

# Dhrumil Shah

(+91) 7600196211 ◊ dhrumilkumars@iisc.ac.in ◊ dhrumil.iisc@gmail.com ◊ : dhrumil.shah488  
/dhrumil488 ◊  dsl.cds.iisc.ac.in/~dhrumil/ ◊  in/dhrumil488/

## AREAS OF INTEREST

Database Systems, OS, Machine Learning, Computer Architecture, Algorithms & Data Structures

## EDUCATION

• <b>Indian Institute of Science, Bengaluru - M.Tech. (Research)</b> Computer Science & Automation Dept. - <b>Advisor:</b> Prof. Jayant Haritsa	<i>August 2018 - Present</i> CGPA: 8.3/10
• <b>Gujarat Technological University - B.E.</b> Computer Engineering	<i>July 2011 - June 2015</i> CGPA: 8.9/10

## RESEARCH WORK

### Cardinality Estimation Problem in Database Query Optimizers (Ongoing)

- Dynamic samples based on query predicates can highly reduce the chances of "Zero Sample" problem and hence, gives closer estimation to actual cardinality after extrapolation.
- To make predicate aware sampling feasible, we introduced **Grid Based Dynamic Sampling** to fetch concentrated and correlated samples for estimation.
- Designed the **modern sampling technique using learned data structures** with the framework of Indexed-Based Join Sampling (IBJS) for a faster & space-efficient estimation process.
- This new estimator beats the state-of-the-art learning models by the **orders of magnitude** for accuracy.

## PROJECTS

### Benchmark for Cardinality Estimation in Industrial Optimizers

Databases Course Project

- Designed a dataset benchmark to challenge the cardinality estimation methods in the industrial optimizers.
- Covered two types of methods: Histogram-based and Independence & Uniformity based estimation.

### Protect the Dynamic Data of a Process in Linux Kernel

OS Course Project

- Developed custom system calls that efficiently protect the dynamically allocated memory regions from malicious code and restore them on corruption.
- Specifically used the concept of "Copy on Write" and modified the Linux Kernel for write-protect fault.

### MINIREL - Minimal Relational Database Engine

Database Project

- Designed & Implemented a simplified single-user relational database system.
- Implemented the logical layer (algebra and schema layer) and the physical layer (files & records, buffer management, catalogs management, page routines).

### CricPredic - Over by Over Cricket win Predictor for ODI match

ML Course Project

- The aim is to accurately predict the winning probability of the team after every over using historical data.
- Used the combination of KNN and ANN models based on input for this regression problem.

### Lightweight Models for Cardinality Estimation of Range Predicates

ML Project

- Implemented very small regression models to deliver fast and accurate estimates for multi-dim. range predicates compared to traditional methods with very small training overheads.

## WORK EXPERIENCE

### Einfochips India Pvt. Ltd., Ahmedabad - Software Engineer

*July 2015 - June 2018*

- Developed **Webpage Automation Framework** (JAVA Selenium) and **Firmware Verification Tool** (Socket Comm.).
- POC: Observed 9% perf. gain for **offloading the computation** from **ARM** to **DSP** core using **RPC**.
- Designed and developed the **cloud-based web application framework** which provides consumer and power consumption data maintenance and reporting services to the power companies.
- Developed **cloud application** to maintain IOT-nodes and their data using MEAN architecture stack.

## TECHNICAL SKILLS

**Programming Languages:** C, C++, PYTHON, JAVA, SQL, NodeJS

**Tools and Frameworks:** Eclipse, PostgreSQL, MSSQL Server, Keras, TensorFlow, MongoDB, GitHub, Jenkins, Latex

## SELECTED GRADUATE COURSES

Databases, Topics in Databases, Operating Systems, Computer Architecture, Design and Analysis of Algorithms, Probability, Machine Learning, Data Analytics