Carnegie Mellon University

Lecture #12

Special Topics:

Self-Driving Database Management Systems

Workload Modeling II

@Kai_Franz // 15-799 // Spring 2022

PREVIOUSLY

- Index recommendation, knob tuning systems
  - Take in some information about workload as input
- Encoding of query plans that captures essential features?

TODAY'S AGENDA

- Query Encoder Model
- Downstream Tasks
- Analysis

CHALLENGES

- Query independence
  - Queries vary within a workload, context does not provide information
- Diverse query structure
  - Representing a query plan tree is non-trivial
- Modeling computational complexity
  - Each operator has its own demand for resources
- Data dependence
  - Need information about index availability, data distribution
- Domain adaptation
  - Adapting to unseen workloads is a challenge

ENCODER

- Structure encoder: autoencoder
- Performance encoder: supervised learning

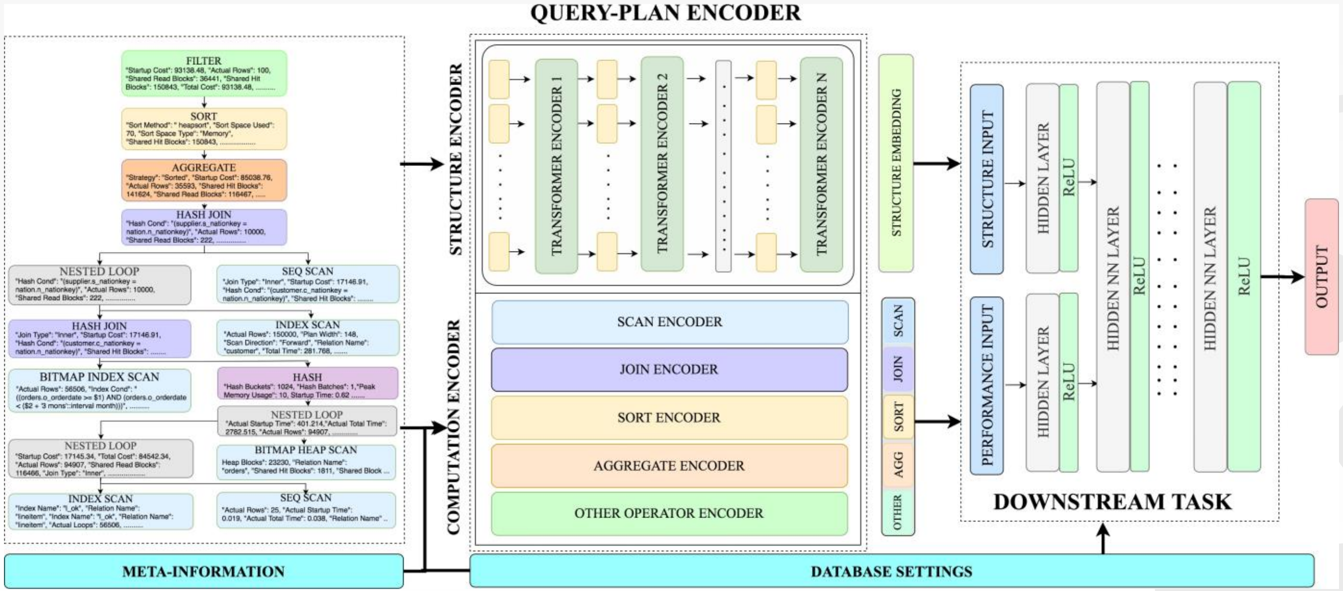

STRUCTURE ENCODER

- First, query plan is flattened via DFS Bracket traversal
- Each node is represented as a triple of subtypes (e.g., Scan-Heap-Bitmap, Join-Merge-Left)

(Filter–, (Sort–, (Aggregate–, (Join-Hash-, (Loop–Nested, (Join-Hash-, (Hash–, (Loop–Nested, (Loop–Nested, Scan-Index-, Scan-Seq-) Scan-Heap-Bitmap)) Scan-Index-Bitmap) Scan-Index-) Scan-Seq-))))

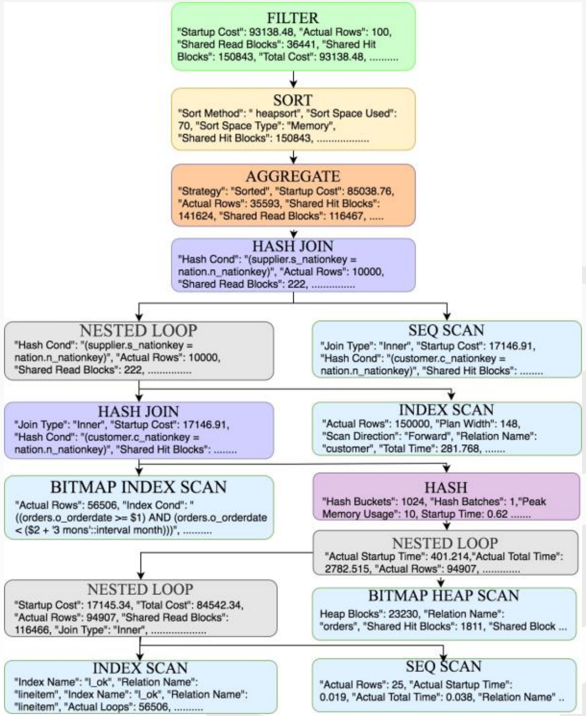

DATABASE WORKLOAD CHARACTERIZATION WITH QUERY PLAN ENCODERS

STRUCTURE ENCODER

- Trained using plan-pair similarity regression
  - Smatch score: degree of overlap between two graphs

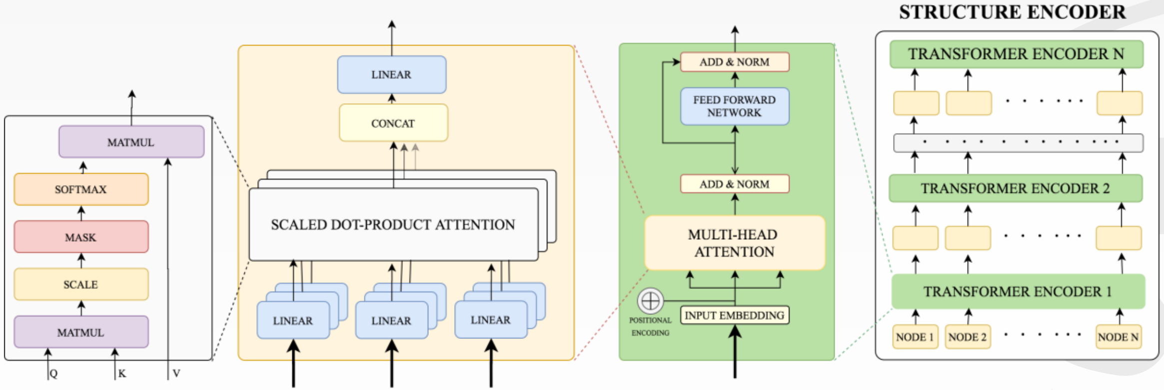

PERFORMANCE ENCODER

- Separate encoders for scan, join, sort, aggregate
- Takes in plan features, meta features, DB settings as input
- To predict total time/cost for plan, add plan feature vectors together for each node type

Operator	Plan Properties or Features
All	Actual Loops , Actual Rows , Local Dirtied Blocks , Local Hit Blocks , Local Read Blocks , Local Written Blocks , Plan Rows , Plan Width , Shared Dirtied Blocks , Shared Hit Blocks , Shared Read Blocks , Shared Written Blocks , Temp Read Blocks , Temp Written Blocks , Parent Relationship , Plan Buffers
Scan	Relation Name, Scan Direction, Index Name, Index Condition, Scan Condition, Filter, Rows Removed, Heap Blocks, Parallel, Recheck Condition
Join	Join Type, Inner Unique, Merge Condition, Hash Condition, Rows Removed by Join Filter, Parent Relationship, Hash Algorithm, Hash Algo, Hash Buckets, Hash Batches, Peak Memory
Sort	Sort Type, Sort Method, Sort Space, Sort Key, Sort Space Type, Sort Space Used, Peak Memory
Aggregate	Strategy, Hash Algo, Hash Buckets, Hash Batches, Parallel Aware, Partial Mode, Peak Memory

Features Type	Feature Attributes
Meta Features	relname, attname, reltuples, relpages, relfilenode, relam, n distinct, distinct values, selectivity, avg width, correlation
DB Settings	bgwriter delay, shared buffers, bgwriter lru maxpages, wal buffers, random page cost, bgwriter lru multiplier, checkpoint completion target, checkpoint timeout, cpu tuple cost, max stack depth, deadlock timeout, default statistics target, work mem effective cache size, effective io concurrency, join collapse limit, from collapse limit, maintenance work mem

PERFORMANCE ENCODER MODEL

- Input: $(f_{node}, f_{meta}, f_{db})$
- Output performance metrics: total cost, total time, startup time
- Evaluation criteria
  - How long it takes to adapt to a new domain
  - Model error after fine-tuning

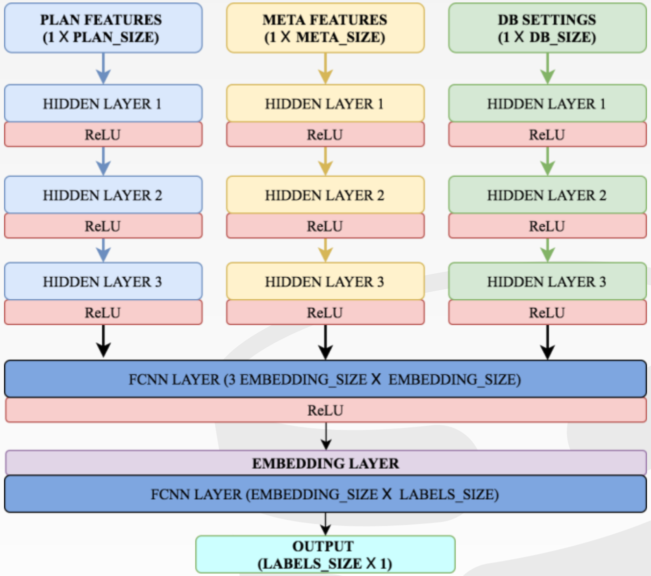

DATABASE WORKLOAD CHARACTERIZATION WITH QUERY PLAN ENCODERS

TODAY'S AGENDA

- Query Encoder Model
- Downstream Tasks
- Analysis

DATASETS

- Crowdsourced query plan dataset
- TPC-H, TPC-DS
- Join Order Benchmark
- Spatial benchmarks
  - Notorious for resource consumption and need for proper tuning
  - Jackpine: spatial queries with multipolygons, lines, and points
  - Open Street Map: Spatial overlap, distance, and routing queries

LATENCY—TPC-DS 100

- Compared to SOTA query latency prediction models: TAM, SVM, RBF, QPP Net

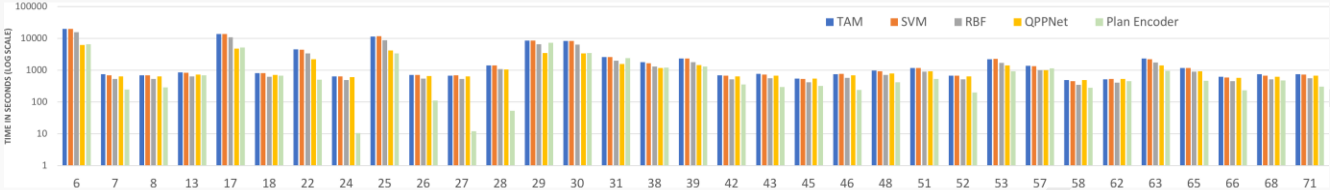

Figure 7: Mean absolute error (MAE) for the 33 TPC-DS query templates with scale factor 100 where Plan Encoder performed better than a baseline.

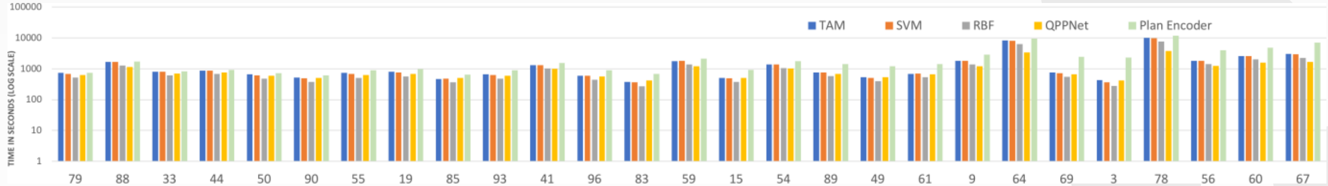

Figure 8: Mean absolute error (MAE) for the 27 TPC-DS query templates with scale factor 100 where Plan Encoder did not performed better than a baseline.

LATENCY—TPC-DS 100

- Features from the performance encoder are dominant
- Features from structure encoder have low importance in latency prediction
- Structure embedding size of 128 or 160 performed best

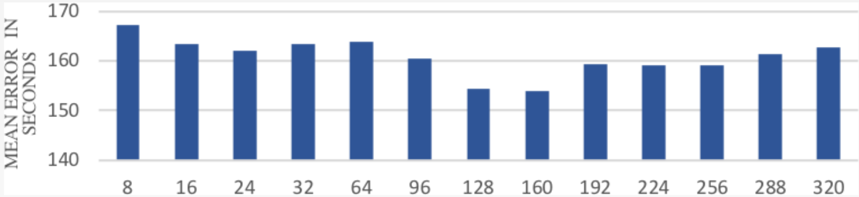

Figure 9: Average of MAEs on 5 test datasets of TPC-DS with scale factor 10 with varying embedding size of structure encoder.

QUERY CLASSIFICATION

- Join Order Benchmark—predict query template and cluster
- Pretrained encoder
- Both structure and performance encoders contribute, but the structure encoder is more important
- Models fine-tuned on 10% and 30% of the dataset performed close to those fine-tuned with the full dataset

Methods	Dev		Test	
	template	cluster	template	cluster
Structure-Only	0.2452	0.4670	0.1946	0.3847
Performance-Only	0.1645	0.2973	0.0977	0.1769
Both	**0.2783**	**0.5573**	**0.2518**	**0.4647**
Both0.1	0.2000	0.4927	0.151	0.334
Both0.3	0.2555	0.5228	0.1843	0.3855

Table 6: Results on Query Classication Accuracy

TODAY'S AGENDA

- Query Encoder Model
- Downstream Tasks
- Analysis

STRUCTURE ENCODER—TRAINING

- Baseline models
  - Sparse autoencoder (Sparse-AE)
  - LSTM with plan-pair similarity regression
- Pretraining datasets
  - Randomly select pairs from crowdsourced plan dataset
- Fine-tuning datasets
  - Randomly generated plan-pairs from TPC-H, TPC-DS, and SPATIAL

STRUCTURE ENCODER—TRAINING

- Pretraining significantly reduces error on TPC-H and TPC-DS
- LSTM and Transformer models work well from scratch on SPATIAL

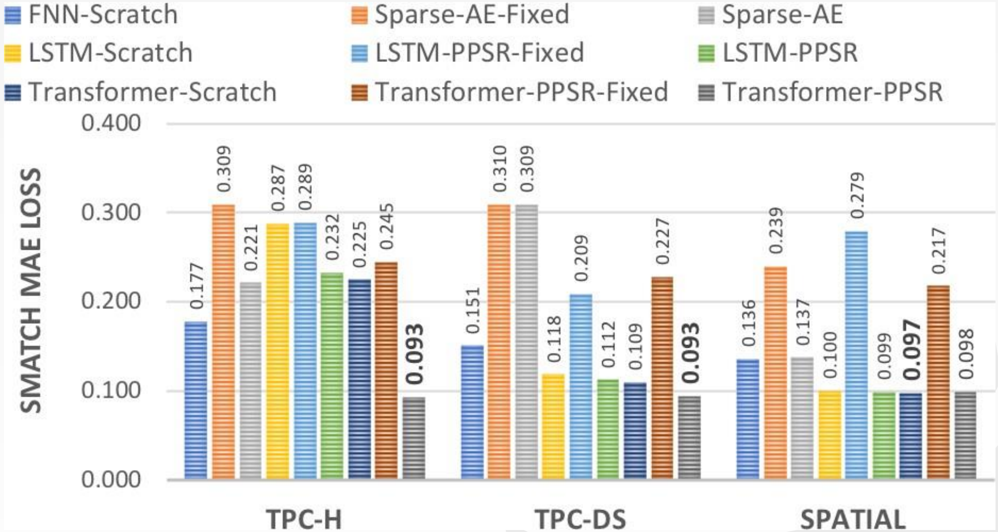

Scratch: no pretraining or fine-tuning
Fixed: pretraining only

STRUCTURE ENCODER—TRAINING

- Pretraining significantly reduces error with less training data for TPC-H and TPC-DS
- Pretraining only improved SPATIAL performance slightly

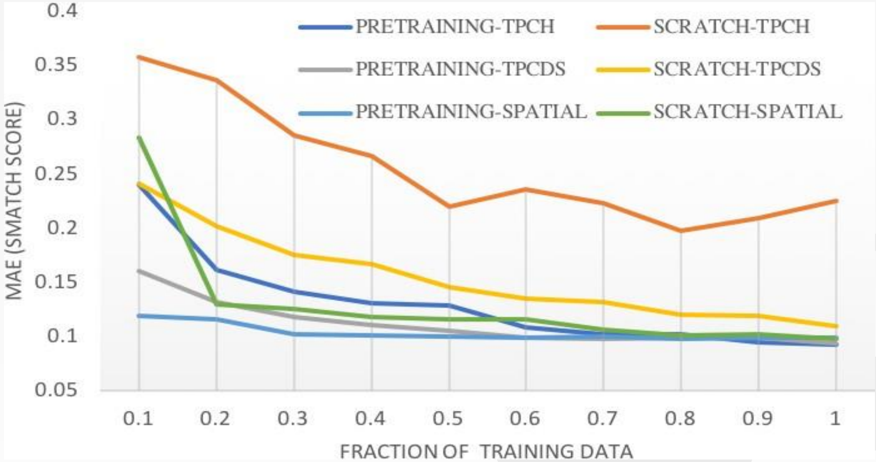

Scratch: no pretraining or fine-tuning
Fixed: pretraining only

PERFORMANCE ENCODER—PRETRAINING

- Pretraining datasets:
  - TPC-H, TPC-DS with scale factors 1, 2, 3, 5
  - 20 randomly generated configuration settings

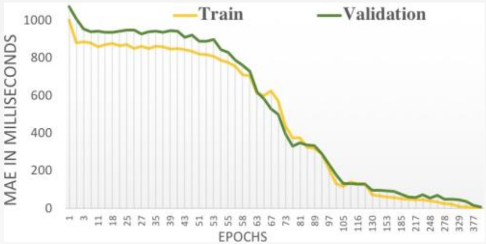

SCAN operator pretraining

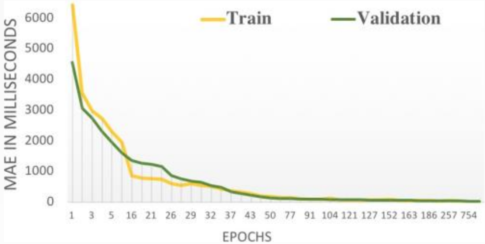

JOIN operator pretraining

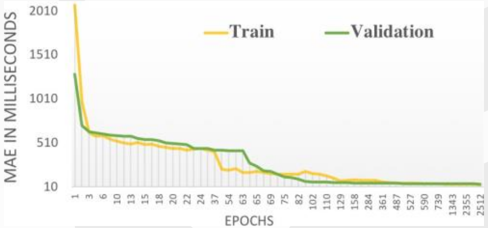

SORT operator pretraining

DATABASE WORKLOAD CHARACTERIZATION WITH QUERY PLAN ENCODERS

PERFORMANCE ENCODER—FINE TUNING

- All pretrained models were fine tuned with new datasets
  - TPC-DS with scale factor 8
  - Spatial—similar dataset
- Pre-trained model with fine tuning always beats scratch

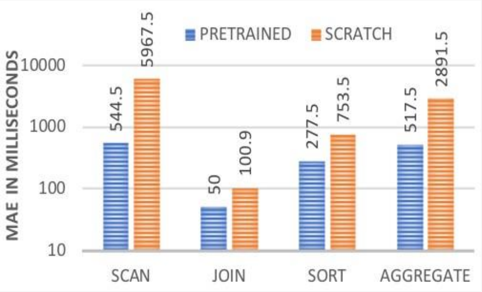

TPC-DS SF-8

Spatial

DATABASE WORKLOAD CHARACTERIZATION WITH QUERY PLAN ENCODERS

PERFORMANCE ENCODER—SCRATCH

- Pre-trained, scratch models evaluated on TPC-DS, Spatial datasets
- Pretrained test accuracy plateaus after 30% of training data
- Scratch model only comes close after 50-70% of training data

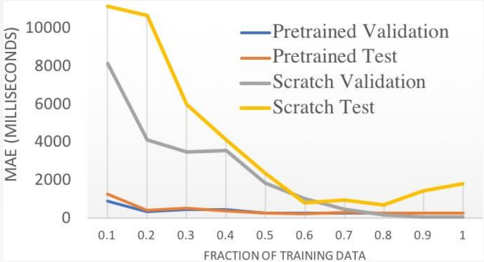

SCAN

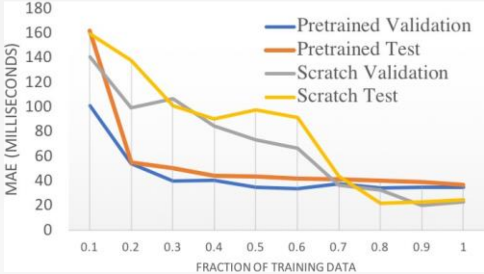

JOIN

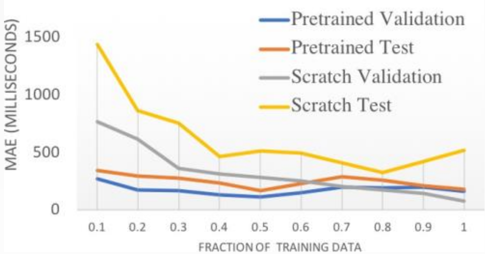

SORT

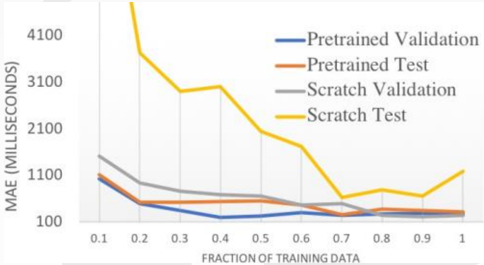

AGGREGATE

DATABASE WORKLOAD CHARACTERIZATION WITH QUERY PLAN ENCODERS

PERFORMANCE ENCODER—MULTI COLUMN

- TPC-DS: three-column DNN performs better except for scan operator
- Spatial workload: three-column DNN always performs better than standard single-column DNN

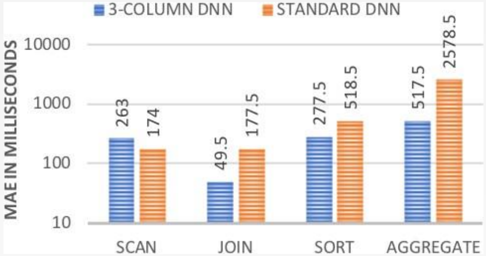

TPC-DS SF-8

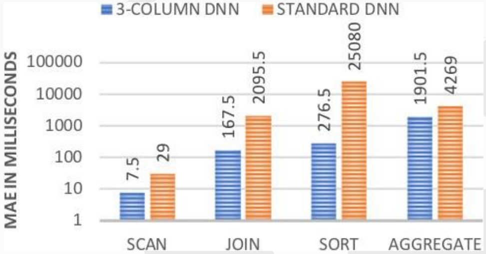

Spatial

DATABASE WORKLOAD CHARACTERIZATION WITH QUERY PLAN ENCODERS

PARTING THOUGHTS

- Pre-trained query plan encoder
- Separate models for structure, performance

NEXT CLASS

Behavior Modeling

REFERENCES

[1] Debjyoti Paul, Jie Cao, Feifei Li, Vivek Srikumar. Database Workload Characterization with Query Plan Encoders. https://doi.org/10.48550/arXiv.2105.12287.