QTune: A Query-Aware Database Tuning System with Deep Reinforcement Learning

Guoliang Li, Xuanhe Zhou, Shifu Li, Bo Gao

Introduction and background

Databases have hundreds of ‘knobs’ that can be tuned to optimise queries.

DBAs cannot tune all knobs and take a long time.

BestConfig and OtterTune require many high quality training examples.

CDBTune only provides coarse-grained tuning

Three levels of tuning

Query level: good latency, bad throughput

Workload level:: bad latency, good throughput

Cluster level: good latency, good throughput

Vectorising queries

Insert	Delete	Update	Select	tbl1	tbl2	tbl3	...	tbl8	Hash_Join	Seq_Scan	Aggregate	...
[0	0	0	1	1	1	1	...	0	0.1401	-0.166	-0.2423	...]

Query type, tables, operation costs (estimate normalised around 0)

There is padding for new tables

To combine, union over flags and sum over operation costs.

Deep Reinforcement Learning

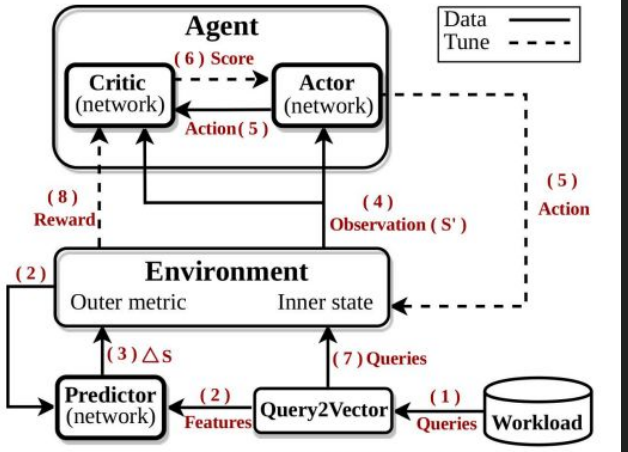

Figure 4: The DS-DDPG Model

Table 1: Mapping from DS-DDPG to Tuning

DS-DDPG	The tuning problem
Environment	Database being tuned
Inner state	Database knobs (e.g., work_mem)
Outer metrics	State statistics (e.g., updated tuples)
Action	Tuning database knobs
Reward	Database performance changes
Agent	The Actor-Critic networks
Predictor	A neural network for predicting metrics
Actor	A neural network for making actions
Critic	A neural network for evaluating Actor

Traditional DRL ignores the effect of the query on the environment.

Double-State Deep Deterministic Policy Gradient (DS-DDPG)

Clustering

It is expensive to compute the continuous knob values for all the of the queries. Instead compute a discrete estimate {-1,0,+1}.

Only consider the top *m* knobs.

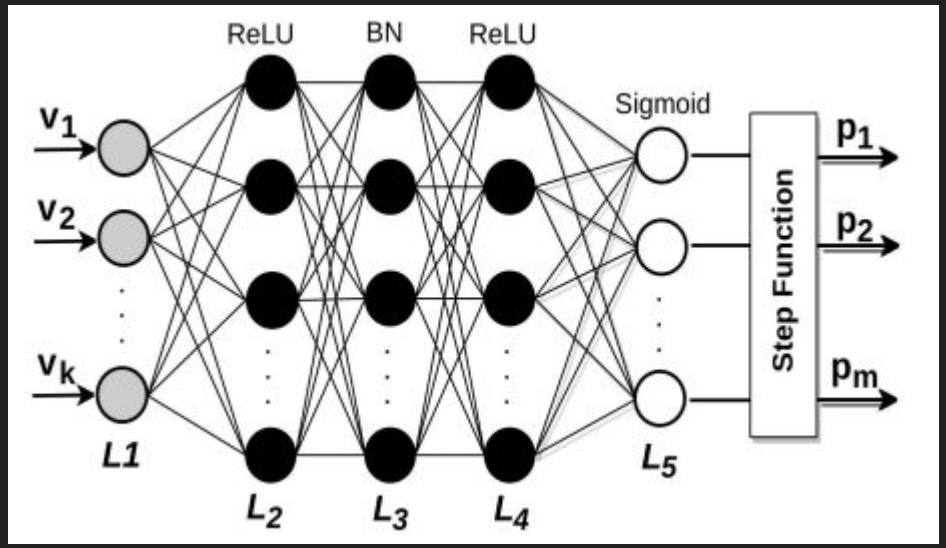

Results

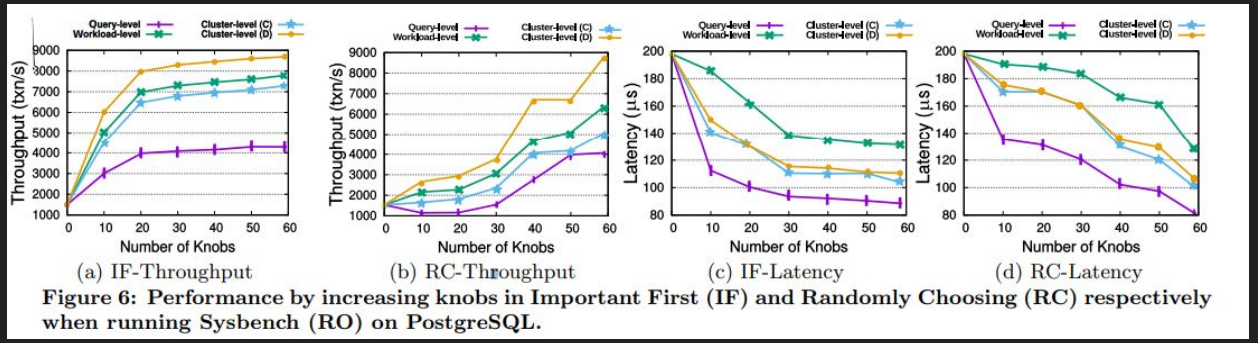

Figure 6: Performance by increasing knobs in Important First (IF) and Randomly Choosing (RC) respectively when running Sysbench (RO) on PostgreSQL.

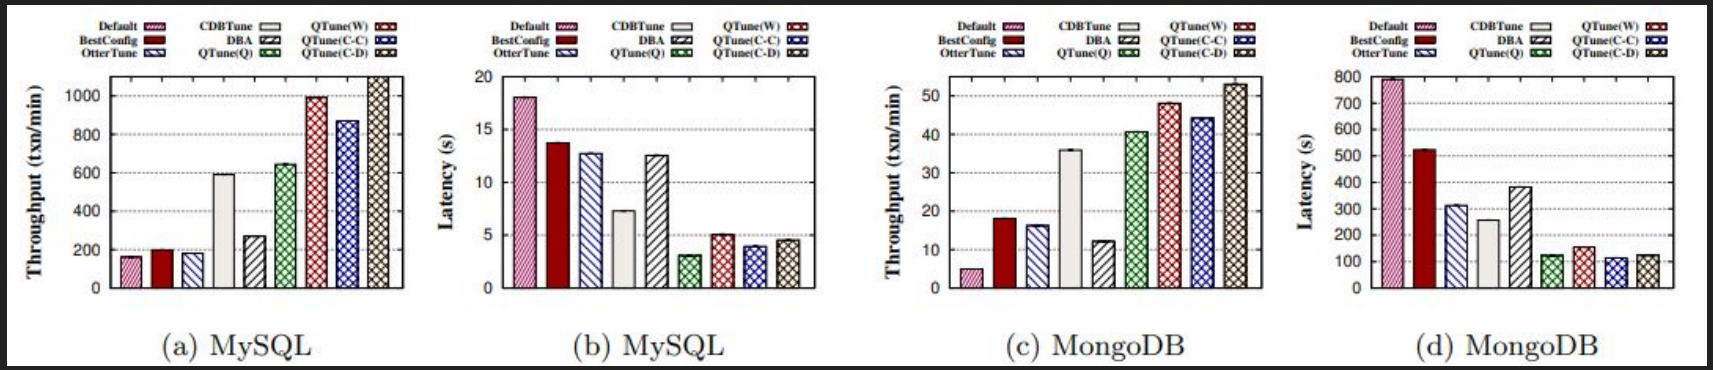

Opinion

Very thorough evaluation.

However, the system is very complicated. There is: actor network, critic network, predictor network, vector2pattern, clustering algorithm.

Training is still on the order of a day.