

Siddharth Bhola

www.siddharthbhola.com | siddharthbhola123@yahoo.co.in | +1 (872) 203 5829

EDUCATION

Northwestern University, McCormick School of Engineering, Evanston, IL

December 2019

Master of Science in Information Technology

Coursework: Statistical Data Mining, Deep Learning, Reinforcement Learning, Information Security, NLP.

University of Pune, Pune, India

May 2014

Bachelor of Engineering, Computer Science.

SKILLS

Languages: C#, Python, Java, R, Scala.

Databases: MySQL, MongoDB, Elasticsearch, Cassandra, Redis.

Frameworks and Tools: Weka, .Net Core, ASP .Net MVC, .Net Web API 2.0, Amazon Web Services (EC2, S3, Redshift, Lambda, DynamoDB, Streams-Firehose), Git, Akka, Kibana, Grafana, Microsoft Power BI, RabbitMQ, Consul, ANTS Profiler.

PROJECTS

Residual Networks for classification of Cifar-10 images.

February 2019

- Implemented a 20-layer deep Convolutional Neural Network with Residual Connections to optimize, with test-accuracy of 92.10% in just 100 epochs.

Data Mining to predict the success of bank telemarketing

November 2018

- Developed a Data Mining solution using Python and Weka to implement KNN and Decision Tree Algorithms to predict and target customers likely to subscribe to a term deposit based on customer's attributes. The solution involves reading the training data files, generating train and validation data sets, training the models, predicting test labels and reporting the training and validation errors.

EXPERIENCE

Senior Software Engineer

Tavisca Solutions, Pune, India

July 2014 - July 2018

Tokenization Proxy for PCI-DSS

- Developed an API in .Net Core to integrate with a data-vault for tokenization of credit-card data.
- Designed and developed a proxy service in .Net Core to abstract credit-card tokenization behind the proxy and make the travel platform PCI compliant.

Revenue Management Engine

- Designed and developed a high-performant revenue management system that provides various pricing techniques and data-analytics for improving e-commerce product revenues, using Finch-Finagle, Scala, Akka, Amazon-Redshift, Amazon-Firehose, and Elasticsearch-Logstash-Kibana.
- Implemented an AWS lambda function in C# for log processing which is integrated with AWS Firehose and Elasticsearch.
- Developed a big data analytics solution using Amazon-Redshift and Microsoft PowerBI to analyze hotel bookings.

OneConnect (Hotel Booking Engine)

- Developed web APIs for a fully managed, cloud-based API platform that eases the complexities of booking travel inventory from multiple suppliers by using ASP .NET Web API 2, C#, Aerospike, Elasticsearch, Redis, and Consul.
- Migrated state store implementation from Aerospike to Cassandra which resulted in 30% improved latency of search API.
- Evaluated the performance profile of the system using ANTS Profiler and reported the bottlenecks in the code. The code changes reduced the memory utilization by 20% and significantly reduced the infrastructure cost.

Clarifi

- Developed a real-time hotel content aggregation and syndication system, along with a tenant management console using technologies such as C#, ASP .NET MVC, MongoDB, Elasticsearch, RabbitMQ, and web technologies such as HTML, CSS, and JavaScript.
- Independently researched and developed an aggregate of the world neighborhood boundaries which allows point-in-polygon test to verify the accurate location of hotels and help find the outliers and remove dirty data.

Software Engineer Intern

Calsoft Incorporation, Pune, India

June 2013 - May 2014

- Developed "Microsoft volume shadow copy service for SQLite database," which is a point-in-time backup mechanism for SQLite database on Windows servers.
- Research paper published in the "International Journal of Advanced Computational Engineering and Networking," with Paper ID – CSMEGOA-16024-021 and ISBN - 978-93-82702-58-0 (Feb 2014).

Software Engineer Intern

InfoStretch Corporation, Pune, India

August 2012 - Sept 2012

- Developed a location-based bus guide application for Pune by using Android SDK, Java, SQLite database.