

PATRICK JOHNSON

patrick@johnsonfamily.me • patrick.johnsonfamily.me

SUMMARY

Software engineer with nearly two decades of experience. Primary experience includes highly scalable and reliable services, infrastructure, system architecture, and automation. Currently working at the intersection of geospatial technologies, machine learning, and environmental impact.

EXPERIENCE

Allen Institute for AI

September 2021 – Present

Senior Software Engineer

- Pioneered the software engineering practice on the [OlmoEarth](#) team. OlmoEarth is an end-to-end platform for scalable planetary intelligence that turns Earth data into decision-ready insights. Work includes distributed data pipelines, STAC-based satellite imagery services, and geospatial ML infrastructure.
 - Technical lead for OlmoEarth Studio, where users annotate data, train models, and run inference.
 - Led design and implementation for computing [custom geospatial embeddings](#).
 - Architected a dataset acquisition service with a unified metadata index across multiple satellite data providers, replacing external API dependencies to improve reliability, performance, and multi-source data access.
- First author on research into high-resolution live fuel moisture content (LFMC) maps for wildfire risk using multimodal earth observation data. Accepted at TerraBytes 2025 and published in PMLR.
- Built AI tools for [wildfire resiliency](#) from the ground up and drove partnerships with research and government collaborators.

Google

July 2007 – May 2021

Senior Software Engineer

- Technical Lead for software to maintain high-quality, accurate models of Google's datacenters and make the data accessible, consistent, and useful.
 - Services are used by planning and execution teams to facilitate efficient management of compute, storage, and network capacity in Google's datacenters.
- Developed backend features in *Google Contacts*, which allows users to intelligently manage their contacts, and *People API*, which provides authenticated developer access to contacts and profile information. Focused on functionality for enterprise users.
 - Improved stability and latency of the *G Suite* directory autocomplete and search system.
 - Co-led a team of six engineers to build an enterprise feature that enables custom directories, allowing administrators flexible control over directory visibility within their organizations.
 - Built a service that provides duplicate contact suggestions and merge functionality.
- Architected automation and tooling for the *Borg* cluster management system, which handles hundreds of thousands of jobs from thousands of applications across Google.
 - Made significant technical improvements to the accuracy and scalability of simulation tooling. This is used to measure system performance under real-world load.
 - Designed and developed a widely adopted test framework for writing system tests.
- Built test infrastructure for the *Caffeine* web indexing system, which when launched in 2010 provided 50% fresher results for web searches than the previous web search index.
- Developed test infrastructure for the *Google Chrome / Chromium* web browser during its pre-launch and early post-launch stages. The browser launched publicly in 2008. Focused on distributed reliability testing: automated crash testing and stack trace analysis by loading millions of popular web pages against new browser builds.
- Hosted two interns and supervised their summer projects.

Parity Computing

December 2005 – July 2007

Software Developer

- Developed text processing software for mining and analyzing scholarly publications, including finding inter-document links and relationships using citations and references.

ISS EarthKAM

June 2006 – September 2006

Software Developer Intern

- Developed software to assist with storing data from the EarthKAM program, allowing middle school students to take photos of specific places on Earth from the International Space Station. Used vector data

from the Johnson Space Center for image targeting.

Teradata, a division of NCR Corporation

June 2005 – December 2005

Software Engineering Intern

- Developed client software for a large-scale data warehouse.

VOLUNTEERING

TEALS

September 2015 – June 2022

Volunteer Teacher / TA

- The goal of the TEALS program is to connect classroom teachers with tech-industry workers to create sustainable Computer Science programs.
- As part of the TEALS program, served as a volunteer teacher and teaching assistant at high schools in Redmond, WA and Easton, WA.
- Classes taught include: Intro to Computer Science, AP Computer Science A, Advanced Projects, and Data Structures and Algorithms.

SKILLS & PROFICIENCIES

- Nearly two decades of professional software engineering experience.
- Substantial experience with distributed systems and multithreaded services.
- Experience with geospatial ML infrastructure, satellite data pipelines, and earth observation platforms.
- Adept at working in and debugging large, complex codebases.
- Effective team collaborator. Able to work with others to set priorities and achieve common goals.

PUBLICATIONS

- High-Resolution LFMC Maps for Wildfire Risk From Multimodal Earth Observation Data. [TerraBytes 2025](#).
- OlmoEarth: Stable Latent Image Modeling for Multimodal Earth Observation. [arXiv 2025](#).

EDUCATION

University of Illinois Urbana-Champaign

Urbana, IL (Online)

- Master of Science in CyberGIS and Geospatial Data Science 2023 – 2025
- Graduate-Level Certificate in Weather and Climate Risk and Data Analytics 2022

University of California San Diego

La Jolla, CA

- Bachelor of Science in Computer Science, *summa cum laude* 2003 – 2007

PROFESSIONAL DEVELOPMENT

Machine Learning, Stanford Online on Coursera

2020

- 11-week course covering the foundations of machine learning, data mining, and statistical pattern recognition.