

# Sankruth Kota

732-422-0689 | [sankruth.kota@gmail.com](mailto:sankruth.kota@gmail.com) | /in/skota121 | github.com/skota121 | skota121.github.io

## EDUCATION

---

**University of Illinois at Urbana-Champaign**

Aug 2017 - May 2021

*Bachelor of Science in Computer Engineering*

- **Relevant Coursework:** Applied Parallel Programming, Operating Systems, Distributed Systems, Applied Machine Learning, Computer Security, Database Systems, Digital Systems (FPGA)

## PROFESSIONAL EXPERIENCE

---

**Oracle**

July 2021 - Present

*Software Engineer (Performance, Scalability and Reliability Team)*

- Improved the performance of Oracle Analytics for Applications (OAX) by bundling groups of requests to be sent out as batches
- Optimized a set of benchmark query execution times (TPCDS) against Oracle Big Data Services clusters on Hive, and Spark by developing a set of optimal tunings to be deployed in future releases as out-of-the-box
- Developed a script to collect metrics such as CPU utilization, disk I/O etc and display them on a dashboard

**Cargill**

Jan 2021 - May 2021

*Software Engineering Intern*

- Built a dashboard for the quant team in **React** and **MongoDB** and deployed this application through **Docker**
- Set up **socket data streaming** to simulate the appearance of real-time data to test our dashboards
- Worked with UI/UX team members to ensure that the dashboard delivered the correct experience for the team

**Oracle**

May 2020 - Aug 2020

*Software Engineering Intern (Performance, Scalability and Reliability Team)*

- Conducted client-side and server-side breakdown of Oracle Analytics for Applications (OAX) to recognize performance and scalability bottlenecks
- Improved OAX load time by 1.8x (5s to 3s) by loading sets of KPI cards sequentially
- Performed scalability tests using **JMeter**, after deploying the improved application on a **WebLogic Server**, to guarantee proportional resource consumption

**Capital One**

*Software Engineering Intern*

May 2019 - Aug 2019

- Focused on detecting changes in data that could prevent machine learning models from making accurate predictions and polluting data (data drift/concept drift)
- Tracked model decay of a **support vector classification** model for data drift and implemented automated model adaptation features to maintain high prediction accuracy
- Developed an ensemble machine learning pipeline, combining model adaptation and statistical detection, using **scikit-learn** to keep machine learning models robust

**John Deere, Intelligent Solutions Group**

*Information Technology Intern*

May 2018 - Aug 2018

- Added the ability to log contextual metadata asynchronously, by implementing **thread-local storage** based on chains of callbacks in **Node.js**
- Implemented a cleansing process by parsing and converting logs to an acceptable format for long-term storage in **Elasticsearch**

## SKILLS

---

**Languages:** Python, Golang, C/C++, Javascript (Node.js), Java, SQL, HTML/CSS, CUDA, System Verilog

**Tools:** Hadoop, Spark, MongoDB, Neo4j, SQL, Data/ML Python libraries, JMeter, Elasticsearch, Logstash, Kibana, Django, Git, Jenkins Pipeline, TDD/BDD