Carnegie Mellon University

D tur σ

Special Topics: Self-Driving Database Management Systems Bao: Making Learned Query Optimization Practical Best Paper of SIGMOD'21

@Lichen Jin // 15-799 // Spring 2022

LAST CLASS

Automatic SQL Rewriting

- \rightarrow Works at changing query at the SQL Language Level
- \rightarrow MCTS to efficiently enumerate and select based on rewriting rules
- \rightarrow Deep Learning model to estimate tree node costs

TODAY'S AGENDA

- \rightarrow Overview: Learned Query Optimization
- \rightarrow Prior Work: Neo
- \rightarrow Bao Introduction
- \rightarrow Enumerating Query Hints
- \rightarrow Contextual Multi-arm Bandit
- \rightarrow Predictive Model: TCNN Cost Model
- \rightarrow Integration and Improvements
- \rightarrow Experiments
- \rightarrow Parting Thought



Overview: Learned Query Optimization



15-799 Special Topics (Spring 2022)

Overview: Learned Query Optimization



SECMU-DB 15-799 Special Topics (Spring 2022)

Source: Ryan Marcus

Prior Work: Neo

 \rightarrow Neural Optimizer

15-799 Special Topics (Spring 2022)

- \rightarrow Complete replacement of default query optimizer
- \rightarrow List alternative plans based on expert rules
- \rightarrow Deep Reinforcement Learning Guided Search



Source: Ryan Marcus

Prior Work: Neo

Promising Results But

Sample Inefficiency

ightarrow Typically takes more than 1 day for pre-train

Brittleness to Workload and Schema Change

 $\rightarrow\,$ The encoding of cardinality estimate needs retrain

Tail Catastrophe

→ Deep RL making wrong estimations due to sample inefficiency

Other state-of-the-art: Similar Issues

SECMU-DB 15-799 Special Topics (Spring 2022)



Bao: Introduction

Bao: Bandit optimizer

By *steering* a traditional query optimizer, Bao:

- \rightarrow Outperforms PG after *1 hour* of training
- → Reduces *99% latency*
- \rightarrow Adapts to changes in workload, schema, and data.



Source: Ryan Marcus

8

SECMU-DB 15-799 Special Topics (Spring 2022)

Bao: Introduction

Three key points:

Offload Alternative Plan Enumeration \rightarrow Query Hint Sets on top of the optimizer

Model Plan Selection as CMAB

- \rightarrow Exploration and Exploitation
- \rightarrow Reduce chances of wrong tail estimation

Deep Estimation Model

- → Tree Convolutional Neural Networks (TCNN)
- \rightarrow Change-adaptive encoding approach



SIGMOD 2021

BAO: MAKING LEARNED QUERY OPTIMIZATION PRACTICA

Query Hints

Slow query.

Run EXPLAIN.

- Loop join plan,
- Low selectivity

Try disabling loop join

Huge improvement



Query Hints

Slow query.

Run EXPLAIN.

- Loop join plan,
- Low selectivity

Try disabling loop join

Huge improvement

Apply this hint globally

... other regressions,

Undo that, need local hints. Now What?

- Opt 1: Apply the hint to every instance of the query
- > Opt 2: Set as default, find regressions, add hints to those queries
- > Opt 3: Give up



Source: Ryan Marcus

EFCMU·DB 15-799 Special Topics (Spring 2022)

Enumerating Query Hints

Query Hint Set: A combination of configuration knobs.

 \rightarrow Example: enable nested loop join

Query plans are enumerated using different query hint sets.

Bao tries to automatically determine the optimal hint to use.



SECMU.DB 15-799 Special Topics (Spring 2022)

Suppose you are in a casino But don't know which bandit wins money



Suppose you are in a casino But don't know which bandit wins money



SECMU-DB

Suppose you are in a casino

You want a mental model to estimate the best



You really try the 'best' for ground truth Used for updating the model



Suppose you are in a casino

The model gradually estimates better on winning





15-799 Special Topics (Spring 2022)

Source: Ryan Marcus

Mapping Back to Query Hints

Exploration-Exploitation Tradeoff

- \rightarrow Losing Exploration: Local optimal, Wrong estimations
- \rightarrow Losing Exploitation: Model cannot converge



15-799 Special Topics (Spring 2022)

Mapping Back to Query Hints

 $\begin{array}{l} \text{Exploration-Exploitation Tradeoff} \\ \rightarrow \text{ Solution: Thompson Sampling} \end{array}$





SECMU.DB 15-799 Special Topics (Spring 2022) Source: Ryan Marcus

CMAB in Bao: Thompson Sampling

An old, well-studied algorithm for balancing exploration and exploitation.

Usual ML training Weight = E[Weight | data] (exploitation)

Pick a random model Weight = sample P(Weight) (exploration)

Sample model weights Weight = sample P(Weight | data) (Thompson Sampling)

Source: Ryan Marcus

CMAB in Bao: Thompson Sampling

An old, well-studied algorithm for balancing exploration and exploitation.

Usual ML training Weight = E[Weight | data] (exploitation)

Pick a random model Weight = sample P(Weight) (exploration) Bao Implementation: Retrain on a random sampled data set sized |E|, randomly drawn from E with replacement

Sample model weights (Thompson Sampling)

Weight = sample P(Weight | data) Othe

Other complex approaches: Bayesian Networks

Source: Ryan Marcus

Second Decision Spring 2022)

Predictive Model: TCNN Cost Model

Tree Convolutional Network: Efficient catching correlation as high-level patterns between operators in the tree structure. \rightarrow Long chain of merge joins without sort.

 \rightarrow Bushy tree of hash operators.



Figure 5: Bao prediction model architecture

Predictive Model: Plan Encoding

Binarize: All internal nodes have two children; *supply with null node* Multiple Children: *Convert to Binary*



Figure 3: Binarizing a query plan tree

BAO: MAKING LEARNED QUERY OPTIMIZATION PRACTICAL SIGMOD 2021

Predictive Model: Plan Encoding

Binarize: All internal nodes have two children; *supply with null node* Multiple Children: *Convert to Binary*

Vectorize: One-hot encoding of operator type, card. Estimator, cost



Figure 3: Binarizing a query plan tree



Figure 4: Vectorized query plan tree (vector tree)

Integration and Improvements

- \rightarrow Generate query plans based on hint sets
- → Select the "best" query hint set predicted by TCNN, execute the corresponding plan
- \rightarrow Add ground truth to Experience data
- \rightarrow Retrain with Thompson Sampling *per n executions* \rightarrow Only keep *most recent k* as Experience data

PostgreSQL Integration

Per-query Activation \rightarrow SET enable_bao TO [on/off]



Active vs. advisor mode

- \rightarrow Active: use bao to select query plan
- \rightarrow Advisor: default; implicitly train the model

Triggered Exploration

- $\rightarrow\,$ Mark query as critical to avoid regression
- $\rightarrow\,$ Weighted retrain on critical misprediction

27

PostgreSQL Integration

Per-query Activation

 \rightarrow SET enable_bao TO [on/off]

Active vs. advisor mode

- \rightarrow Active: use bao to select query plan
- \rightarrow Advisor: default; implicitly train the model

Triggered Exploration

- \rightarrow Mark query as critical to avoid regression
- \rightarrow Weighted retrain on critical misprediction

Source Code: <u>https://github.com/learnedsystems/BaoForPostgreSQL/tre</u> e/master/pg_extension

Secmu∙db

15-799 Special Topics (Spring 2022)

switch (arm) {

case 0:

enable_hashjoin = true; enable_indexscan = true; enable_mergejoin = true; enable_nestloop = true; enable_seqscan = true; enable_indexonlyscan = true; break;

case 1:

enable_hashjoin = true; enable_indexonlyscan = true; enable_indexscan = true; enable_mergejoin = true; enable_seqscan = true; break;

case 2:

enable_hashjoin = true; enable_indexonlyscan = true; enable_nestloop = true; enable_seqscan = true; break;

#ifndef BAO_CONFIGS_H #define BAO_CONFIGS_H

#include "c.h"

#define BAO_MAX_ARMS 26

// Each Bao config variable is linked to a Postgre
// See the string docs provided to the PG function
static bool enable_bao = false;
static bool enable_bao_rewards = false;
static char* bao_host = NULL;
static int bao_port = 9381;
static int bao_num_arms = 5;
static bool bao_include_json_in_explain = false;
#endif

Experiment: Settings

OLAP Workload:

- \rightarrow Tail Dominate Pattern
- \rightarrow Dynamic Changes
- → IMDb (Default)

Machine Types:

	Size	Queries	WL	Data	Schema
IMDb	7.2 GB	5000	Dynamic	Static	Static
Stack	100 GB	5000	Dynamic	Dynamic	Static
Corp	1 TB	2000	Dynamic	Static ^a	Dynamic

^{*a*}The schema change did not introduce new data, but did normalize a large fact table.

 \rightarrow GCP N1-2; N1-4 (Default); N1-8; N1-16

Baseline Databases:

- → PostgreSQL (Default); Comsys [Oracle]
- \rightarrow Bao integrated in corresponding DBMS

Concurrency: 1 (Default), 2, 4 Bao configuration: n=100, k=2000 Model:

3 Layer convolution, 2 Linear, Relu, Batch Norm; Adam Batch=16; epoch=100 with early stop

Time-split training: Model only updated with data collected from previous time split

Experiment: Results



(a) Across our three evaluation datasets, Bao on the PostgreSQL engine vs. PostgreSQL optimizer on the PostgreSQL engine.



(b) Across our three evaluation datasets, Bao on the ComSys engine vs. ComSys optimizer on the ComSys engine.

500 600 Bao 500 400 PostgreSQL Cost (cents) (m 400 300 1 200 300 200 200 100 100 0 Λ n1-2 n1-4 n1-8 n1-16 n1-2 n1-4 n1-8 n1-16 Machine Type Machine Type

(a) Across four different VM types, Bao on the PostgreSQL engine vs. PostgreSQL optimizer on the PostgreSQL engine.



(b) Across four different VM types, Bao on the ComSys engine vs. ComSys optimizer on the ComSys engine.

Significant Lower Cost and Runtime

Tail Latency Analysis

IMDB: Significant Lower Tail Latency



Analysis over Runtime

IMDB-PostgreSQL: Model converges in a few hours



Figure 10: Number of IMDb queries processed over time for Bao and the PostgreSQL optimizer on the PostgreSQL engine. The IMDb workload contains 5000 unique queries which vary over time.

BAO: MAKING LEARNED QUERY OPTIMIZATION PRACTICAL SIGMOD 2021

Per-query Analysis

IMDB-PostgreSQL: Only 3 of 105 has minor regression Mostly with significant improvement



Varying # of Arms (Hint Sets)

IMDB-PostgreSQL: 48 arms ordered by observed benefits Only a few arms is sufficient for optimization



BAO: MAKING LEARNED QUERY OPTIMIZATION PRACTICAL SIGMOD 2021

Varying Concurrency and Memory

IMDB-PostgreSQL: Good performance when I/O bound High CPU usage contention when CPU bound



Adaptivity Analysis

Bao better adapts to changes compared to state-of-the-art



Figure 14: Comparison of number of queries finished over time for Bao, Neo [51], DQ [40], and PostgreSQL for a stable query workload (left) and a dynamic query workload (right).

BAO: MAKING LEARNED QUERY OPTIMIZATION PRACTICAL SIGMOD 2021

Modeling Evaluation

- \rightarrow TCNN is necessary for Query Latency Predict
- \rightarrow Bao avoids regression when model not converged
- \rightarrow GPU time linear to the experience window size



(a) Random forest (RF) and linear models (Linear) used as Bao's model. "Best hint set" is the single best hint set. IMDb, N1-16 VM, on PostgreSQL.

(b) Median Q-Error (0 is a perfect prediction) of Bao's predictive model vs. the number of queries processed. IMDb workload on N1-16 VM using PostgreSQL engine. (c) Simulated and observed time to train Bao's performance prediction model (GPU) based on the sliding window size (number of queries used during each training iteration).

Model Regret (Loss) Analysis

$$R_q = \left(P(B(q)(q)) - \min_i P(HSet_i(q)) \right)^2$$



(b) Physical I/O regret

BAO: MAKING LEARNED QUERY OPTIMIZATION PRACTICAL SIGMOD 2021

Bao on Distributed Databases

After SIGMOD submission: Applying Bao to distributed systems



Source: Ryan Marcus

SECMU-DB 15-799 Special Topics (Spring 2022)

PARTING THOUGHTS

- → Support searching on larger enumeration space (different query hints at different points) with approaches like MCTS
- → Inject learnt query optimizer changes to lower level like JIT bytecode with minimum recompiling
- → Why Bao eliminates tail latency better? Just the model is better sampled and hence more accurate?
- → Not suitable for OLTP optimizing? (Simple queries, no tail patterns)

PERMUTABLE COMPILED QUERIES: DYNAMICALLY ADAPT COMPILED OUERIES WITHOUT RECOMPILING

NEXT CLASS

→ Query Optimization II

 \rightarrow P2 Discussion