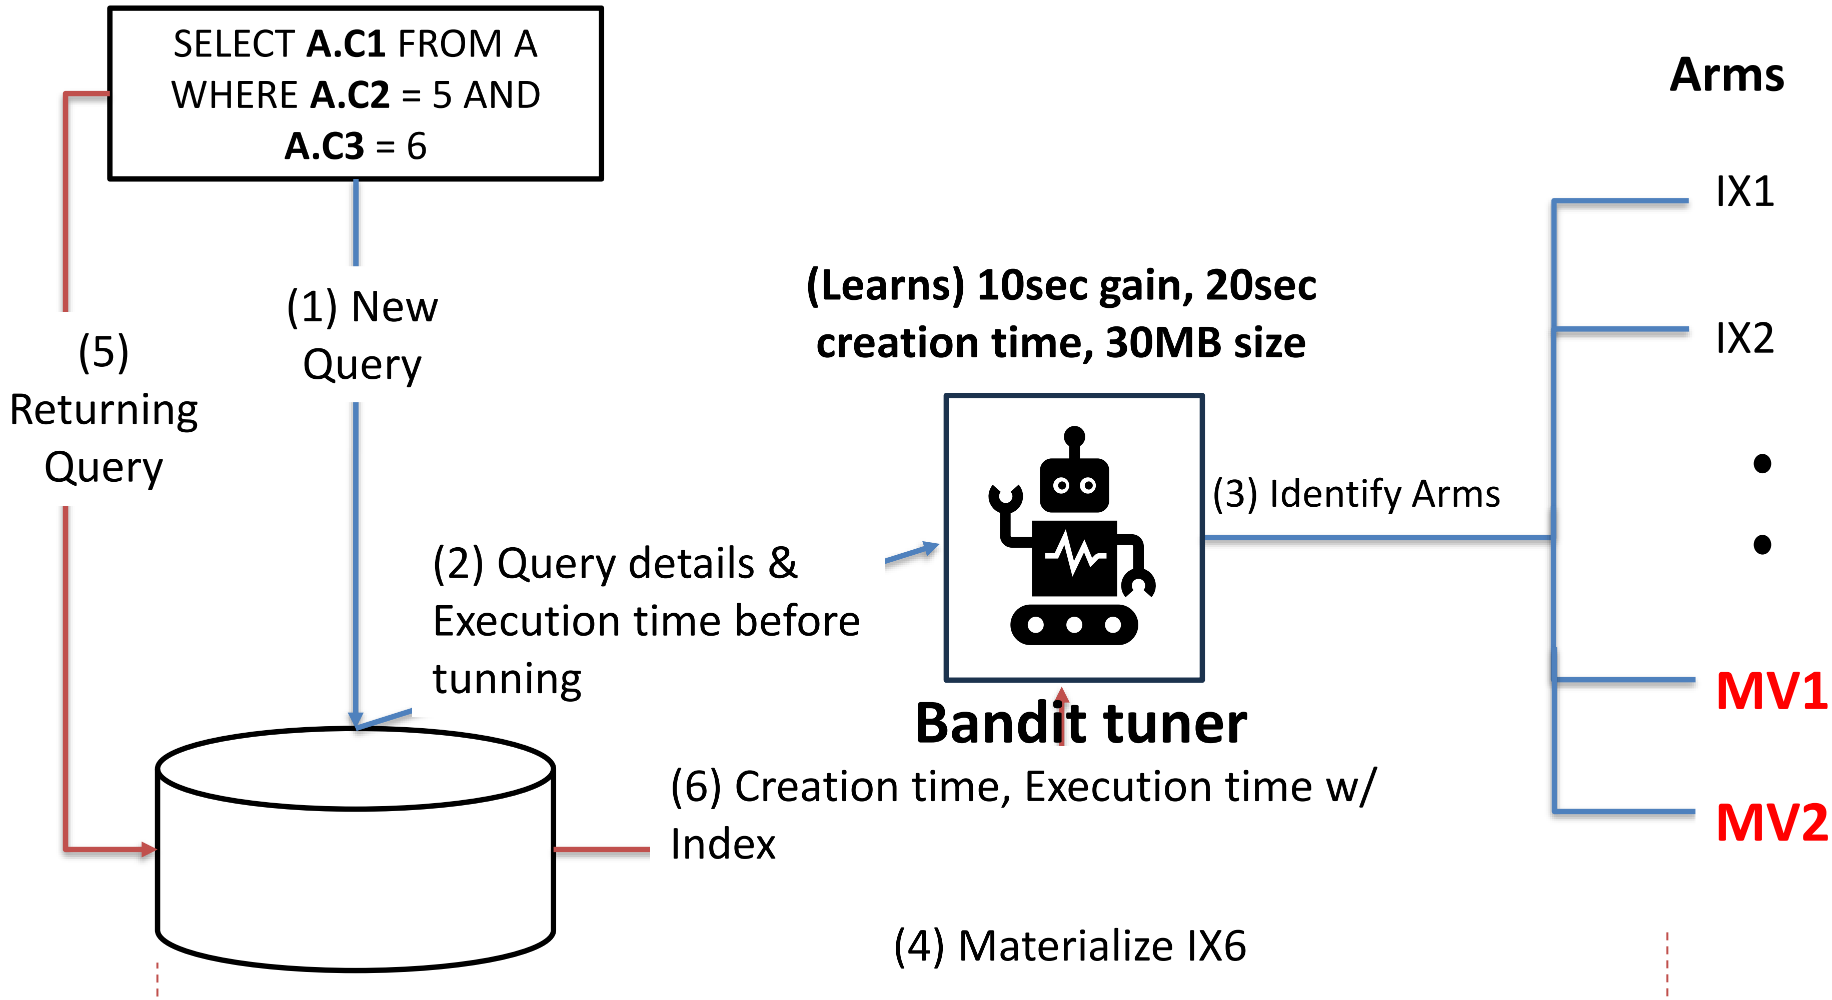

Design too complex, too large action space

HMAB: Hierarchical Multi-armed Bandit Architecture for Integrated Physical Design Tuning

[VLDB'22]

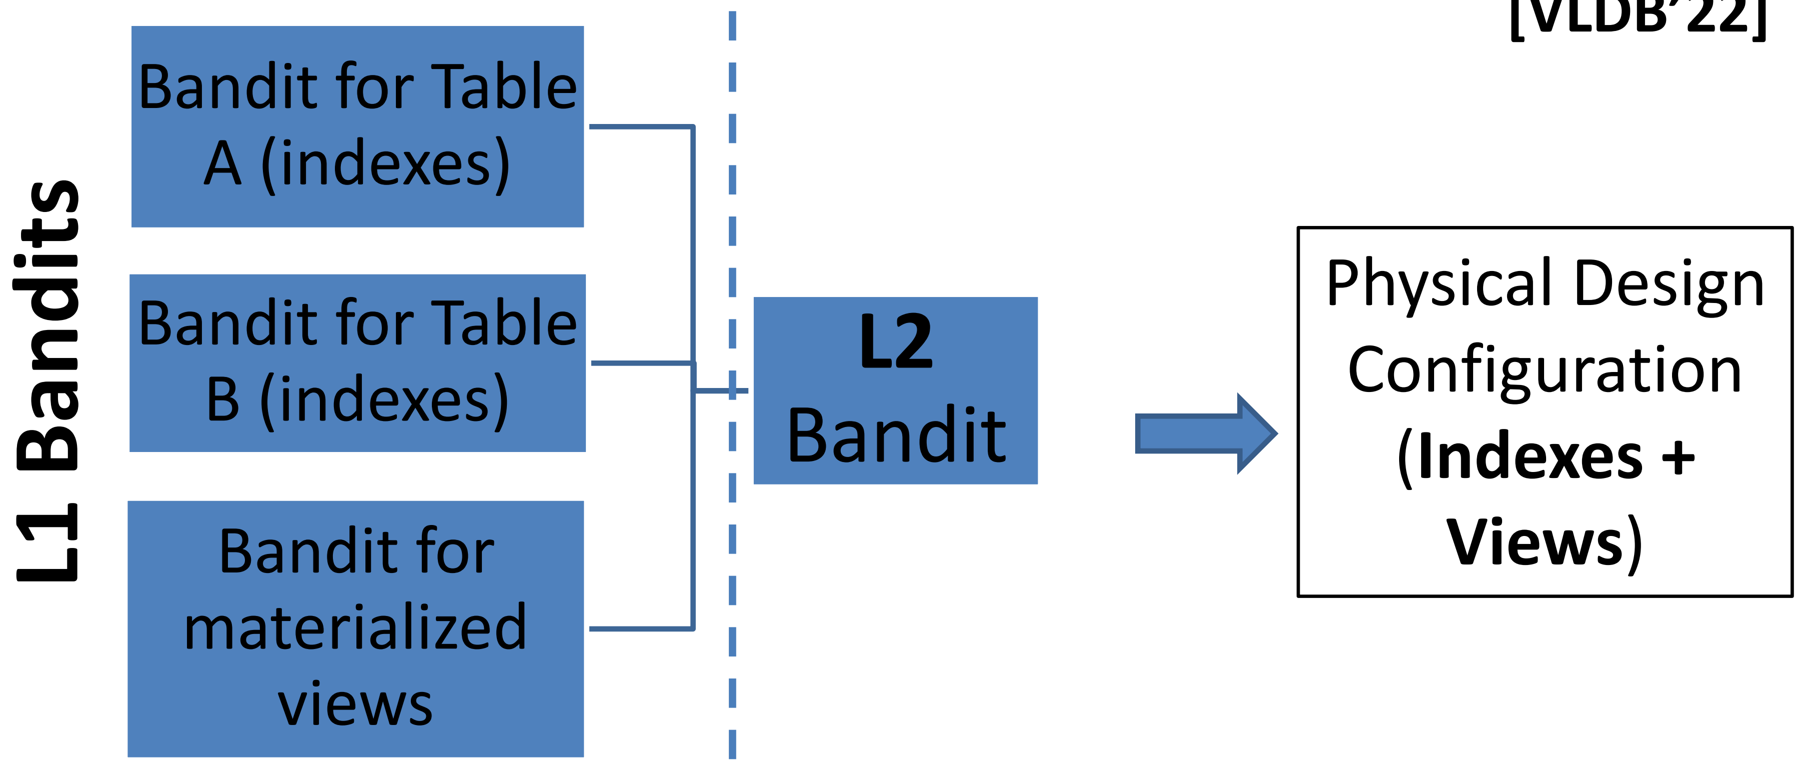

Smaller bandits for faster convergence – divide and conquer

Knowledge sharing via central bandit – global optimality

[VLDB'22] HMAB: Self-Driving Hierarchy of Bandits for Integrated Physical Database Design Tuning. M. Perera, B. Oetomo, B. Rubinstein, and R. Borovica-Gajic.

HMAB in Action

[VLDB'22]

Setting: TPCH, TPCH skew, TPC DS, IMDb datasets; static (repetitive) vs random (ad hoc) queries, MAB vs PDTool, 25 rounds, tuning indices and materialised views

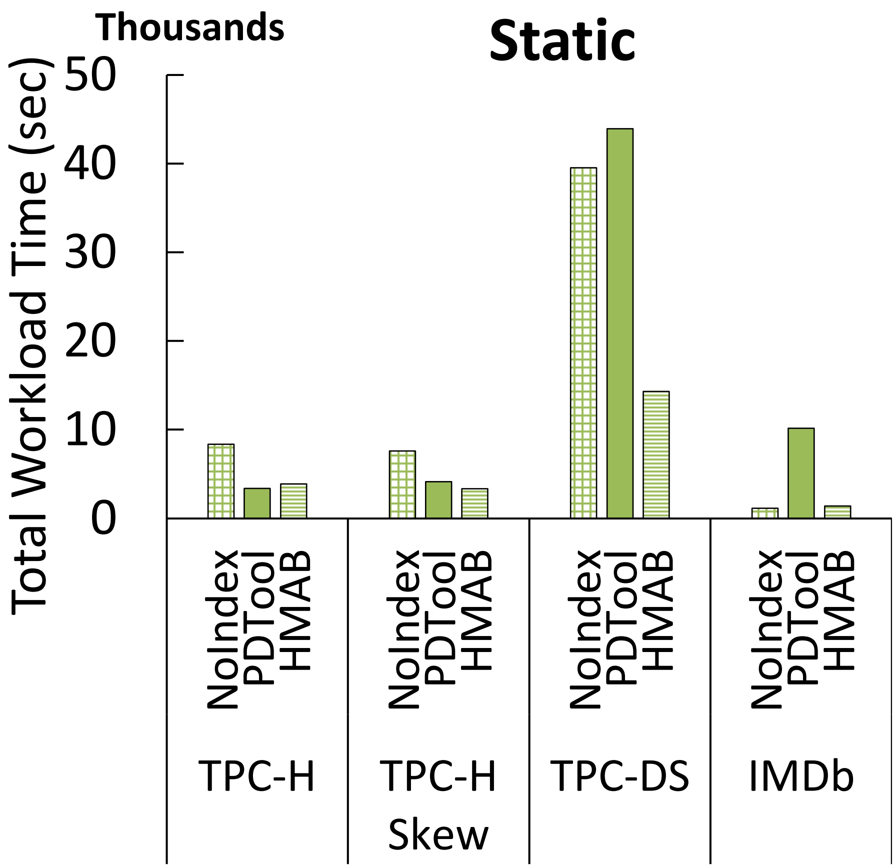

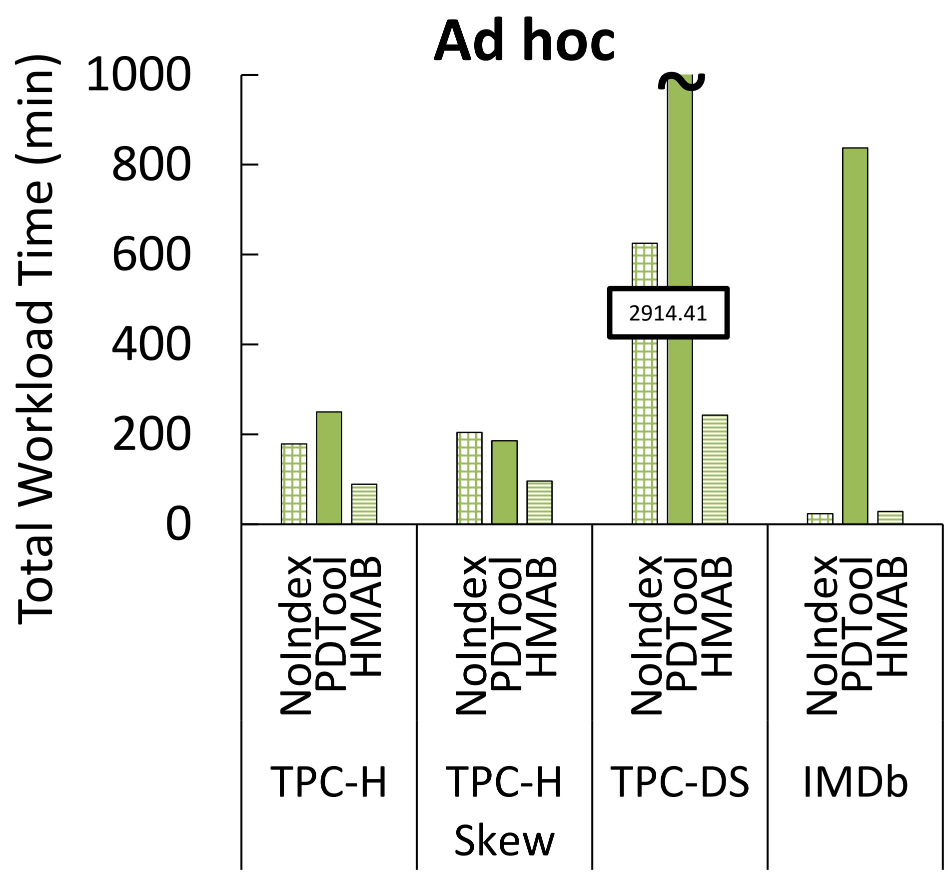

Up to 96% speed-up, and 67% on average

Dealing with complexity (HTAP) [TKDE’23]

Setting: CH-BenCHmark under static workloads, MAB vs. PDTool, 25 rounds

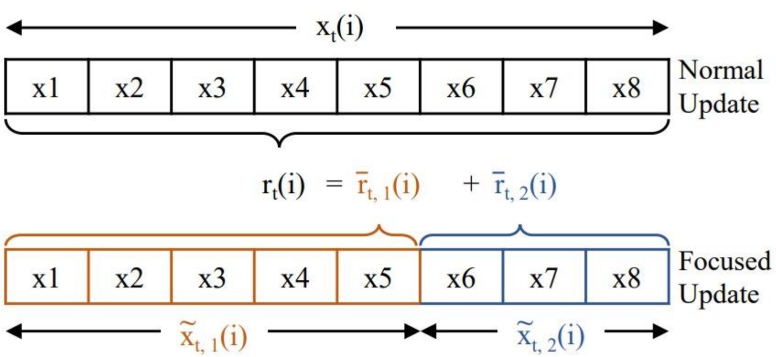

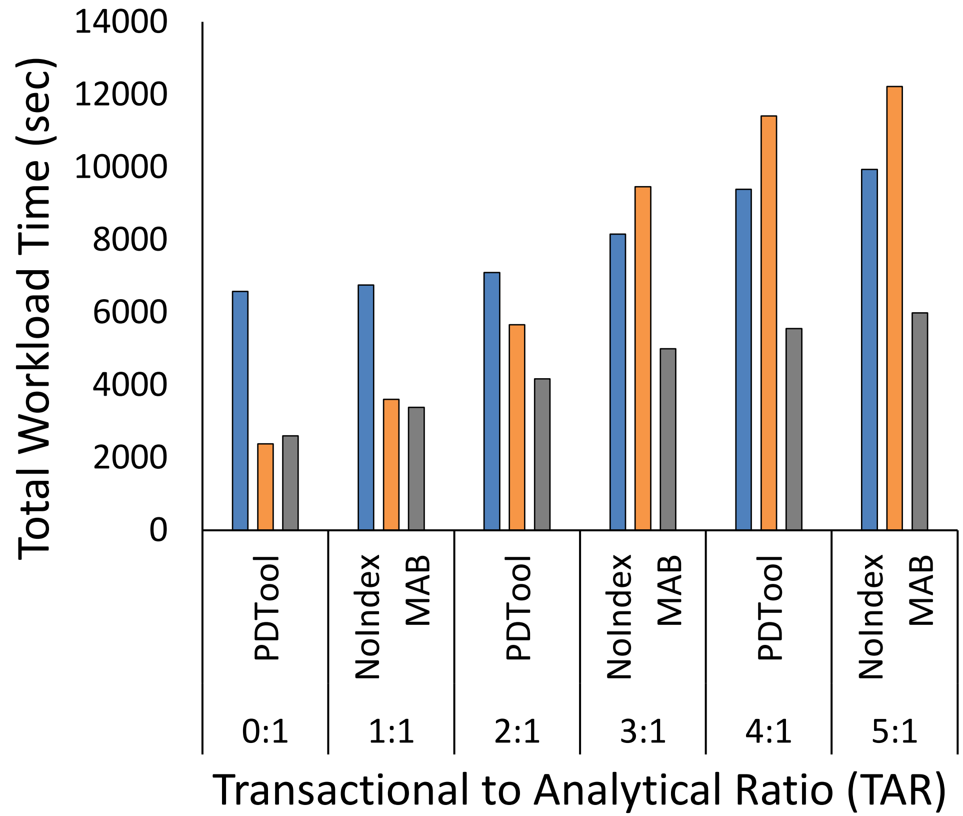

MAB with focused updates to support HTAP

New bandit flavor with better regret bounds

[TKDE’23] No DBA? No regret! Multi-armed bandits for index tuning of analytical and HTAP workloads with provable guarantees. M. Perera, B. Oetomo, B. Rubinstein, R. Borovica-Gajic.

MAB Summary

- (H)MAB is a lightweight MAB solution for *(integrated)* physical database design tuning
- HMAB is the first learned solution to work in the combined space of indices and views
- (H)MAB successfully tackles tuning challenges: optimizer *misestimates, unpredictable and HTAP* workloads
- Up to 40% and 70% average improvement for integrated view and index tuning under static and random settings compared against a SOTA commercial tuning tool
- **Extensions**: bandit warm up [ICDM’21], bandits under latent reward scaling [ICDM’24]

Outline

- Select physical design structures

  [ICDE'21, ICDM'21, VLDB'22, TKDE'23, ICDM'24]

- **Tune the layout of physical design structures**

  [SIGMOD'23, ICDE'24]

- Prefetch data ahead of time

  [VLDB'24]

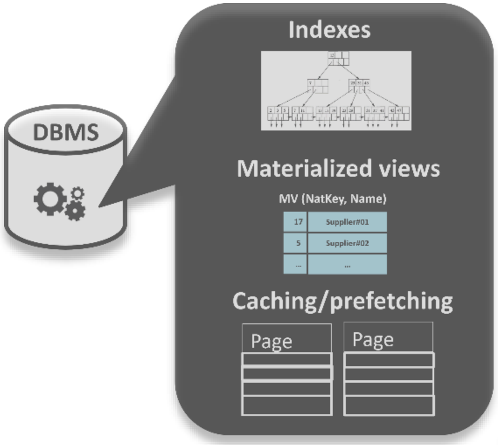

Classic vs learned index layout

Classic B+ tree

Learned index

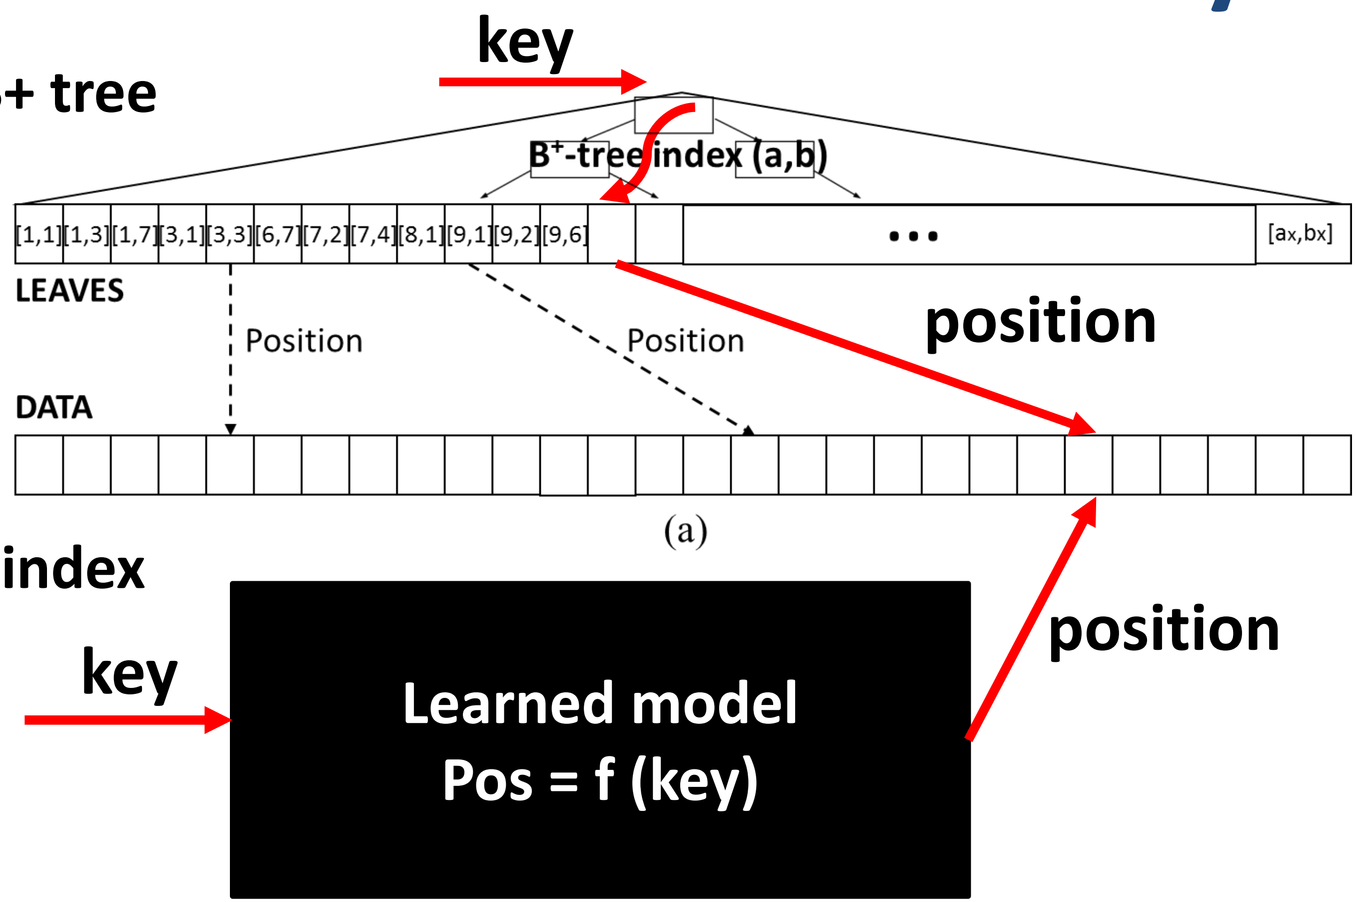

Learned indexes promise lower memory footprint and faster lookup

(M) Learned indexes …

Google Scholar | learned index

Articles

About 9,460,000 results (0.?1 sec)

Any time
Since 2024
Since 2023
Since 2020
Custom range...

Sort by relevance
Sort by date

Any type
Review articles

include patents
include citations

Create alert

The case for learned index structures [PDF] acm.org
T Kraska, A Beutel, EH Chi, J Dean... - Proceedings of the 2018 ..., 2018 - dl.acm.org
... The remainder of this paper is outlined as follows: In the next two sections we introduce the general idea of **learned indexes** using B-Trees as an example. In Section 4 we extend this ...
Save Cite Cited by 1125 Related articles All 23 versions

ALEX: an updatable adaptive learned index [PDF] acm.org
J Ding, UF Minhas, J Yu, C Wang, J Do, Y Li... - Proceedings of the ..., 2020 - dl.acm.org
... on "**learned indexes**" has changed the way we look at the decades-old field of DBMS **indexing**. ... In this paper, we present a new **learned index** called ALEX which addresses practical ...
Save Cite Cited by 302 Related articles All 10 versions

Learned index: A comprehensive experimental evaluation [PDF] vldb.org
Z Sun, X Zhou, G Li - Proceedings of the VLDB Endowment, 2023 - dl.acm.org
... of new **learned** indexes for researchers. We compare state-of-the-art **learned indexes** in the ... , and provide findings to select suitable **learned indexes** under various practical scenarios. ...
Save Cite Cited by 24 Related articles All 2 versions

Why are learned indexes so effective? [PDF] mlr.press
P Ferragina, F Lillo... - ... on Machine Learning, 2020 - proceedings.mlr.press
... This is especially known in the context of **indexing** data ... that **learned** indexes are provably better than classic **indexes**, ... and time occupancy of those **learned indexes**. Our general result ...
Save Cite Cited by 44 Related articles All 15 versions

Related searches

learned index **structures**	**updatable** learned index
spatial learned index	learned index **scheme in storage**
learned index **alex**	learned index **string keys**
scalable learned index	learned index **lisa**

RadixSpline: a single-pass learned index [PDF] acm.org
A Kipf, R Marcus, A van Renen, M Stoian... - Proceedings of the third ..., 2020 - dl.acm.org
... While this is a very promising result, existing **learned** structures are ...), a **learned index** that can be built in a single pass over the data and is competitive with state-of-the-art **learned index** ...
Save Cite Cited by 183 Related articles All 6 versions

Cdfshop: Exploring and optimizing learned index structures [PDF] acm.org
R Marcus, E Zhang, T Kraska - Proceedings of the 2020 ACM SIGMOD ..., 2020 - dl.acm.org
... models (**learned index** structures) can achieve low lookup ... model **indexes** (RMIs), a type of **learned index** structure. This ... of RMIs and why **learned index** structures can greatly accelerate ...
Save Cite Cited by 61 Related articles All 2 versions

Are learned indexes disk ready?

[SIGMOD'23]

Normalized throughputs on the FB dataset

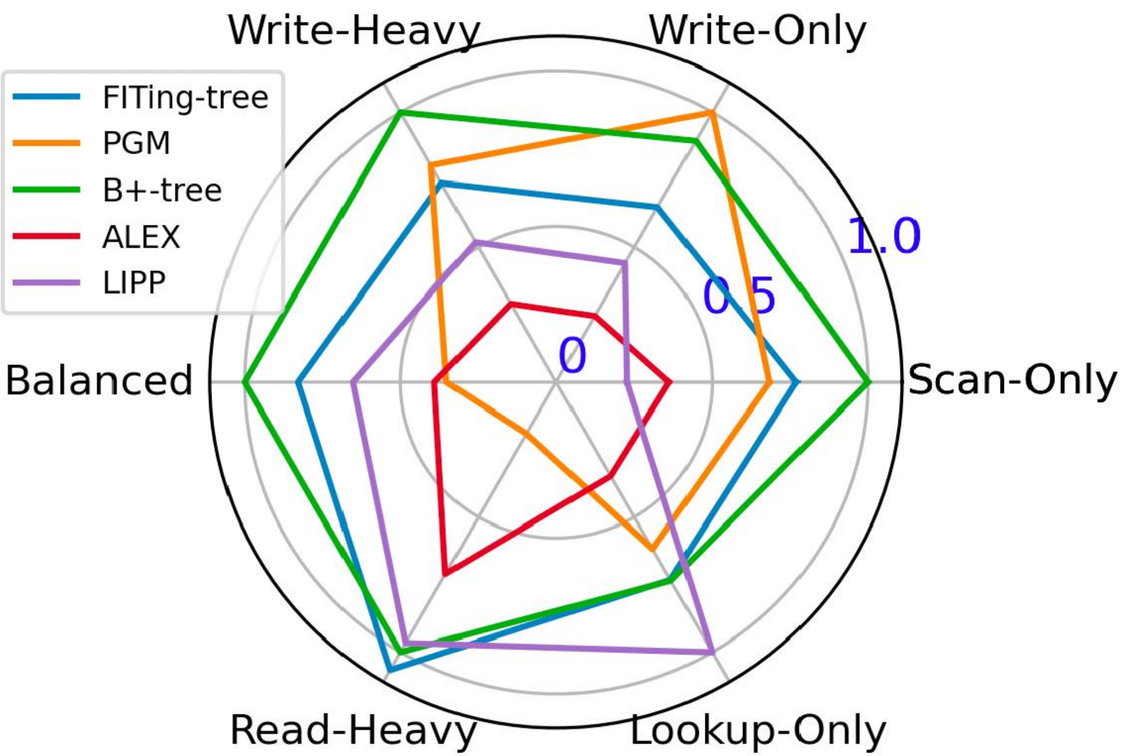

B+tree (still) the best choice when disk resident

[SIGMOD'23] Updatable Learned Indexes Meet Disk-Resident DBMS - From Evaluations to Design Choices. H. Lan, Z. Bao, S. Culpepper, and R. Borovica-Gajic.

Where does time go?

[ICDE'24]

#blocks fetched (for reads)

	# Inner Nodes	# Inner Blocks	# Total Blocks (L)	# Total Blocks (S)
FITing-tree	5	3	4.2	5
PGM	6	3.9	5.2	5.6
ALEX	7.7	6.5	8.1	10.6
LIPP	1.8 (18.8)	-	3	24
B+-tree	4	3	4	4.5

Lookup

Scan

Latency breakdown (for writes)

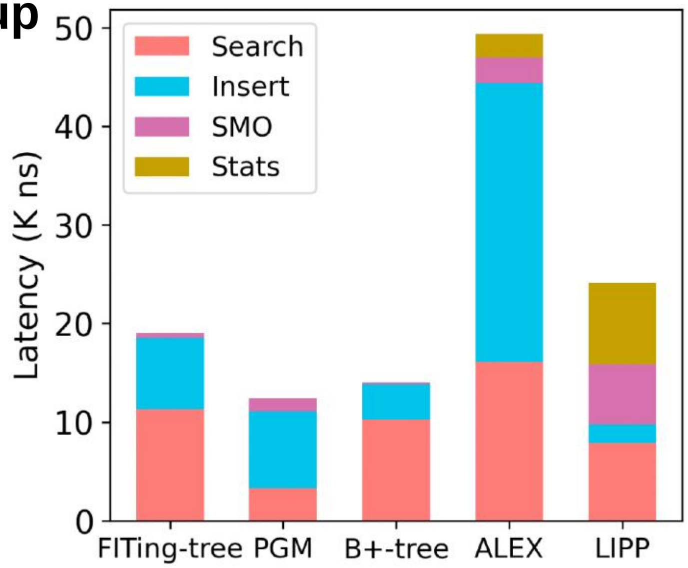

- **Challenge 1:** A learned index cannot guarantee to reduce **I/O costs** when searching data on disk.
- **Challenge 2:** Most learned indexes suffer from large **insertion overheads**.

[ICDE'24] A Fully On-disk Updatable Learned Index. H. Lan, Z. Bao, S. Culpepper, R. Borovica-Gajic and Y. Dong.

Design principles for effective on disk learned index

[ICDE'24]

Challenge 1. A learned index cannot guarantee to reduce **I/O costs** when searching data on disk.

Challenge 2. Most learned indexes suffer from large **insertion overheads**.

P1. Reducing the Tree Height of the Index

P2. Model-based Operations (Search and Insert)

P3. Lightweight Structure Modification Operations

P4. Improve Scan Performance

P5. Support Duplicate Index Keys

AULID: an updatable learned index on disk

Simple Yet Effective

AULID Index Layout

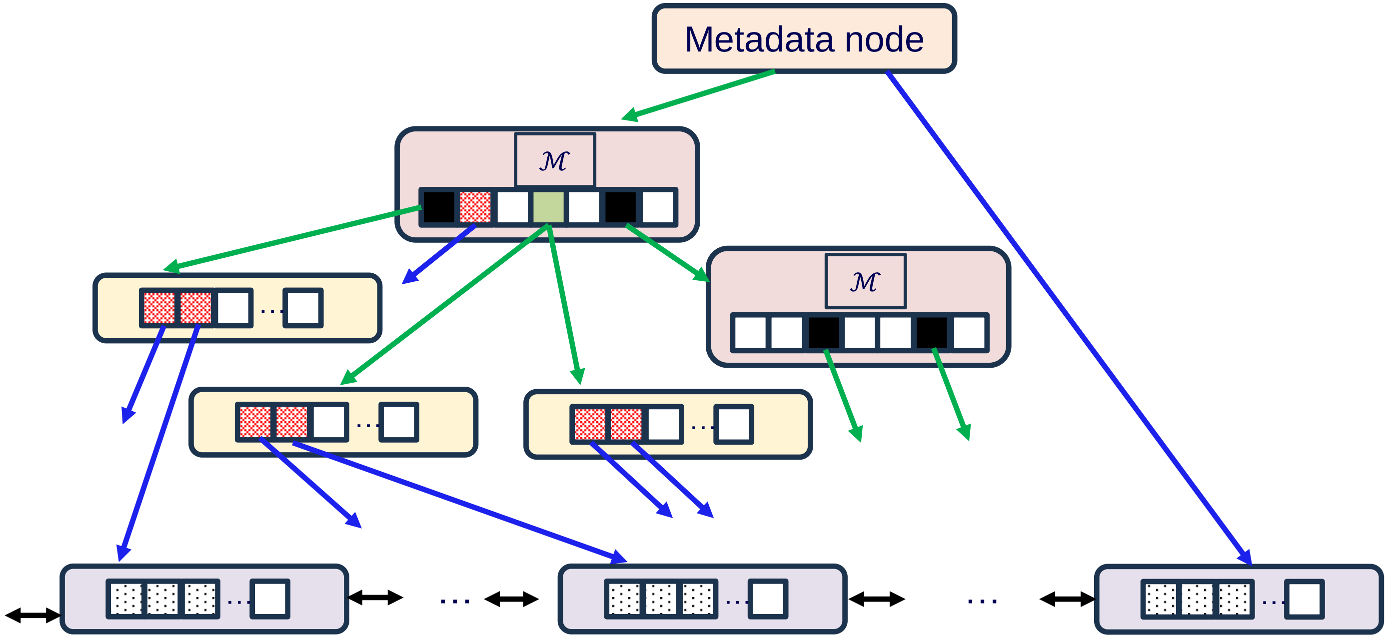

Bring the best of both worlds

AULID Index Layout

Leaf Node Layer

AULID Index Layout

Leaf Node Layer

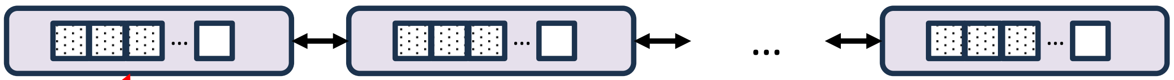

Leaf node

Set the size equal to the block size

AULID Index Layout

Leaf Node Layer

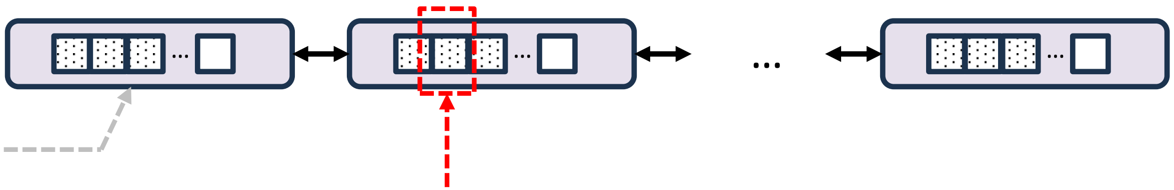

Leaf node

Set the size equal to the block size

Data slot – store the indexed value

AULID Index Layout

Leaf Node Layer

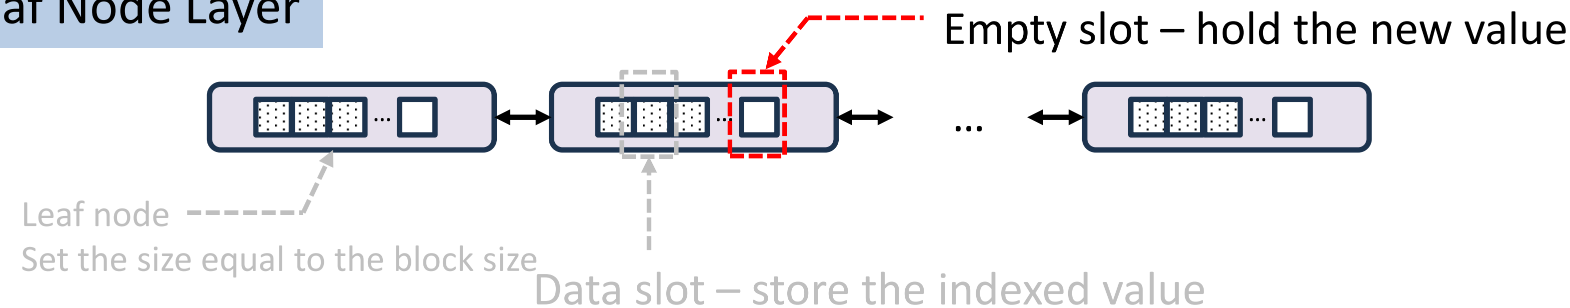

AULID Index Layout

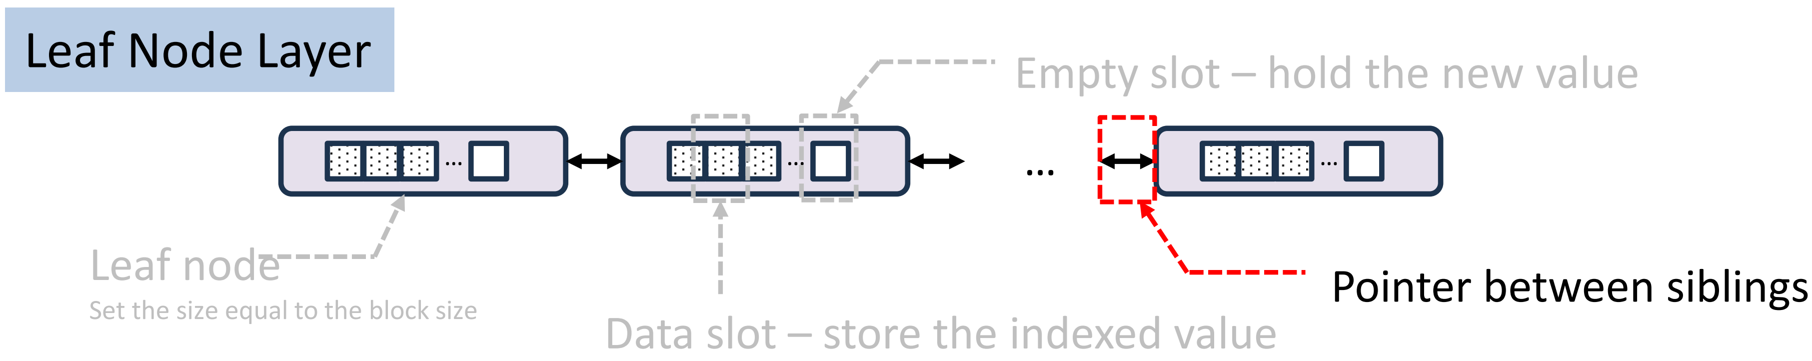

AULID Index Layout

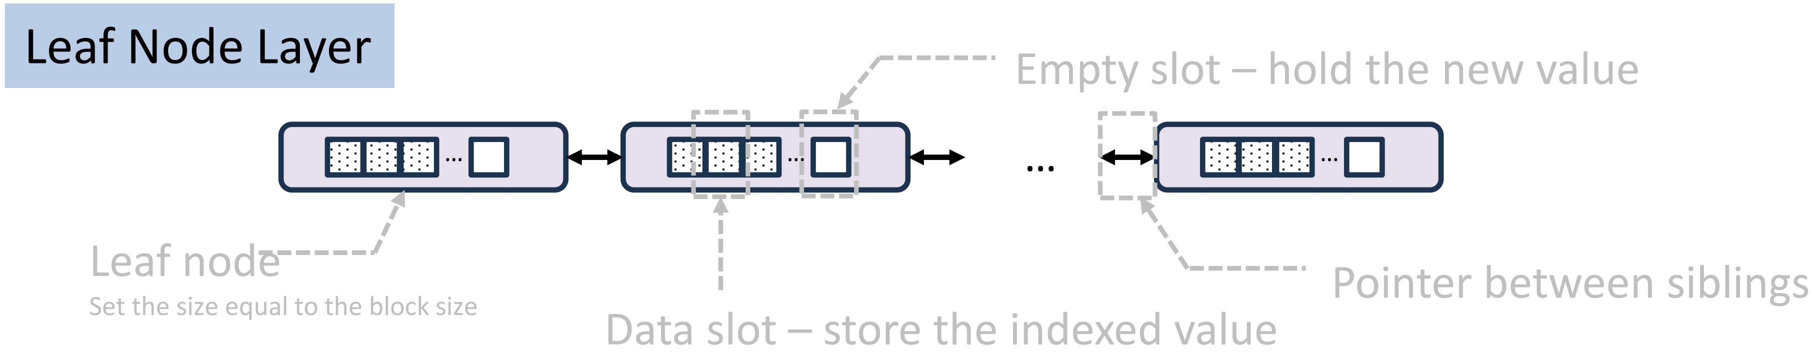

Benefits

- Low overhead for **scan** operations in fetching the *next* item (**P4**).

AULID Index Layout

Leaf Node Layer

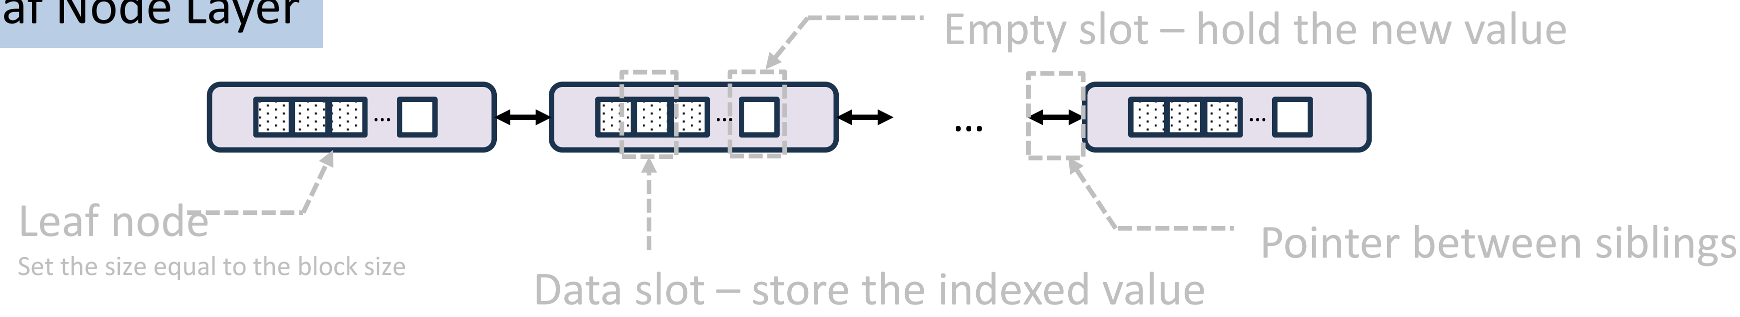

Benefits

- Low overhead for **scan** operations in fetching the *next* item (**P4**).
- Low **insertion** overhead and **SMO** overhead (**P3**).

AULID Index Layout

Inner Node Layer

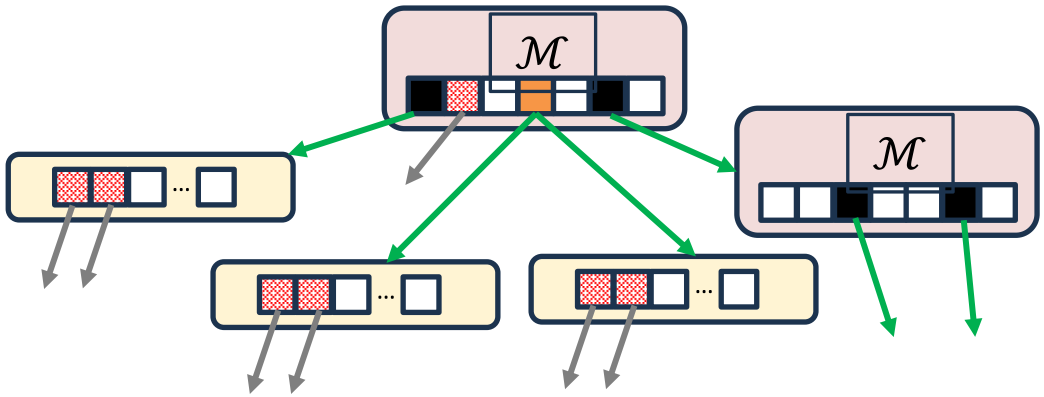

AULID Index Layout

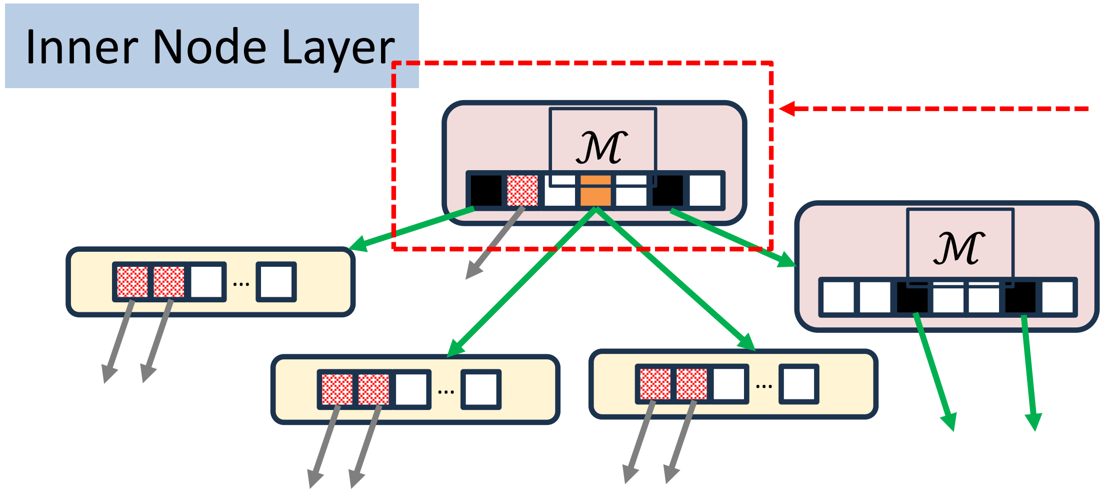

Mixed inner node

- Can hold different **slot** types
- Use a **model** to determine which slot to be accessed next

AULID Index Layout

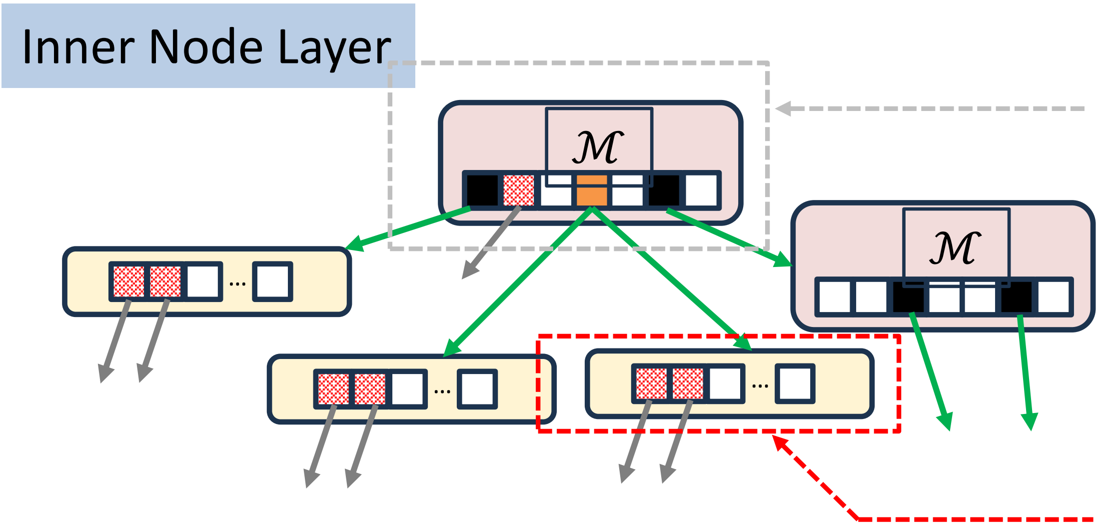

Mixed inner node

- Can hold different **slot** types
- Use a **model** to determine which slot to be accessed next

Packed inner node

- Hold the **pointer** to the leaf node and the **maximum key** in the indexed leaf node

AULID Index Layout

Inner Node Layer

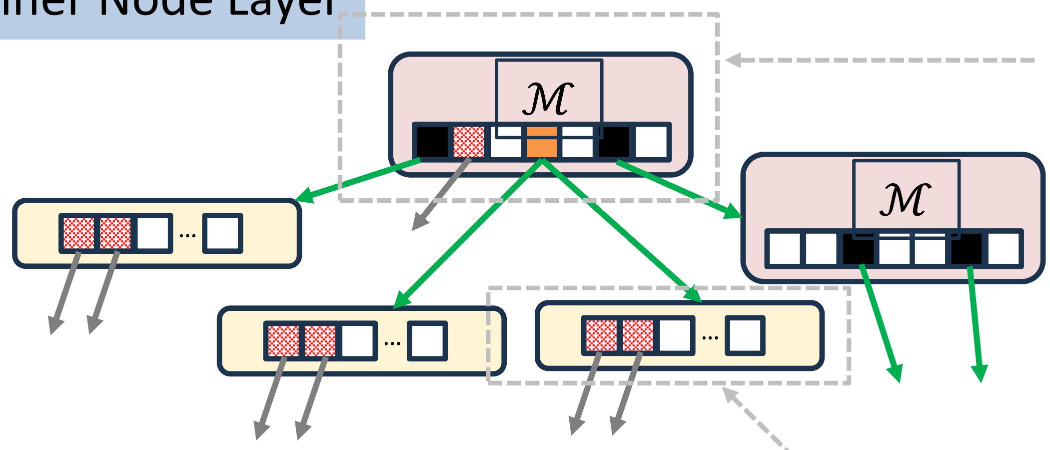

Mixed inner node

- Can hold different **slot** types
- Use a **model** to determine which slot to be accessed next

Packed inner node

- Hold the **pointer** to the leaf node and the **maximum key** in the indexed leaf node

Benefits

- Reducing the **tree height** of the index (**P1**).

AULID Index Layout

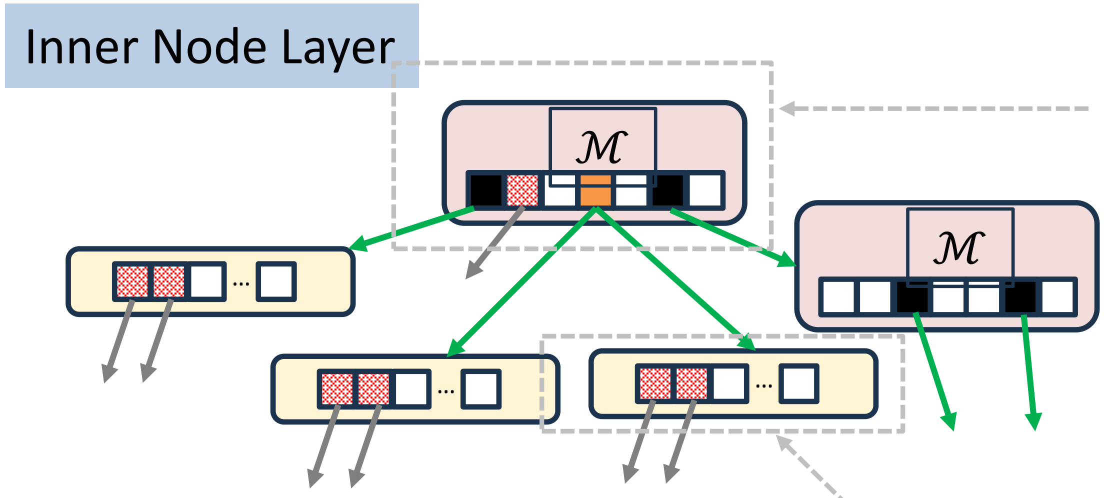

Mixed inner node

- Can hold different **slot** types
- Use a **model** to determine which slot to be accessed next

Packed inner node

- Hold the **pointer** to the leaf node and the **maximum key** in the indexed leaf node

Benefits

- Reducing the **tree height** of the index (**P1**).
- **Model-based** operations (search and insert) (**P2**).

AULID Index Layout

Inner Node Layer

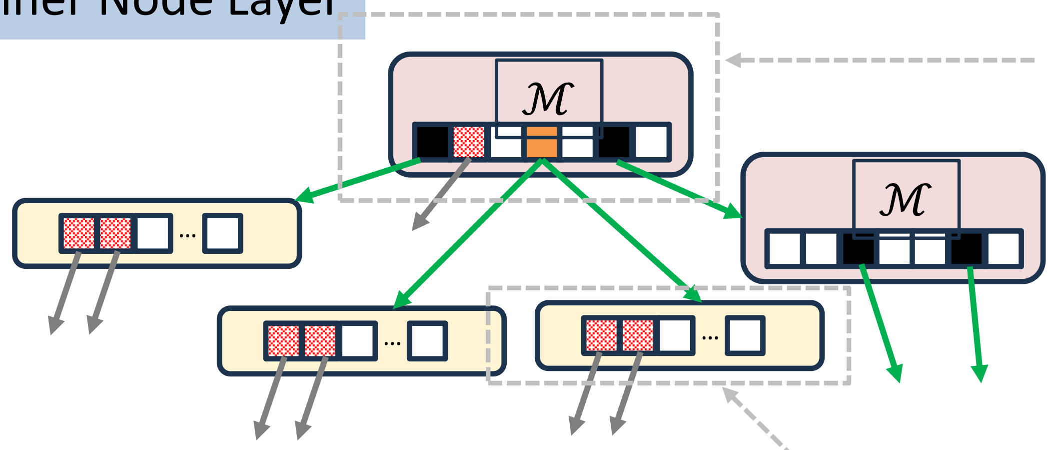

Mixed inner node

- Can hold different **slot** types
- Use a **model** to determine which slot to be accessed next

Packed inner node

- Hold the **pointer** to the leaf node and the **maximum key** in the indexed leaf node

Benefits

- Reducing the **tree height** of the index (**P1**).
- **Model-based** operations (search and insert) (**P2**).
- Low **SMO** overhead in inner nodes (**P3**).

AULID in action

[ICDE'24]

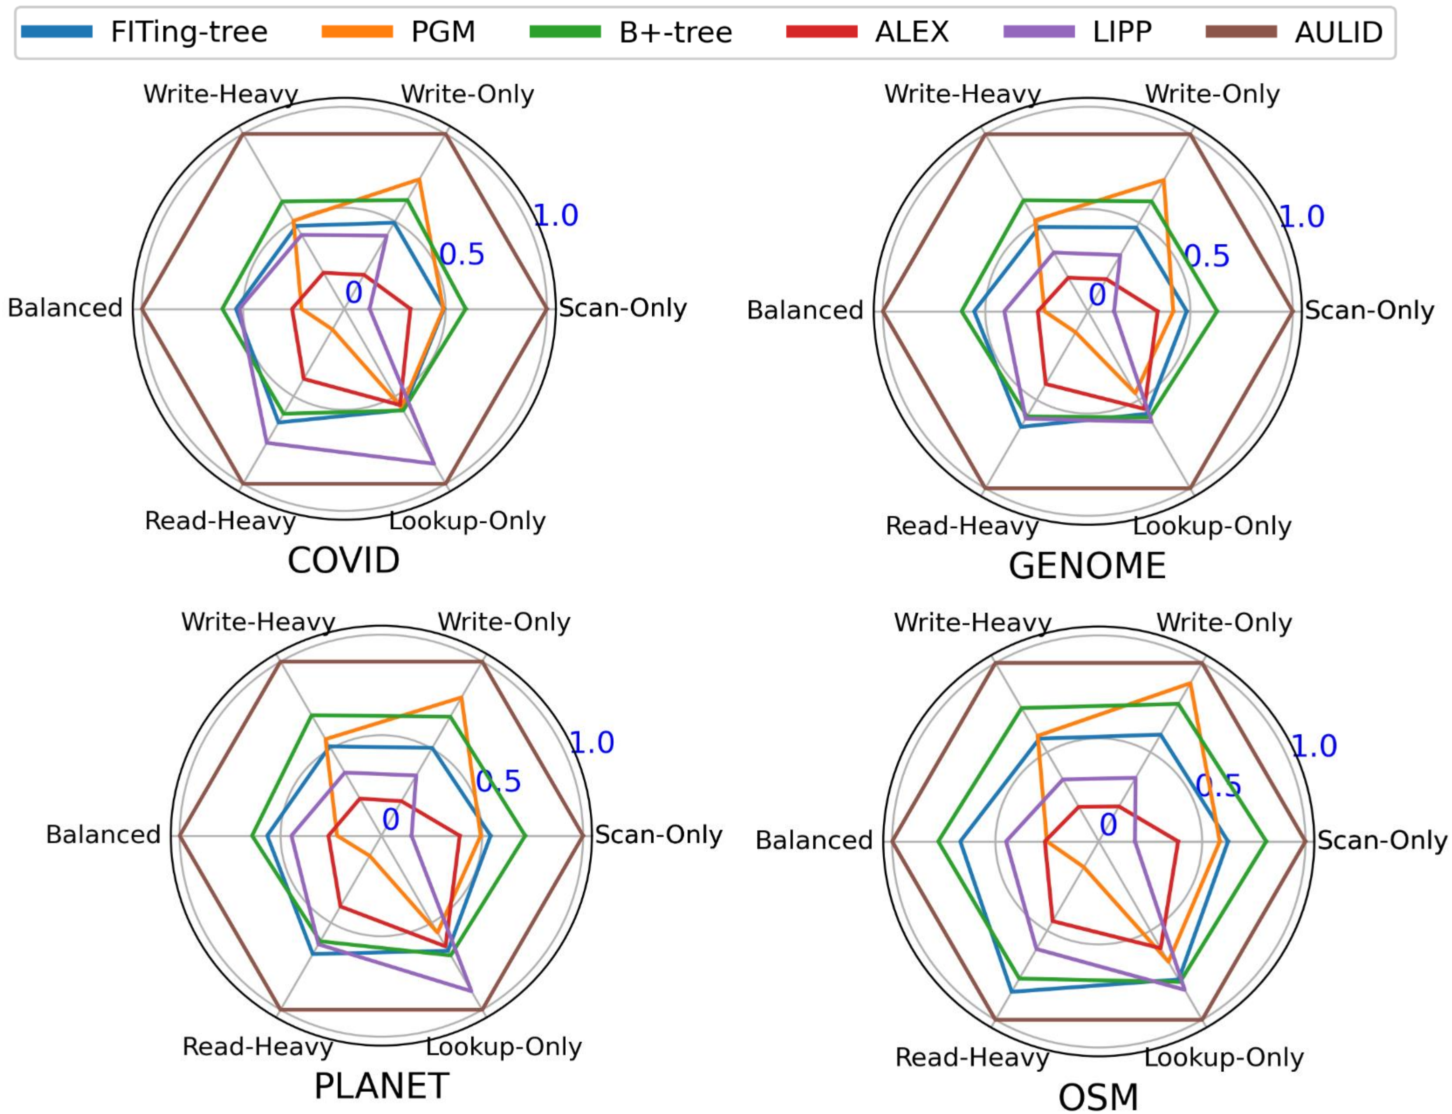

Consistently outperforming baselines across a range of workloads and data sets

AULID Learned Index Summary

- We identify the **challenges** when applying the learned indexes on disk and propose new design **principles**
- **AULID** adopts the principles with the carefully designed index layout and operations
- AULID **significantly outperforms** the SOTA across a range of workloads and datasets

Outline

- **Select physical design structures**

  [ICDE’21, ICDM’21, VLDB’22, TKDE’23, ICDM’24]

- **Tune the layout of physical design structures**

  [SIGMOD’23, ICDE’24]

- **Prefetch data ahead of time**

  [VLDB’24]

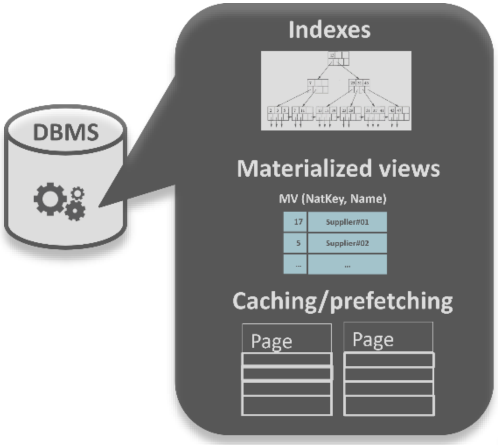

From data to rapid insights with interactive data exploration

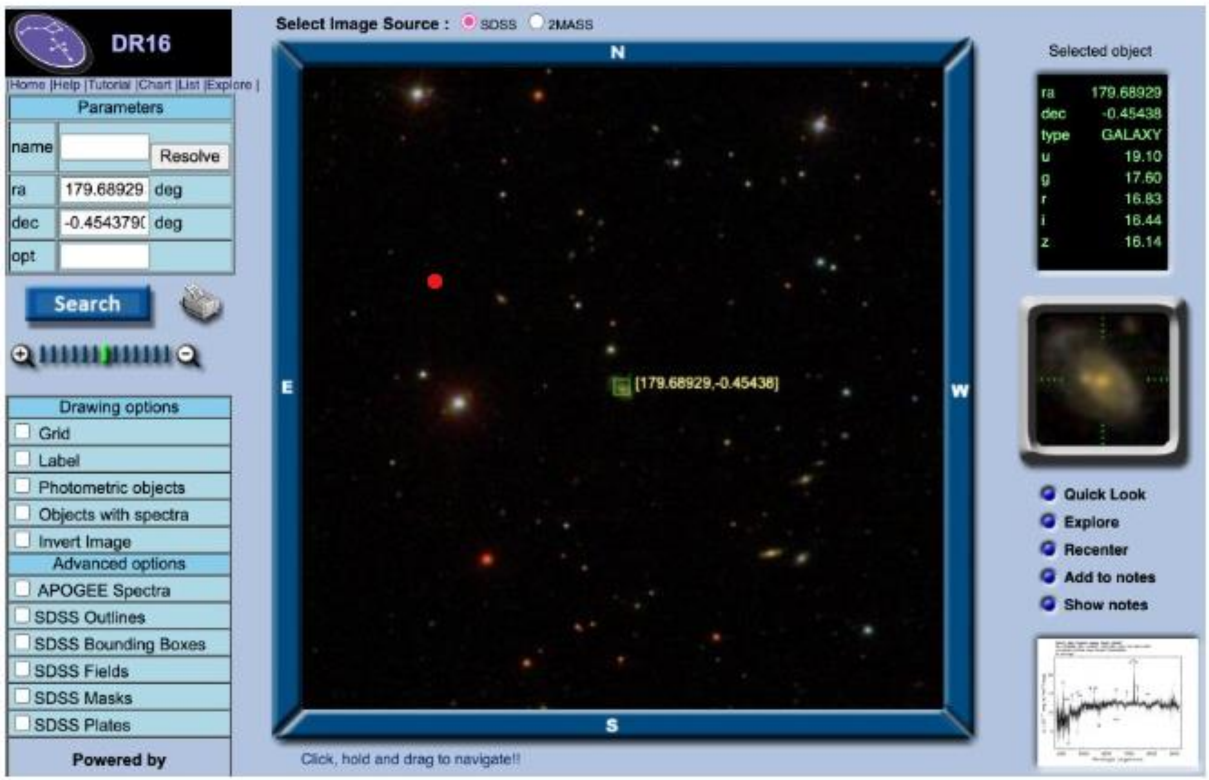

SDSS navigate tool [1]

- Need for interactive data exploration with sub-second latency
- Fast retrieval of large amounts of (scientific) data

Support interactivity with hands-free semantics-driven prefetching

Prefetching in the current landscape

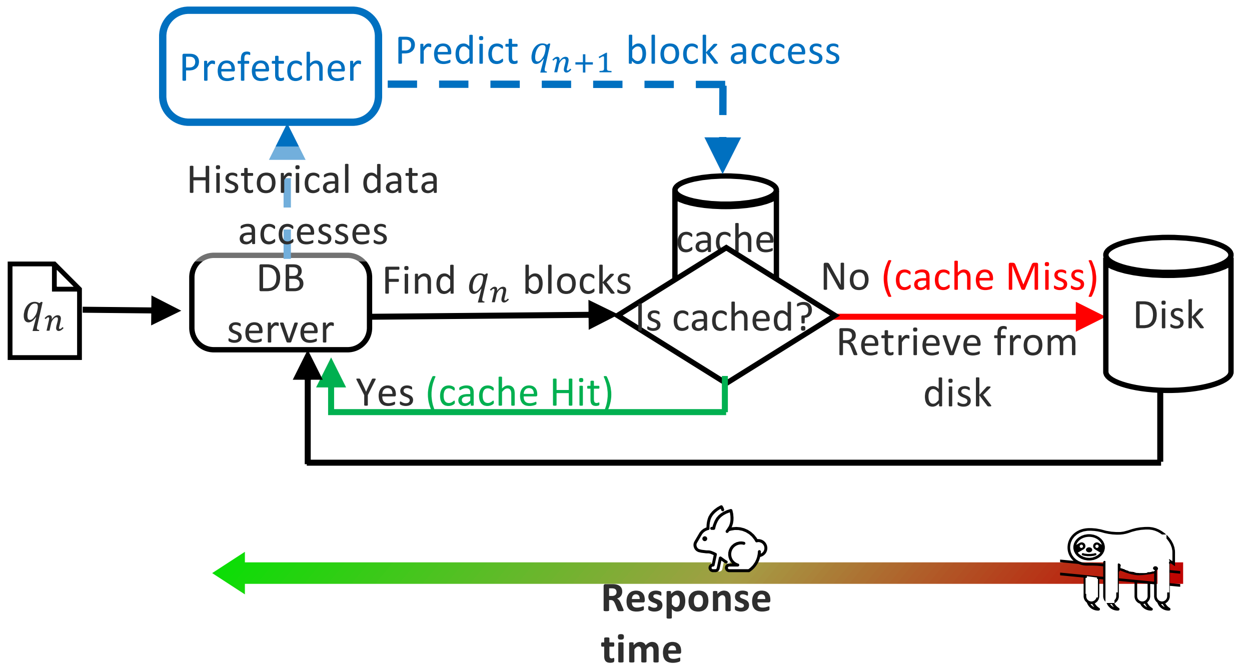

- ☒ Not suitable for SQL workloads
- ☒ Work with block addresses: No data semantics
- ☒ Not adaptive

Prefetching as timeseries forecasting

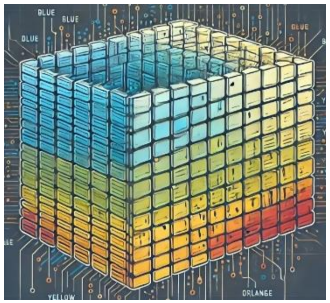

Data semantics is important

There is usually an inter-dependency among values stored in the data blocks accessed together

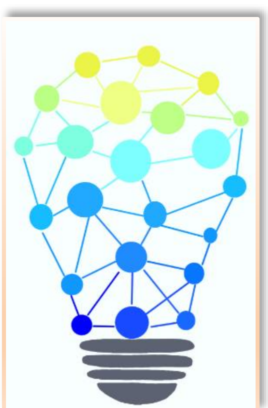

Prefetching → time series forecasting

Results observed from the queries in the previous time steps form the upcoming queries

SeLeP Overview

[VLDB'24]

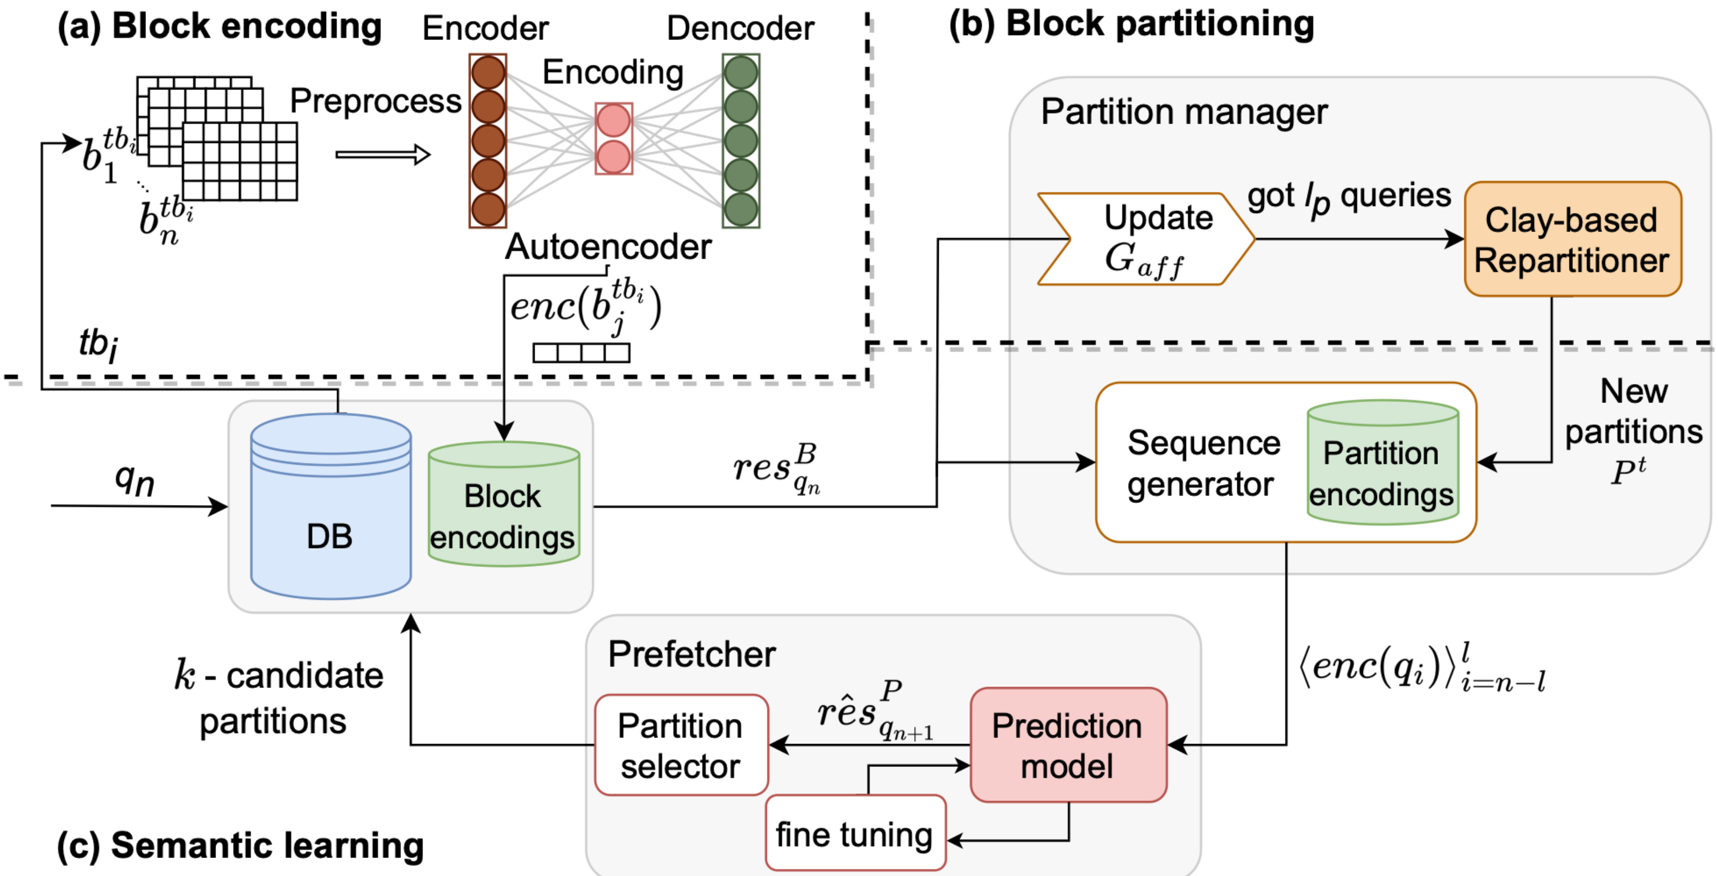

[VLDB'24] SeLeP: Learning Based Semantic Prefetching for Exploratory Database Workloads. F. Zirak, F. Choudhury, and R. Borovica-Gajic.

Block encoding

- Block can contain hundreds of values
- Need a concise block representation which captures the distinctive characteristics of the data
  - Encode blocks into vectors and aggregate them to form query encodings

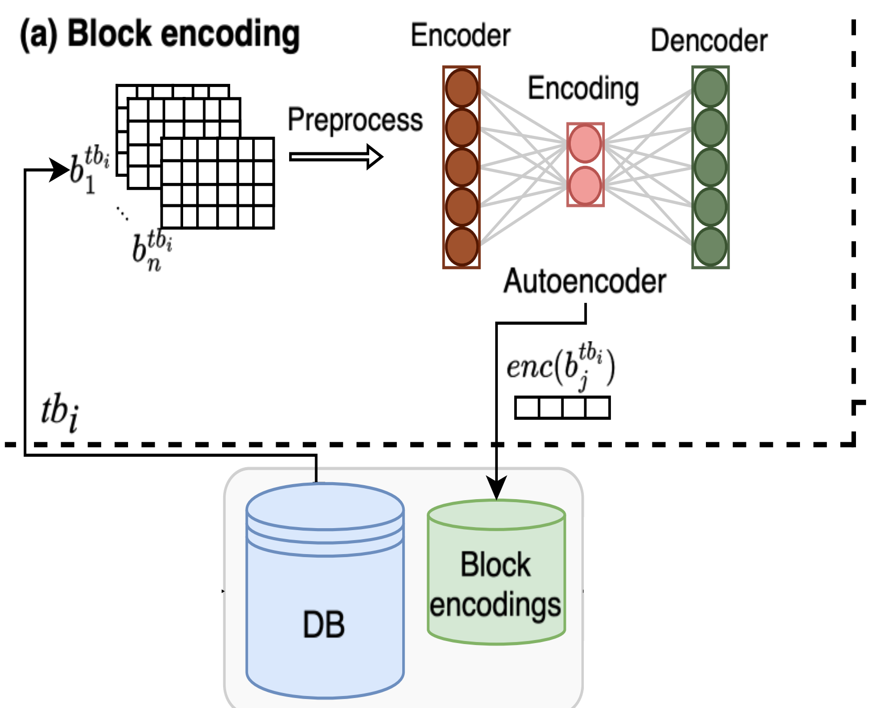

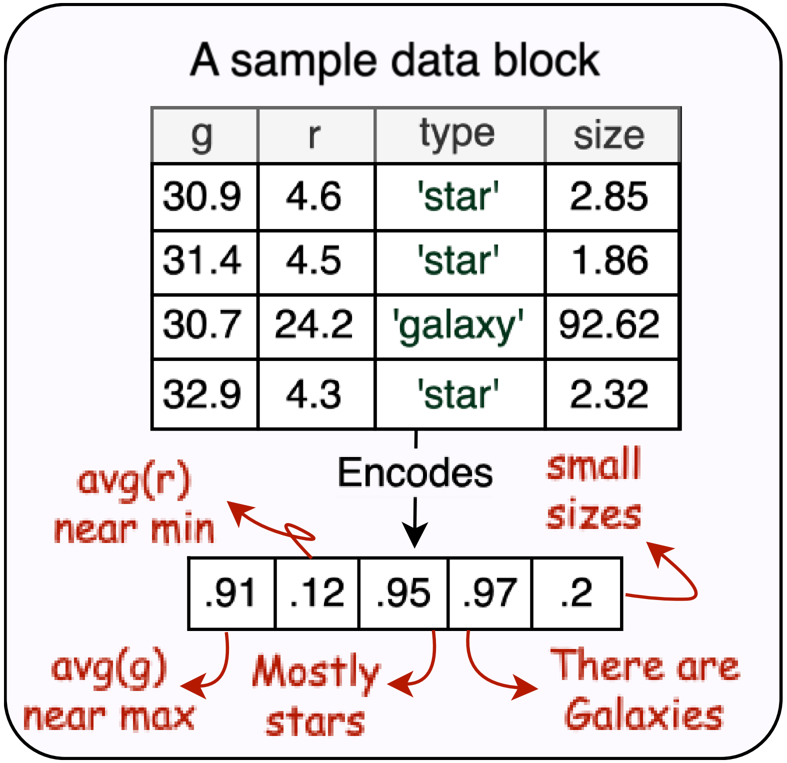

Block partitioning

- **Classification problem:**
  Having the sequence of last l query encodings, predict and fetch blocks that will be accessed next

Large dataset → **Substantial number of labels**

- Group blocks frequently accessed together into partitions

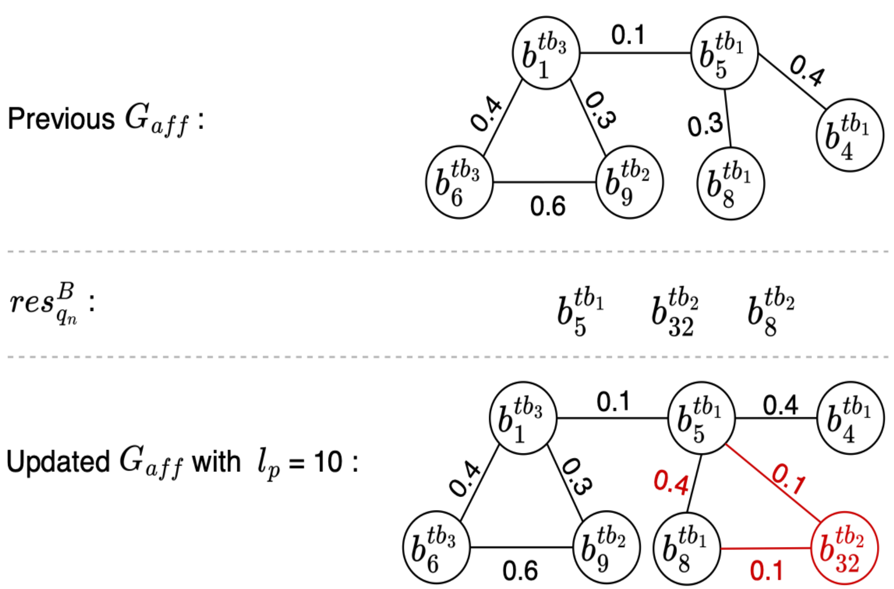

Graph partitioning on affinity graph

tb_1	0	0	0	0	0
tb_2	.89	.21	.79	.68	.24
tb_3	0	0	0	0	0
tb_4	.4	.13	.33	.19	.81

Partition encodings

Semantic Learning

- Learn partition access pattern from a sequence of query encodings and fine tune the model with new workloads

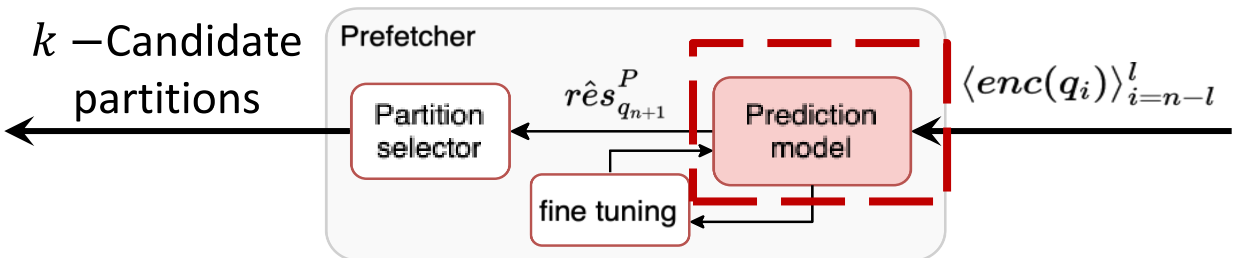

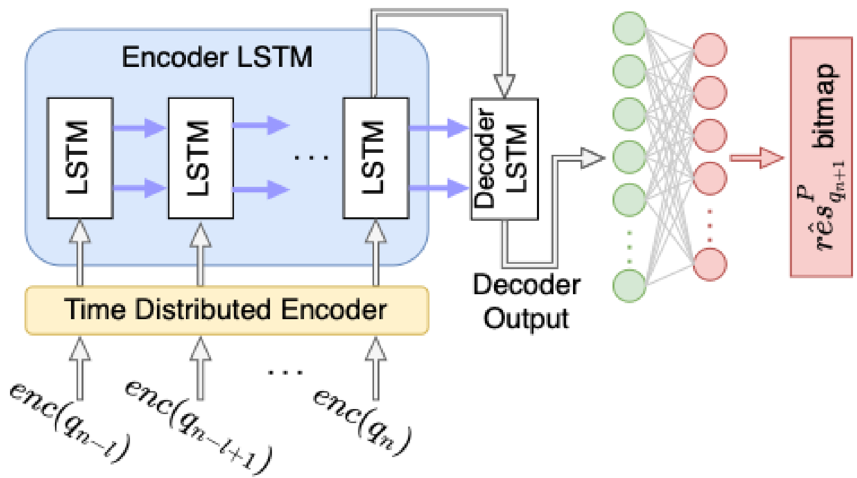

Query encodings

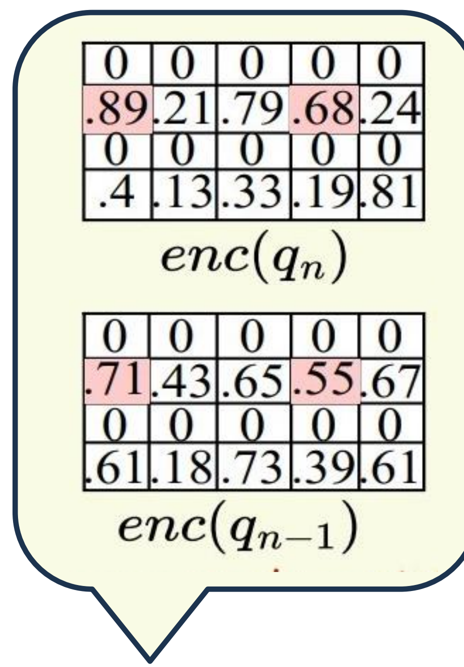

Accessed partitions contain blocks with galaxies and high g average

SeLeP in Action

[VLDB'24]

Setting: 16GB SDSS DR7, prefetch size = $k \times 128$ block, 4GB cache

Queries: multi-table join SQL workloads

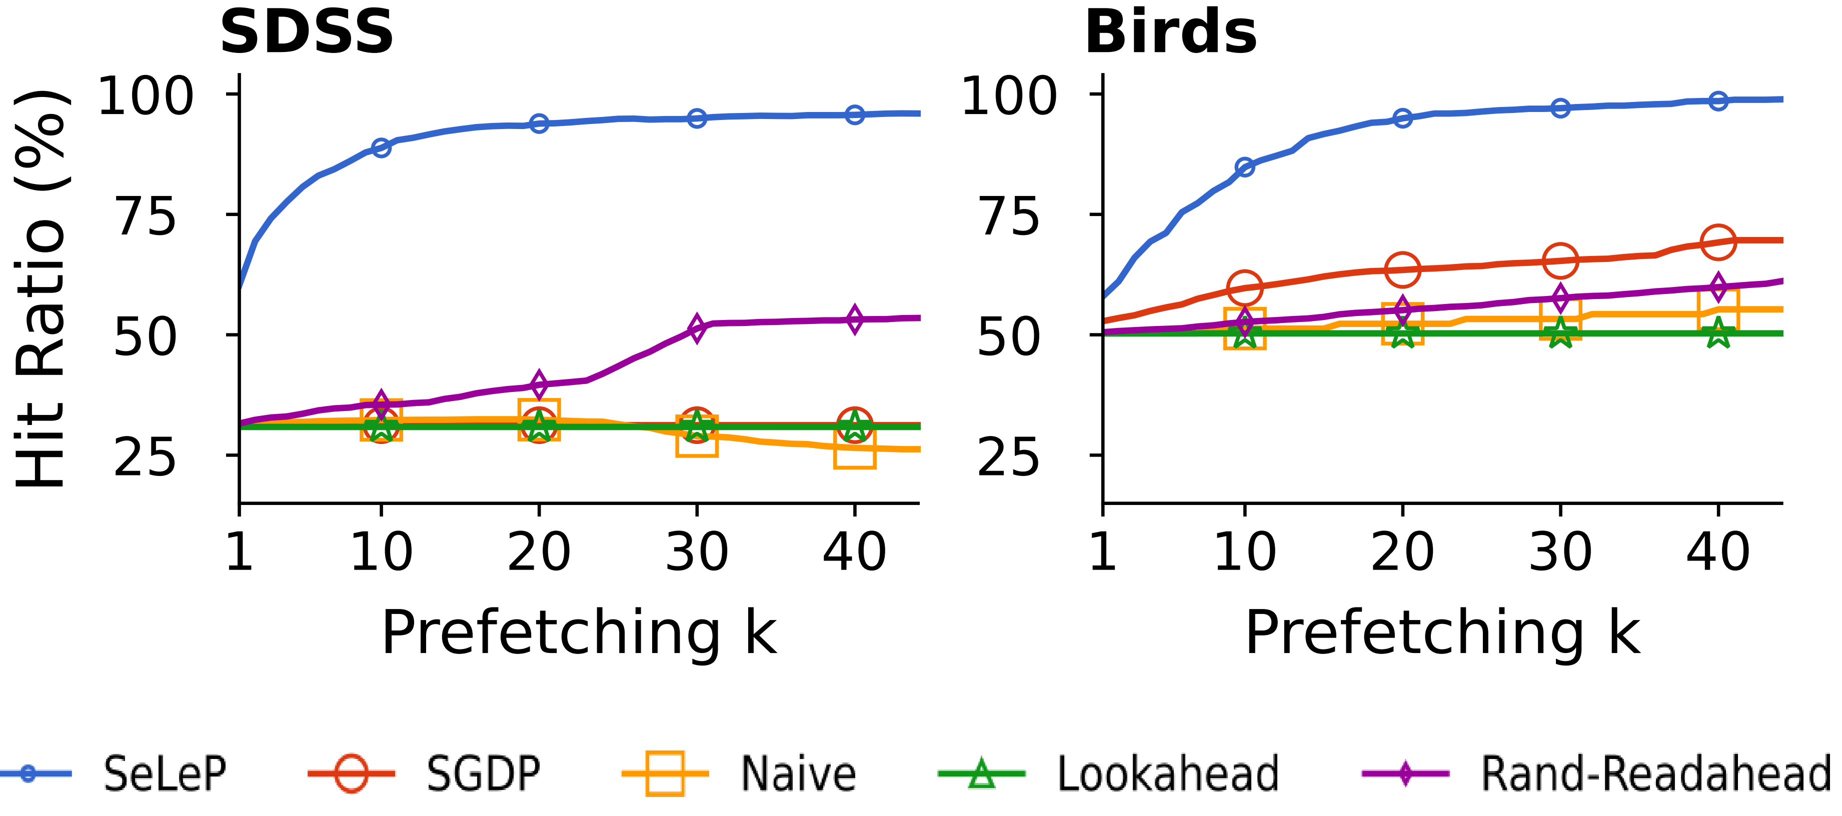

95% average hit ratio, outperforming SOTA by 40%

I/O Reduction with SeLeP

[VLDB'24]

Setting: 16GB SDSS DR7, prefetch size = $k \times 128$ block, 4GB cache

Queries: multi-table join SQL workloads

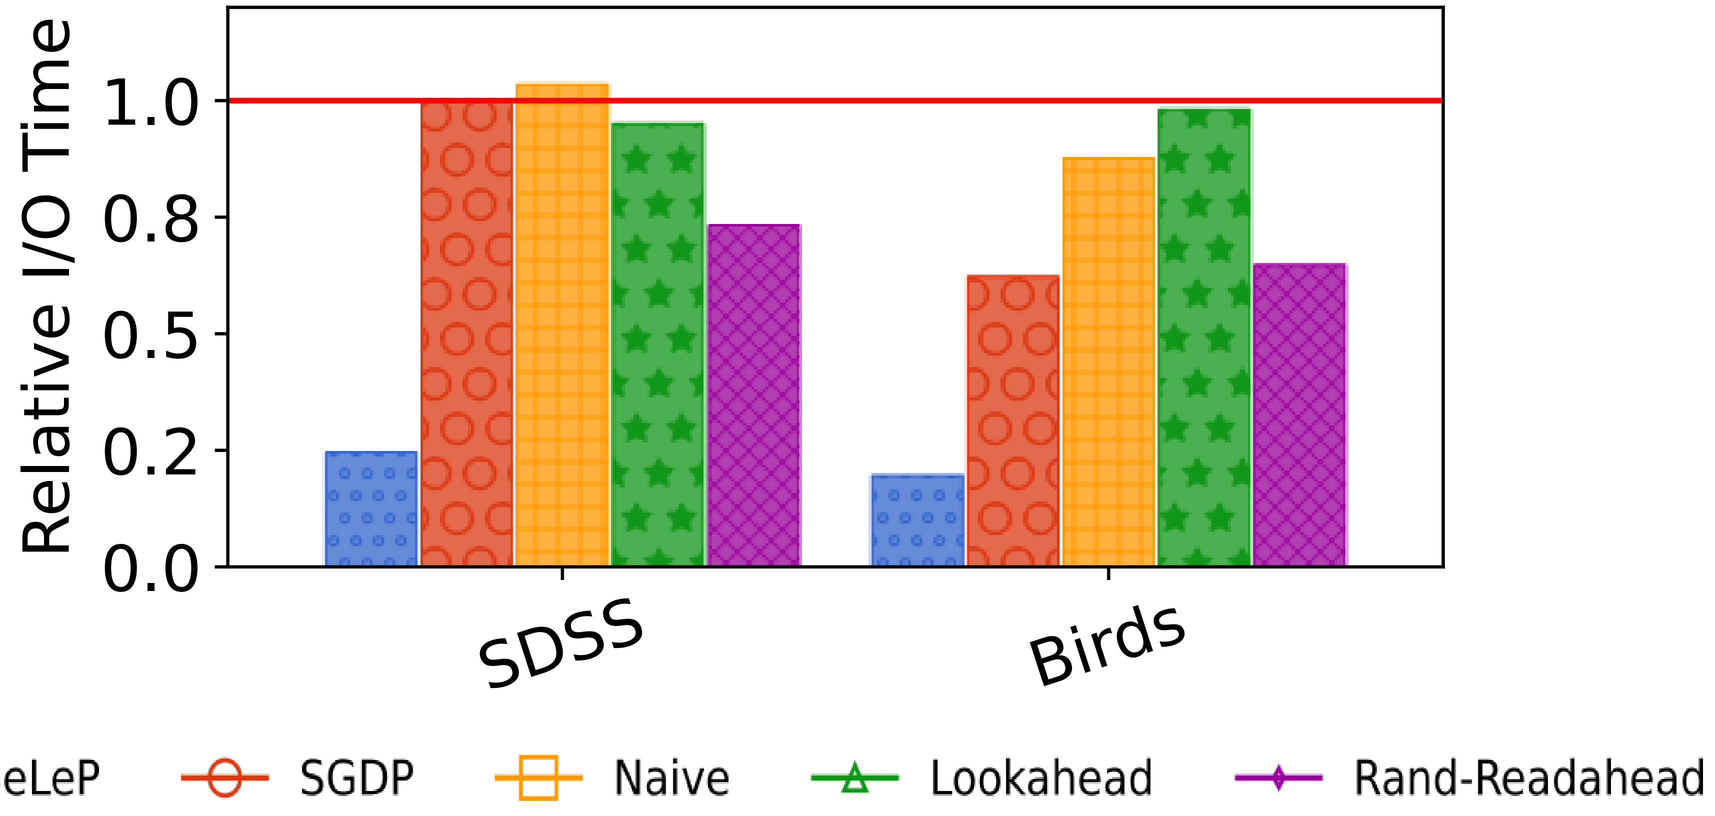

80% average I/O reduction, outperforming SOTA by 45%

SeLeP Adaptivity

[VLDB'24]

Setting: 16GB SDSS DR7, prefetch size = $k \times 128$ block, 4GB cache
Queries: Shifts at sequence number = {2000, 4000, 6000} with novel query templates and access to unseen data

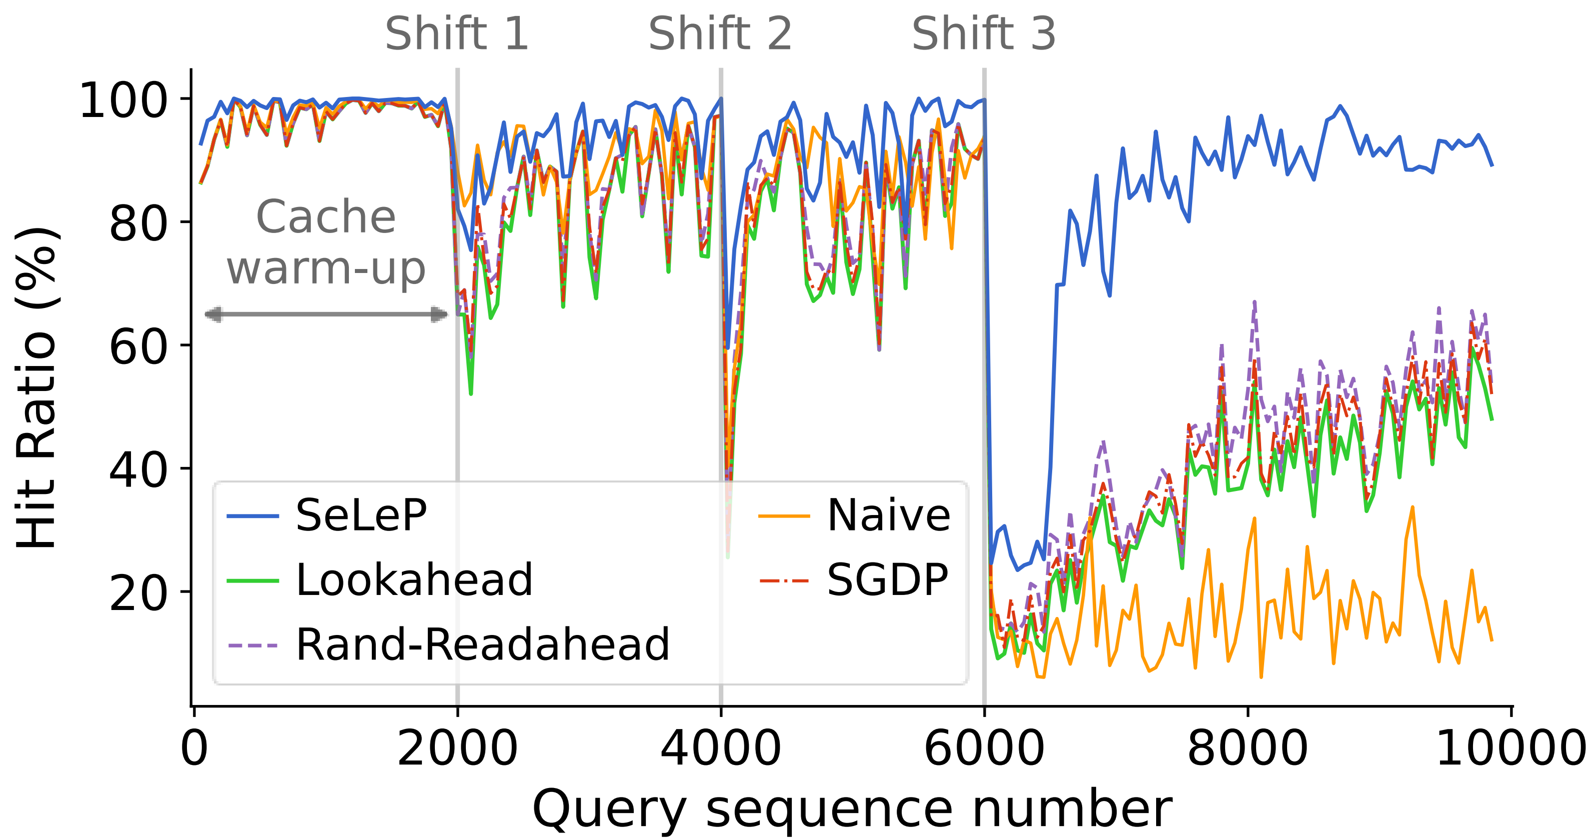

Graceful adaptation to unpredictable workloads

SeLeP Summary

- Prefetching can substantially reduce I/O time, but the existing SOTA prefetchers ignore data semantics and cannot deal with ad hoc workloads
- SeLeP can benefit all types of exploratory workloads by leveraging **data semantics**
- SeLeP improves hit ratio up to 40% and reduces I/O time up to 45% compared to SOTA prefetchers

AI-Powered Databases

https://ai-db-uom.github.io/

Custom-tailored (AI-driven) databases that automatically learn from user interactions with the database to optimize its performance

Learn from .....

1. Queries
[SIGMOD'12]
[CACM'15]
[ICDE'21]
[ICDM'21]
[VLDB'23]

2. Data
[ICDE'15]
[VLDBJ'18]
[ADC'20]
[SIGMOD'23]
[ICDE'24]

3. Hardware
[VLDB'16]
[ADMS'17]
[CACM'19]

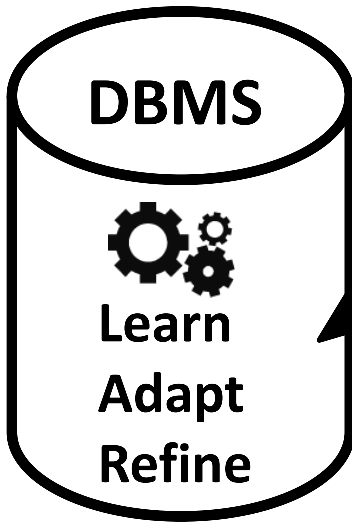

Fast responses

Physical design tuning
[ICDE'21, ICDM'21, VLDB'23, TKDE'23]
- PD tuning via Multi-Armed Bandits
- Tailored for HTAP, ad hoc workloads
- Provable performance guarantees

Learned Indexes
[ADC'20, SIGMOD'23, ICDE'24, VLDB'25]
- Updateable (on disk) learned indexes
- Indexing via function interpolation
- Spatial learned indexes

Caching/prefetching
[DEBS'22, VLDB'24]
- Semantics driven prefetching
- Tailored for (ad hoc) data exploration

Where to go from here?

AI-Powered Databases: From a passive data retrieval provider to a co-pilot for insight discoveries

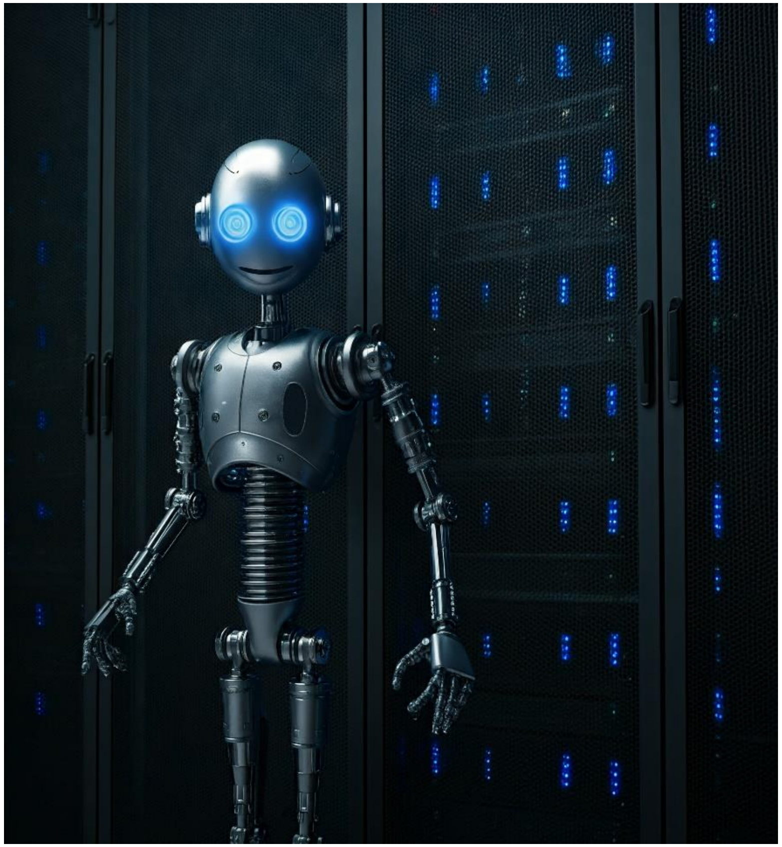

Credit: generated with Gemini

- What makes some data interesting?
  - Measures of interestingness
  - Novelty discovery
- Can we predict user intent?
  - Long term and short-term goal recognition
- Can we recommend queries that lead to insights?
  - Query formulation (LMs)
  - From text to SQL and back

Special thanks to the AI-powered DB team

- My students/postdocs:
  - Malinga Perera
  - Bastian Oetomo
  - Hai Lan
  - Farzaneh Zirak
  - Guanli Liu
  - David Adams
  - Lankadinee Rathuwadu
  - Dinuka de Zoysa
- Collaborators:
  - Benjamin Rubinstein
  - Zhifeng Bao
  - Farhana Choudhury
  - Jianzhong Qi
  - Lars Kulik
  - Nir Lipovetzky
  - Christopher Leckie
  - James Bailey

And many more (external) collaborators...

THE UNIVERSITY OF MELBOURNE

Questions?

Personal website: https://renata.borovica-gajic.com/

Project website: https://ai-db-uom.github.io/

Email: renata.borovica@unimelb.edu.au

Looking for PhD students!

**This work is supported by the Australian Research Council Discovery Early Career Researcher Award DE230100366, and L'Oréal UNESCO FWIS'23 Fellowship.*

THANK YOU!