

ANDREW R. SUMNER

Laurel, MD 20707 ♦ 301-741-2262 ♦ andrew.r.sumner@gmail.com ♦ [linkedin.com/in/andrew-sumner/](https://www.linkedin.com/in/andrew-sumner/) ♦ andrewsumner.com

► EDUCATION

University of Maryland, College Park

B.S., Computer Science | B.S., Physics | Minor in Technology Entrepreneurship

Graduated May 2021

GPA: 3.52

► SOFTWARE ENGINEERING SKILLS

Programming Languages: Python, Typescript, Javascript, SQL, Go, HTML/CSS, Bash, Ruby

Frameworks/Libraries: Node.js, Next.js, React, NumPy, Pandas, SciPy, Tensorflow, Keras, PyTorch, Scikit-learn, PySpark

Tools: Github Copilot, AWS (Lambda, S3, Kinesis, DynamoDB, SQS, RDS, Step Functions, EMR, Cloudwatch), Github, Git, New Relic, Snowflake, Elasticsearch, PagerDuty, Confluence, Jira, Splunk, Jenkins, Docker, MySQL, Linux, Postman, PostgreSQL, Supabase, Vercel, Tailwind CSS, Selenium, Grafana, Datadog

Miscellaneous: Agile/Kanban, microservices, back-end development, data processing, REST APIs, GraphQL APIs

Notable Certifications: AWS Cloud Practitioner

► EXPERIENCE

Campbell & Company, Quantitative Software Engineer

December 2024 – Present

- Automated simulation and reporting processes using Python and SQL for assessing quantitative investing risk target parameters, reducing the time spent by quantitative researchers on repetitive tasks by 60%
- Contributed to AI working group by testing internally hosted LLM for code generation using various models

Capital One, Senior Associate Software Engineer

August 2021 – October 2024

Team #2 (August 2022 – October 2024)

- Developed a real-time marketing decisioning platform for assigning customers to best-fit marketing campaigns. Grew this platform from early concept stage to production deployment, now used with 11+ large-scale campaigns targeting millions of customers/accounts. Daily throughput has exceeded 100 million records and is rapidly scaling.
- Built 6 lambdas (Typescript, Python), DynamoDB, S3, SQS, Kinesis, and other AWS components as part of this platform. Our infrastructure writes downstream to Snowflake and Salesforce Marketing Cloud.
- Utilized Github, Splunk, New Relic, Snowflake, ServiceNow, PagerDuty, Jenkins, Confluence, Jira, Github Copilot
- Maintained other applications that assembled marketing customer profiles, facilitated welcome email delivery, and delivered marketing data. These apps utilized Python, PySpark, AWS (Lambda, EMR, ECS), Java, Docker

Team #1 (August 2021 – August 2022)

- Contributed extensively to development, first production deployment, and successful customer pilot of self-service cashier's check kiosk for future use in 300+ bank branch and café locations, serving over 100 million customers
- Implemented backend REST APIs (AWS lambda, Node.js, Javascript), middleware to interface between backend APIs and kiosk hardware (Node.js, Typescript), frontend mobile webapp (React), database (MySQL, RDS), and deployed these components on AWS infrastructure via Jenkins
- Achieved 99.999% reliability via testing of prototype ATM replacement kiosk, and created backend Python lambda

Capital One, Technology Incubator Intern

September 2020 – December 2020

- Improved machine learning model accuracy by 10% for speech recognition tool used by company helplines
- Developed bash scripts and set up Docker container for containerized speech recognition application

Capital One, Software Engineering Intern

June 2020 – August 2020

- Collaborated with fraud investigators to build a full-stack application for investigating suspicious card transactions
- Developed GraphQL APIs using Go and AWS Lambda to load/search transaction data from Elasticsearch database
- Implemented filtering, sorting, and searching functionality within Angular UI to quickly discover transactions

► PROJECTS

ParkScout.io, a mapping tool for planning trips to National Parks

November 2024 – Present

- Developed and deployed web app using React, Next.js, Tailwind CSS, Leaflet, Vercel, and Supabase/PostgreSQL
- Scraped and tabulated national park data (including GIS data) using Selenium and pandas
- Generated park descriptions with LLMs by feeding scraped data to the OpenAI API

Predicting Land Use from Satellite Imagery using TensorFlow

May 2021

- Collected satellite image data using Google Earth API, labeled images according to land use categories, trained a convolutional neural network in Tensorflow to predict the classifications, and achieved 85% accuracy on test subset

Panorama Image Stitching Pipeline

April 2021

- Assembled panorama image pipeline in python, implementing corner detection, adaptive non-maximal suppression of corners, SIFT algorithm, homography feature matching, RANSAC algorithm, and image warping

Classifying Galaxy Imagery using PyTorch

April 2021

- Built convolutional neural network with PyTorch for classifying 18,000 galaxy images, achieving 70% accuracy

Data Analysis of National Park Visitation (found at <http://asumner1.github.io>)

May 2019 – July 2019

- Scraped and analyzed data in python to understand national park visitation and identify undervalued parks