VIKRAM SUBRAMANIAN

vnsubram@uwaterloo.ca
vikramsubramanian.github.io/
vikram-n-subramanian/
vikramsubramanian

SKILLS

Languages: Python, C++, SQL, Bash, Learning: Java & Go

Technologies: Docker, AWS, Git, Jenkins, Flask, RESTful Services, Microsoft SQL Server, MySQL, MongoDB, NumPy, Pandas, Elasticsearch, Logstash, Kibana & Grafana, Linux, Jira, Agile Development

INTERNSHIPS

SOFTWARE ENGINEERING INTERN — Santa Clara, CA (remote) · May 2021 to Aug. 2021
Nvidia

- Built an extensive data pipeline using Python, Beats (Golang), Logstash (Ruby), Elasticsearch and Kibana to automatically detect new logs produced by deep learning models, process them, ship them from customer machines, upload them to Elasticsearch and automatically create Kibana visualizations.
- Built a custom Elastic Beat (a data shipper from Elasticsearch) to process Nvidia's native DL-Log format using Go.

SOFTWARE ENGINEERING INTERN — Ottawa, ON (remote) · Sep. 2020 to Dec. 2020
Wind River

- Used data provided by the GNU Linker to map dependencies between different files in a Make build and create a stack trace of calls to find serialization points and optimizations in a build.
- Created scripts to automatically convert GNU Make rules to Ninja (Google Chromium's build system) build rules using data captured from the LD_PRELOAD mechanism of the GNU Linker.
- Built a data pipeline to collect, process and visualize dependency data from different Yocto Project (an embedded Linux distribution) builds using MongoDB, Python and Flask in an attempt to find bottlenecks and differences between builds.

SOFTWARE ENGINEERING RESEARCH INTERN (PART-TIME) — Waterloo, ON · May 2020 to Aug. 2020
BlackBerry Research

- Lead a team of 3 students to build a tool in Python that can automatically apply code patches to forks/modified versions of code- a task where the 'git apply' command usually fails (as it searches for a verbatim match).
- Developed a novel method using fuzzy searching and program slicing to identify the same code block in 2 different versions of the same code and achieved a ~31% increase in apply percentage compared to the 'git apply' command.
- Main author of research paper to be submitted to the 2022 ACM/IEEE International Conference on Software Engineering- bit.ly/3nURxhm

SOFTWARE ENGINEERING INTERN — Waterloo, ON · Jan. 2020 to Apr. 2020
Thomson Reuters

- Designed and developed a REST API backend for an NLP model using Flask, Docker, SQLite and AWS. Demoed the product to the customer (a 6 person internal team) and provided support till the end of my tenure.
- Built a multithreaded data pipeline in Python to extract over 25 million rows of legal data from a customer's Microsoft SQL Server, run Flair and Spacy NLP models to remove personal information and then visualize them using Pandas.

DEVOPS ENGINEERING INTERN — Waterloo, ON · May 2019 to Aug. 2019
Sandvine

- Built a pre-merge CI pipeline for 4 repositories to run a battery of tests affecting 50+ commits/week using Python, Groovy, Jenkins, JIRA API and GitLab API.
- Created Python scripts to generate docker-compose files from configuration settings provided by Sandvine's native build system and substantially reduced the number of instructions the user had to provide.

RESEARCH PUBLICATIONS AND AWARDS

Research Publications

I have been working with Prof. Mei Nagappan at SWAG Lab, University of Waterloo and have worked on multiple projects-

- Main author- Analyzing first contributions on GitHub: what do newcomers do. Published by IEEE Software journal- bit.ly/2M3QH3C
- Main author- An empirical study of the first contributions of developers to open source projects on GitHub. Published in The 2020 International Conference on Software Engineering- bit.ly/38Raxbq
- Second author- A Curated Archive for Software Engineering Research Tools. Published in SIGSOFT Softw. Eng. Notes- bit.ly/3bNOJzn

Winner of the 2020 ACM Microsoft Student Research Competition for my work titled "An empirical study of the first contributions of developers to open source projects on GitHub" at ICSE2020- The most premier conference for software engineering research- bit.ly/3oUXP14

Overall winner at Hack the North 2019 (1500+ participants) for our project titled 'Complete'- bit.ly/3BjUsal

EDUCATION

University of Waterloo — 2018 to 2023
Software Engineering, Honours
Selected Coursework: Operating Systems, Algorithms and Data Structures, Database management, Sequential Programs (Compilers), Software Engineering Principles (OOD), UI/UX Design (Java), Data Abstraction and Implementation (C++, OOP), Programming Principles (C).