SUB-OPTIMAL JOIN ORDER IDENTIFICATION WITH L1-ERROR

SIGMOD PODS 2024

Santiago Chile

article

code

Yesdaulet Izenov[2], **Asoke Datta**[1], **Brian Tsan**[1] and **Florin Rusu**[1]

[1]Dept. CSE, University of California, Merced, CA, USA

[2]Dept. CS, Nazarbayev University, Astana, Kazakhstan

Q-error – the standard metric for quantifying the error of individual cardinality estimates – has been widely adopted as a surrogate for query plan optimality in recent work on learning-based cardinality estimation. However, the only result connecting Q-error with plan optimality is an upper-bound on the cost of the worst possible query plan computed from a set of cardinality estimates – there is no connection between Q-error and the real plans generated by standard query optimizers. Therefore, to identify sub-optimal query plans, we propose a learning-based method having as its main feature a novel measure called **L1-error**. Similar to Q-error, L1-error requires complete knowledge of true cardinalities and estimates for all the sub-plans of a query plan. Unlike Q-error, which considers the estimates independently, L1-error is defined as a permutation distance between true cardinalities and estimates for all the sub-plans having the same number of joins. Moreover, L1-error takes into account errors relative to the magnitude of their cardinalities and gives larger weight to small multi-way joins. Our experimental results confirm that, when L1-error is integrated into a standard decision tree classifier, it leads to the accurate identification of sub-optimal plans across four different benchmarks. This accuracy can be further improved by combining L1-error with Q-error into a composite feature that can be computed without overhead.

1 Problem

Can we classify the sub-optimality of a query plan based on true and estimated cardinalities?

2-way joins (CARDINALITY only)					
index ρ	1	2	3	4	5
sub-plan	$cn \bowtie mc$	$k \bowtie mk$	$mc \bowtie t$	$mk \bowtie t$	$mc \bowtie mk$
true Y	388	41.8K	2.6M	4.5M	34.9M
est. $\hat{Y}$	973	20	1.5M	2.7M	13.8M
Q-error q	2.51	2,092	1.70	1.70	2.53

3-way joins (CARDINALITY only)					
index ρ	1	2	3	4	5
sub-plan	$cn \bowtie mc \bowtie t$	$cn \bowtie mc \bowtie mk$	$k \bowtie mk \bowtie t$	$k \bowtie mk \bowtie mc$	$mc \bowtie t \bowtie mk$
true Y	388	1,588	41.8K	148.6K	34.9M
est. $\hat{Y}$	973	8,739	20	104	2.7M
Q-error q	2.51	5.50	2,092	1,428.38	12.69

4-way joins (CARDINALITY only)			
index ρ	1	2	3
sub-plan	$cn \bowtie mc \bowtie mk \bowtie k$	$cn \bowtie mc \bowtie t \bowtie mk$	$k \bowtie mk \bowtie t \bowtie mc$
true Y	4	1,588	148.6K
est. $\hat{Y}$	1	8,739	104
Q-error q	4	5.50	1,428.38

5-way joins (CARDINALITY only)	
index ρ	1
sub-plan	$cn \bowtie mc \bowtie t \bowtie mk \bowtie k$
true Y	4
est. $\hat{Y}$	1
Q-error q	4

Plan	Join order	$\mathcal{C}(\mathcal{P}, Y)$	$\mathcal{C}(\mathcal{P}, \hat{Y})$
$\mathcal{P}_{opt}$	$cn \bowtie mc \bowtie t \bowtie mk \bowtie k$	2,364	10.7K
$\mathcal{P}_{pg}$	$k \bowtie mk \bowtie t \bowtie mc \bowtie cn$	232.2K	144

Query plan enumeration over the search space corresponding to JOB query 2c.

2 Q-error

$$q_i = \max\left(\frac{\hat{Y}_i}{Y_i}, \frac{Y_i}{\hat{Y}_i}\right) \quad (1) \qquad \frac{\mathcal{C}(\mathcal{P}_{pg}, Y)}{\mathcal{C}(\mathcal{P}_{opt}, Y)} \leq q^4 \quad (2)$$

where Y_i and $\hat{Y}_i$ are true and estimated cardinality of a single sub-plan. $q = \max_{i \subseteq X}\{q_i\}$ and X is the set of all sub-plans. The cost ratio is called *P-error*.

The bound states that, given the estimated and true cardinalities of a query, the cost ratio of any query plan and the optimal query plan is at most q^4.

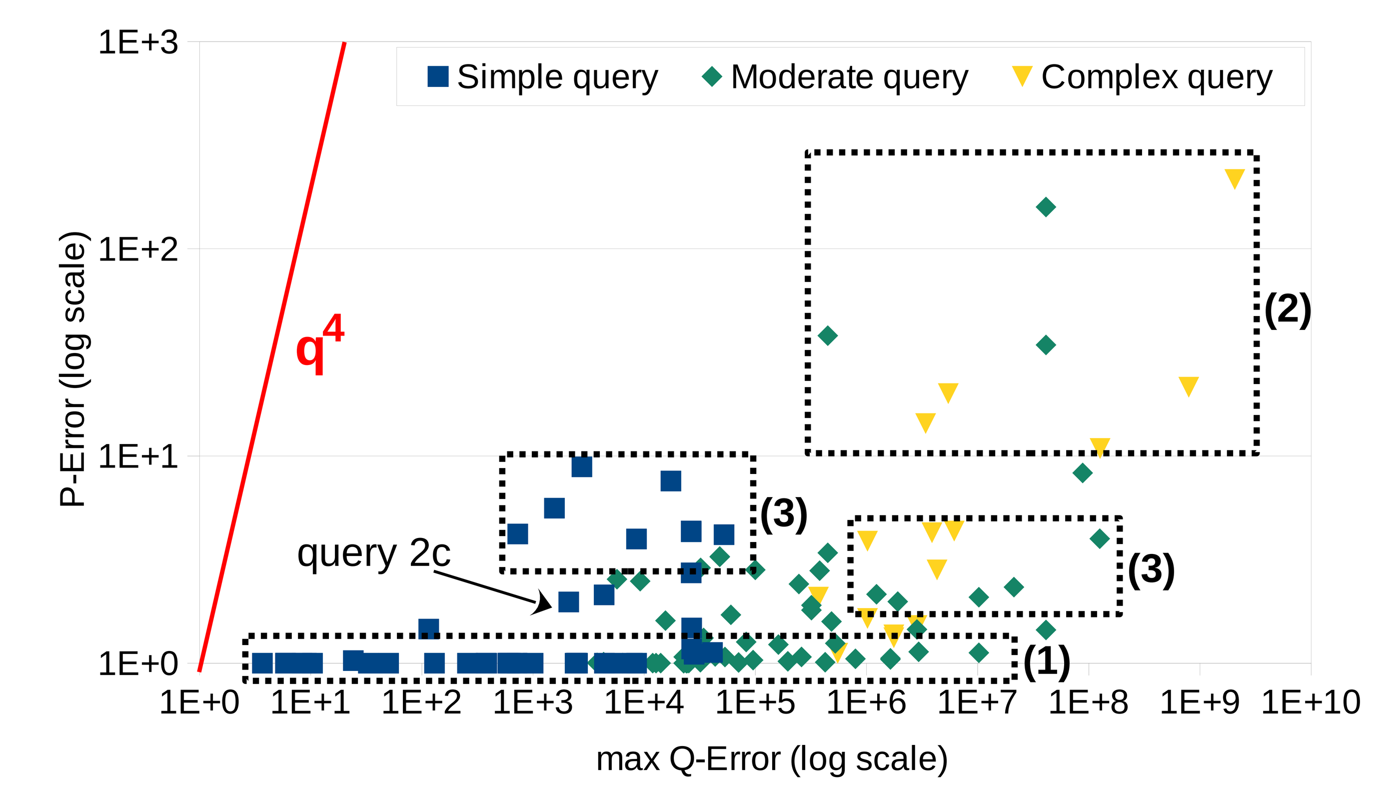

P-error and Q-error are computed for all 113 JOB queries. Queries complexity: 45 Simple (4-9 joins), 53 Moderate (10-19 joins), and 15 Complex (20-28 joins).

Observations:

- loose as a bound — large Q-error and small P-error
- no observable relationship — between Q-error and P-error
- reversed relationship — smaller Q-error with larger P-error and vice versa

3 L1-error

Spearman's footrule (a.k.a. L1-distance):

$$\text{L1}^k(\rho, \hat{\rho}) = \sum_{i}^{d} |\rho(i) - \hat{\rho}(i)|$$

where $\rho(i)$ and $\hat{\rho}(i)$ are the indexes of i-th sub-plan in the position vectors. In the case of identical sub-plan arrangements, $\text{L1}^k(\rho, \hat{\rho}) = 0$.

Our Objectives – Weighted L1-error:

- Significantly over and underestimating cardinality should be associated with greater penalties. Conversely, mispositioning sub-plans with similar cardinalities in the position vectors should carry fewer and relatively similar penalties.
- Cardinality misestimations that occur early in the position vector should attract greater penalties. This is particularly beneficial for plan search algorithms, which are more likely to choose a sub-plan from the first half of the position vector.

2-way join (CARDINALITY only)					
ρ	1	2	3	4	5
$\hat{\rho}$	2	1	3	4	5
Swap cost	1.0	107.84	62.36	1.73	7.71
Monotonic weight	1.0	108.84	171.2	172.93	180.64

Impact weight, W (2-way join)	1	2	3	4	5
1	1.0	107.84	6,724.56	11,659.61	89,854.74
2		1.0	62.36	108.12	833.26
3			1.0	1.73	13.36
4				1.0	7.71
5					1.0

3-way join (CARDINALITY only)					
ρ	1	2	3	4	5
$\hat{\rho}$	3	4	1	2	5
Swap cost	1.0	4.09	26.35	3.55	234.69
Monotonic weight	1.0	5.09	31.44	34.99	269.68

Impact weight, W (3-way join)	1	2	3	4	5
1	1.0	4.09	107.84	382.87	89,854.74
2		1.0	26.35	93.55	21,954.43
3			1.0	3.55	833.26
4				1.0	234.69
5					1.0

4-way join (CARDINALITY only)			
ρ	1	2	3
$\hat{\rho}$	1	3	2
Swap cost	1.0	397.0	93.55
Monotonic weight	1.0	398.0	491.55

Impact weight, W (4-way join)	1	2	3
1	1.0	397.0	37,138.0
2		1.0	93.55
3			1.0

Sub-plan weights used by the plan enumeration algorithms for JOB query 2c.

4 Experiments - L1-error classifier accuracy on test data using PostgreSQL cardinality estimates

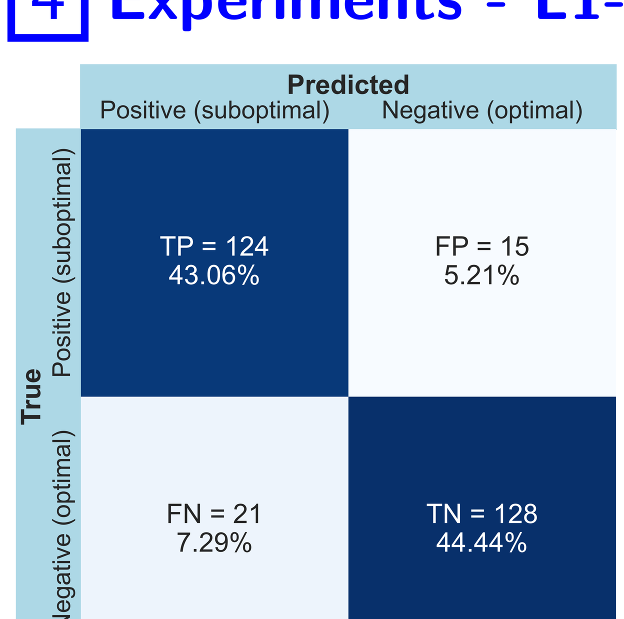

DSB benchmark. Exhaustive enumeration.
1152 train and 288 test queries.

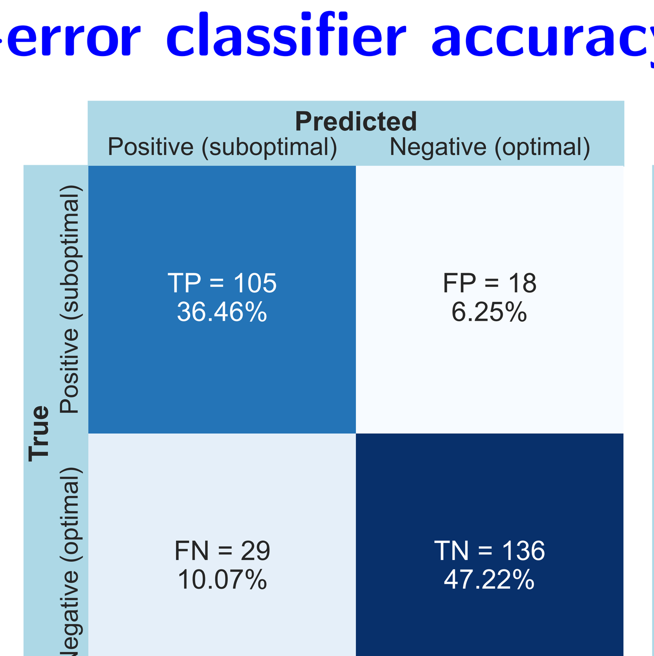

DSB benchmark. Greedy enumeration.
1152 train and 288 test queries.

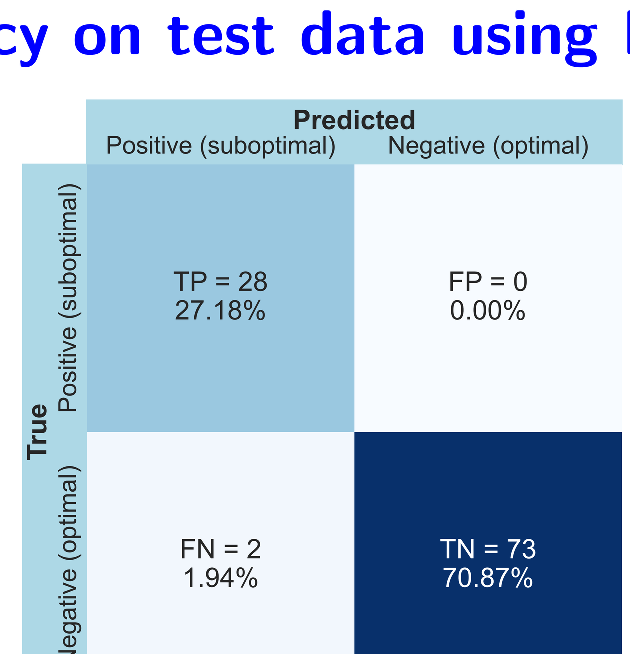

JCCH benchmark. Exhaustive enumeration.
408 train and 103 test queries.

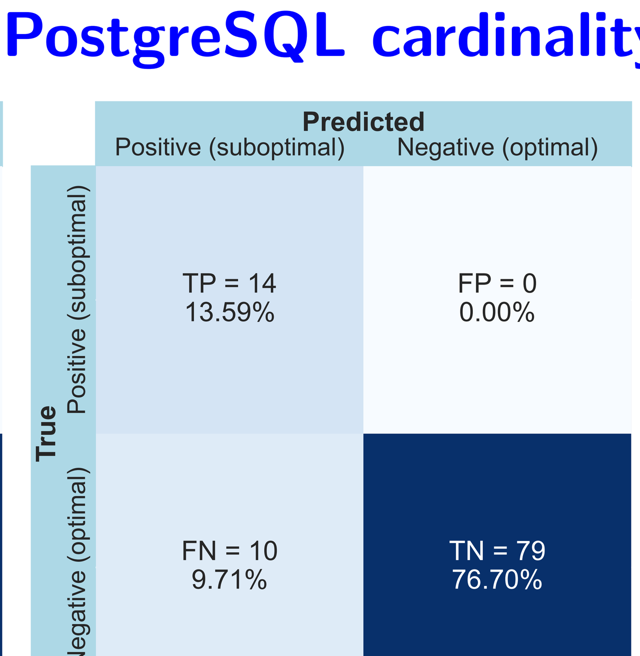

JCCH benchmark. Greedy enumeration.
408 train and 103 test queries.

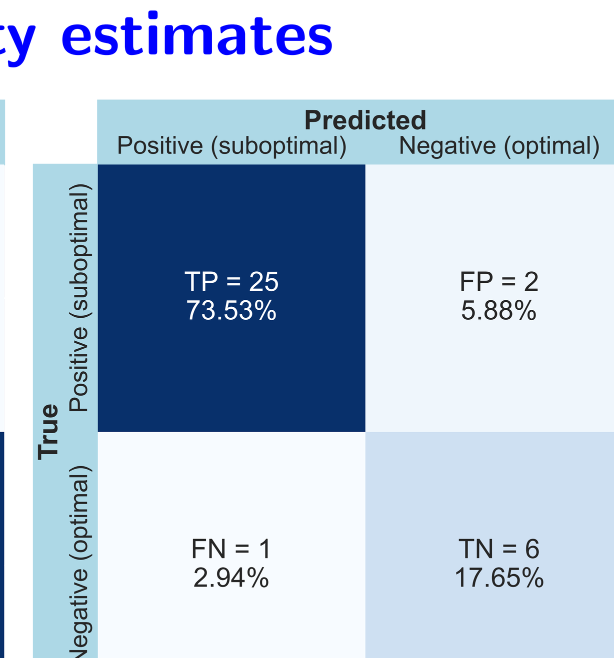

JOB benchmark. Exhaustive enumeration.
79 train and 34 test queries.

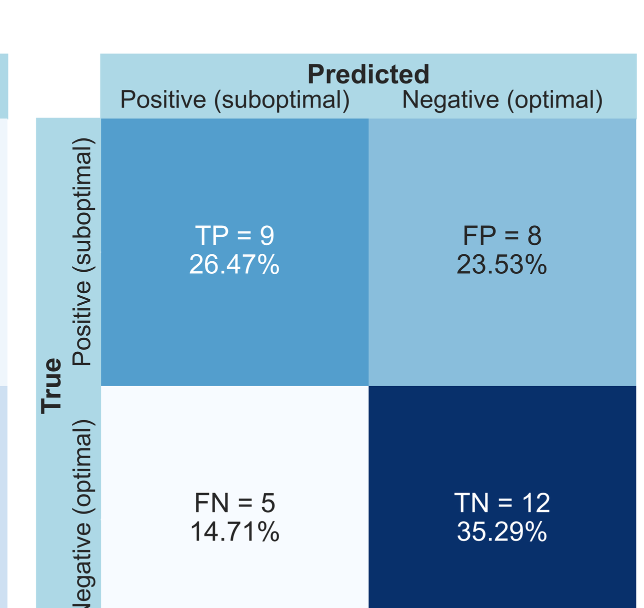

JOB benchmark. Greedy enumeration.
79 train and 34 test queries.