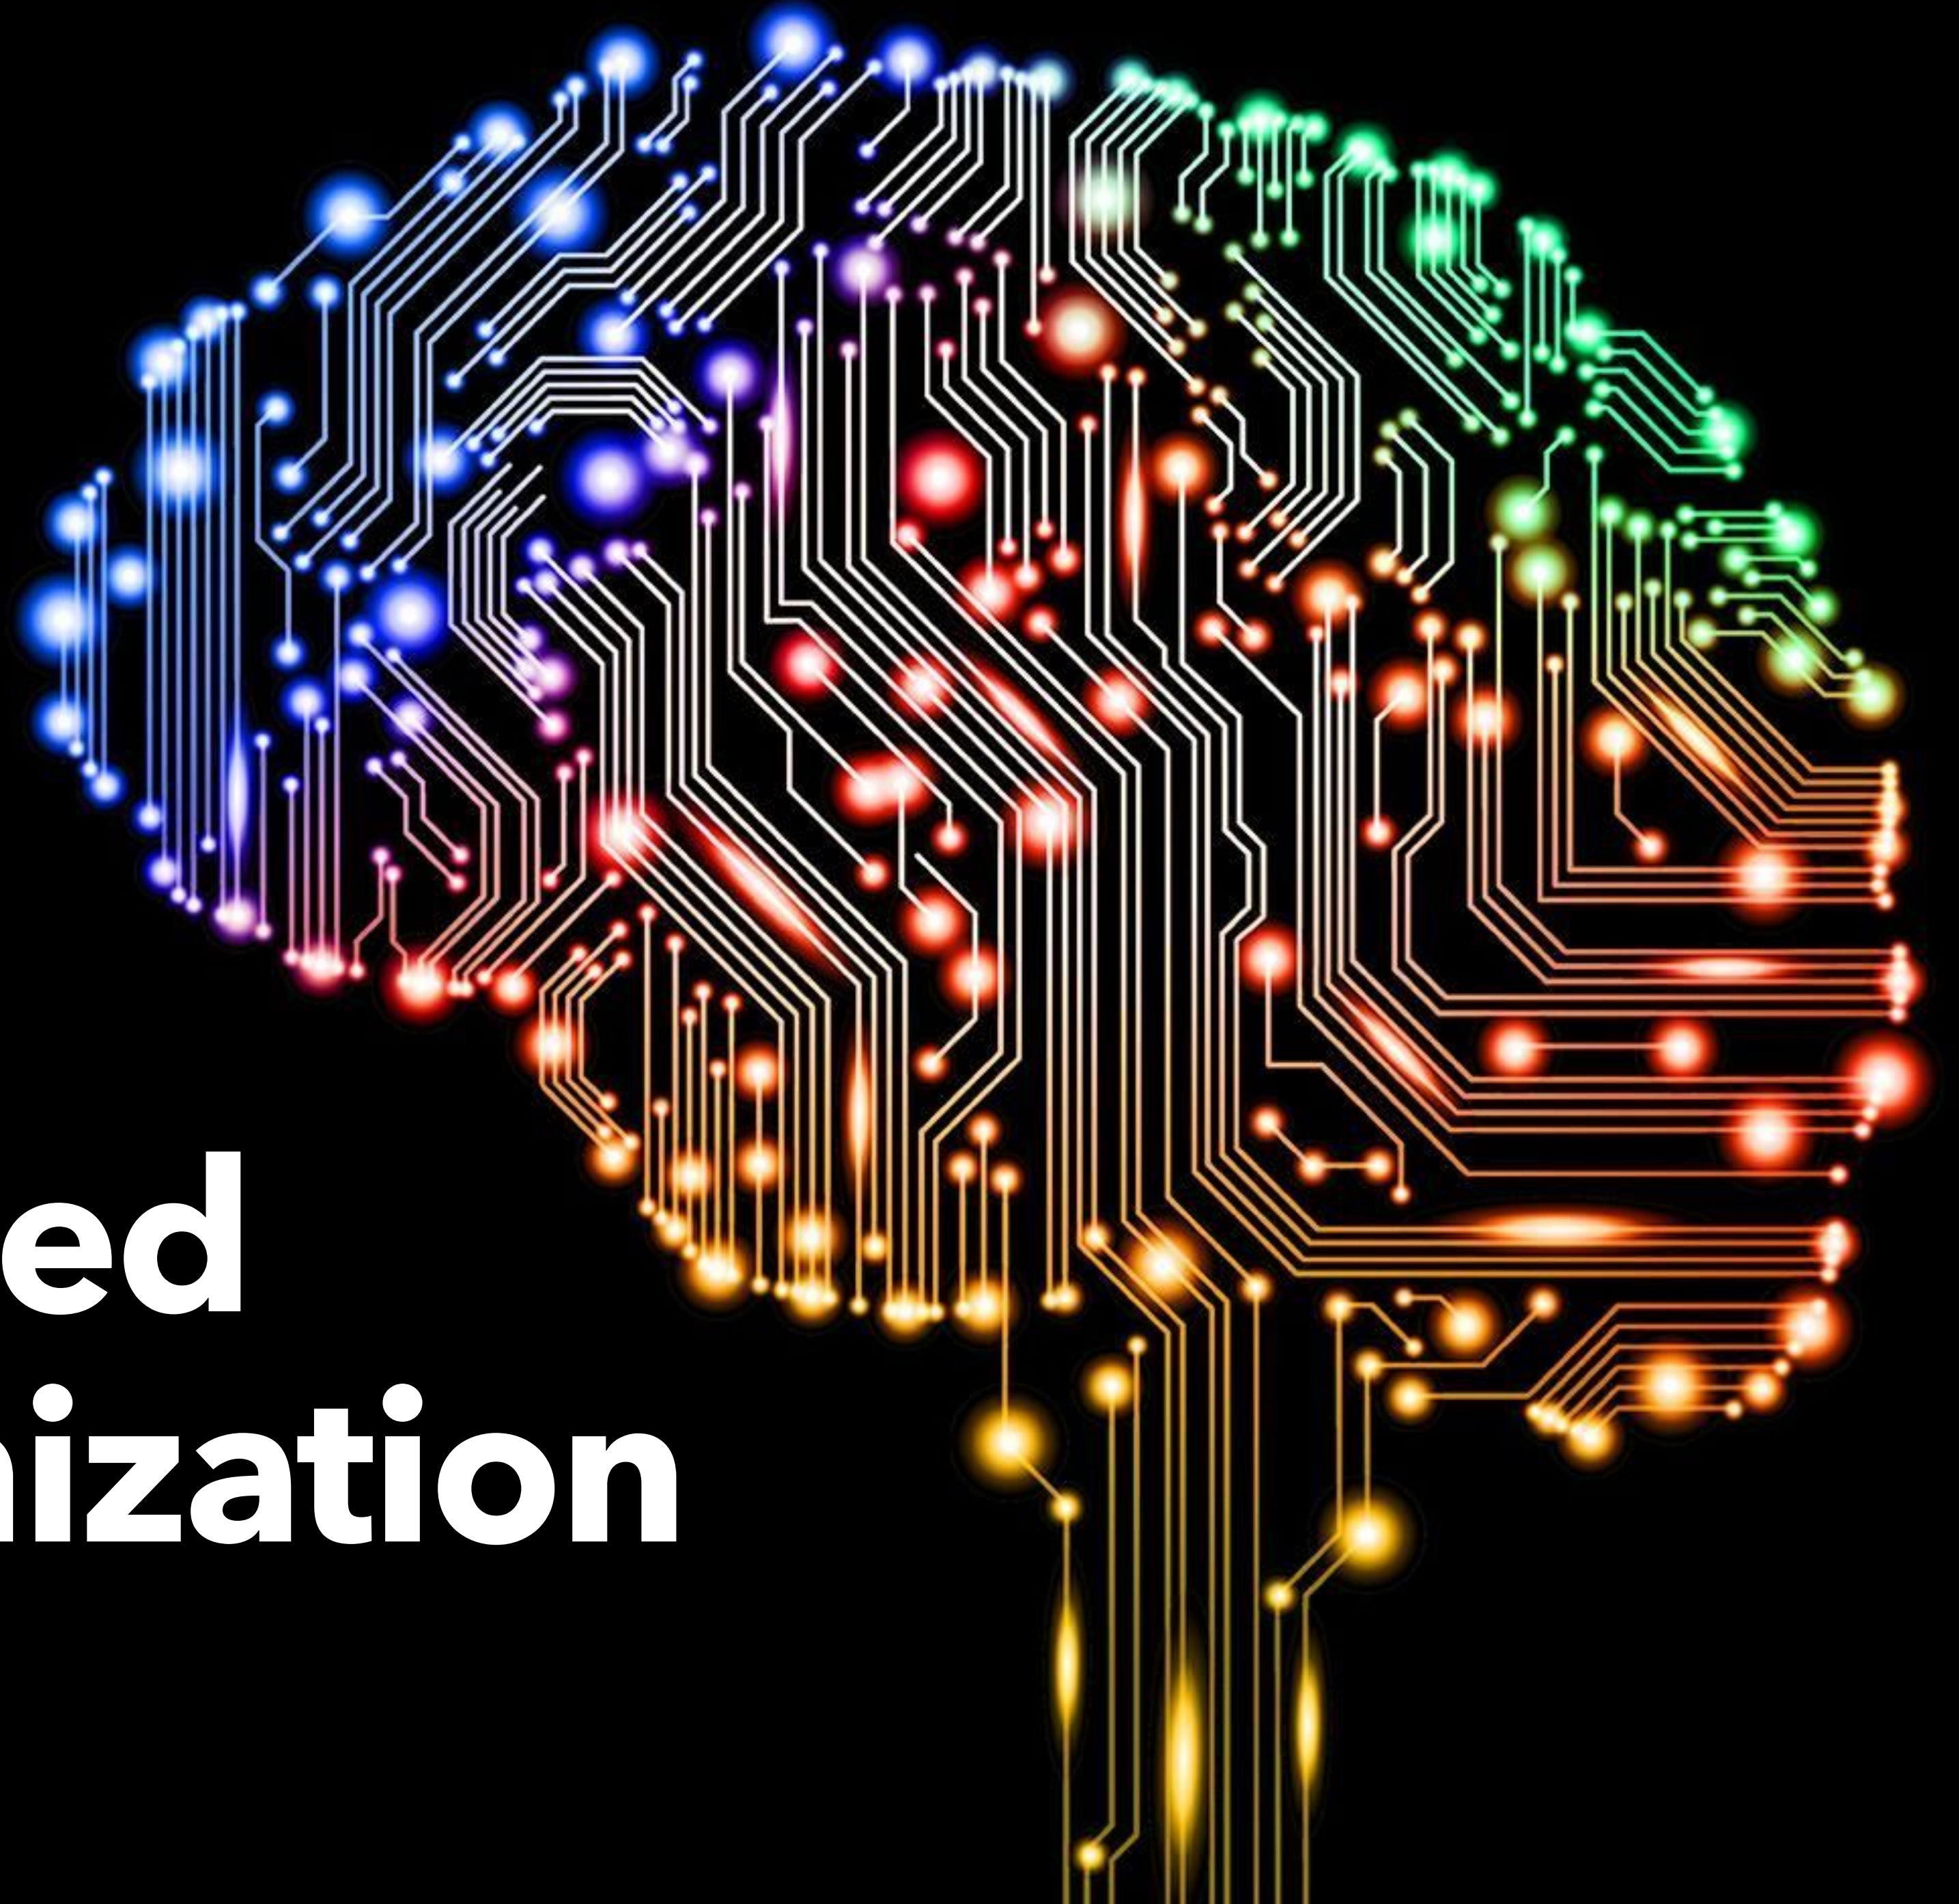



Zoi Kaoudi - 13.02.2023



Learning-based Query Optimization

What are we still missing?

Why query optimization?



Apache Wayang




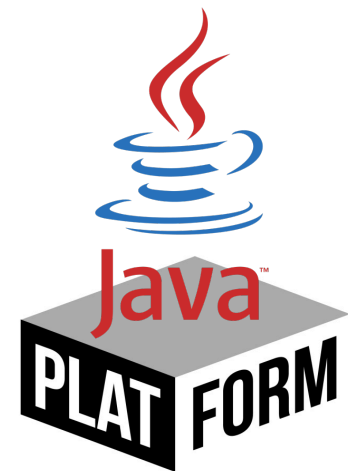
Decoupling applications from underlying processing platforms

Applications

- Platform-agnostic code
- SQL queries



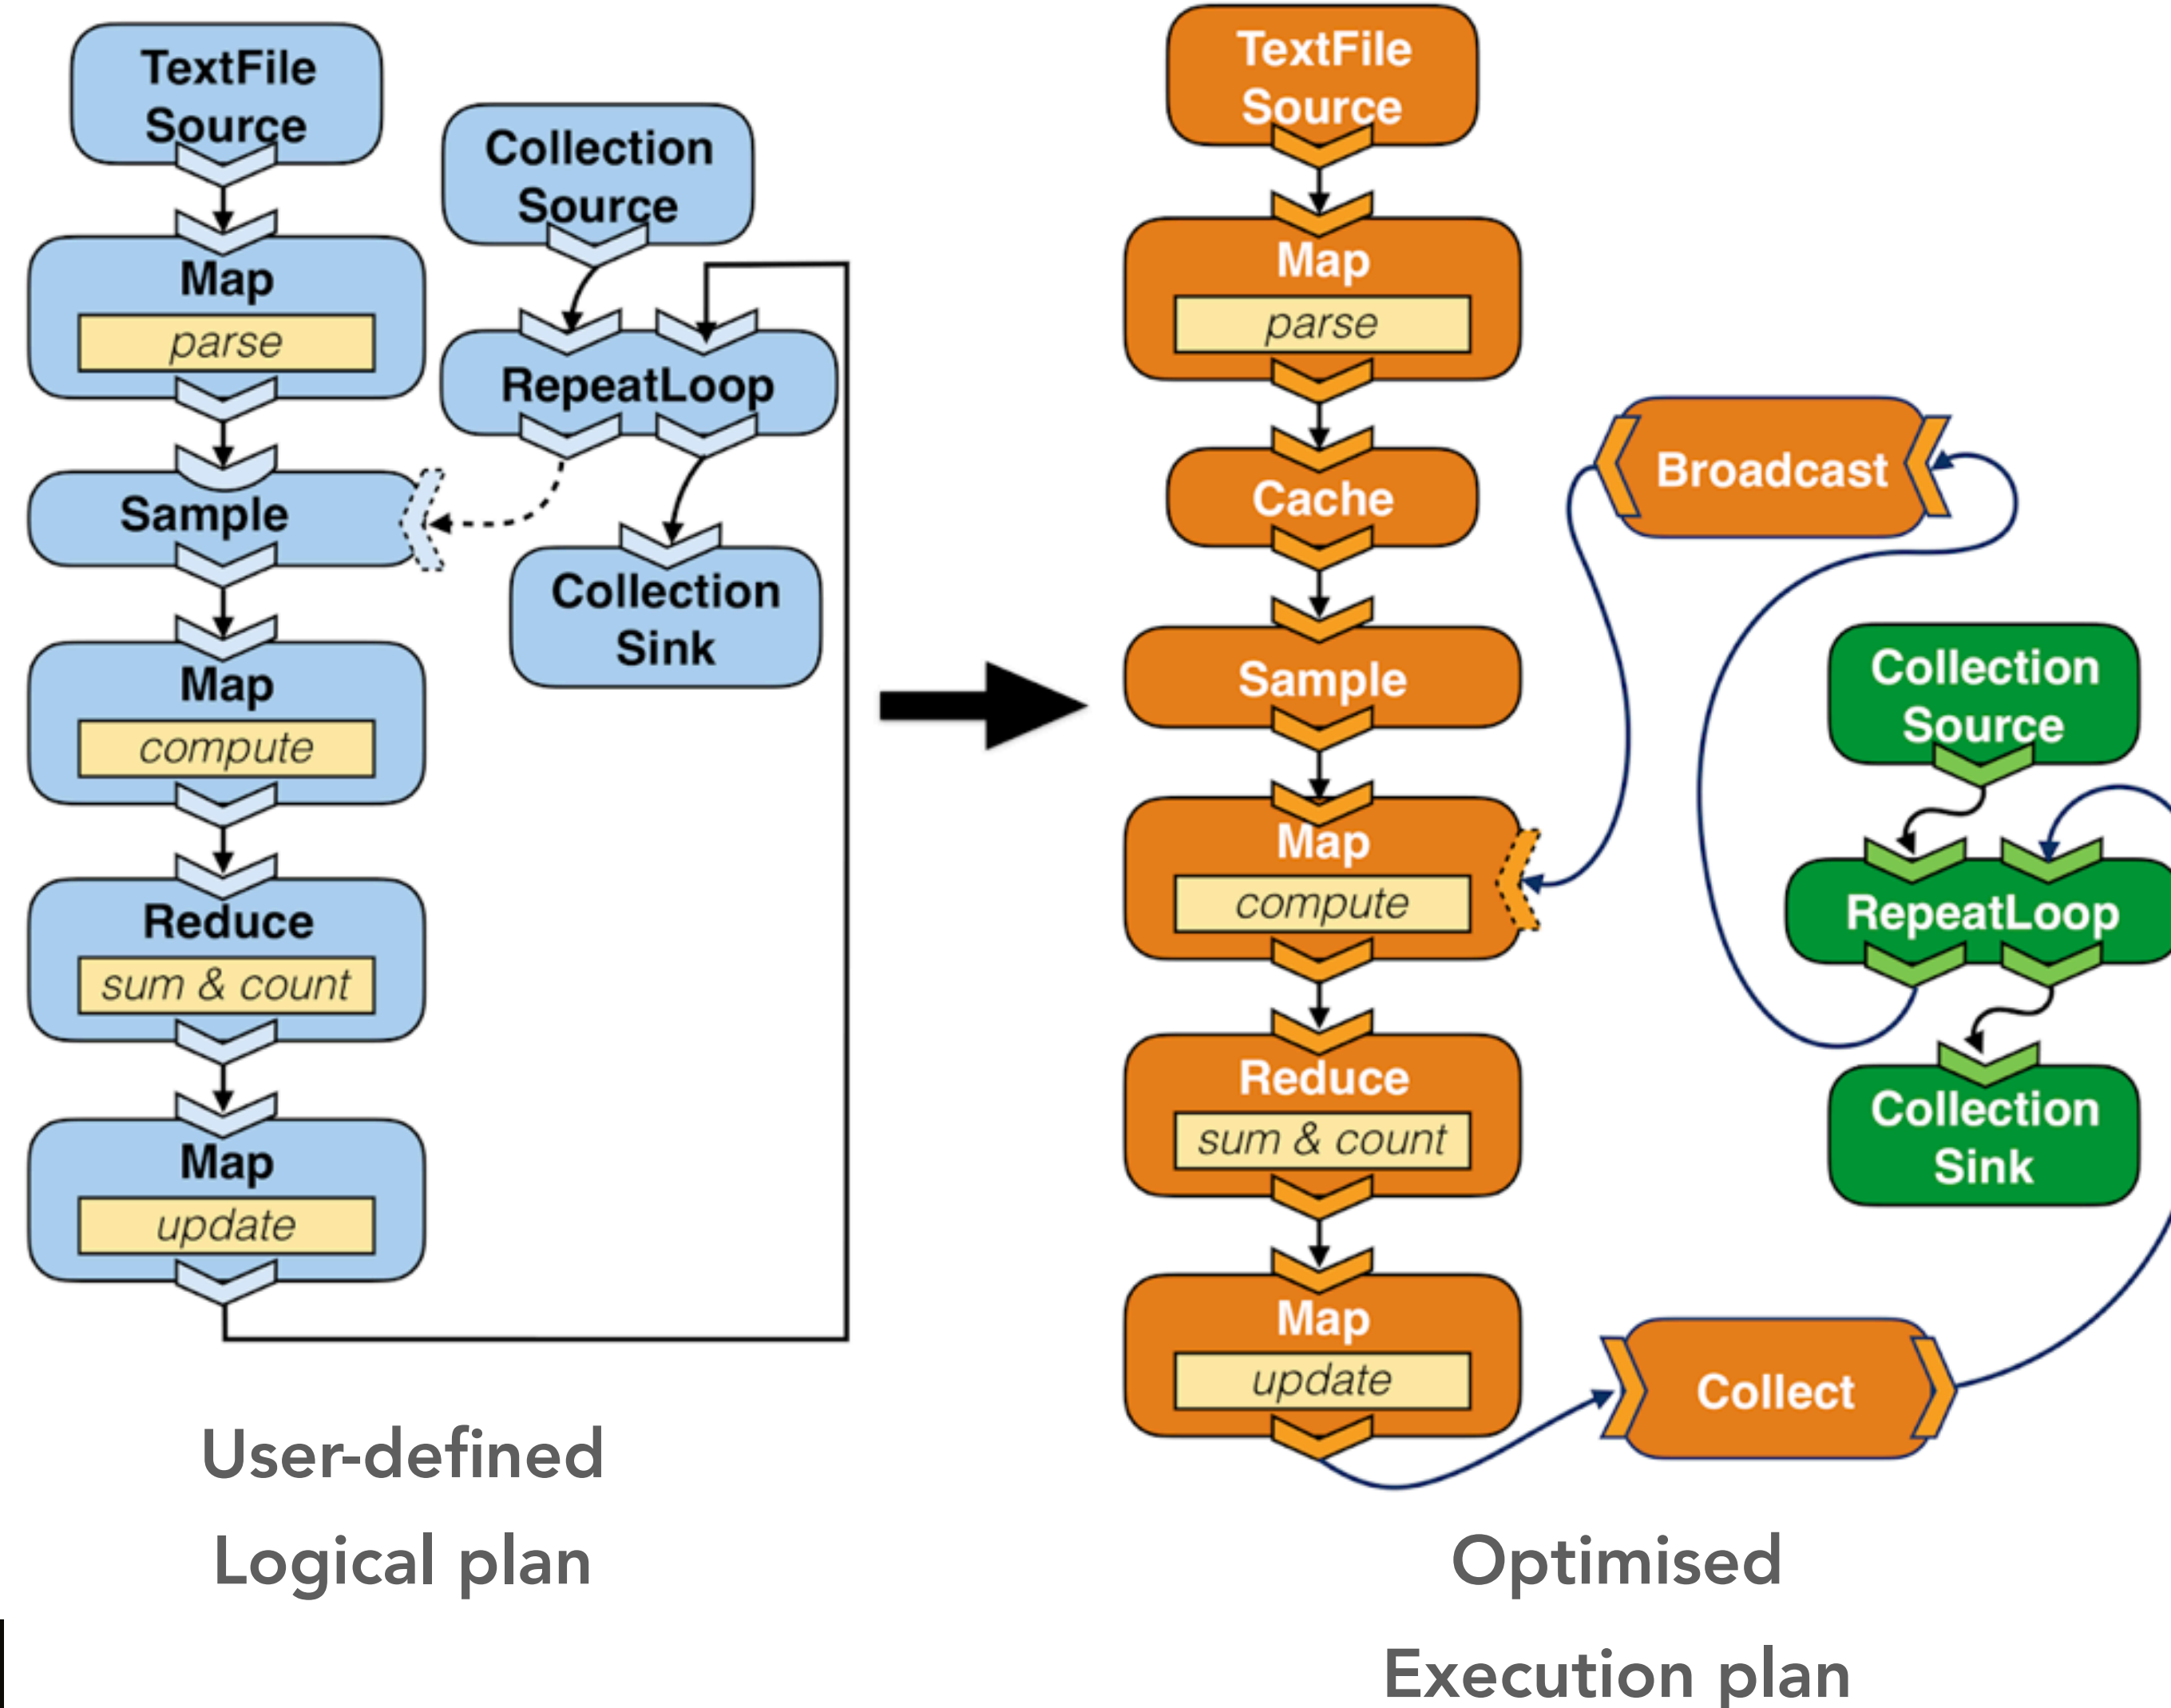
Platforms



Apache Wayang

Decoupling applications from underlying processing platforms

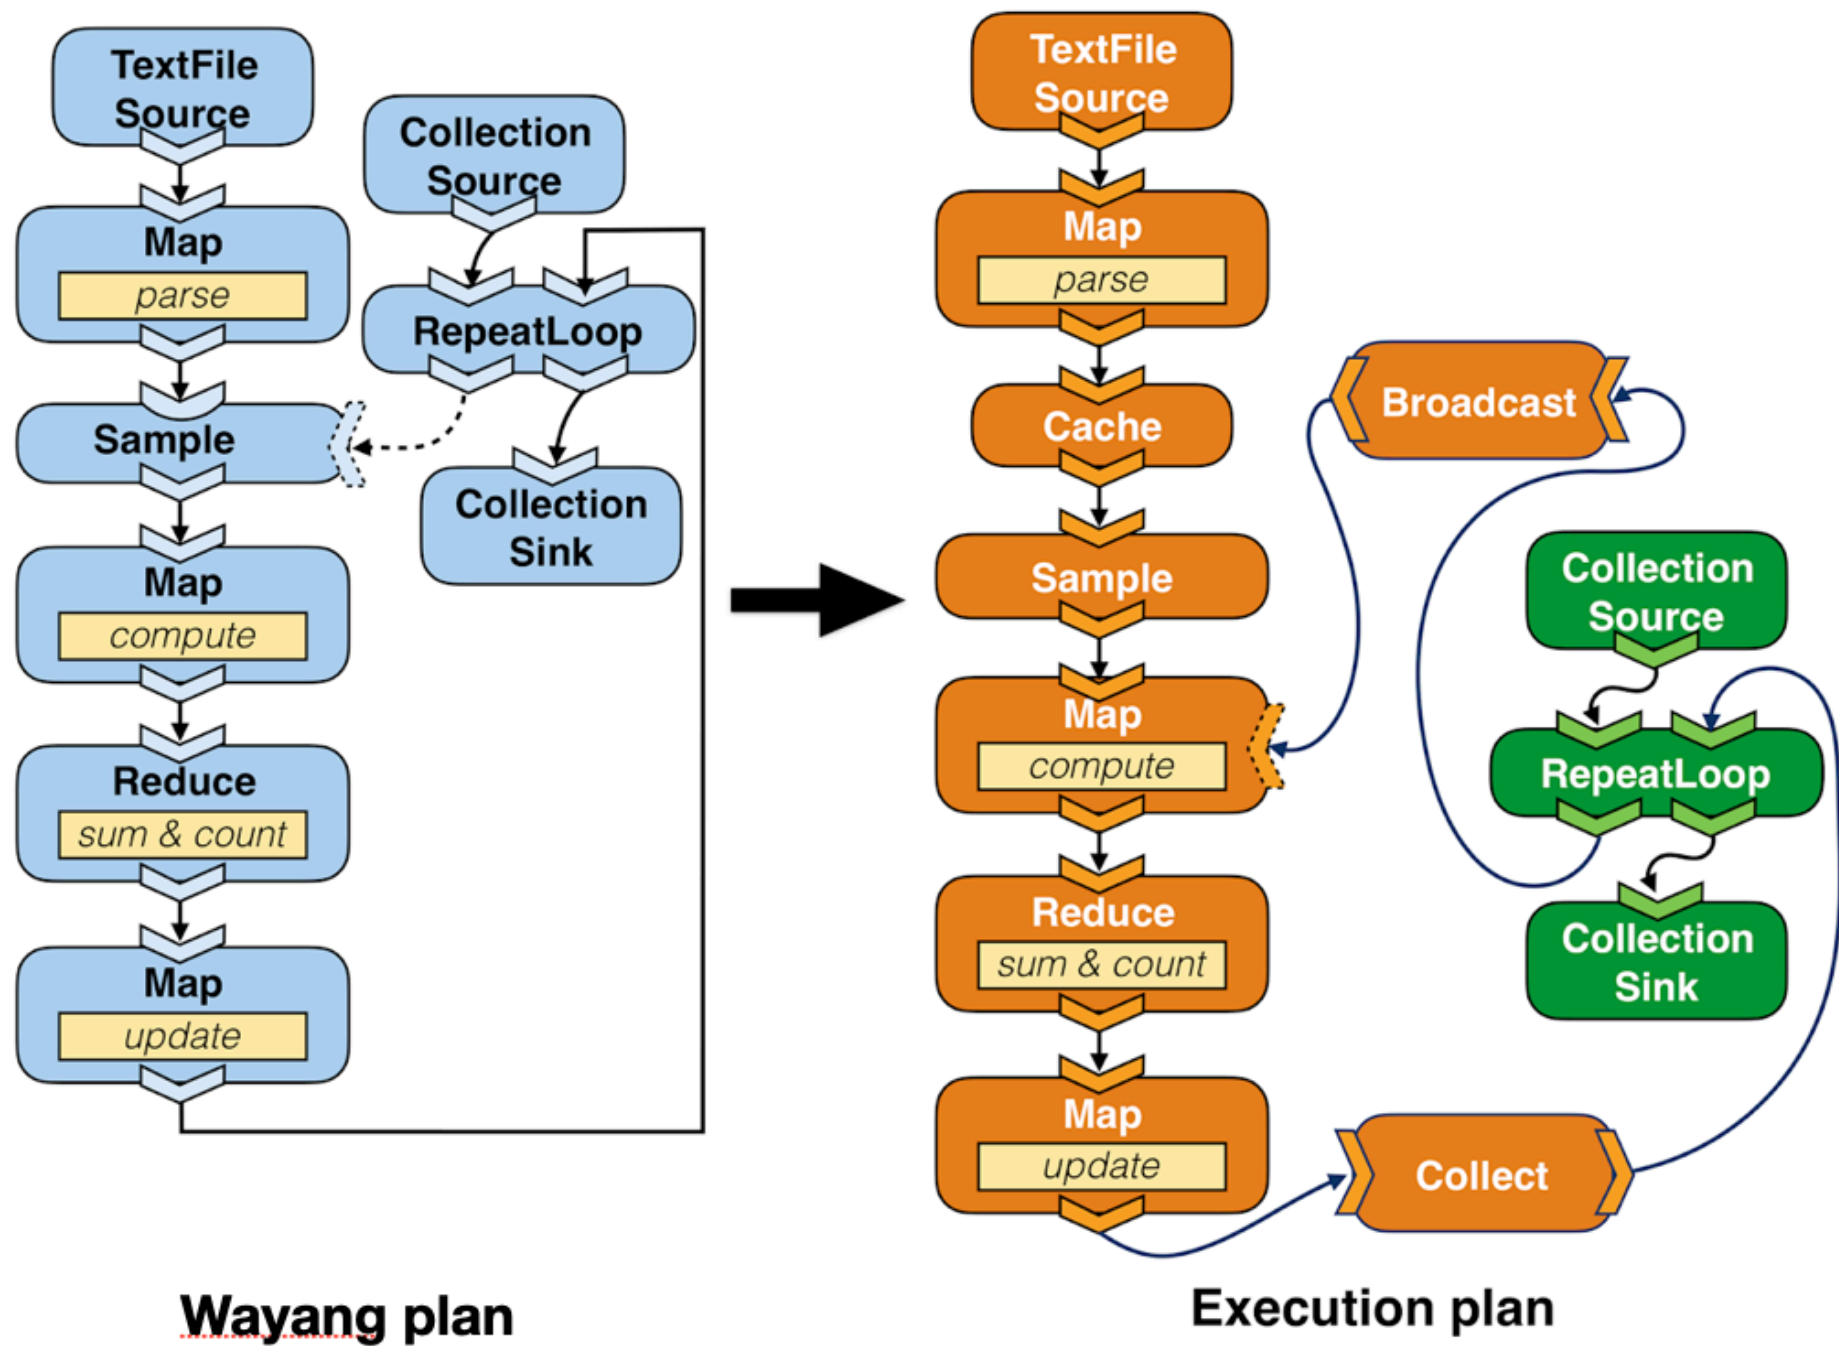
- Wayang operator
- Spark execution operator
- JavaStreams execution operator
- Input/Output
- UDF
- Data flow
- Broadcast data flow




Apache Wayang

Decoupling applications from underlying processing platforms

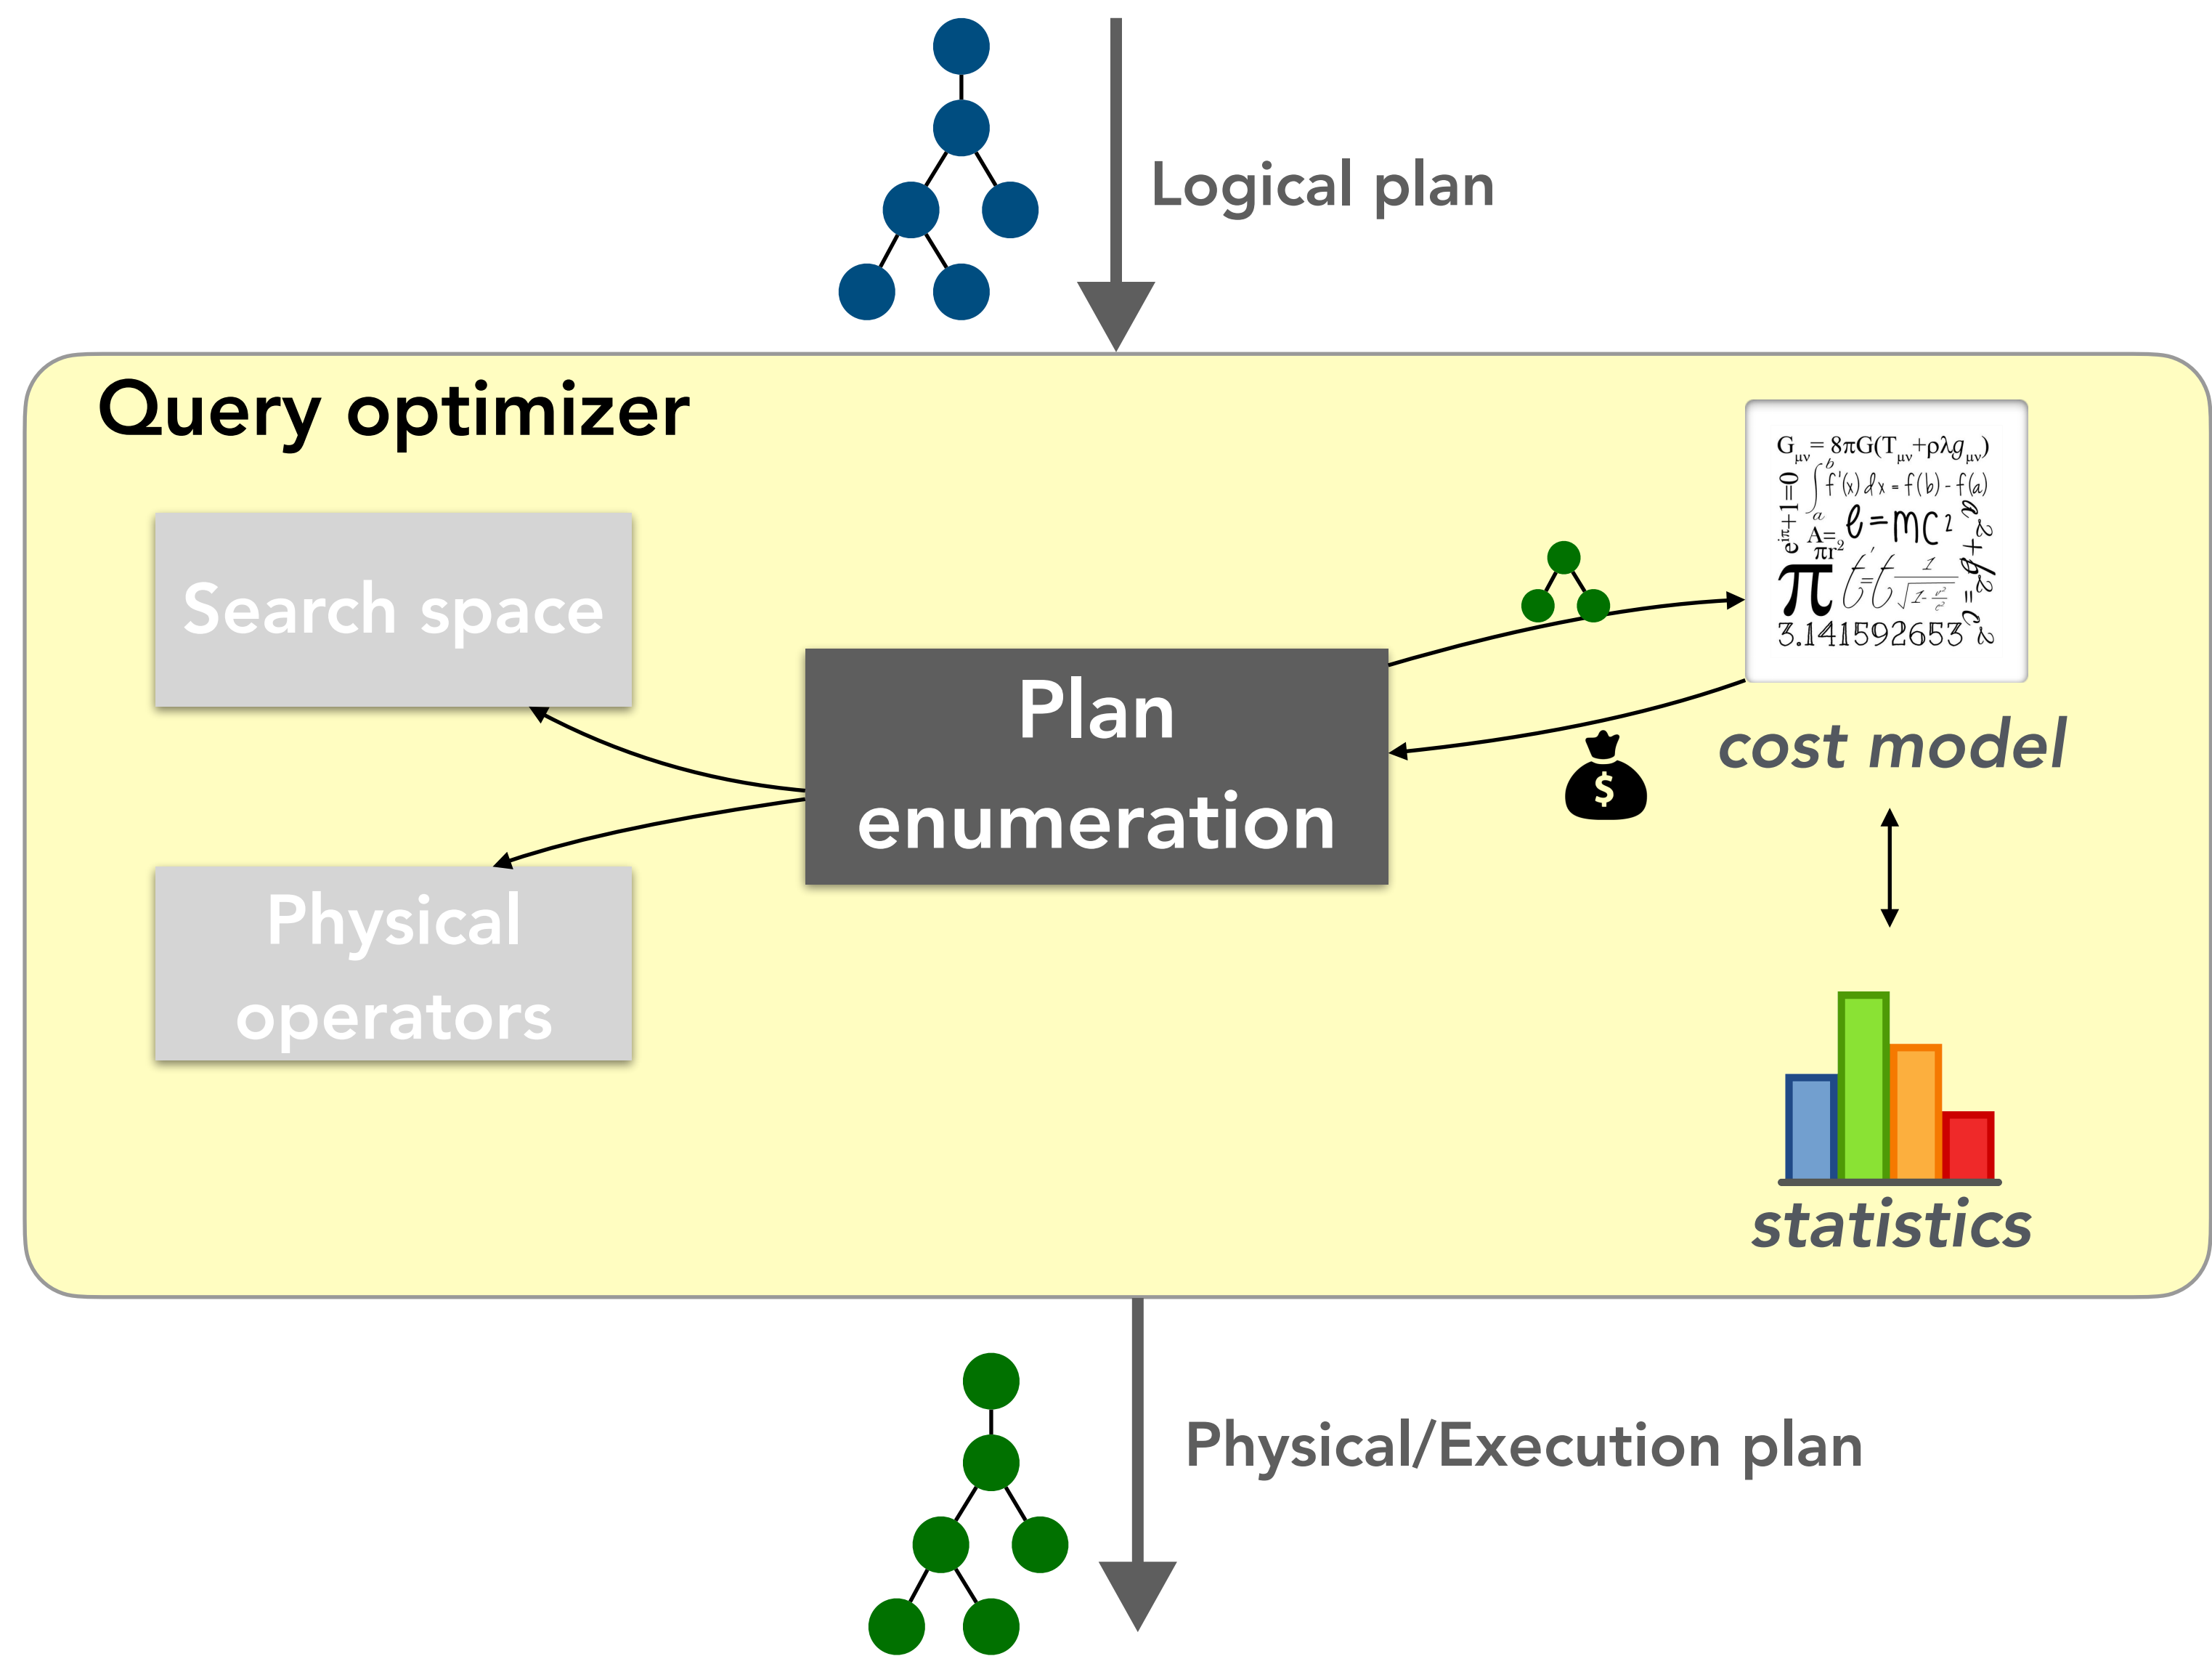
Wayang operator Spark execution operator JavaStreams execution operator
Input/Output UDF → Data flow ⇢ Broadcast data flow



SGD



Cost-based query optimization process



Cost model tuning is important ...

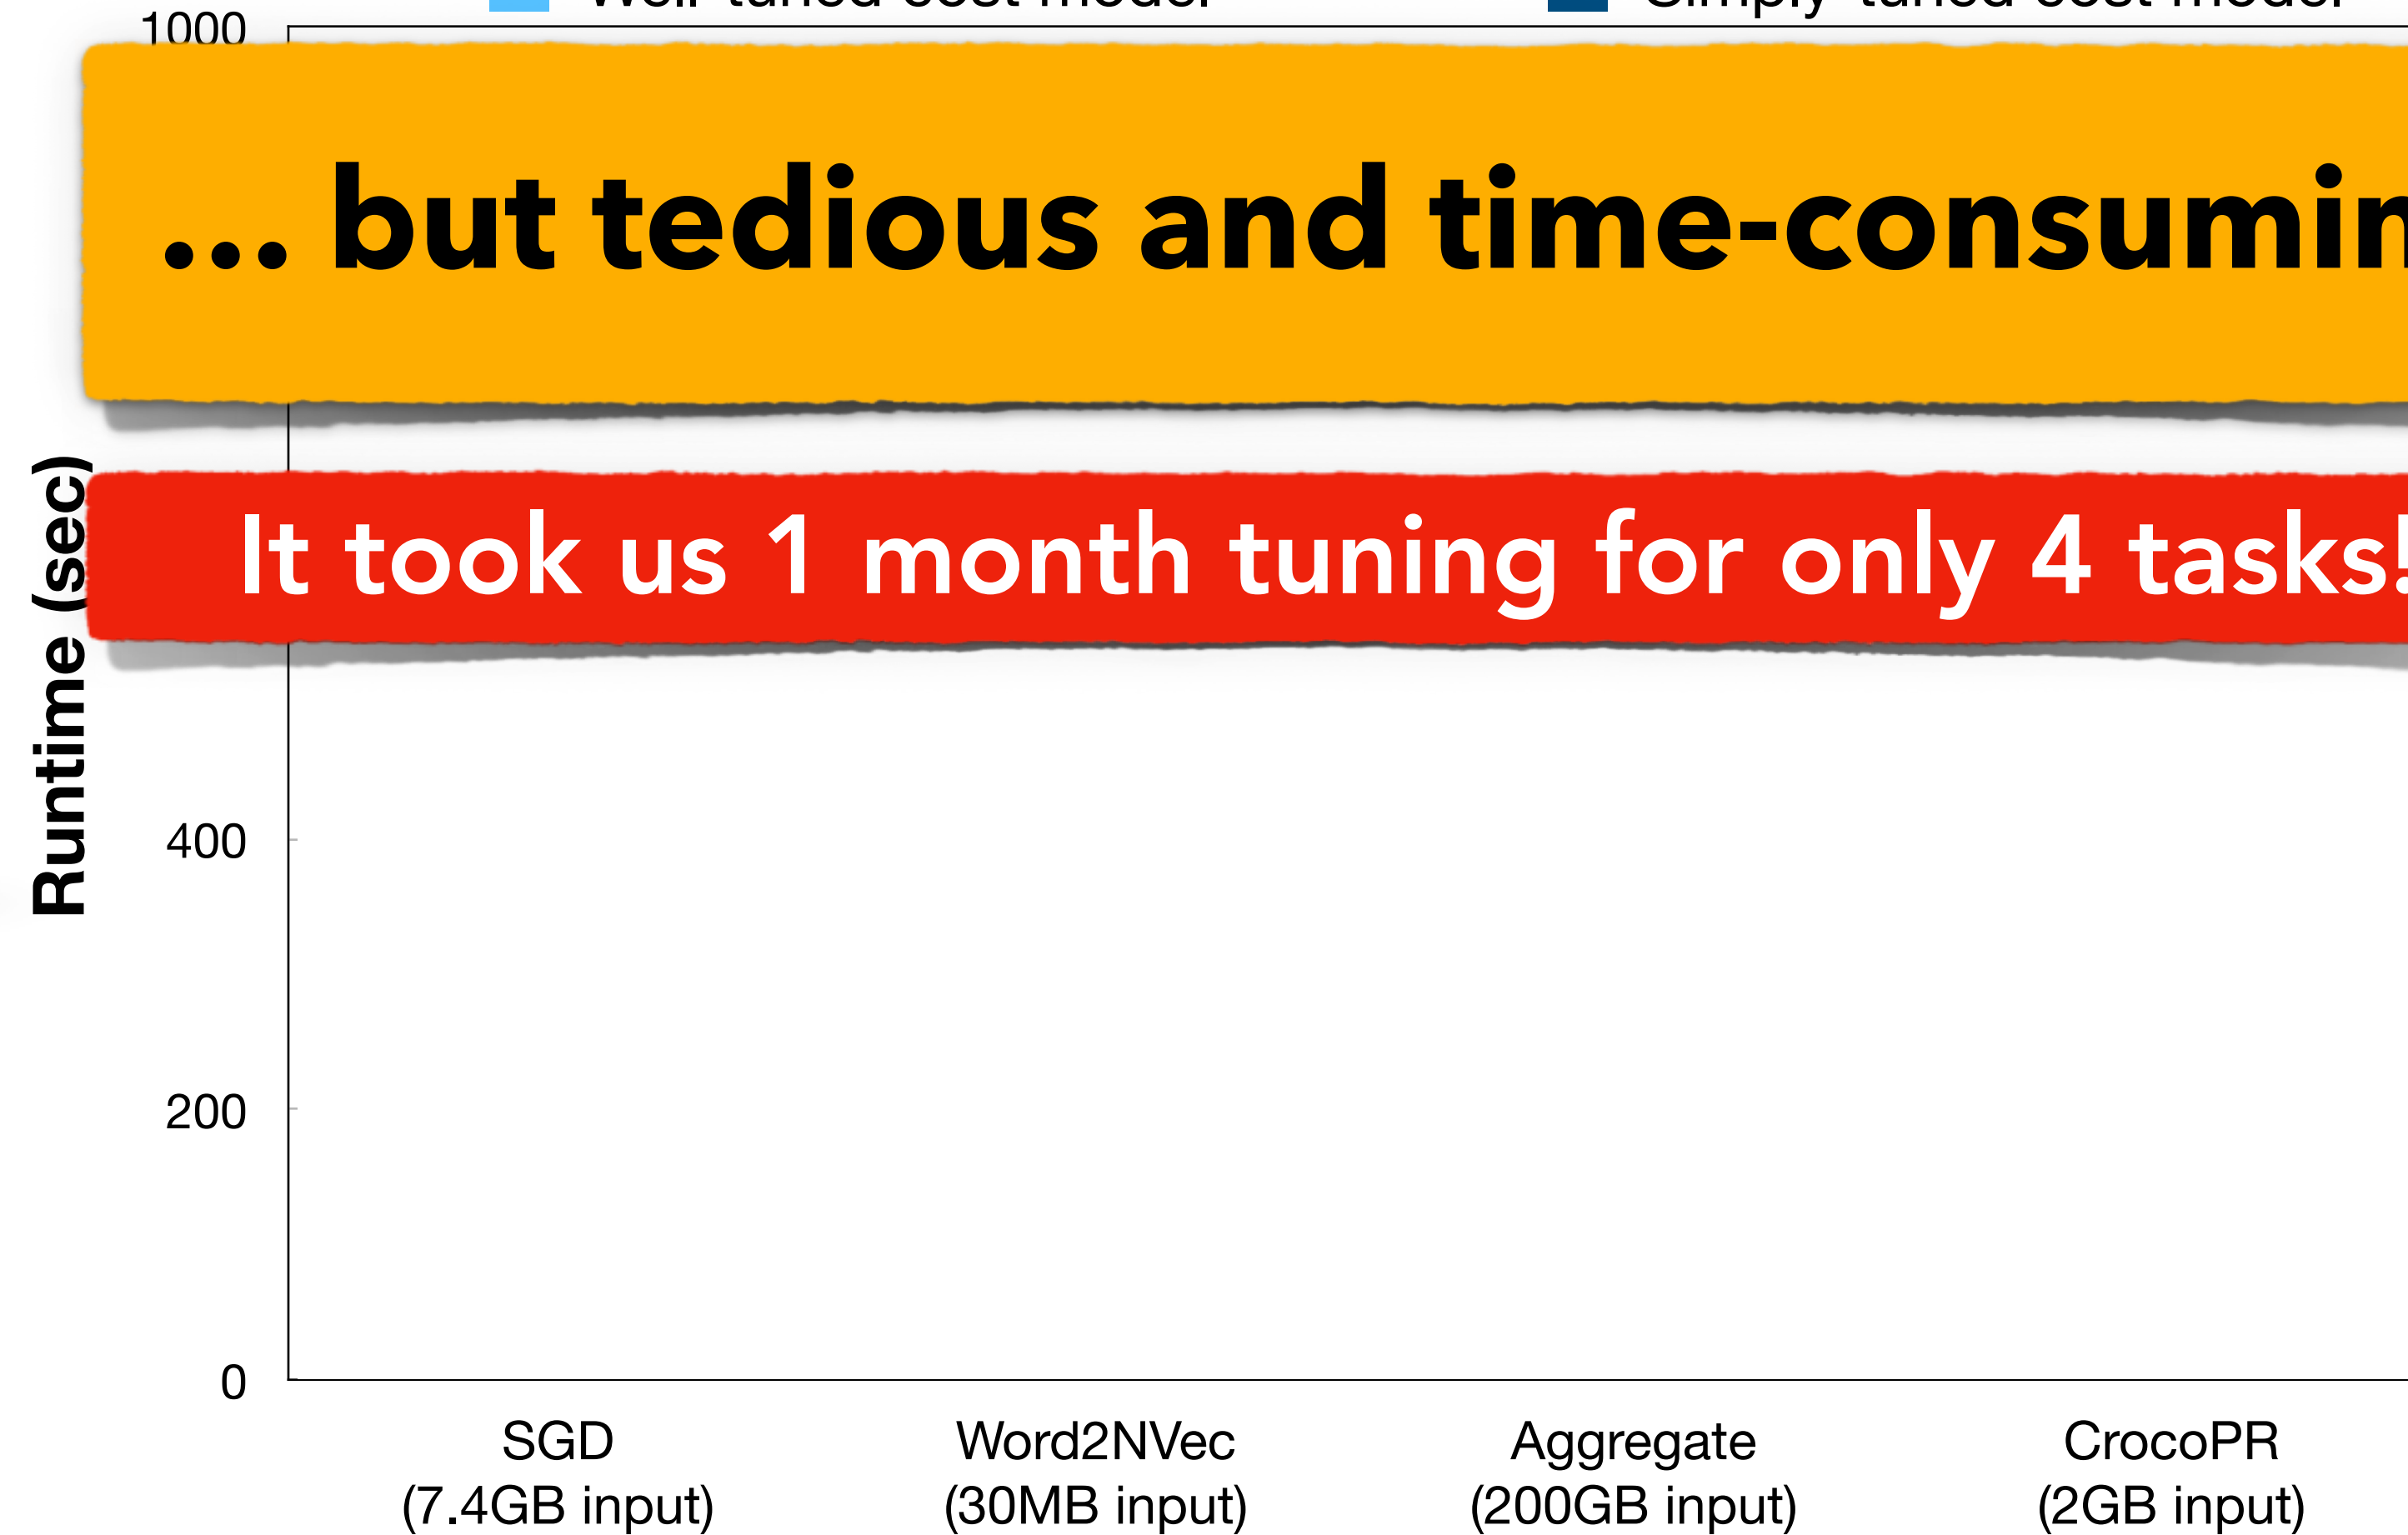


Well-tuned cost model Simply-tuned cost model

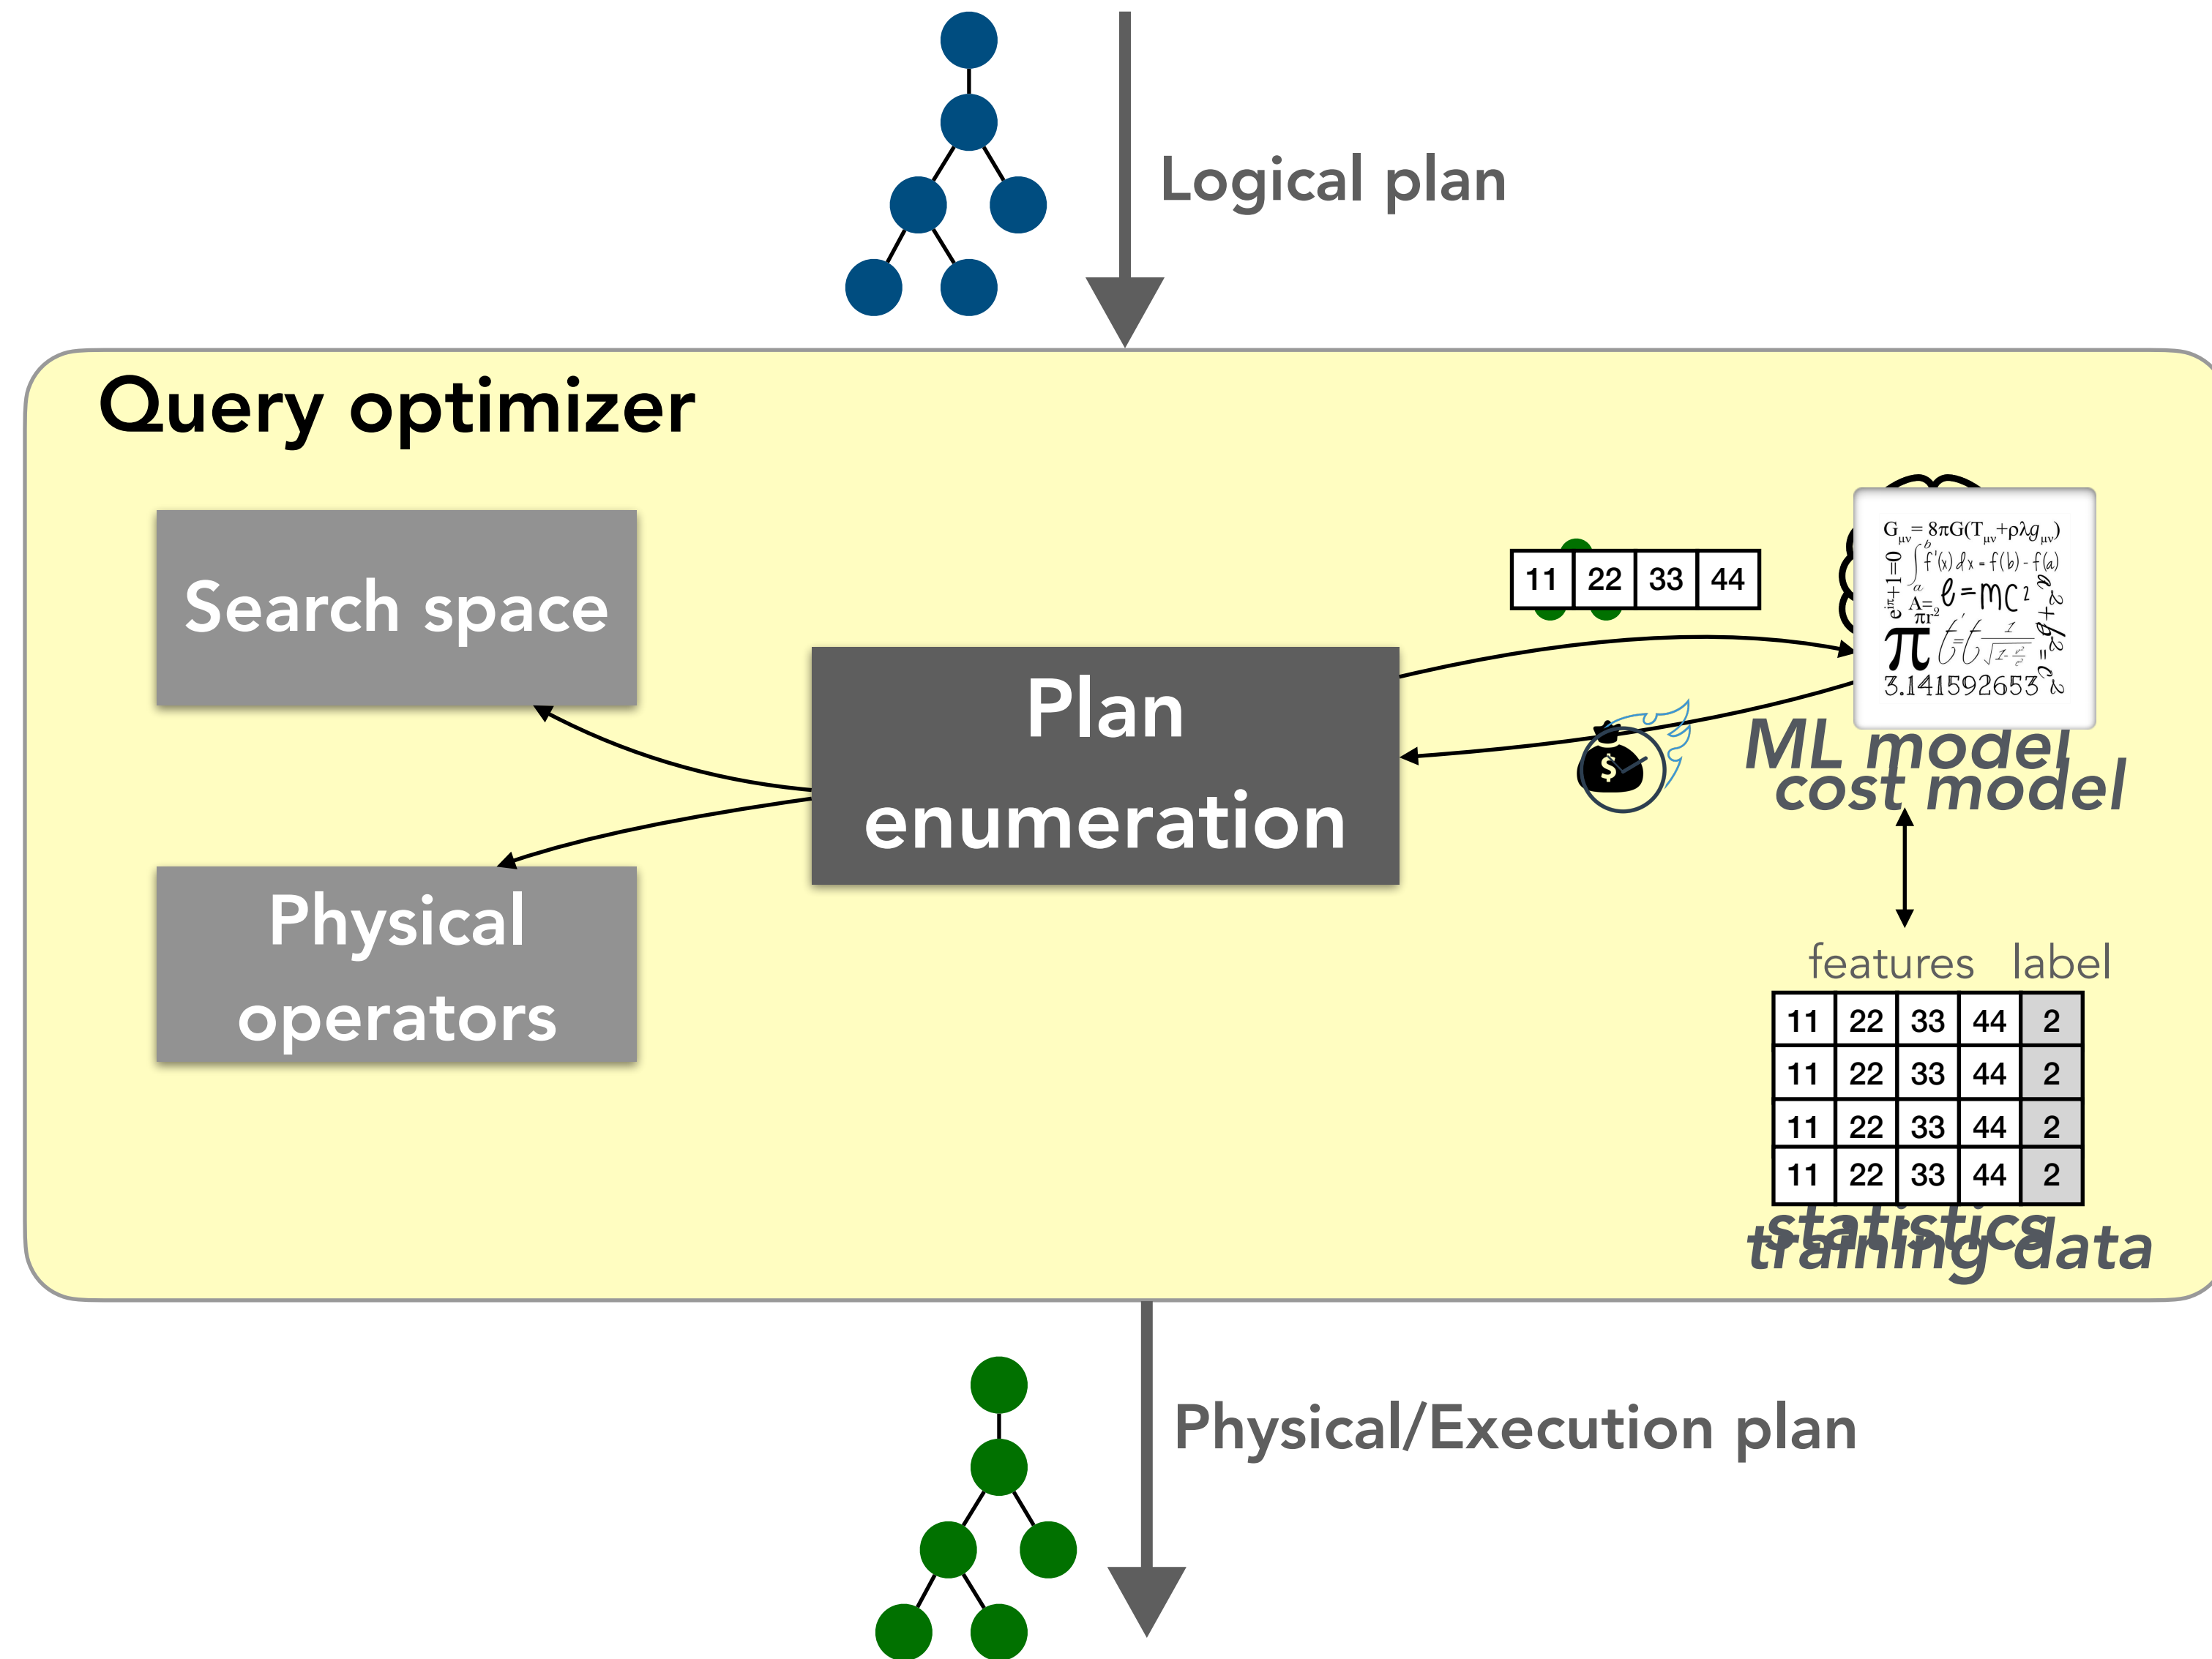
... but tedious and time-consuming!

It took us 1 month tuning for only 4 tasks!

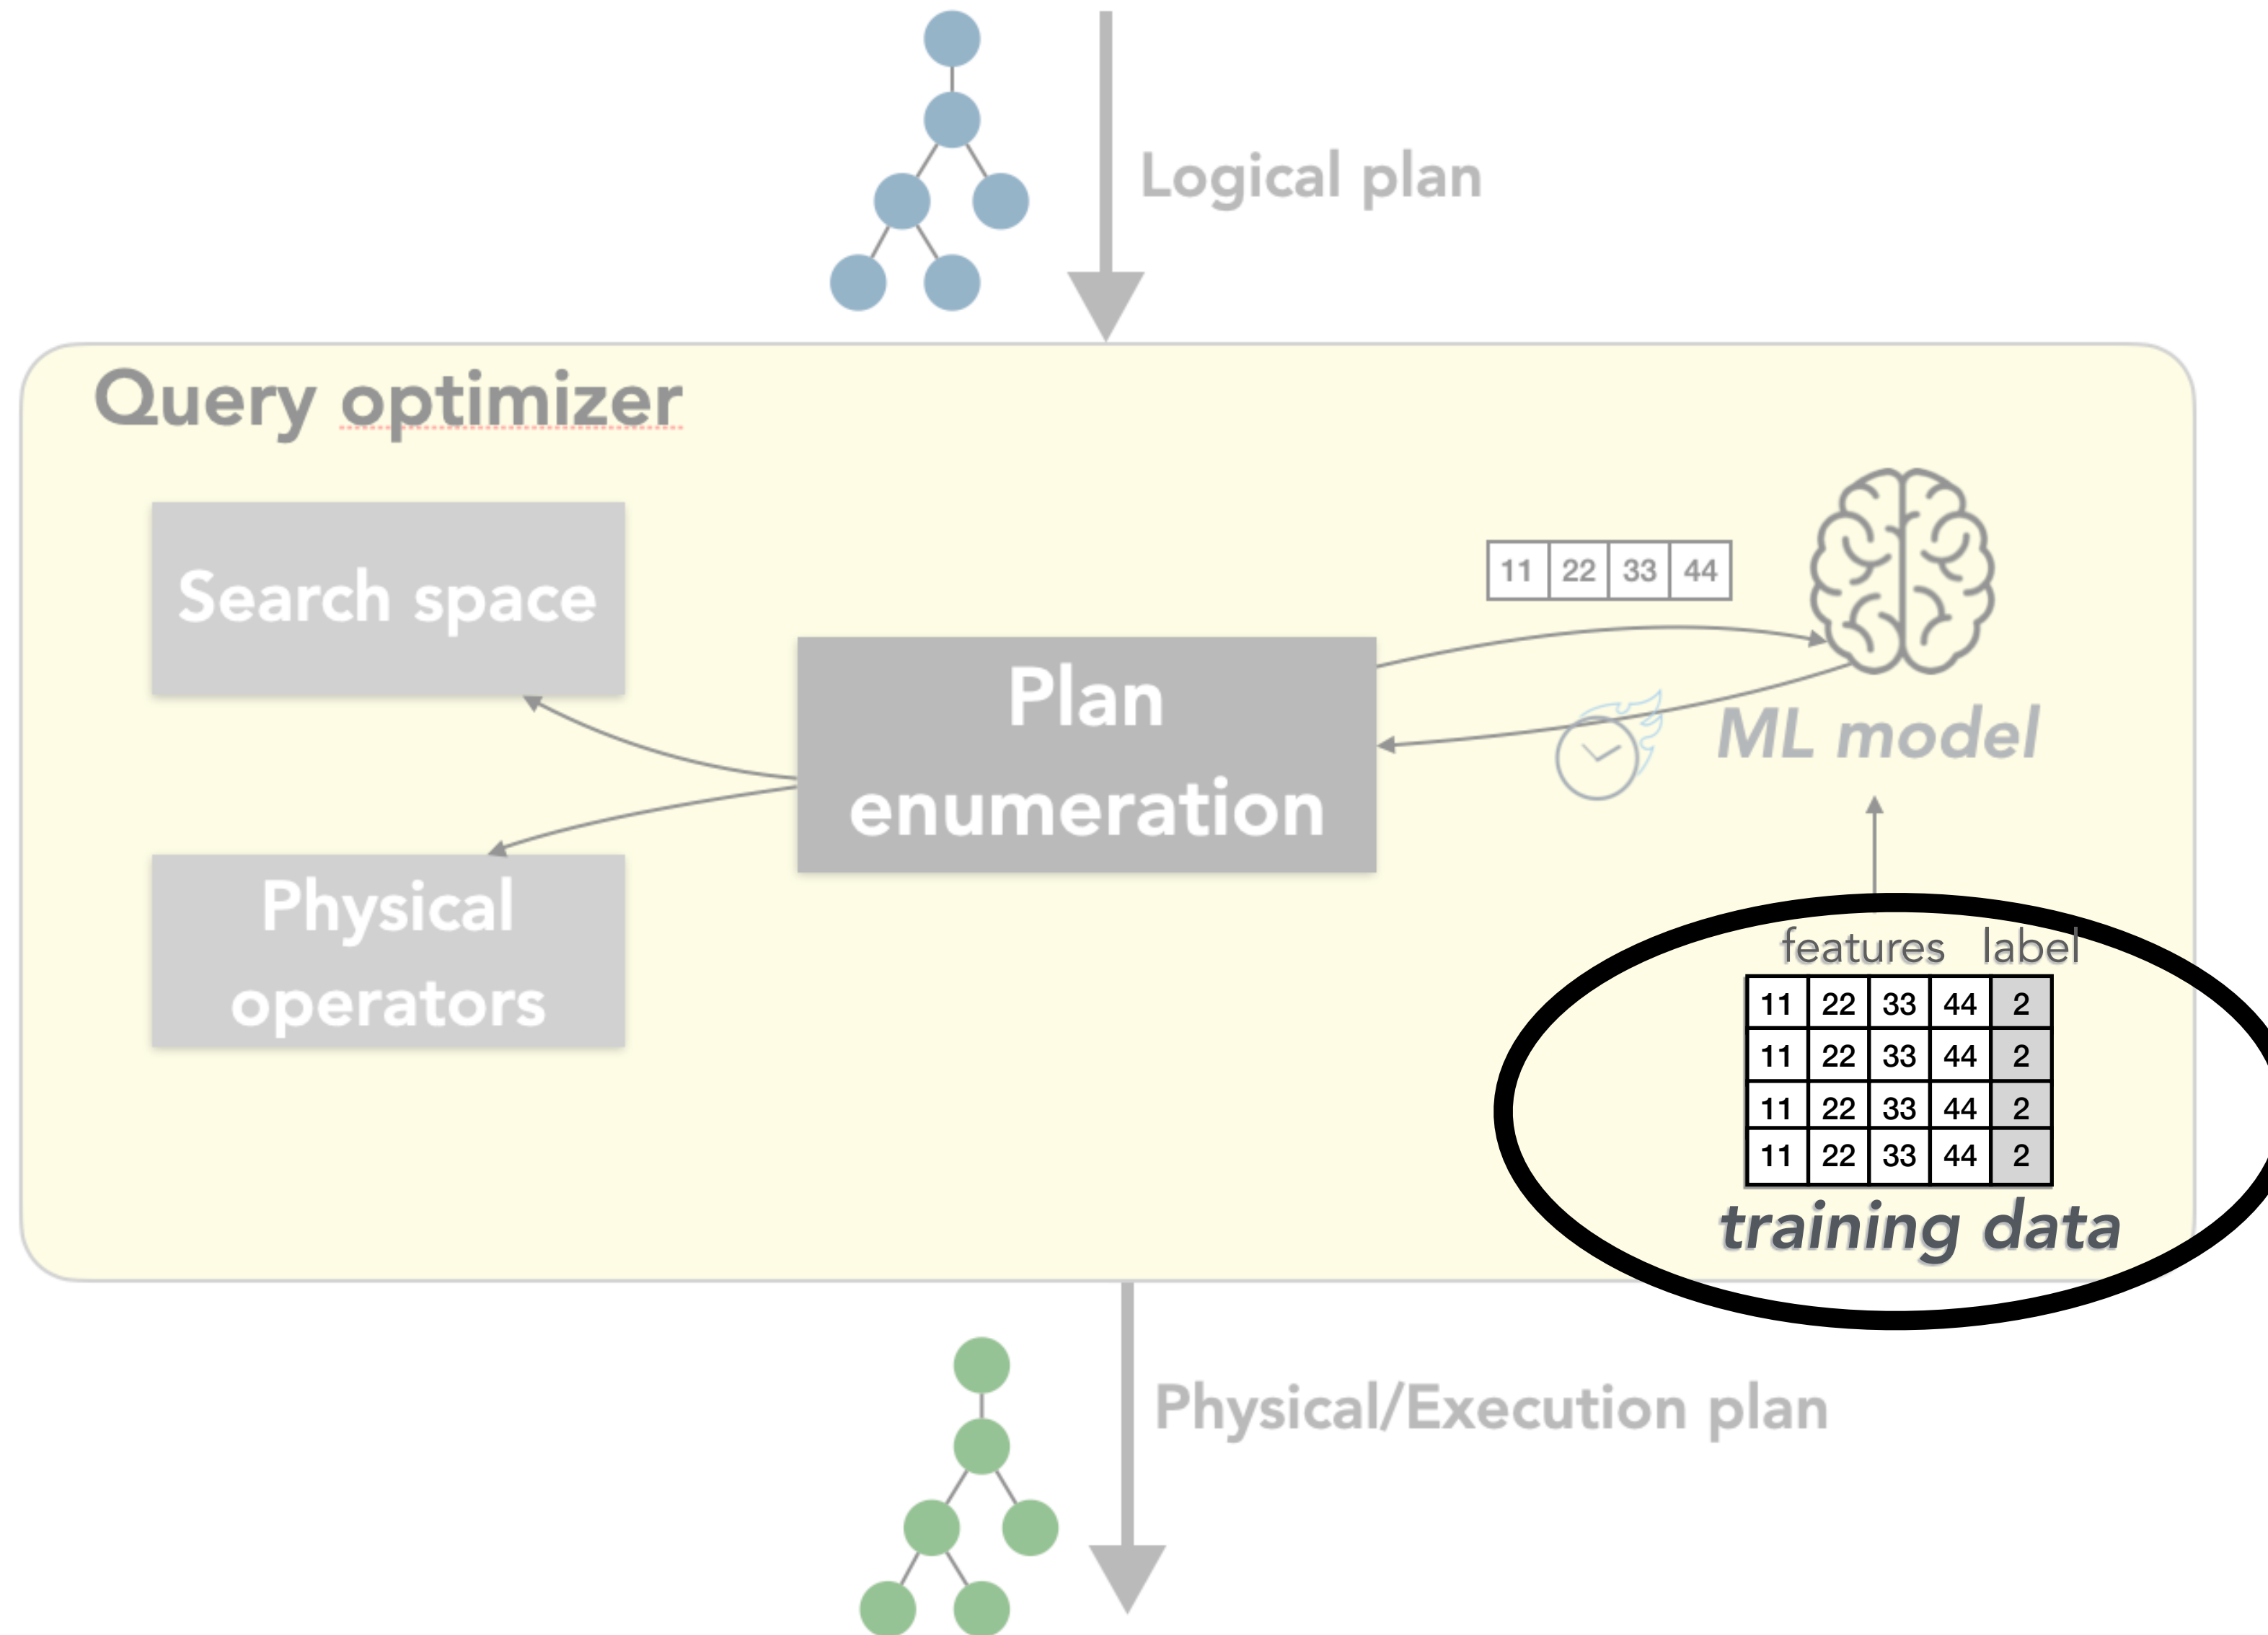
Using real cardinalities



ML to the rescue



Learning-based query optimization



**ML models are
data-hungry**



Naive solution to generate training data

Time-consuming

Extrapolated cost of 10,000 plans with 1TB input data > 6 months*

P1
P2
P3
...

P1	8s
P2	3s
P3	36s
...


(1) Manually create THOUSANDS of (good and BAD) plans

(2) Extract features for each plan

(3) Execute ALL plans to collect labels (e.g., exec time)

*On four-node quad-core cluster

Proposed solution to generate training data



	F1	F2	F3
P1
P2
P3
...

	F1	F2	F3	L
P1	8s
P2	3s
P3	36s
...

real
estimated

(1) Generate diverse THOUSANDS of plans based on initial BAD workload


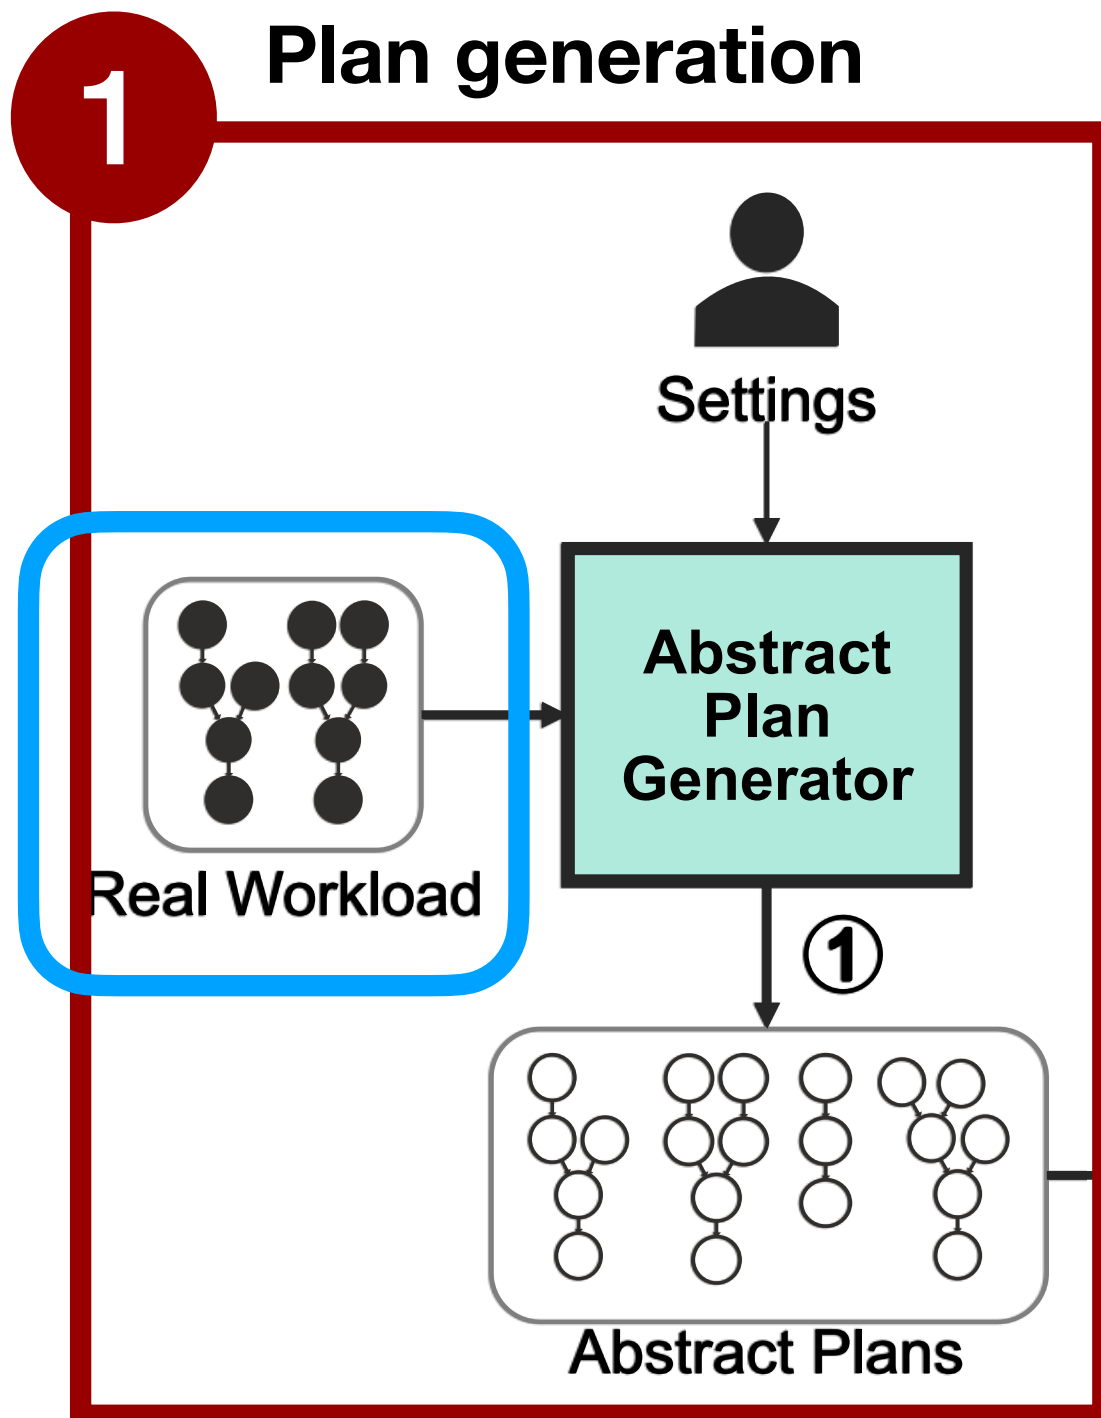
(2) Extract features for each plan

(3) Execute all plans sampled at once (e.g., forecast time) test

DataFarm: Training data generator for learning-based QO

White-Box

Data-driven



Explainable ML process



Easy to customise and debug



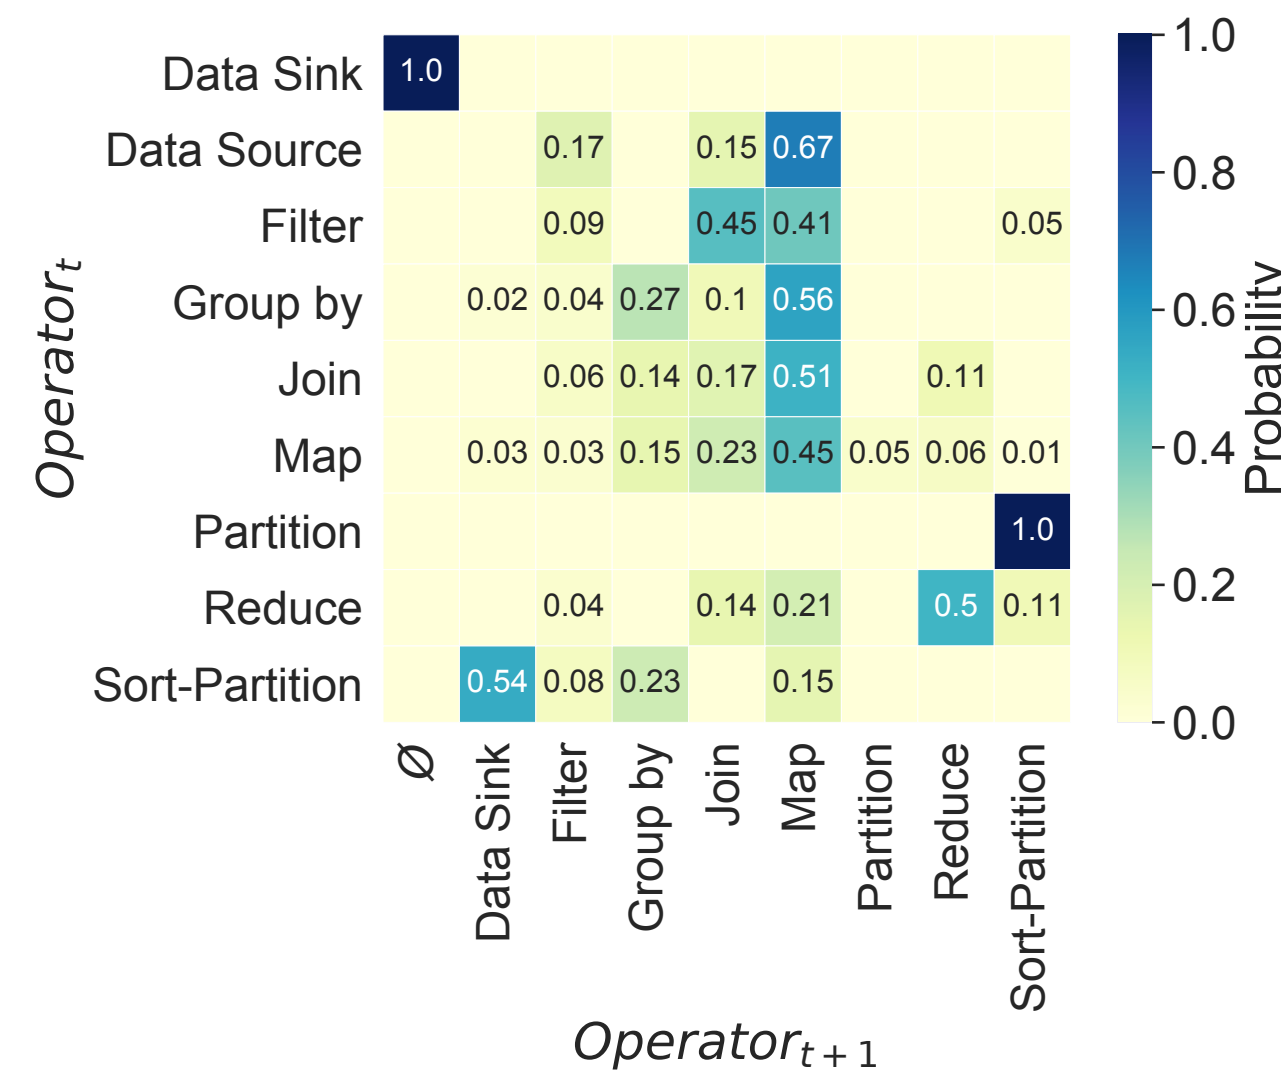
Expand your Training Limits! Generating Training Data for ML-based Data Management. **SIGMOD 2021**: 1865-1878
DataFarm: Farm Your ML-based Query Optimizer's Food! – Human-Guided Training Data Generation – **CIDR 2022**
Farming Your ML-based Query Optimizer's Food. **ICDE 2022 (best demo award)**

DataFarm: Training data generator for learning-based QO


◆ Learns real execution patterns



Children Transition Matrix

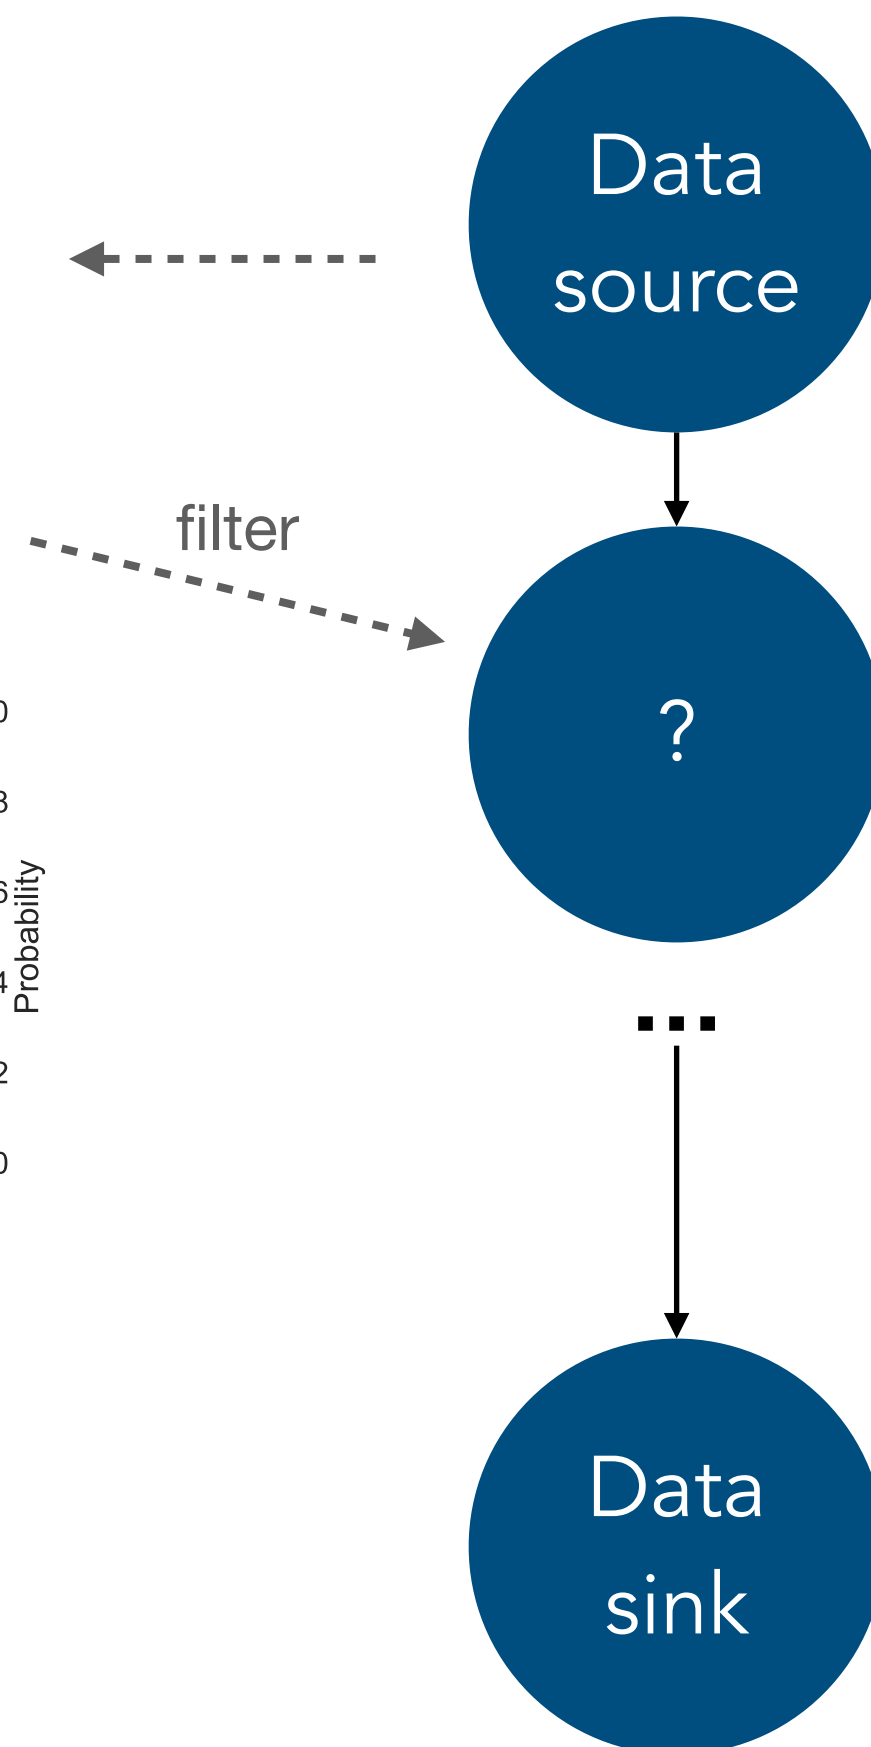


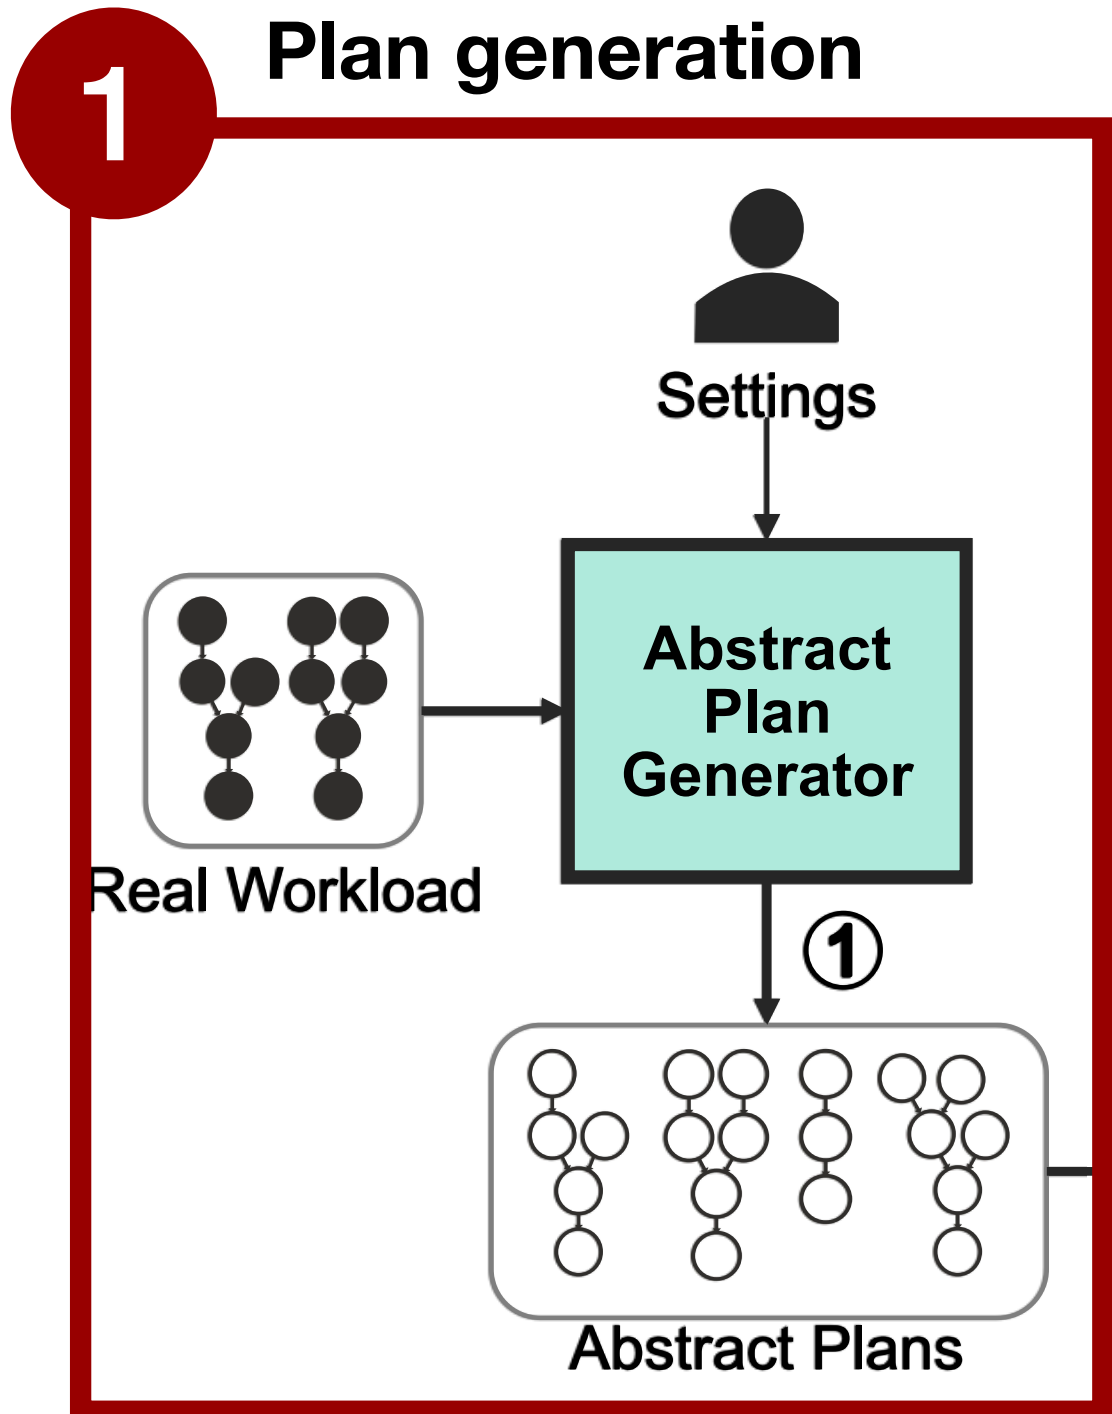


Parent Transition Matrix

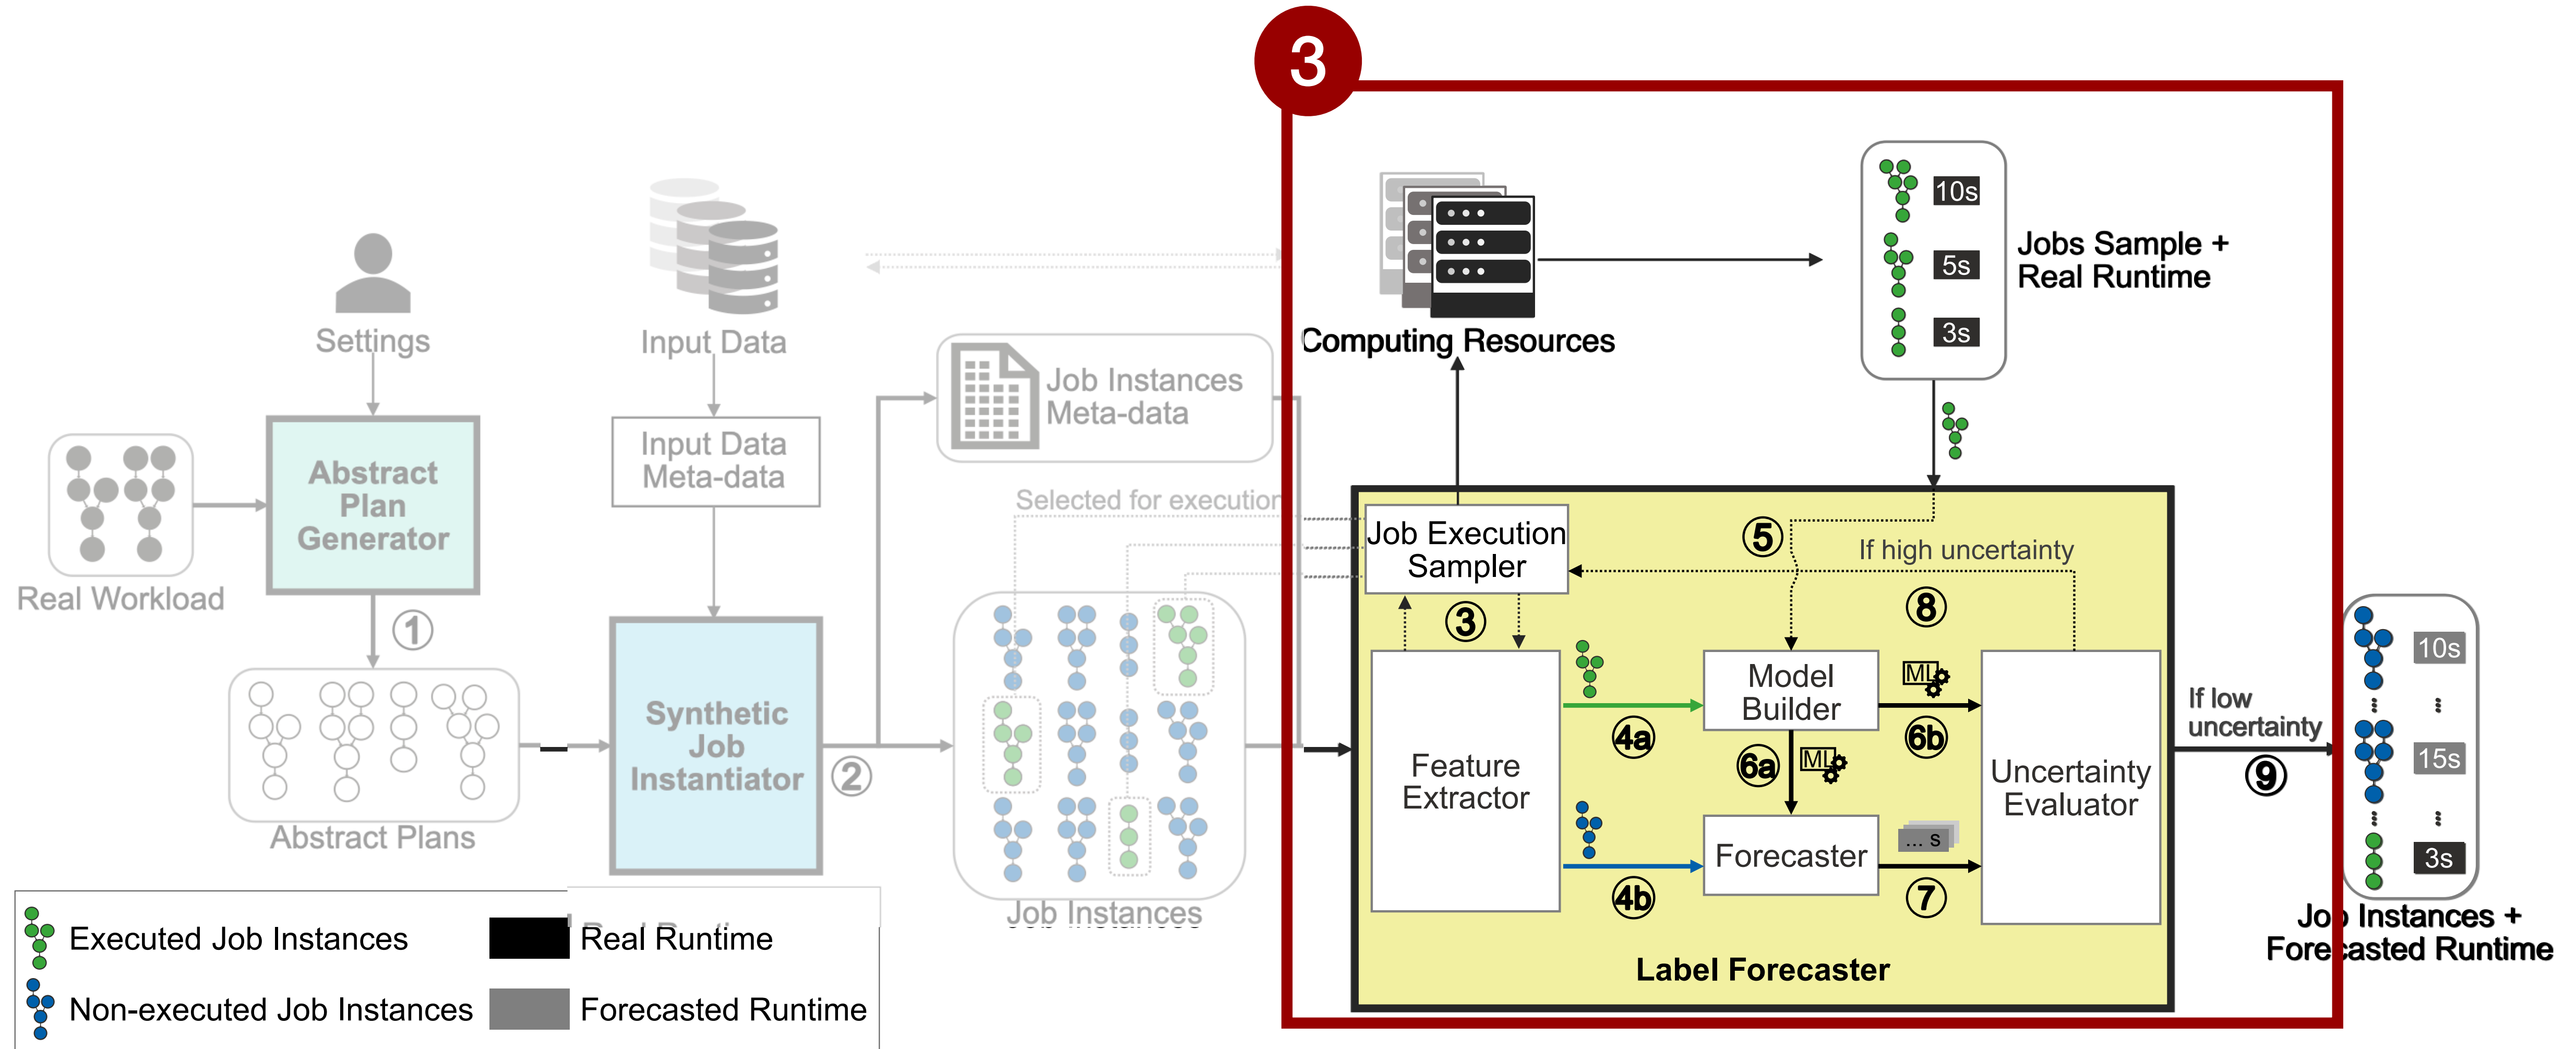


DataFarm: Training data generator for learning-based QO

- ◆ **Learns** real execution **patterns**
- ◆ **Generates** new **representative** plans

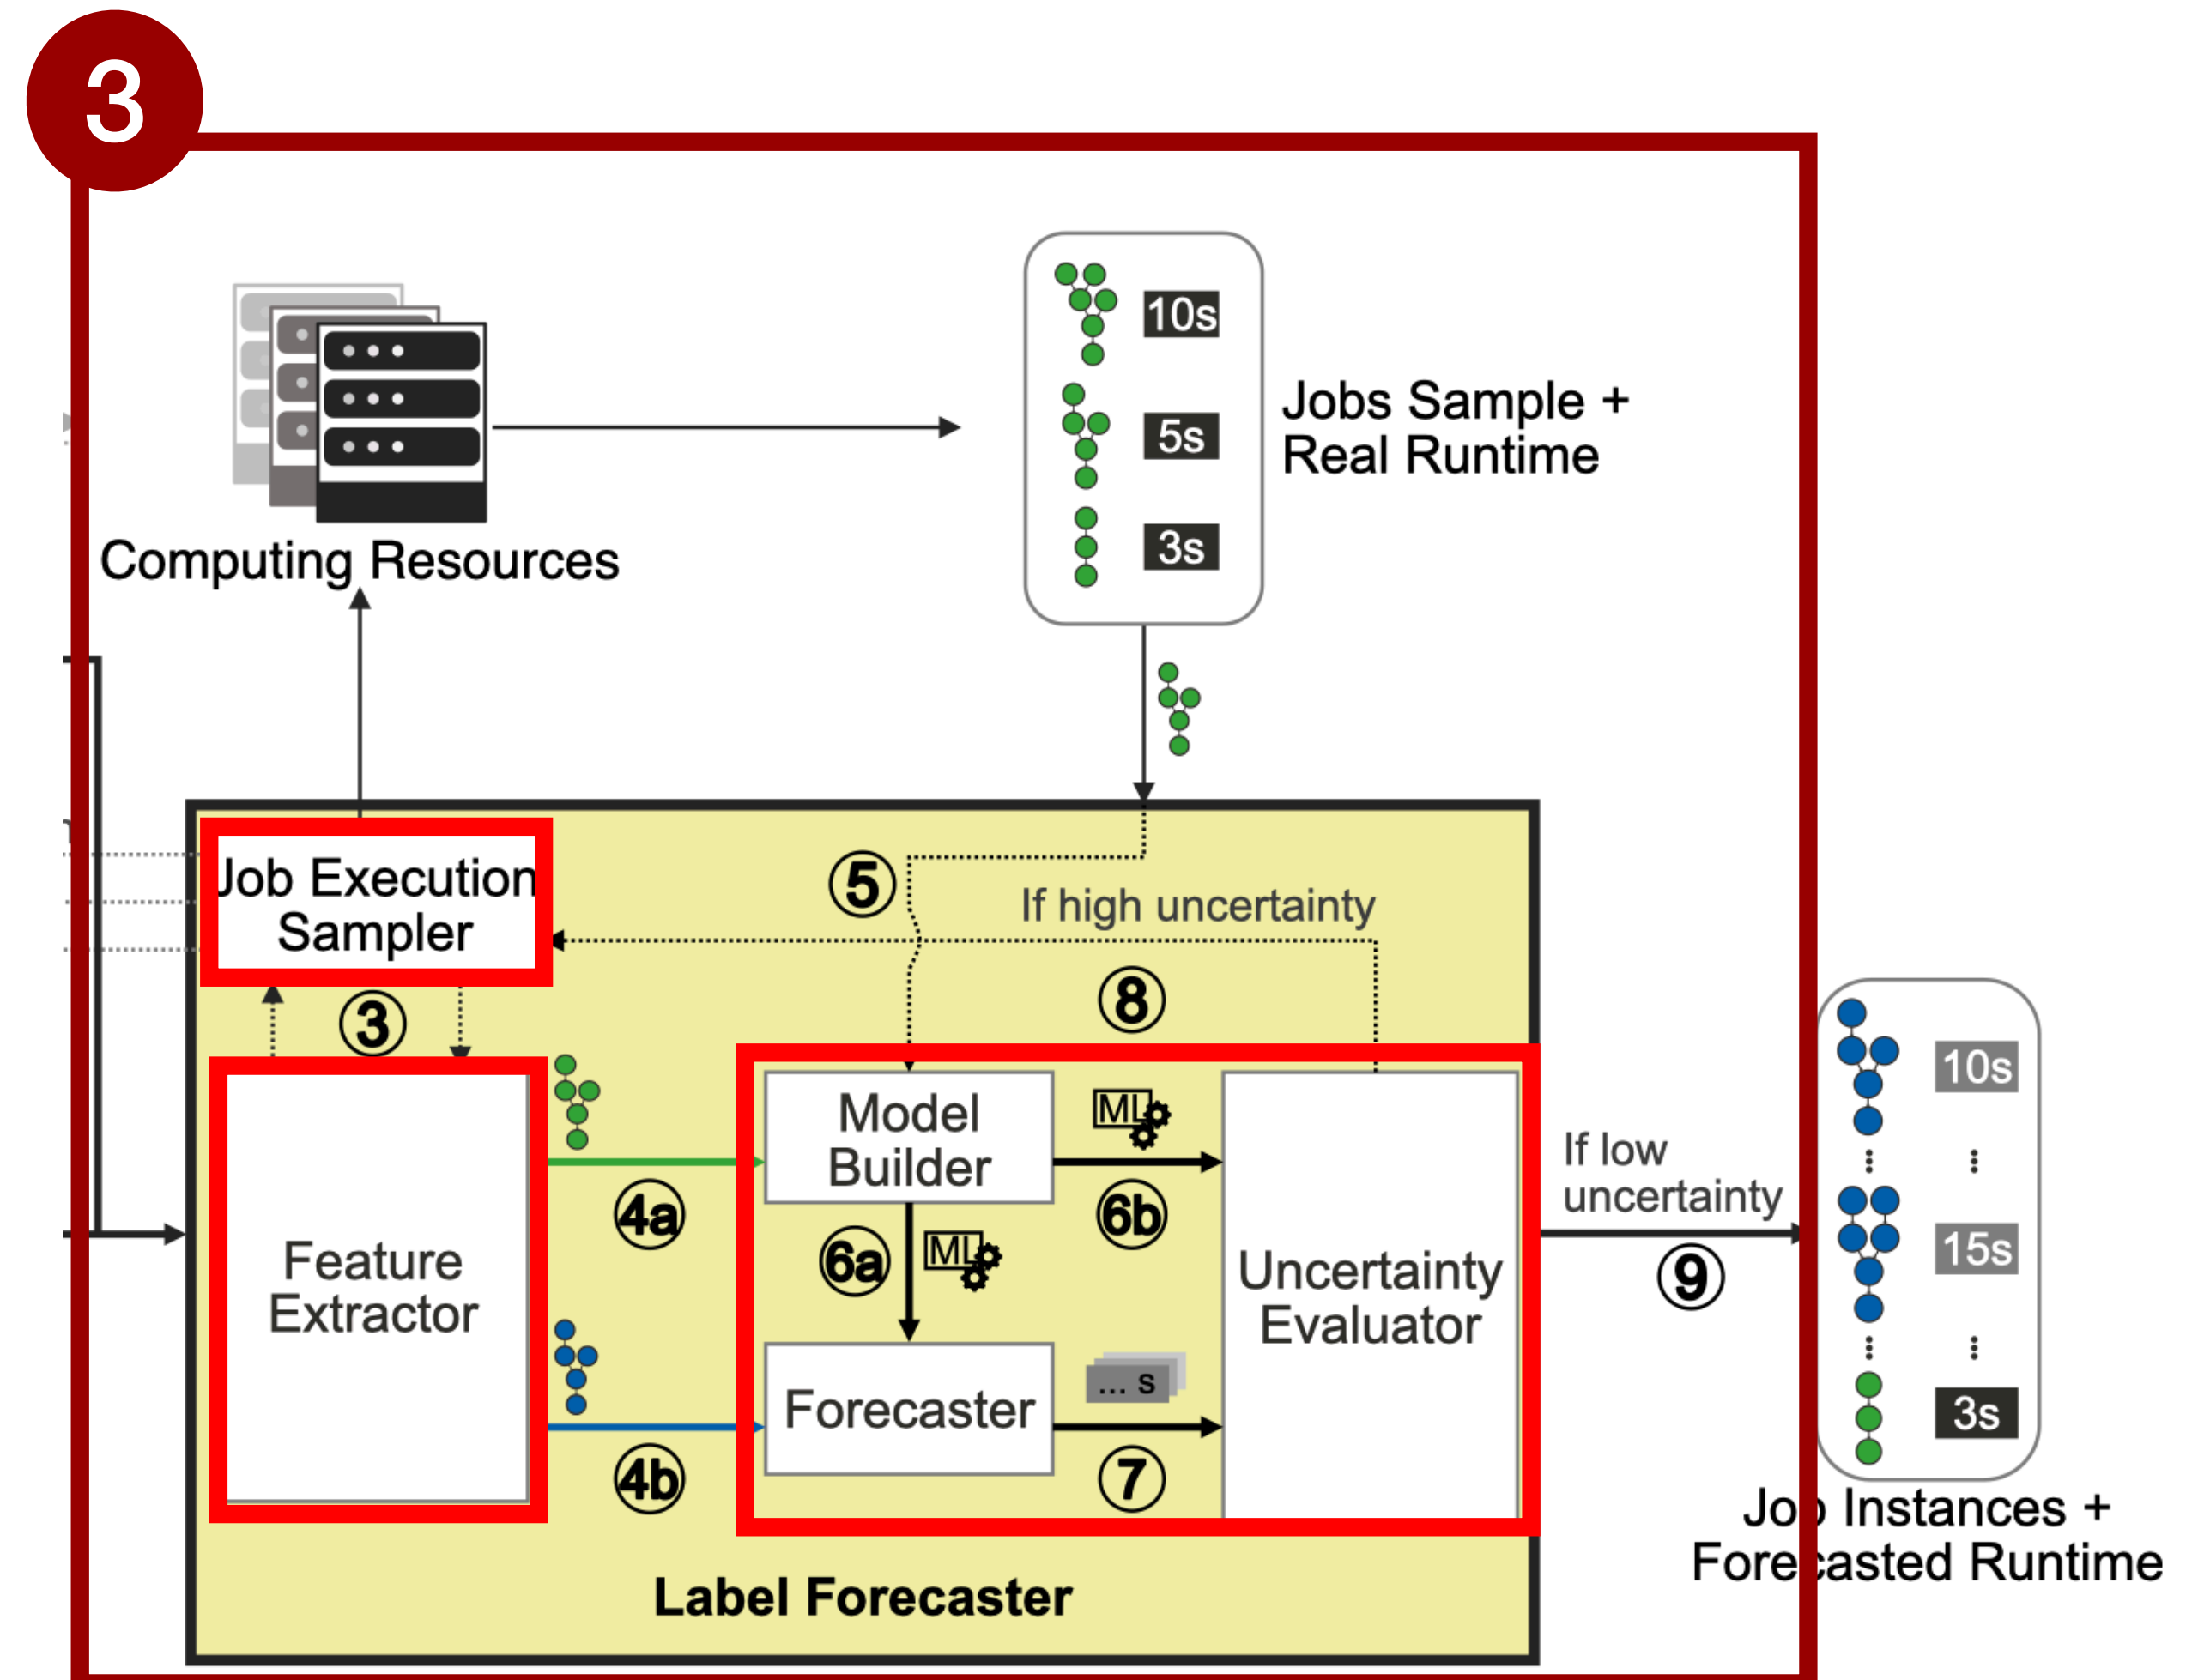


Focus on Label Forecaster



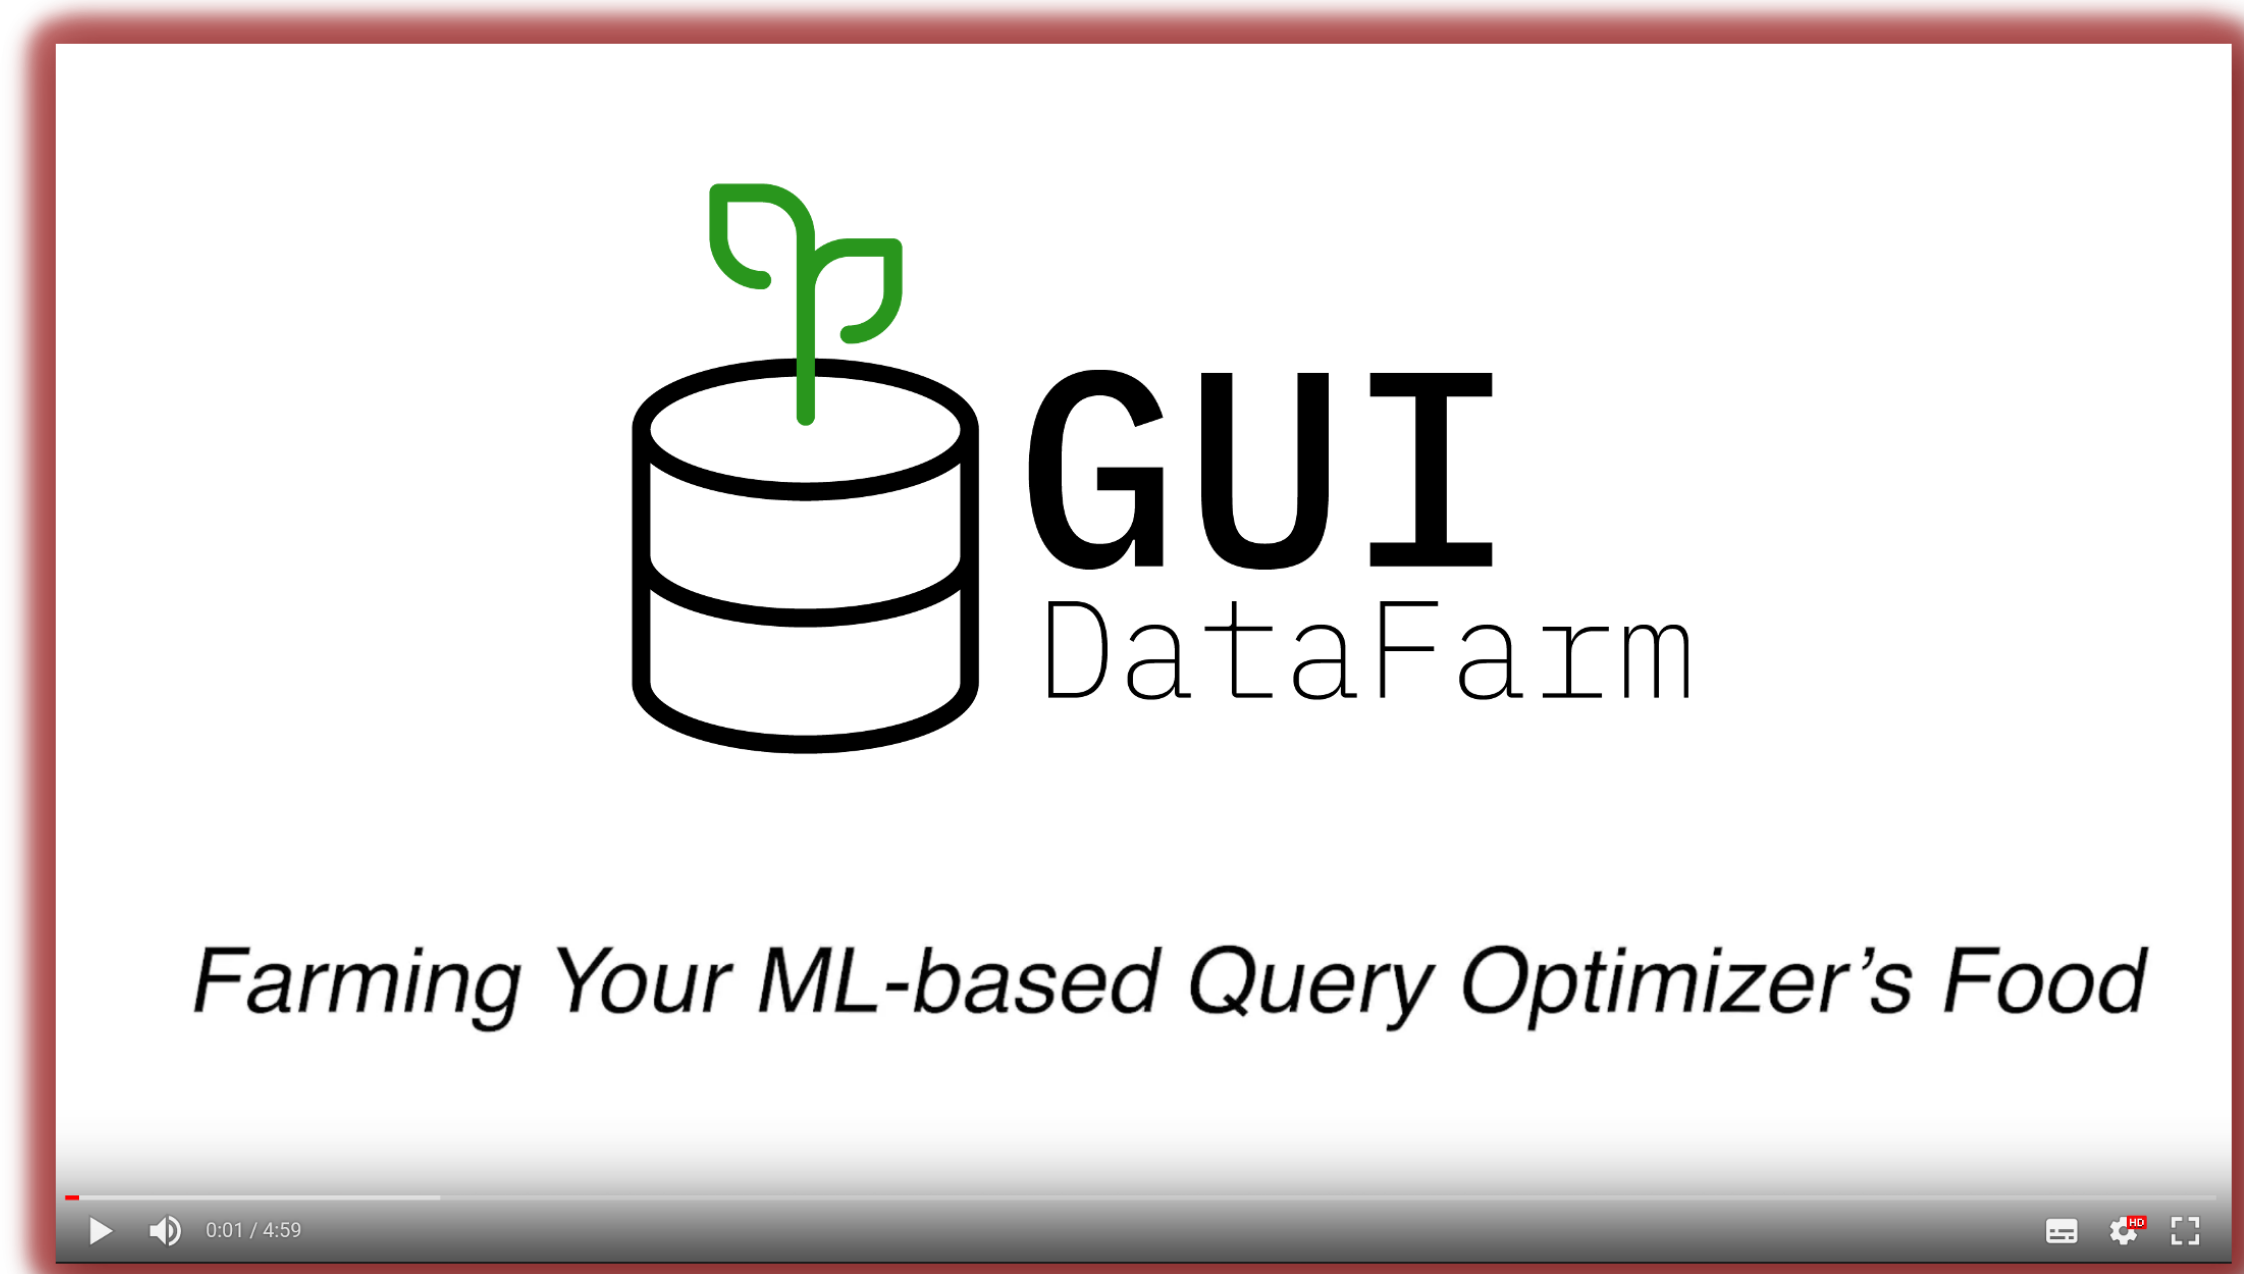
DataFarm: Label Forecaster

- ◆ Characterize jobs with **interpretable** features
- ◆ Find the smallest set of **representative** jobs to execute
- ◆ Predicts the **labels** and **uncertainty** for the non-executed jobs
- ◆ **Incrementally** execute more jobs and improve the model



GUIDataFarm: Human-in-the-loop data generation

Job Explorer (Job0v1, '1GB')

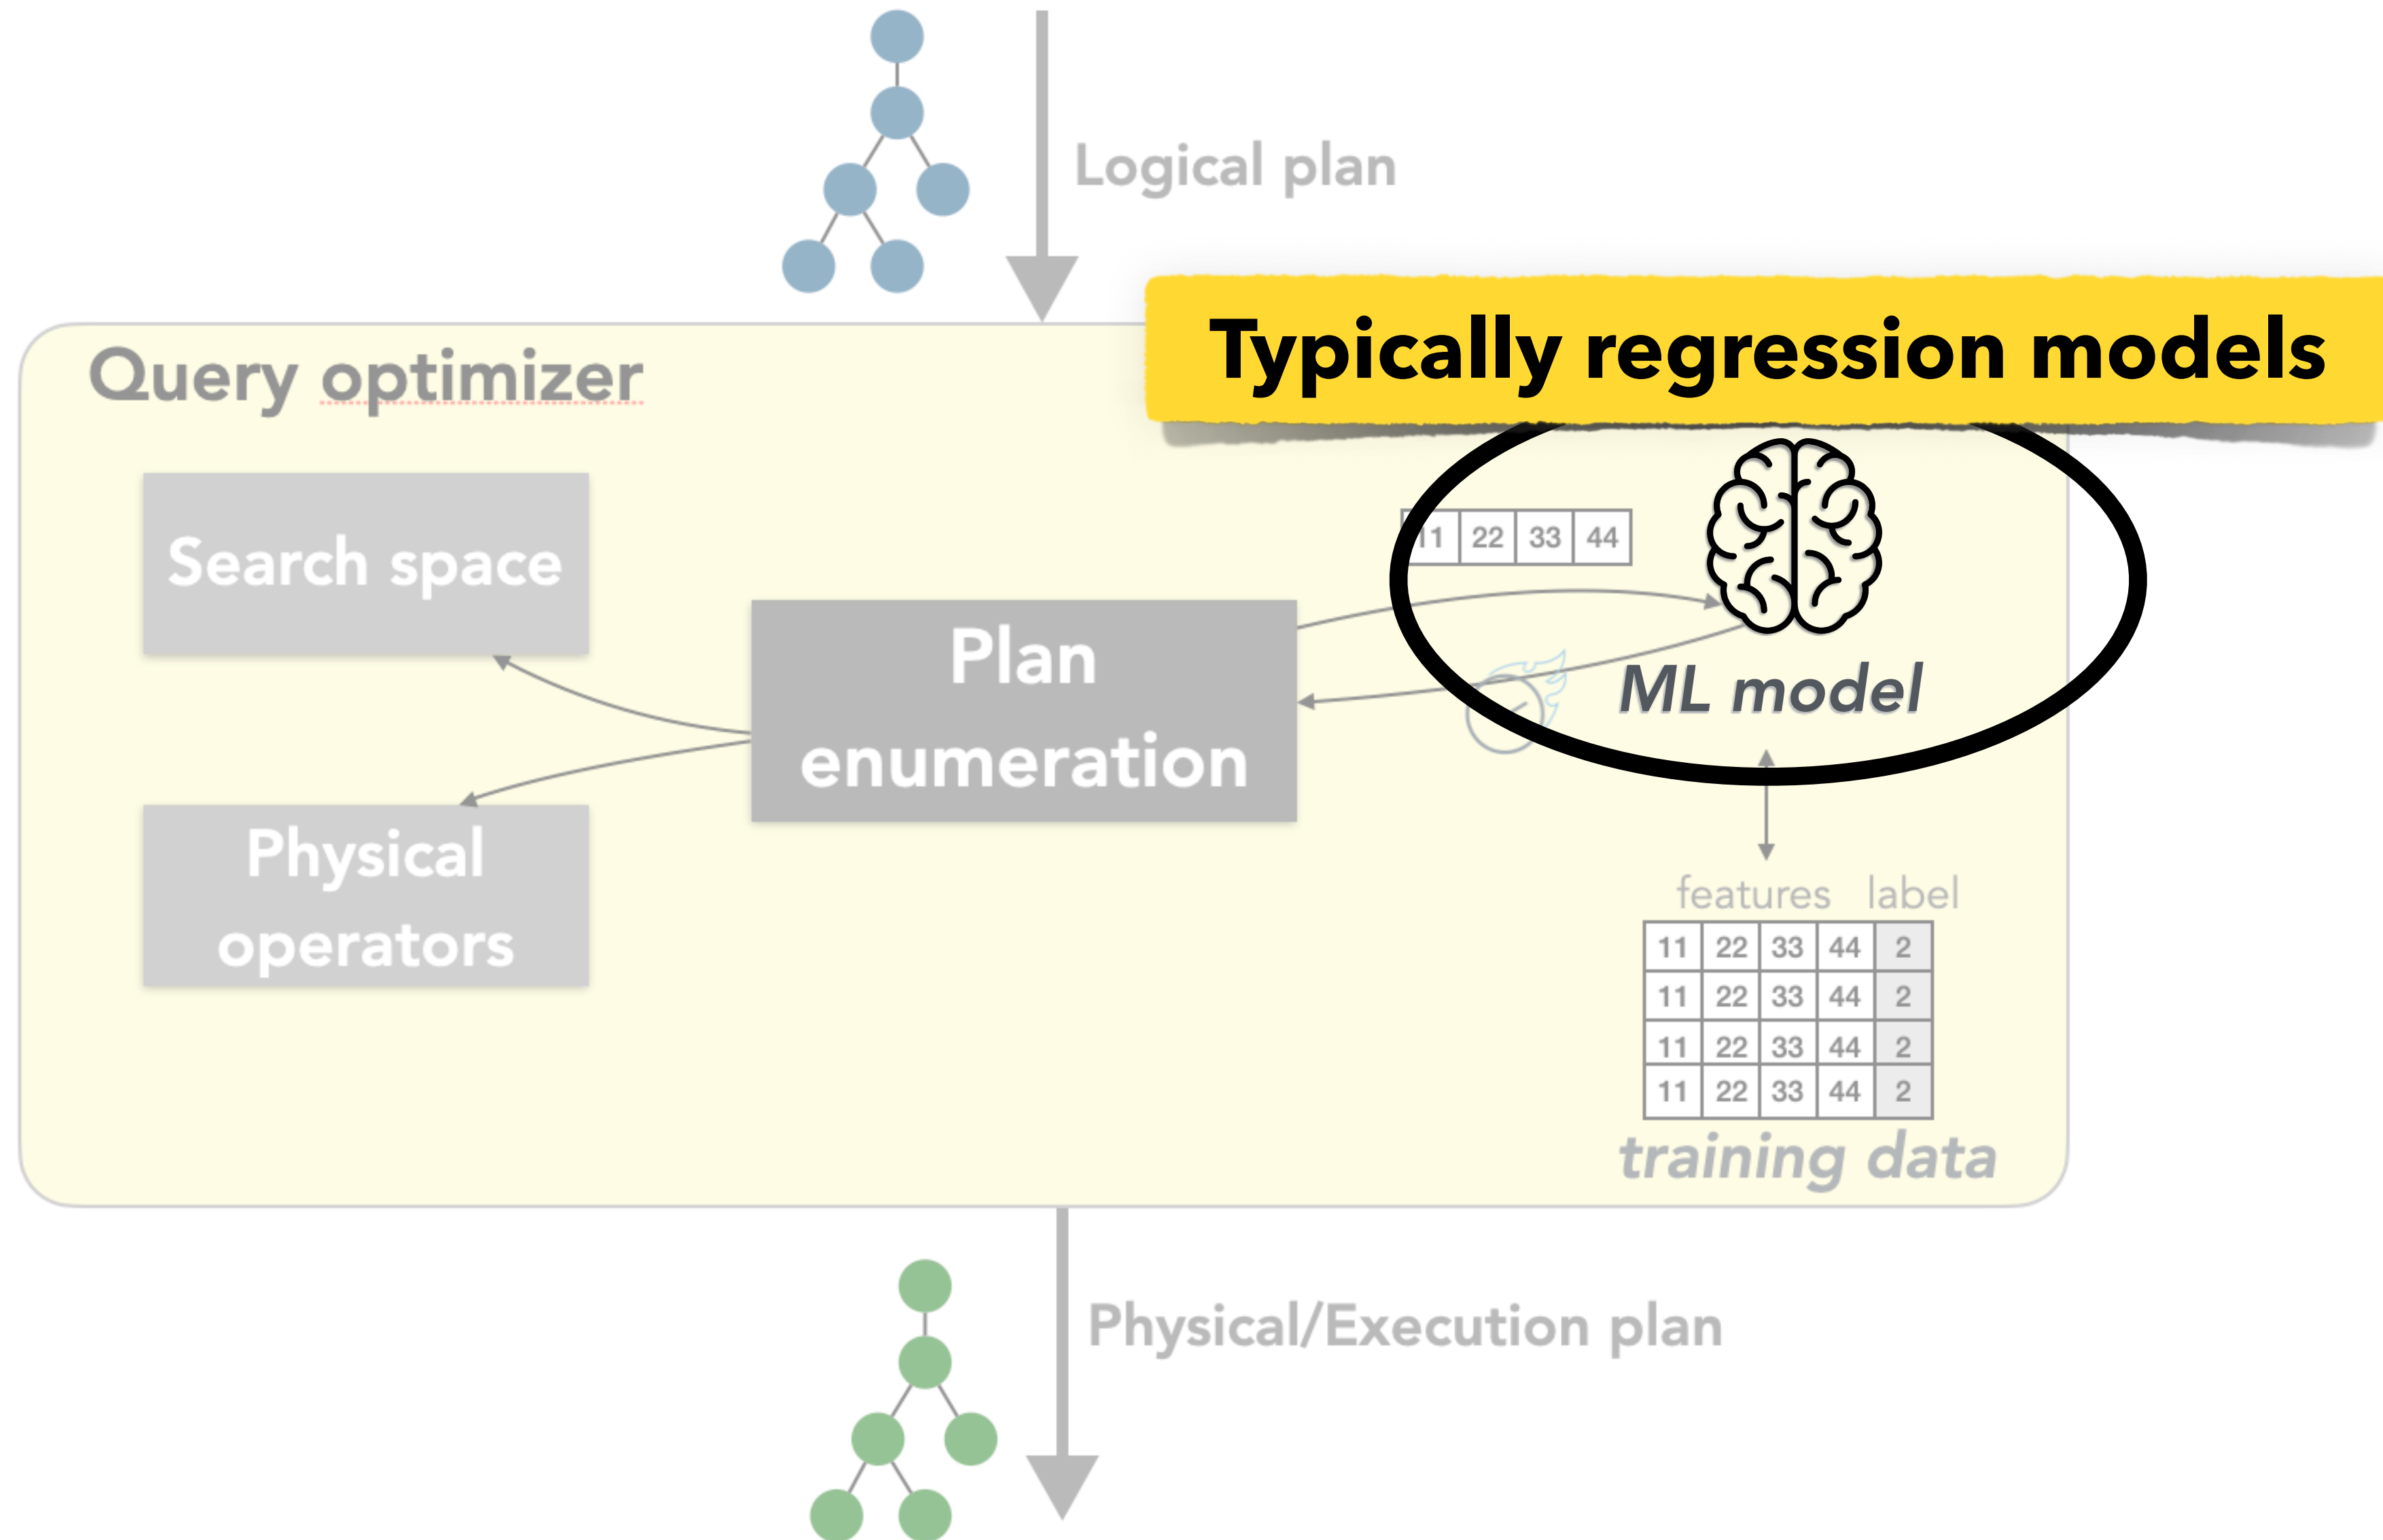


Github



Video demonstration

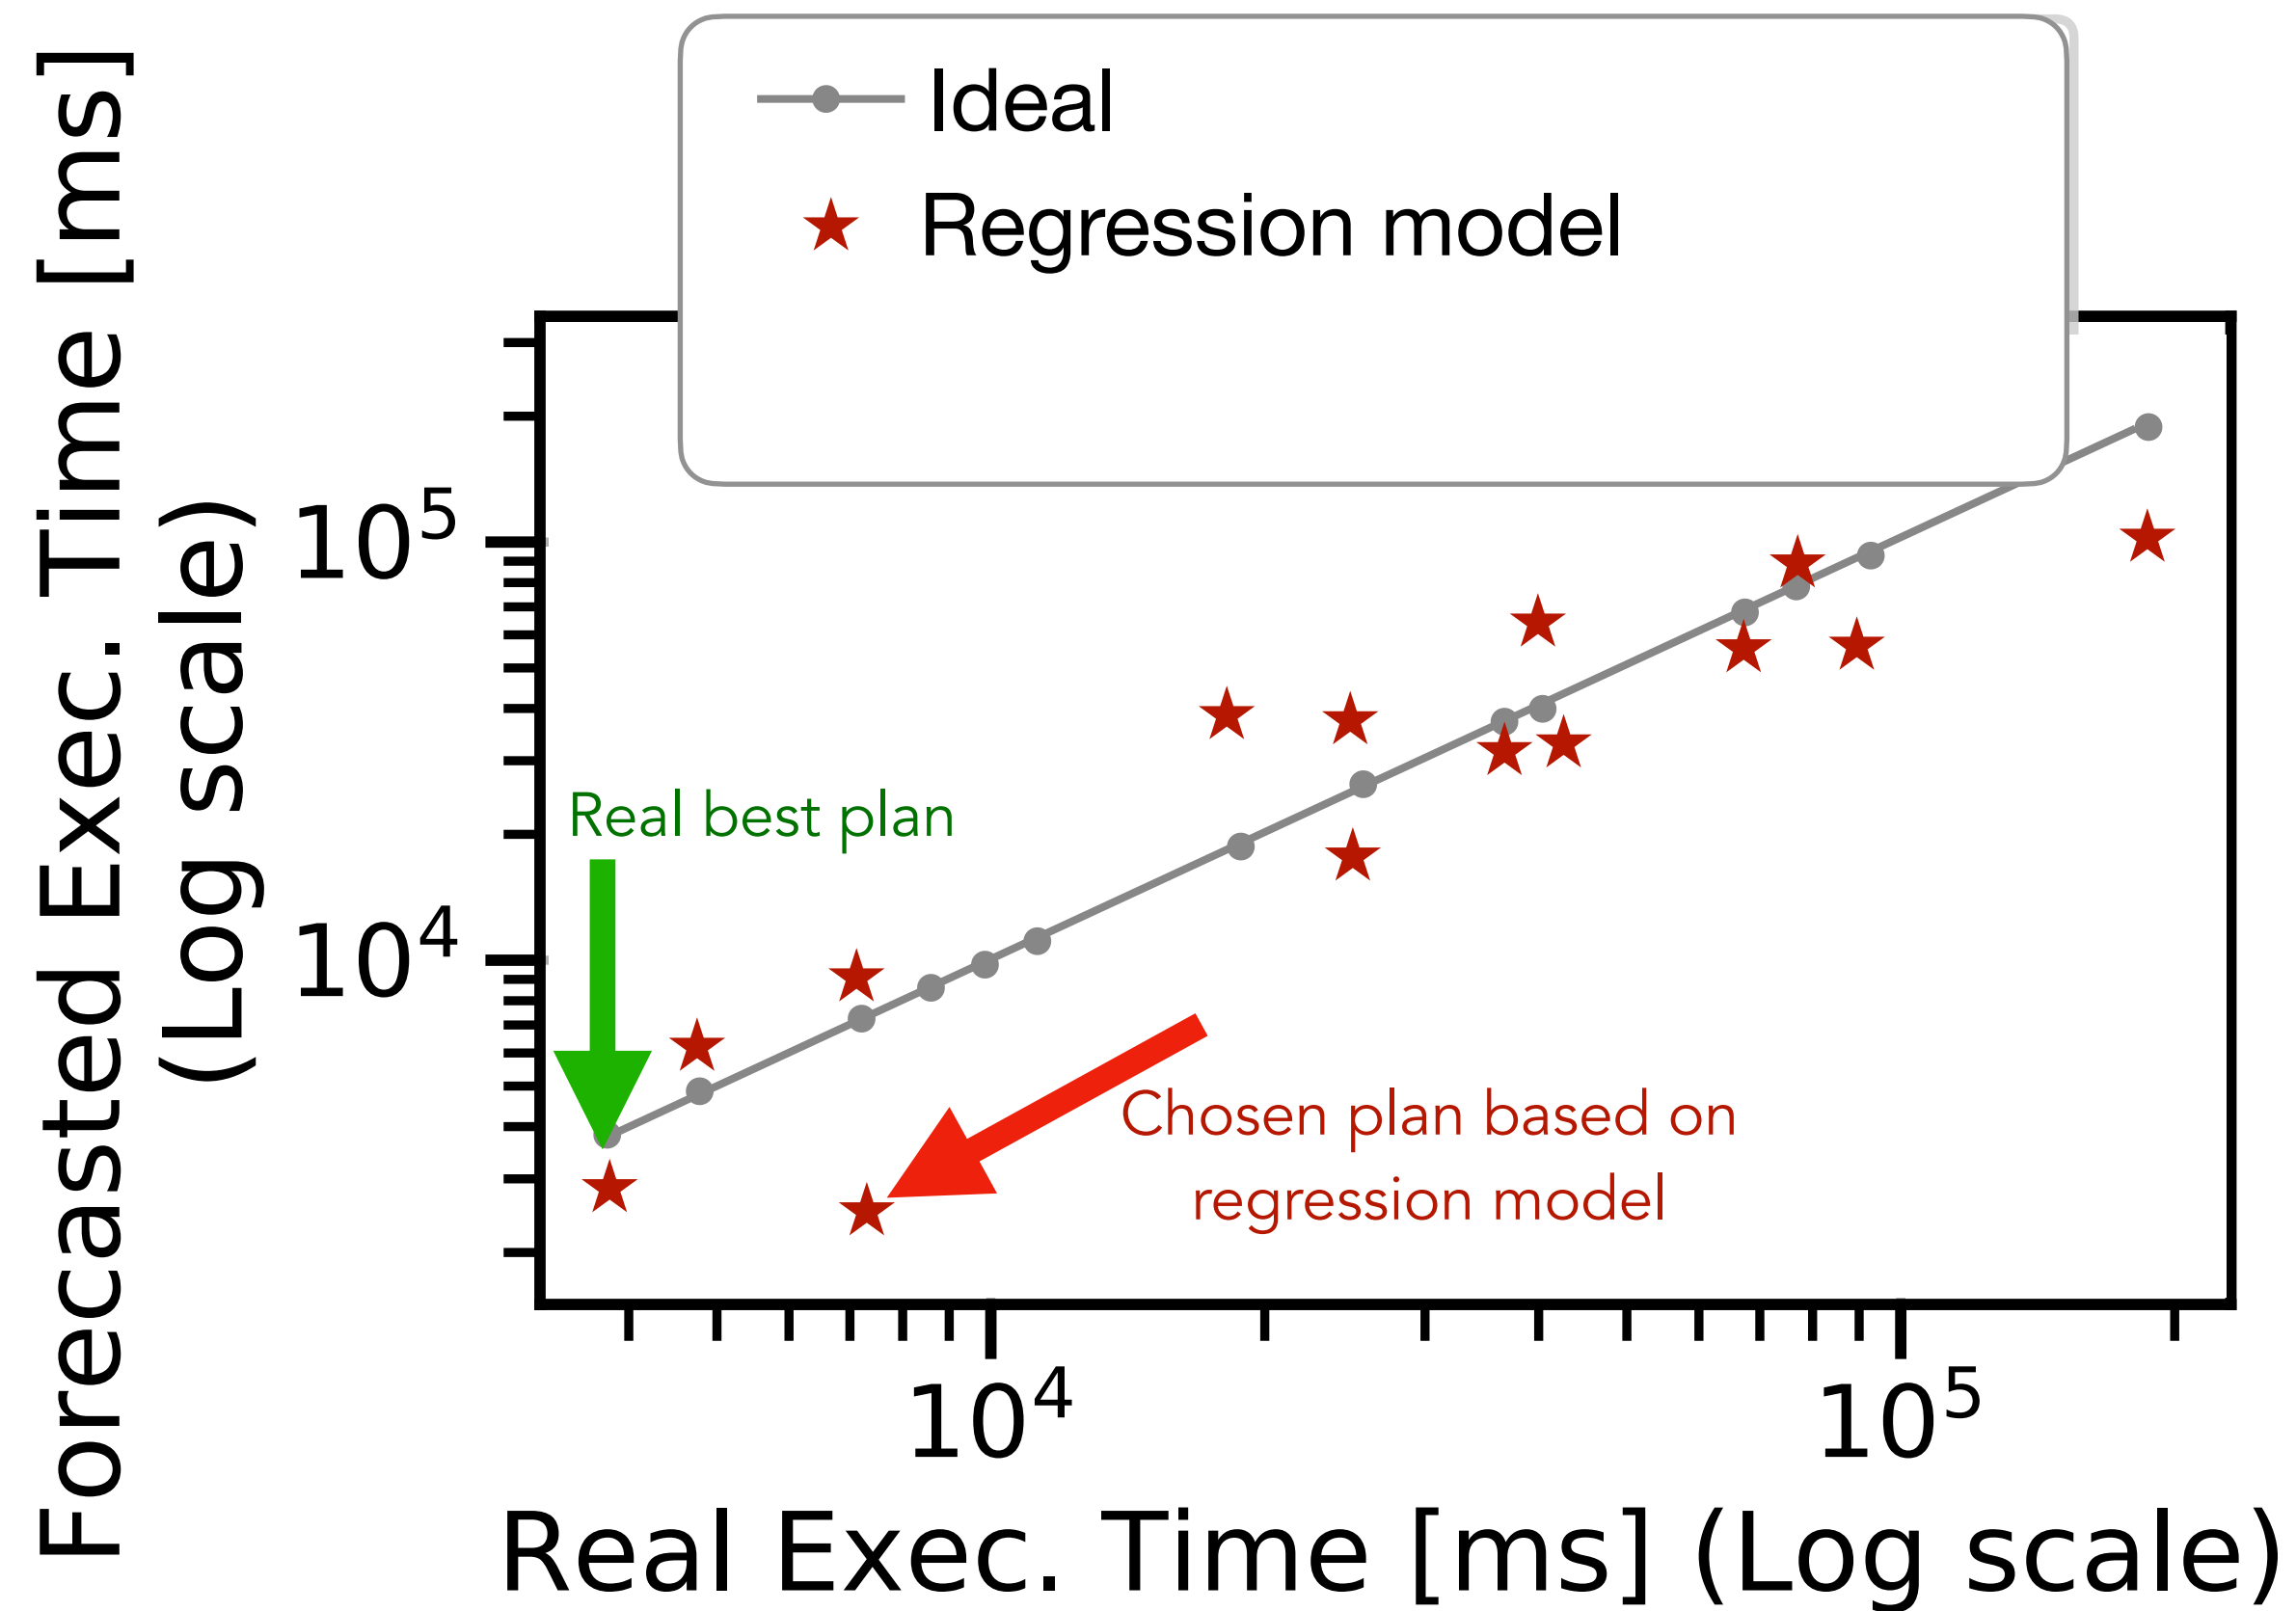
Learning-based query optimization



**Regression: hard to
get completely right**



Effect of regression models errors



~ 2x worse execution time

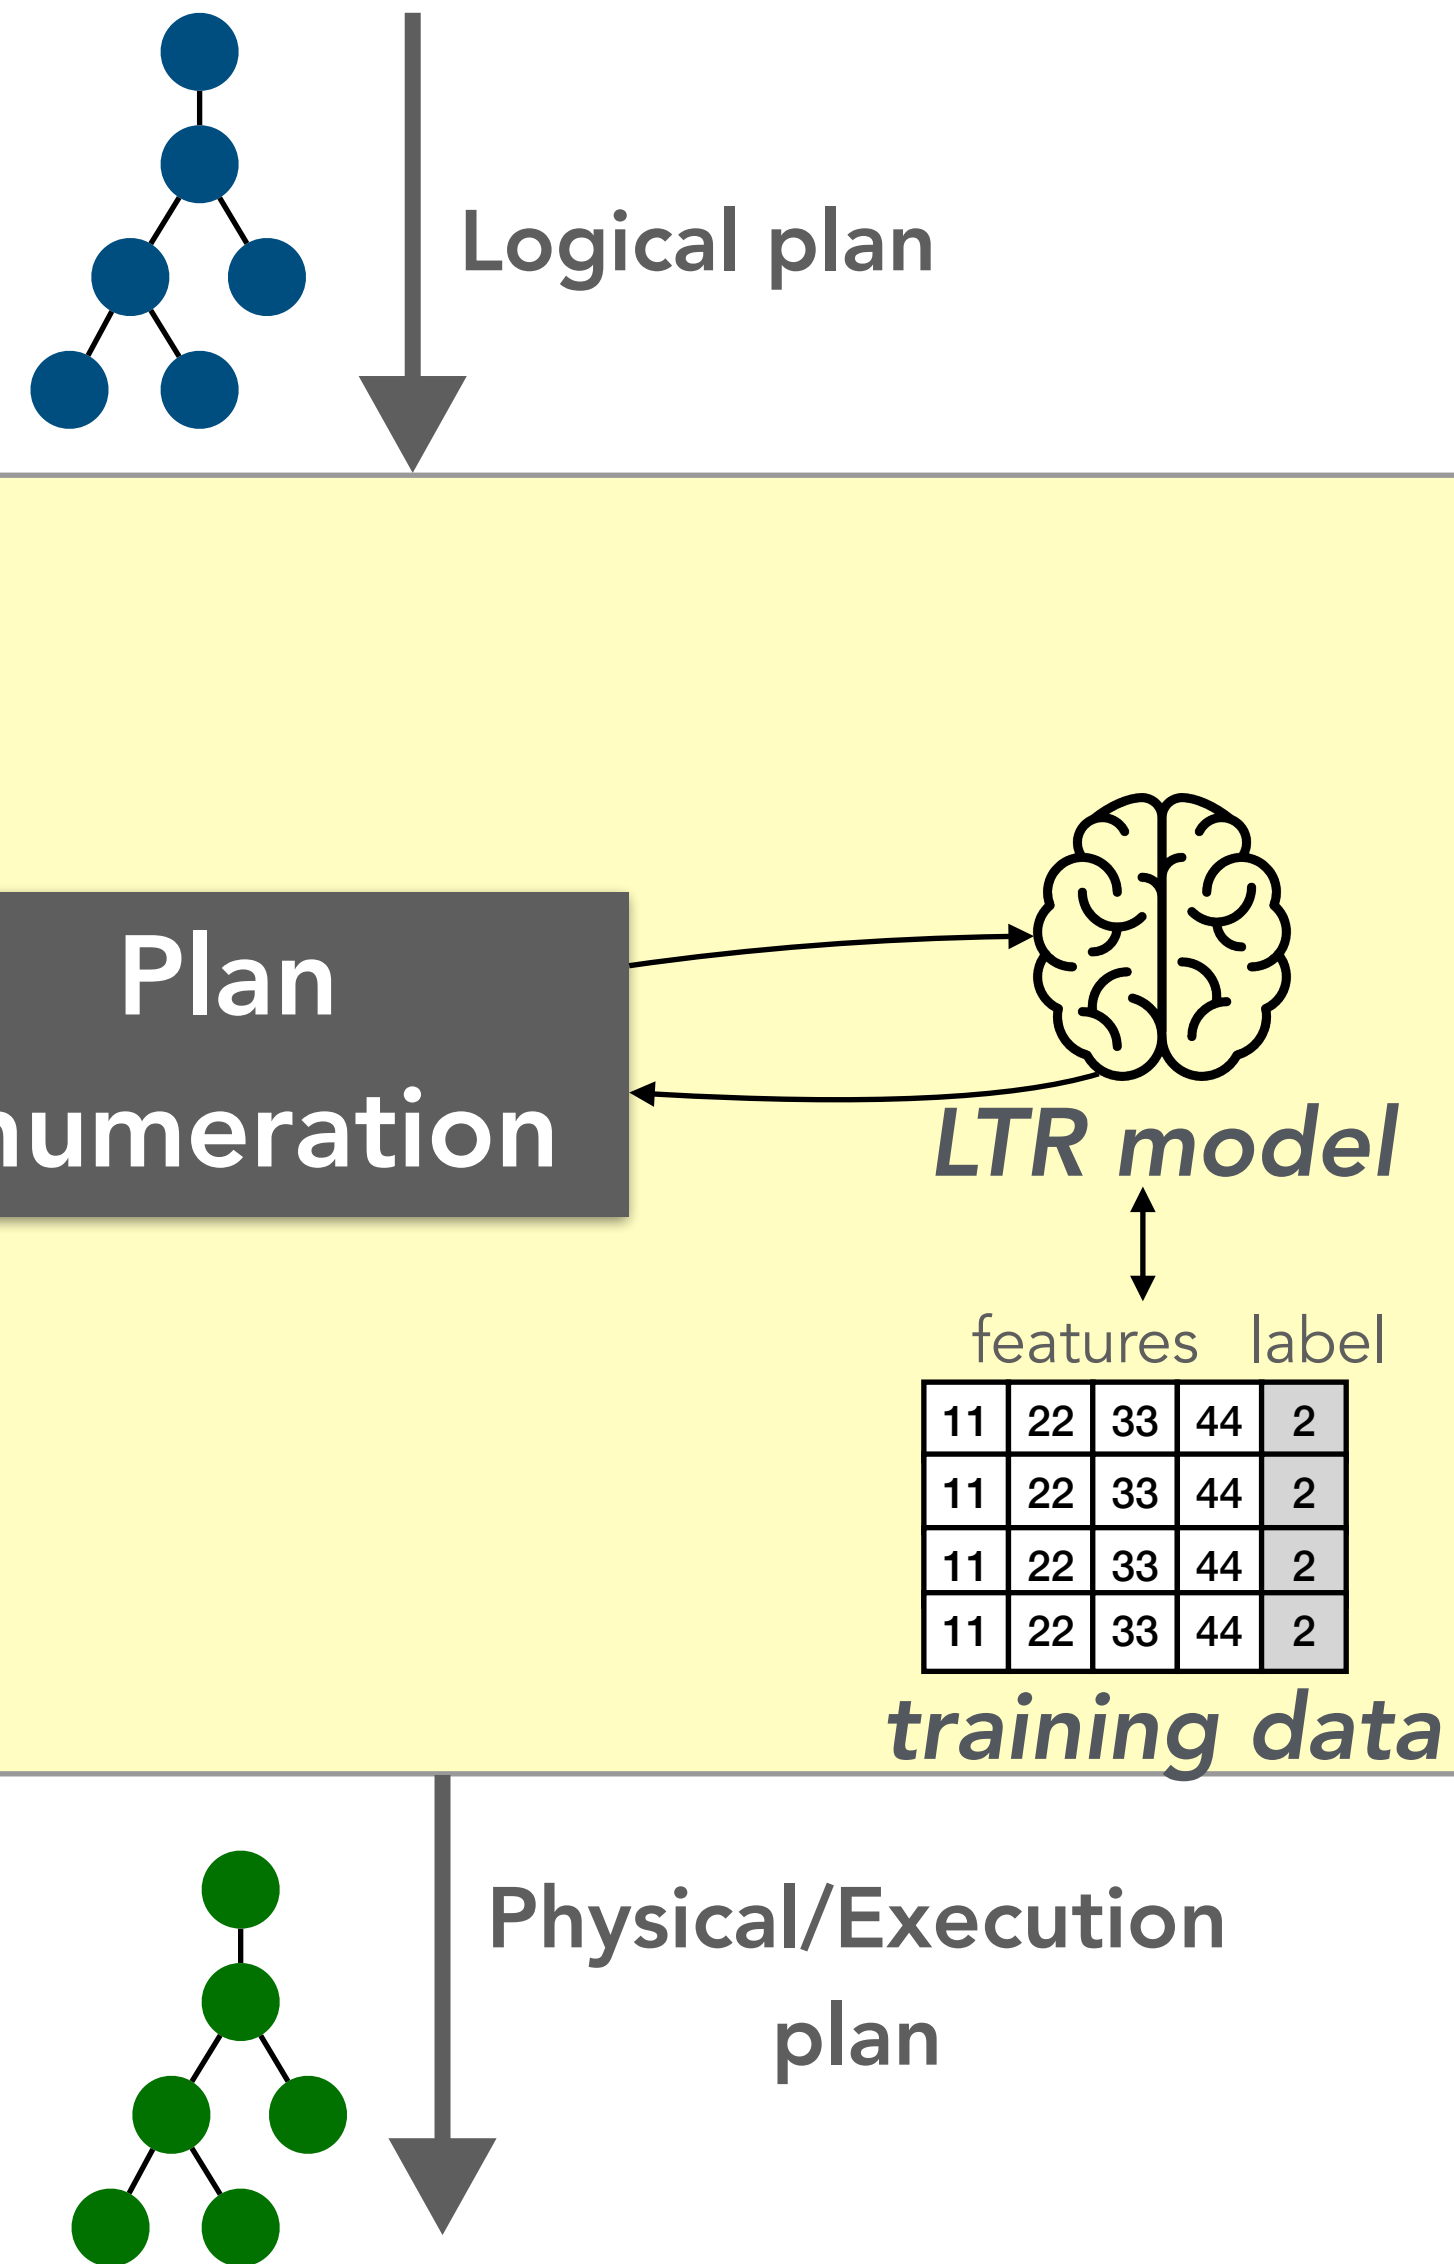
What really matters in QO is the **relative order** of the plans



Can we leverage
learning-to-rank
models?



Learning-to-rank (LTR) in query optimization




◆ Questions

◆ Type of LTR

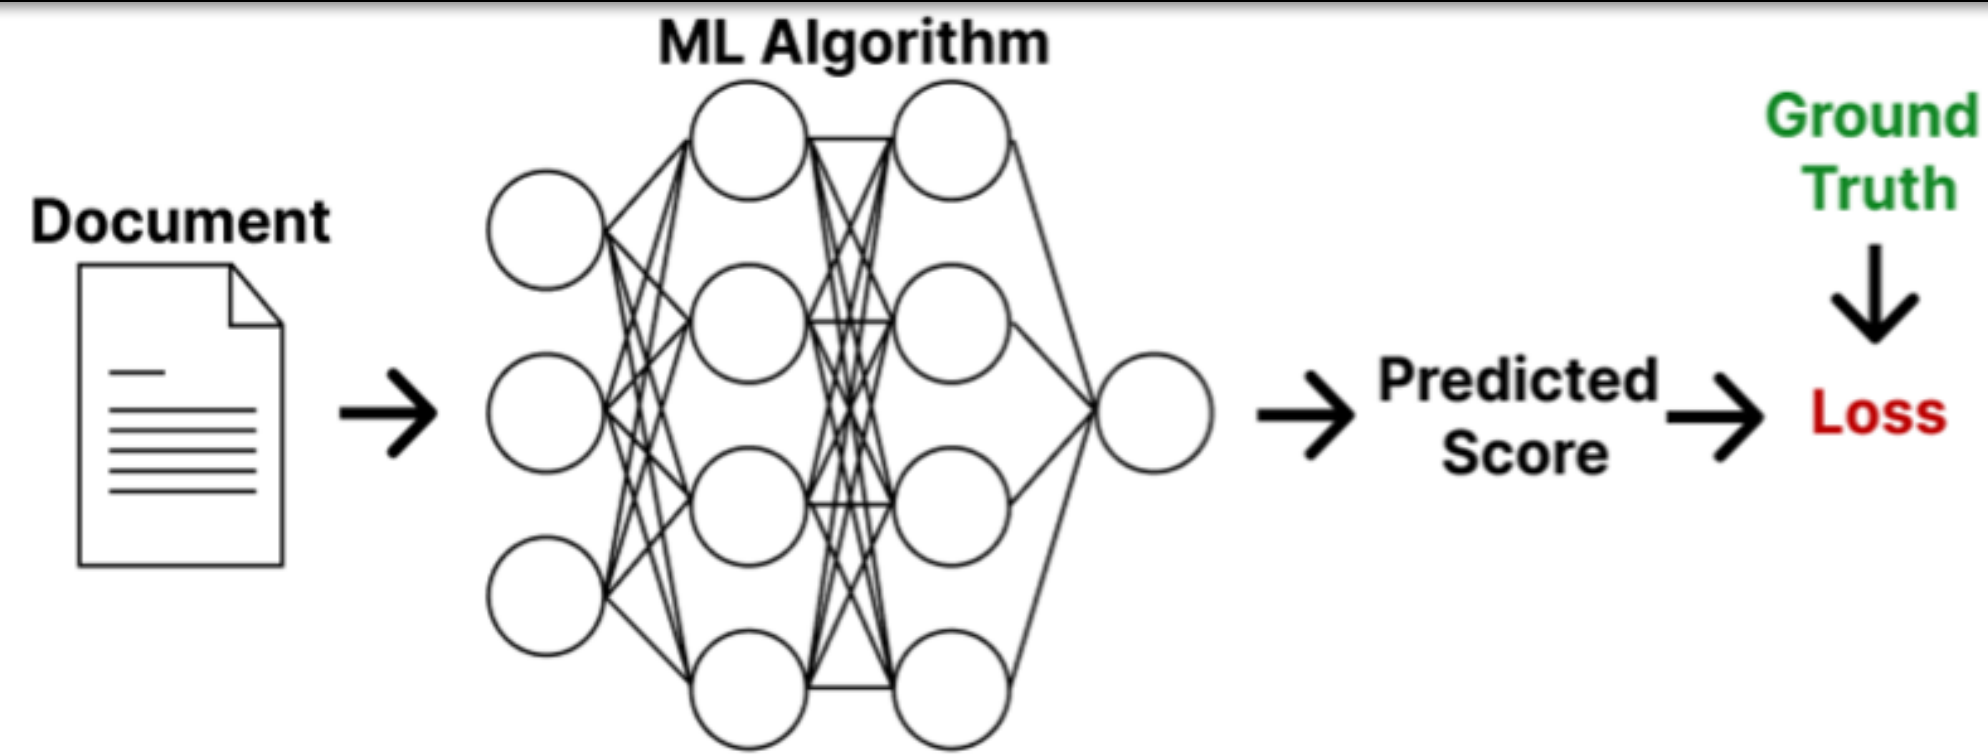
◆ Model architecture and features

◆ Ranking scores from training data




Learning-to-rank (LTR) approach

Pointwise




No comparison among items

Pairwise



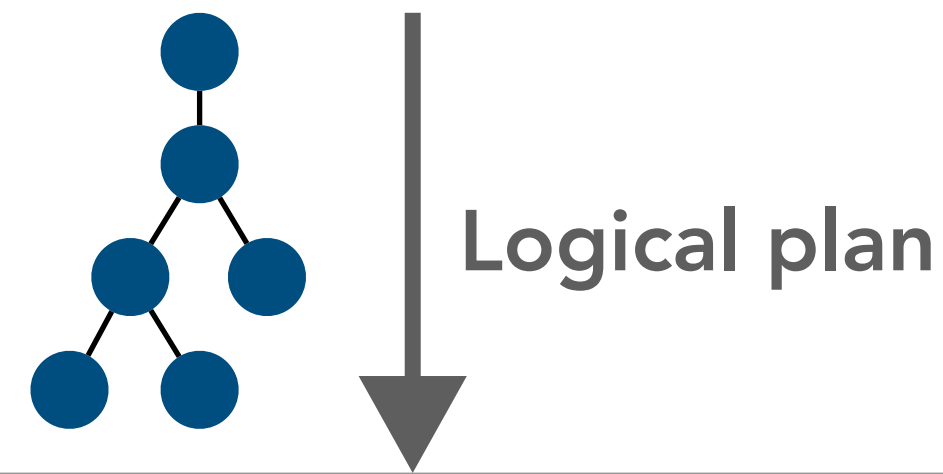
Too expensive
Assumes pairs are i.i.d

Listwise



Emphasises the ranking objective

Learning-to-rank (LTR) in query optimization




◆ Questions

◆ Type of LTR

◆ Model architecture and features

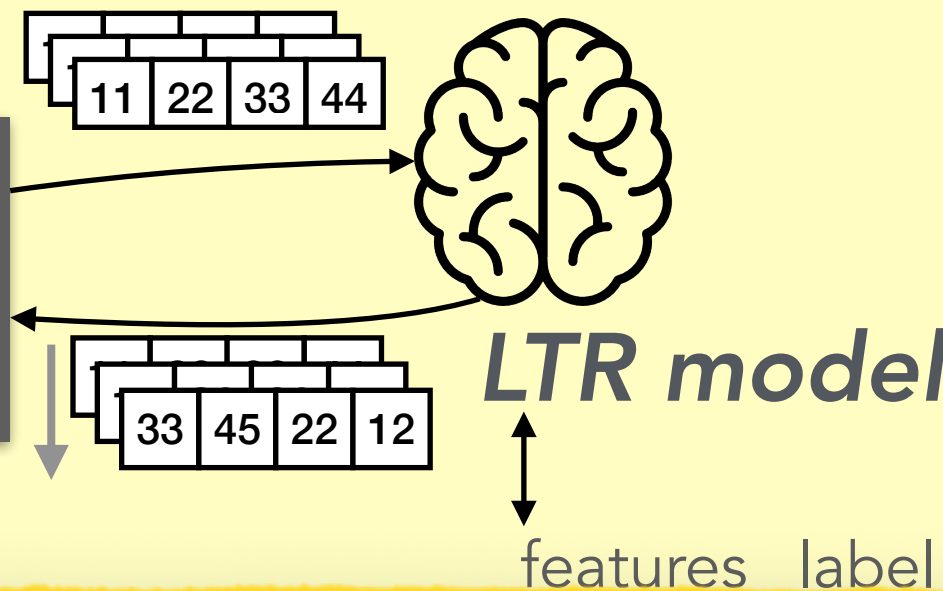
◆ Ranking scores from training data



Query optimizer

Search space

Plan enumeration




11 22 33 44

33 45 22 12

Physical operators

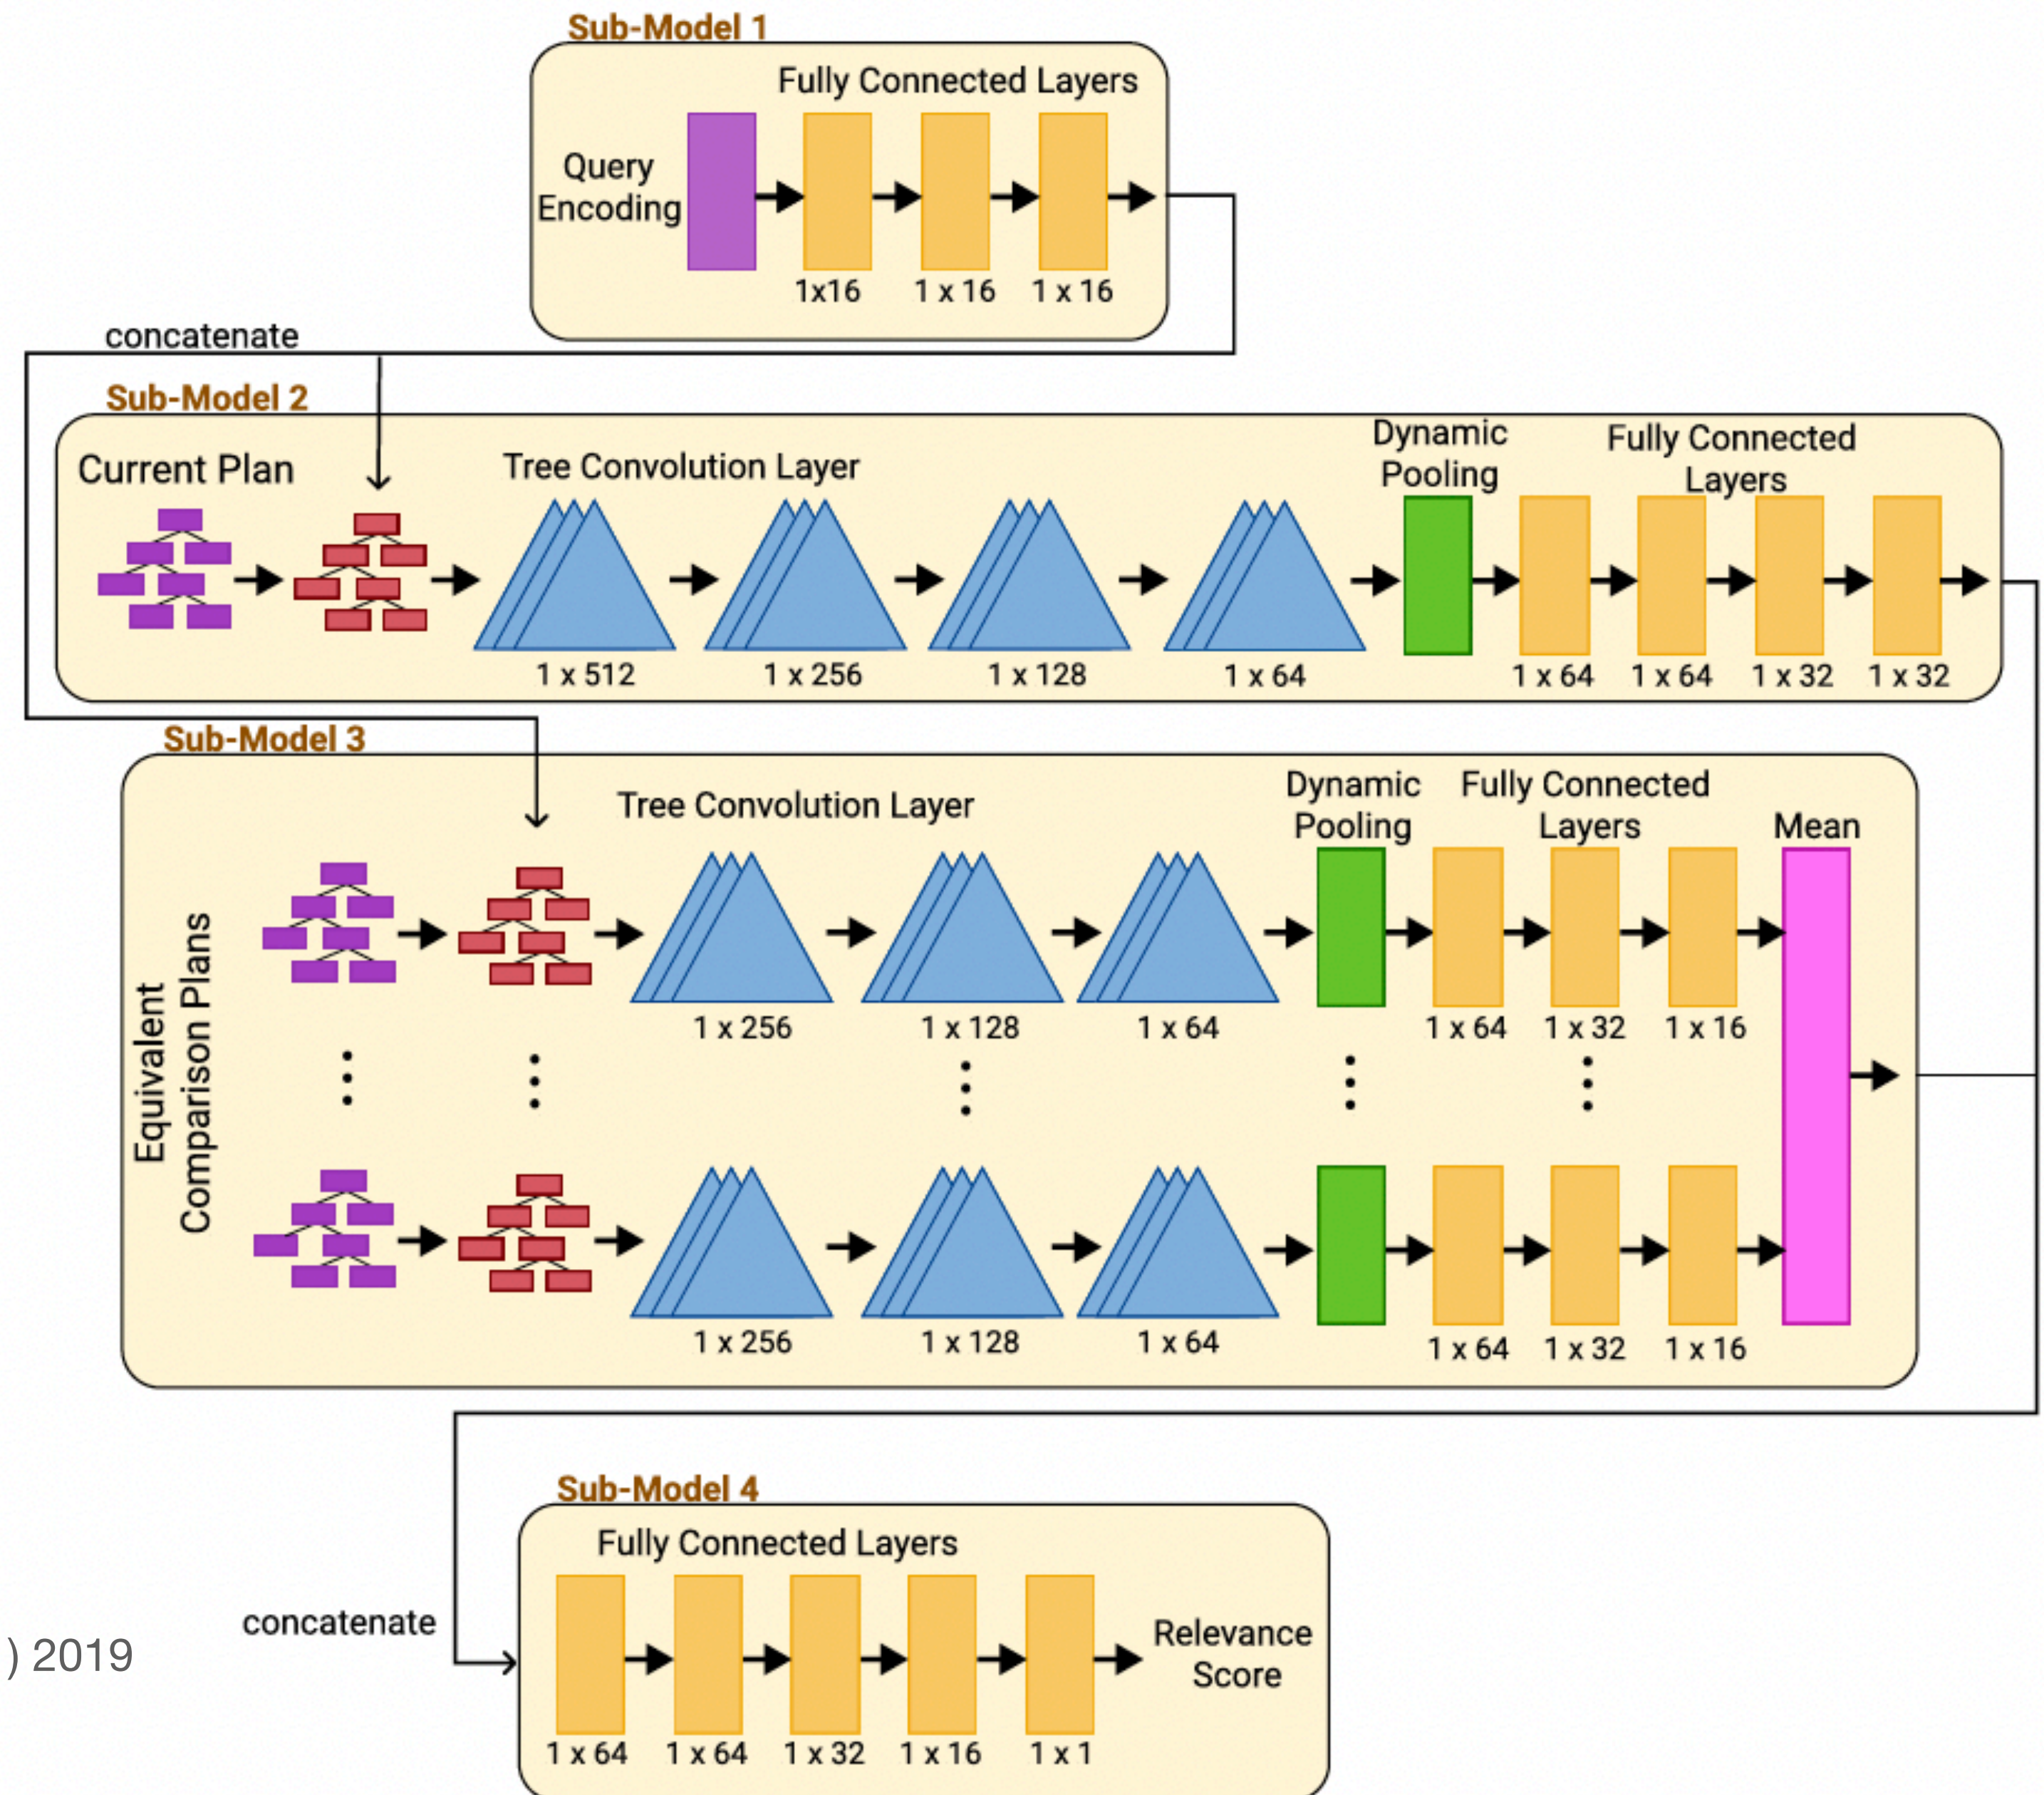
Now plan enumeration needs to consider lists and not pairs



LTR model architecture

Inspired by FATE [1] and Neo [2]

Each plan against the rest equivalent plans




[1] K. Pfannschmidt et al.: Deep architectures for learning context-dependent ranking functions. CoRR abs/1803.05796 (2018)]

[2] R. Markus et al.: Neo: A Learned Query Optimizer. In: PVLDB 12(11) 2019

Featurization

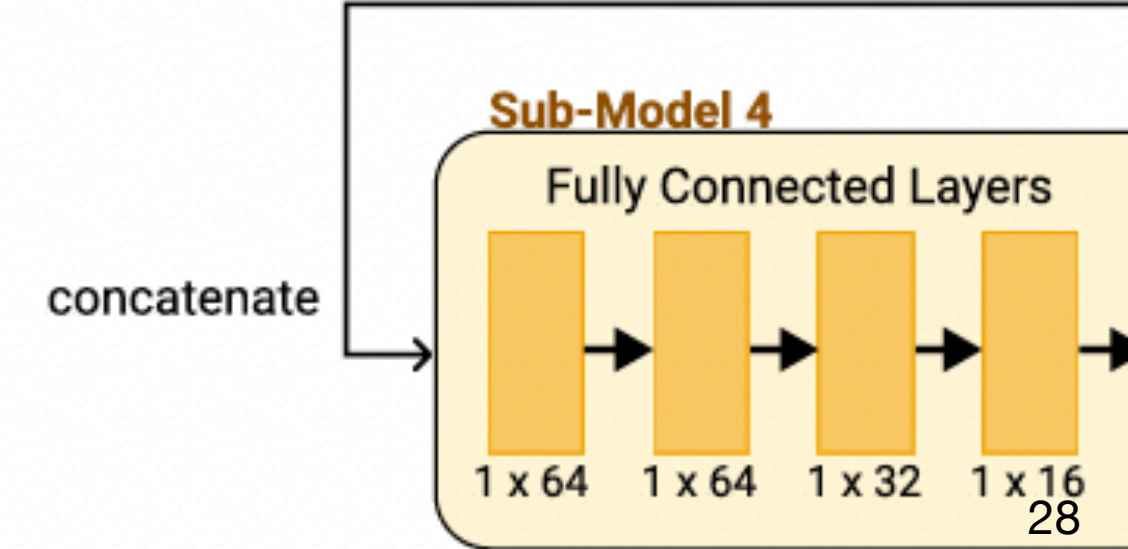
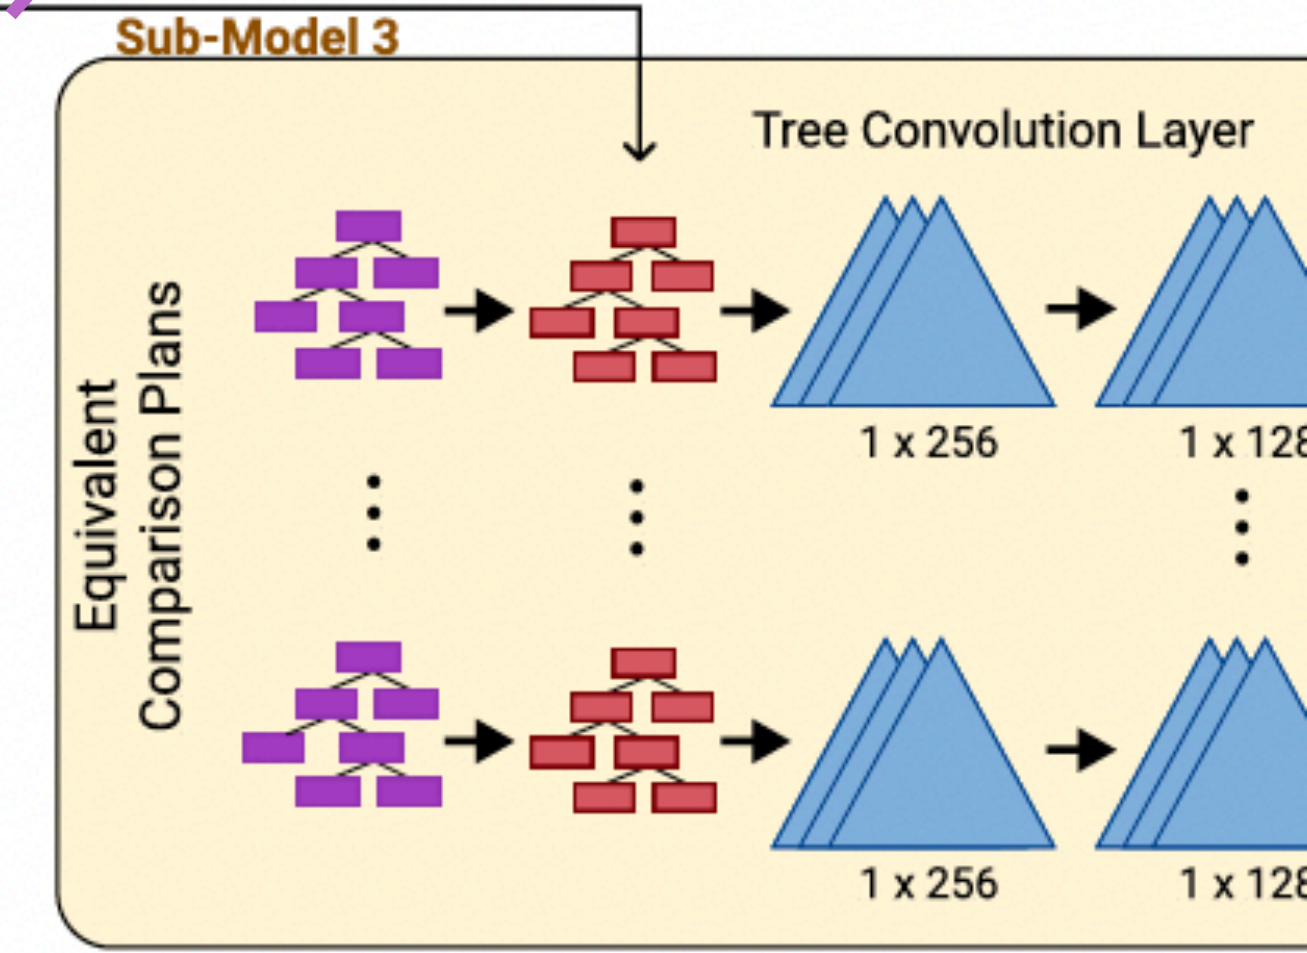

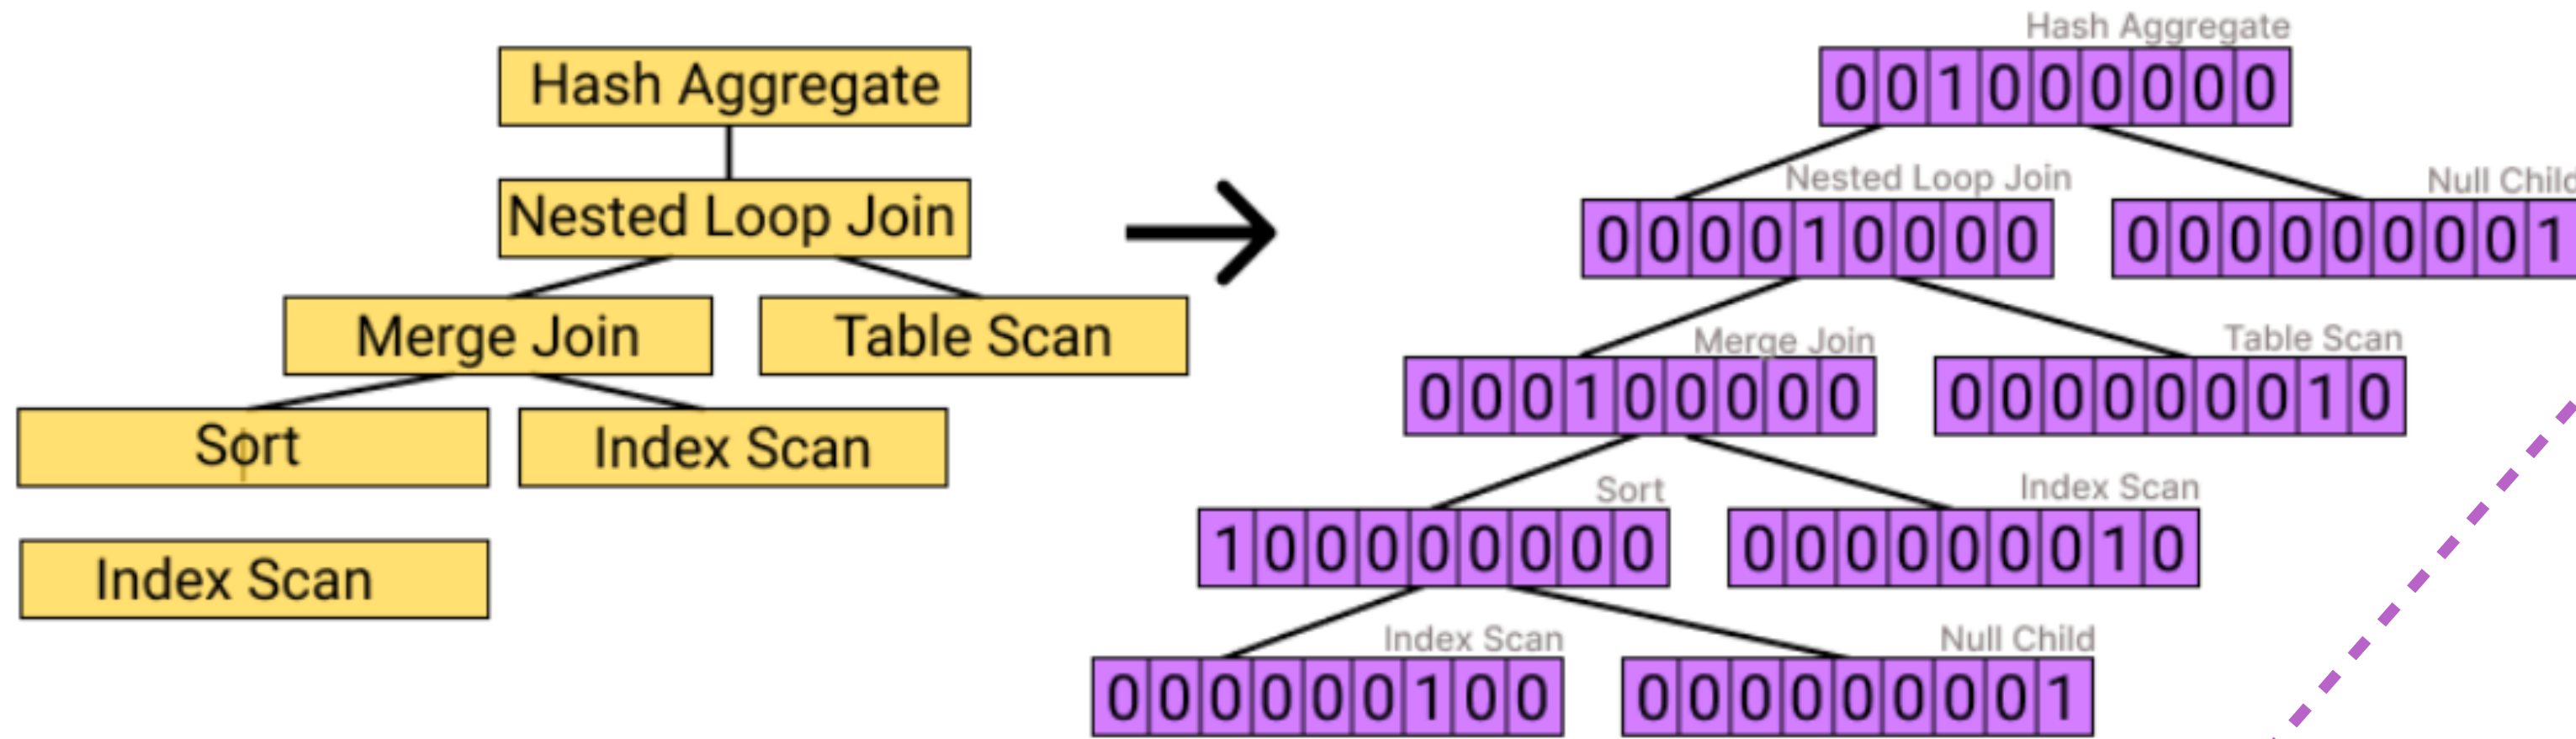
Query encoding (logical)

Has Order By	Has Group By	Number of Joins	Estimated Number of Rows	Max. Relation	Min. Relation
--------------	--------------	-----------------	--------------------------	---------------	---------------



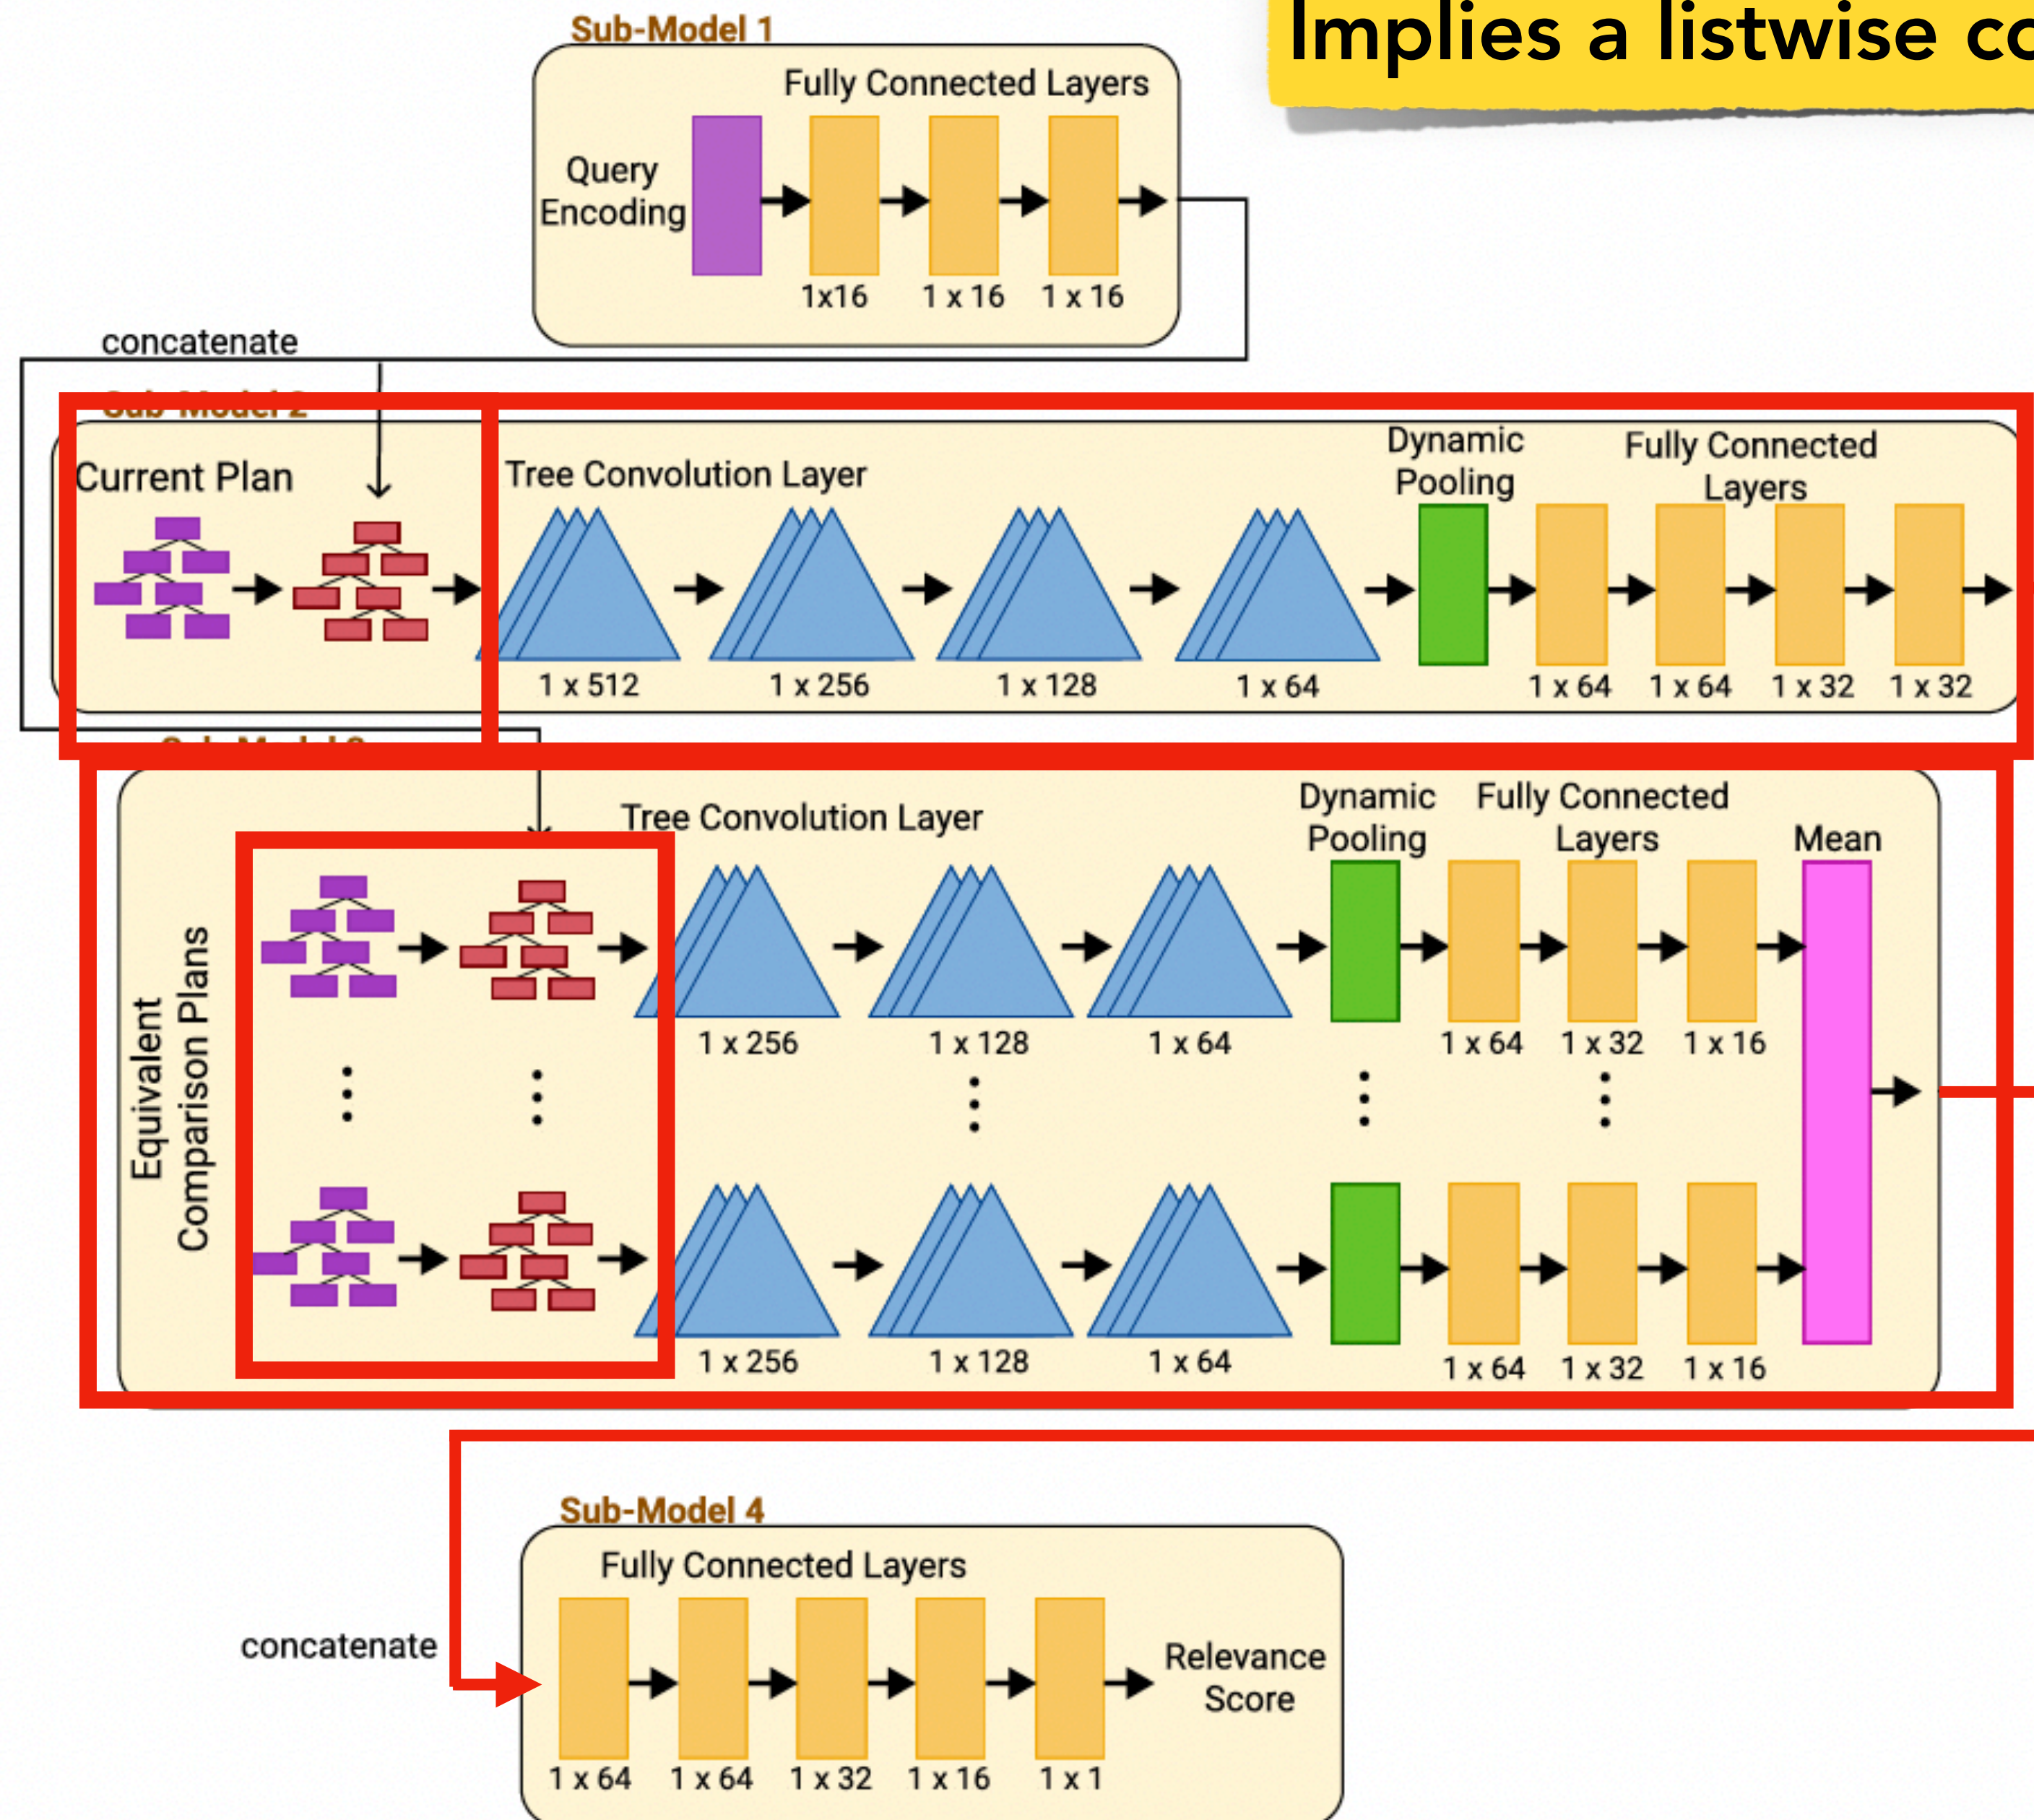
Plan encoding (physical)

Sort	Stream Aggregate	Hash Aggregate	Merge Join	Nested Loop Join	Hash Join	Index Scan	Table Scan	Null	Estimated Rows
------	------------------	----------------	------------	------------------	-----------	------------	------------	------	----------------




LTR model architecture

Implies a listwise comparison of plans



Learning-to-rank (LTR) in query optimization



Query optimizer


Search space

Plan enumeration

Physical operators

11 22 33 44

33 45 22 12




LTR model

features label

11	22	33	44	2
11	22	33	44	2
11	22	33	44	2
11	22	33	44	2

training data




Physical/Execution plan

Questions

◆ Type of LTR

◆ Model architecture and features

◆ Ranking scores from training data



Scoring training plans

- ◆ Option 1: Use execution time, sort, extract rank
- ◆ Information loss on how two plans differ



- ◆ Option 2: Devise a **score function**: given the runtime of a plan and a maximum score value, output its score w.r.t. a list of plans and their runtimes

- ◆ **Local linear** score function

- ◆ **Global** agglomerative function




Execution time

	features	label		
11	22	33	44	78
11	22	33	44	23
11	22	33	44	99
11	22	33	44	7

training data

Preliminary Results



Experimental setup

◆ Datasets: TPC-H and JOB

◆ Queries:

Query types	1 Join	2 Joins	3 Joins	4 Joins	5 Joins	6 Joins	7 Joins
Short (< 2s)	8	8	8	8	8	8	8
Medium ($\geq 2s$ & $< 30s$)	2	8	8	8	8	8	8
Long ($\geq 30s$)	0	2	4	7	8	8	5

◆ ~25K execution plans produced by DataFarm [1]

◆ Baselines:

◆ DB X (cost-based optimizer)

◆ Neo [2]

◆ Bao [3]

◆ baseline model (pairwise LTR)

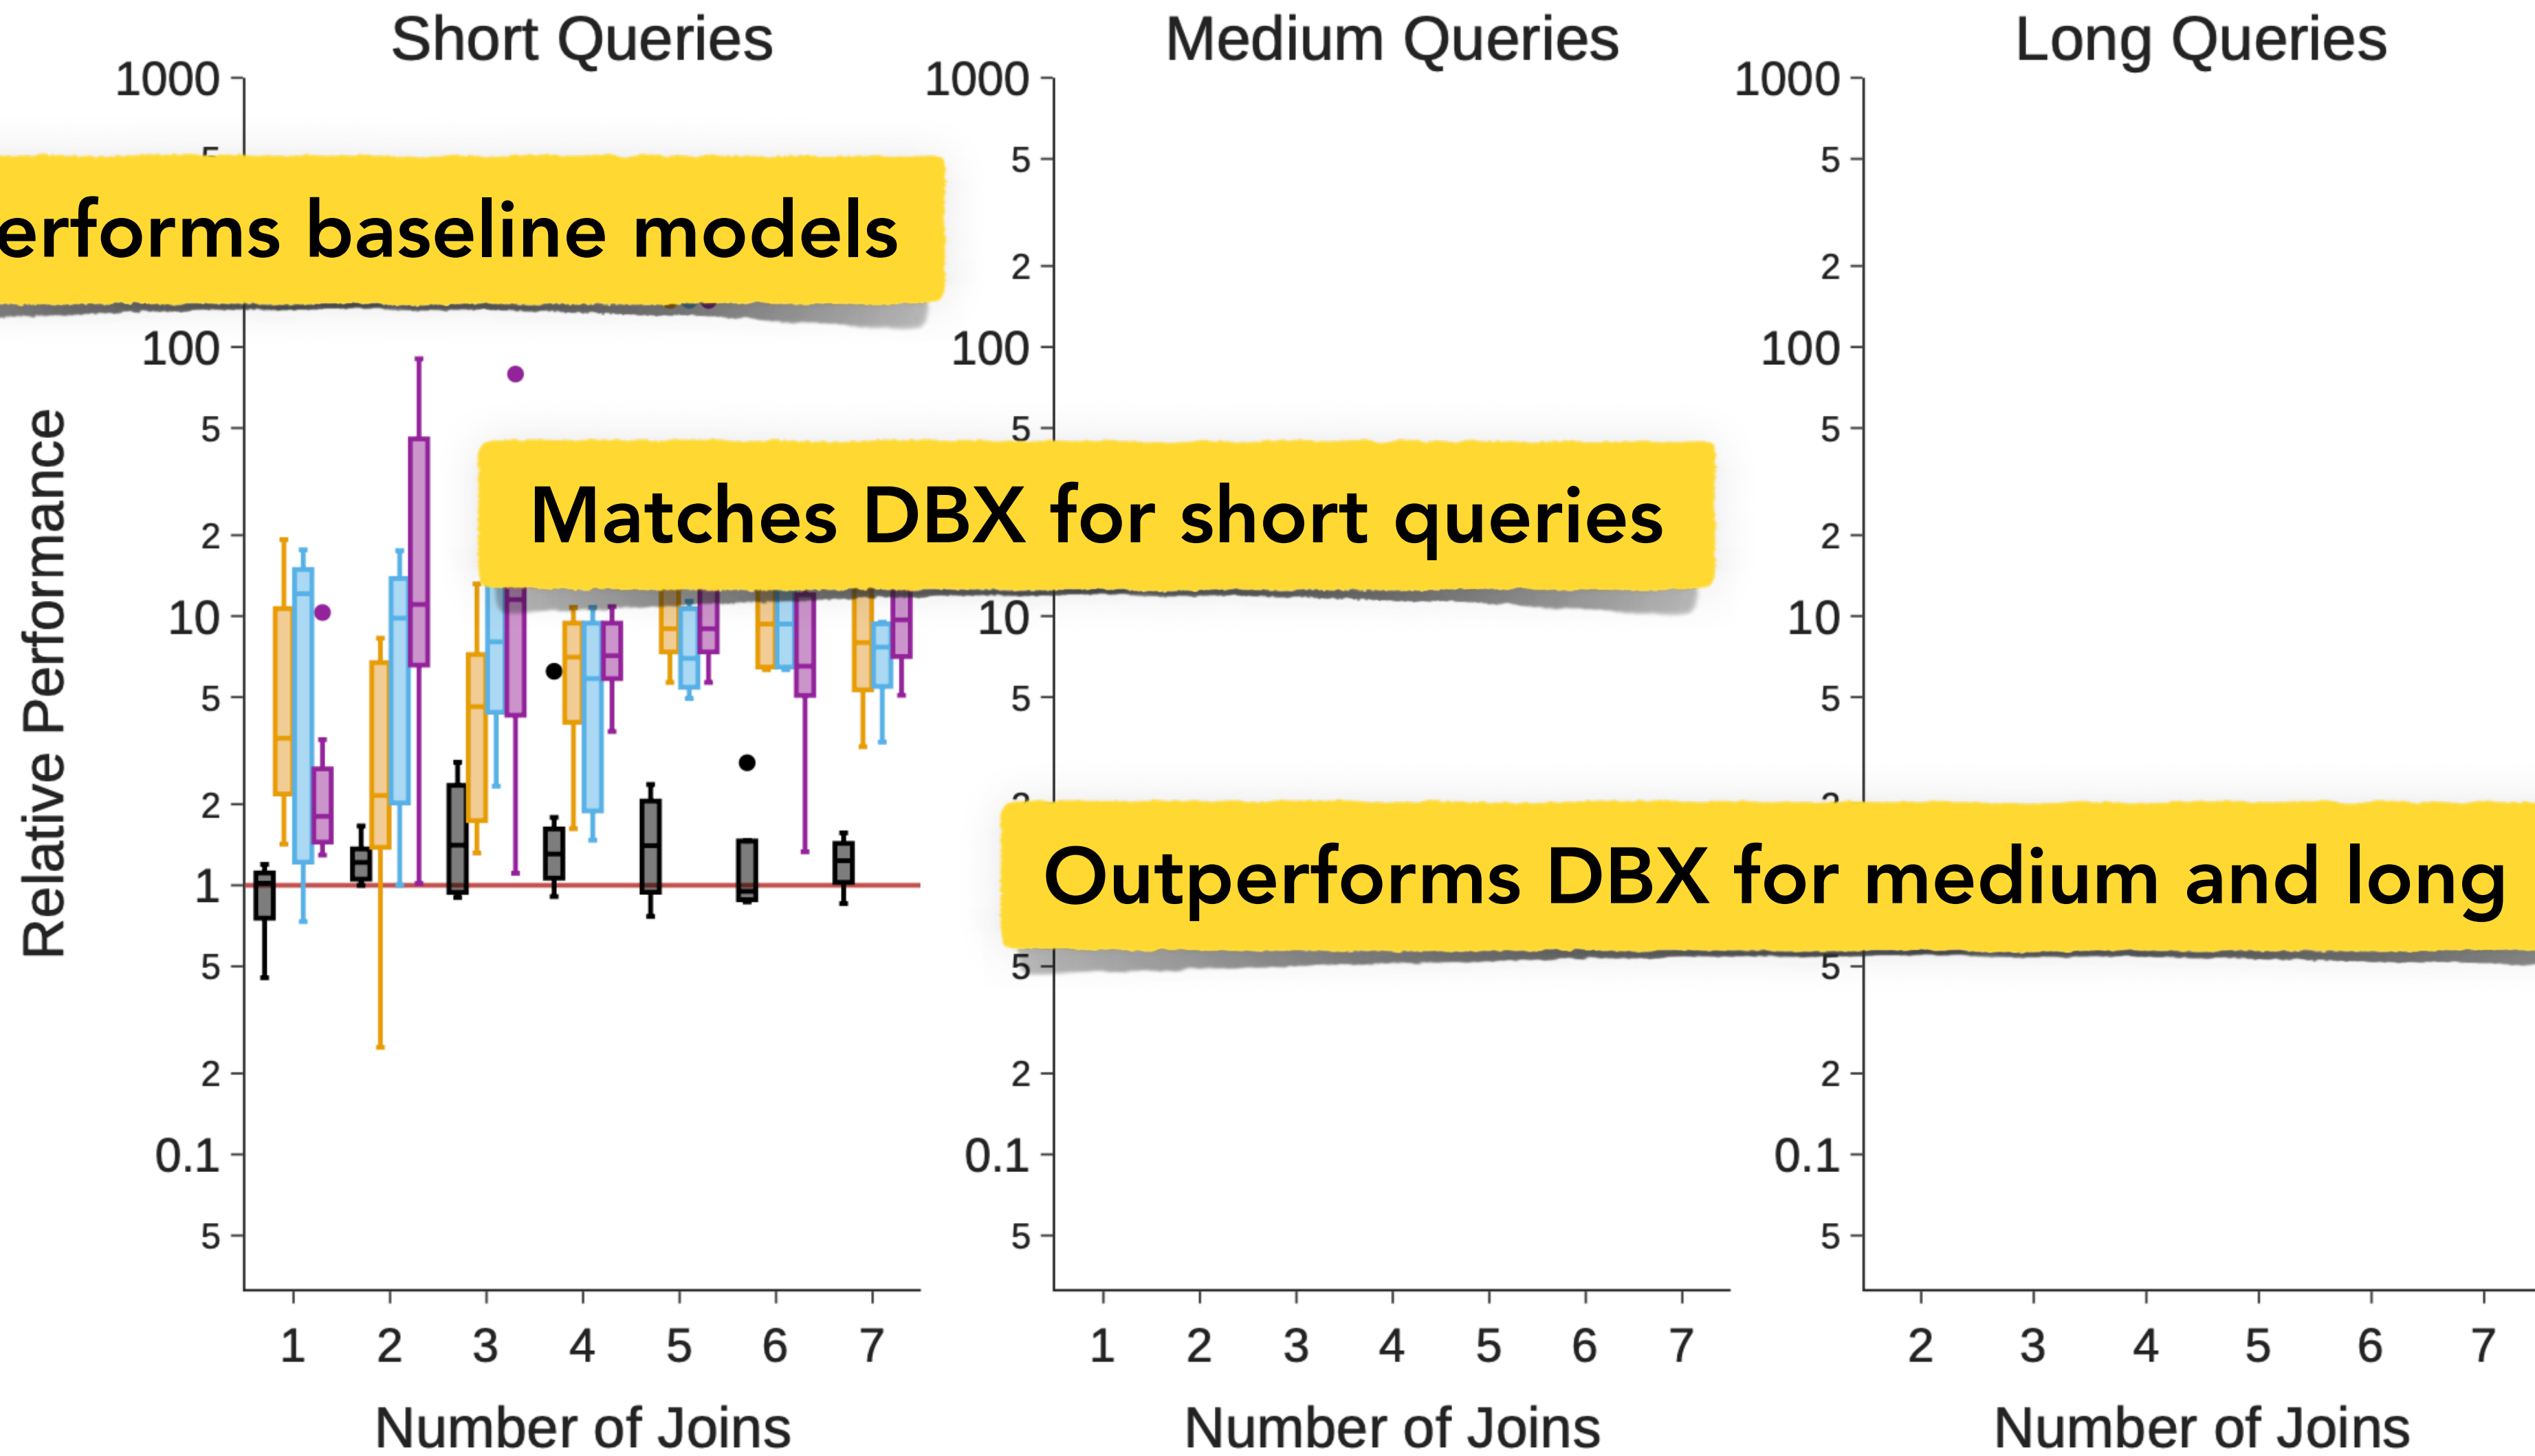
[1] F. Ventura et al: Expand your Training Limits! Generating Training Data for ML-based Data Management. SIGMOD 2021

[2] R. Markus et al.: Neo: A Learned Query Optimizer. PVLDB 12(11) 2019

[3] R. Markus et al.: Bao: Making Learned Query Optimization Practical. SIGMOD 2021

Performance on TPC-H (1 GB)

■ LTR model ■ baseline model ■ Bao ■ Neo — DB X Baseline



Outperforms baseline models

Matches DBX for short queries

Outperforms DBX for medium and long queries

Learning-based Query Optimization

What are we still missing?

Training data generation



Learning-to-rank methods



The logo for GUI DataFarm, featuring a green plant growing out of a database cylinder. To the right is a QR code. The text "GUI DataFarm" is written in a large, bold, sans-serif font.

Github

