Vishnu Prasad S
Male, Indian, 26 Years

Gangabagh
M.M.Street,
Kodaikanal-624101,
Mobile: +917899077904
vishnu667@outlook.com
vishnu667@gmail.com

Personal Details

Date of Birth : 08th October 1991.
Father's Name : Sankar A.
Languages Known : English, Tamil and Malayalam.

Areas of Experience

Web Designing	: Html5, CSS3, nvd3.js, d3.js
Languages Known	: Scala, Python, JavaScript
ML libraries	: Scikit-learn, Scipy, Tensor flow, Keras, Mahout, Spark-ml, Nltk, Spacy, Gensim
Client Side Web Frameworks	: Ember.js
Server Side Frameworks	: Django, Flask, PlayFramework, Http-kit, Netty.io, Undertow.io, RestExpress, spray.io
DataBases	: Cassandra, Hbase, Redis, RocksDb.
Big-Data	: Spark, Druid.io , Kafka, Zookeeper

Achievements

Github profile : https://github.com/vishnu667
Blog : http://vishnu667.github.io

Module Lead / Data Scientist – Happiest Minds (April 2015 - Current)

About: Happiest Minds is an IT company based in Bangalore. The company primarily focuses on providing services on disruptive technologies.

Worked on R&D Projects for Happiest Minds

Image analytics : Image analysis using CNN.Using CNN to find faces in an image and tag the person to the location. Identifying objects in an image and classifying them to get meaningful statistics on campus (number of vehicles in the area & people movement).

- Preprocessing & cleaning images to train the classifier.
- Training a Bounding Box regressor using Faster R-CNN, YOLO & SSD.
- Training image classifiers using tensor flow.
- Combining models to identify objects in a given image.
- Exporting and using the models on Raspberry Pi for edge analytics.

Anomaly Detection : A scalable domain agnostic Anomaly Detection solution. The System is built upon a distributed stack. Using Akka, (Hbase + Phoenix) later replaced by ElasticSearch, Zookeeper and Spark.

- Building a ETL layer for Data Ingestion from various sources to Kafka.
- Performing Aggregates and predictions on the Edge before pushing the data for reducing workloads at the consumer end.
- Implement PMLL Models for deploying the algorithms across various platforms.

Vega-http: Low Latency High performance Http server Built in C. Used mainly as a Jump Server for Data Pipelining, Due to its very low memory map, worked well on embedded devices running linux.
(http://www.happiestminds.com/wp-content/uploads/2016/05/Vega-HTTP-The-High-performance-web-server.pdf)

Senior Software Engineer – Reduce Data (Jun 2013 - Mar - 2015)

About: Programmatic Demand Side Platform – Reduce Data (https://www.reducedata.com/) . Reduce Data is a programmatic demand side platform that helps advertisers buy media using RealTime Bidding (RTB) across various exchanges. Reduce Data helps advertisers to buy specific audiences across devices such as Desktops, Tablets and Mobile Phones. Reduce Data is a realtime platform which use realtime machine learning algorithms and large scale realtime data processing systems, built to eliminate media waste and improve the return on investment (ROI). Reduce Data processes over 6-15 billion Bid Requests Every Day.

Worked on Projects for Reduce Data (https://www.reducedata.com/**)**

- High Performance, High Availability and Low Latency (< 30 ms internal processing time) DSP Ad Server using Java 7, JBoss Netty, Java Springs, REST Express, DI, Disruptor Framework, CQengine, Guava cache, Google protobuf and Apache Kafka.
- RealTime Bid Computation Engine using RealTime performance data.
- Programmatic Optimization and Predictive Analysis with realtime performance data and Machine Learning algorithms using Vowpal Wabbit and Apache Mahout.
- RTB (RealTime Bidding) and SSP (Supply Side Platform) integration with various SSPs (Rubicon, Mopub, Nexage, Appnexus, Facebook Exchange).
- Agency White Label Features (Multi Agency, Multi Account and Multi User management) with custom logo, custom css and custom cname domain URL configuration.
- User Retargeting using smart pixel. Storing and replicating userprofile at RealTime across servers using Kafka and Rocksdb.
- Audience Targeting Feature integrated with various Audience data providers (Semcasting, Bluekai).
- Real Time Data Processing and Analytics using Zookeeper quorum, Distributed kafka, Distributed HBase and Multiple Reporters.
- PreProcessing Large data sets with apache spark for machine learning.

Strengths

- Adaptable
- Good Team Player
- Cooperative & Open Minded
- Good Communication Skills
- Versatile
- Quick Learner

I, Vishnu Prasad S , hereby declare that the particulars furnished above are true, correct and complete to the best of my knowledge.

Place: Bangalore
Date : 7-Aug-2018 **Vishnu Prasad S**