

Gautam Singaraju

Engineering Leader, Product Builder, Machine Learning, Ph.D.

(925) 384-7950

gautam@singaraju.com

<http://gautamsingaraju.github.io>

LinkedIn: [gautams](#)

Oracle Digital Assistant, CA — *Engineering Leader, Architect Machine Learning And Natural Language Understanding*

February 2016 - Present

Oracle Digital Assistant product helps customers build virtual agent assisted conversational interfaces to deflect call traffic resulting in lower total cost of ownership.

Leader building out Machine Learning and Natural Language Understanding Features architectures, engineering and mentoring team. Wrote code for core features: Intent classification (features Sentence Embeddings, ML classification models, Train-Serve architecture), Data Curation Tools and xAI (explainability). Optimized the Machine Learning Models for Fast Inference with Minimal dependency (Apache Spark, Java-C++ JNI bridge). Explored Knowledge Graph Embeddings for improved classification. Software and data compliance and reviewing product user guide for the features built. Involved in computer vision stack.

Oracle Digital Assistant's customer references:

<http://www.oracle.com/us/technologies/mobile/mobile-customer-reference-lookbook-3634213.pdf>

GE Digital, CA — *Senior Architect*

October 2014 - January 2016

GE Digital Platform provided analytics solutions across GE. As an architect working for Power & Water and Oil & Gas teams, designed systems for ingesting, analyzing asset time-series data. The analytics identify generator failure detection from sensor data.

Electronic Arts, CA — *Technical Lead*

March 2012 - October 2014

Electronic Arts Data Platform provides an unified analytics platform across Video Game Studios for large scale of gaming data.

First engineer and later team lead, built the multi-tenant platform and analytics to multi-million dollar games (FIFA, NFL, NHL, Battlefield); running into PetaByte scales.

SKILLS

Product Engineering Leader, Machine Learning, Cloud and Distributed Architecture, Business Analytics

PATENTS

Justia Patents list:

<https://patents.justia.com/search?q=gautam+singaraju>

EDUCATION

UNC Charlotte, Ph.D.

2003 - 2009

Dissertation: Towards Sender Accountability on Email Infrastructure using Sender Identity and Reputation Management.

UNC Charlotte, MS

2001 - 2003

Thesis: Testbed for evaluation of intrusion detection systems.

JNTU, Bachelors

2001 - 2003

Project: Signature Analysis Using Neural Networks.

Ask.com, CA— Senior Engineer

June 2010 - March 2012

Ask Data Team provided analytics for Ask Search engine.

Built analytics for identifying the top queries at the search properties across different channels and countries, Search Engine Monetization.

Built analytics to report top queries given a keyword using HBase.

UNC Charlotte, NC— Postdoctoral Researcher

August 2009 - May 2010

Assessing Botnet Population Dynamics and Relationships. Profiling Malicious Domain Behavior From Top-Level Domain Data.

Technical Staff Intern, CA— Postdoctoral Researcher

May 2007 - August 2007

Performance analysis of iSCSI on VMware ESX Servers. Worked on Disaster Recovery performance scenarios with VMware ESX server.

UNC Charlotte, NC - Graduate Assistant

Jan 2002 - May 2009

Research: Email Security, Data Synchronization techniques, Applied Cryptography, Intrusion Detection Networks, Honeypots.

Taught graduate level courses on Information Security.

Selected Publications

Entropy-Based Measurement of IP Address Inflation in the Waledac Botnet, FloCon 2011.

Concord: A Secure Mobile Data Authorization Framework for Regulatory Compliance, USENIX LISA November 1, 2008.

Tracking Email Reputation for Authenticated Sender Identities, CEAS, April 1, 2008.

RepuScore: Collaborative Reputation Management Framework for Email Infrastructure, USENIX LISA 2007.

A Testbed for Quantitative Assessment of Intrusion Detection Systems Using Fuzzy Logic, IWIA, January 1, 2004.

SOFTWARE

Java, Apache Spark ML, Kafka, TensorFlow, Python, Kubernetes, Big Data.

Selected Blog

[Introduction to Embedding in Natural Language Processing](#)

[Caffe Deep Learning: Dog Breed Alexnet Model](#)

[Object Classification using Deep Learning and Raspberry PI](#)

Volunteer

System Administration for Non-profit organizations

Researcher with George Mason University

1. Botnet population dynamics
2. Profiling malicious domain names from top level domain registration data.

Program Committee Member for International Conference on Operations Research and Enterprise Systems 2017