

RISHAVPREET SINGH

New Grad Software Engineer | Data & AI Products, Backend Systems, AWS

+1(604)771-5800 | rishavpreetsingh30@gmail.com | [linkedin.com/in/rpss30](https://www.linkedin.com/in/rpss30) | rishavpreetsingh.site | github.com/rpss30

EDUCATION

University of British Columbia

Sept 2021 – May 2026

Bachelor of Science | Computer Science (Minor Data Science)

Vancouver, BC

TECHNICAL SKILLS

Languages:	Python, SQL, Java, C, C++, TypeScript/JavaScript, R
Data & AI:	Pandas, NumPy, SciPy, scikit-learn, PyTorch, PySpark, Groq APIs, Altair, Plotly
Testing & Tools:	Pytest, JUnit, Mocha, Chai, Git/GitHub, Linux, GDB, Valgrind, Test-Driven Development
Backend & Web:	FastAPI, Flask, SQLAlchemy, PostgreSQL, React.js, Node.js, Streamlit, Tailwind CSS, REST APIs
Cloud & DevOps:	AWS (Lambda, S3, EC2, ECR, ECS Fargate, RDS, ALB, API Gateway, CloudFront, Cognito, Secrets Manager, CloudWatch), Docker, Docker Compose, GitHub Actions

WORK EXPERIENCE

Aeroqube

April – Aug, 2025

Software Engineer Intern | Cloud, AI Products & Web Applications

Vancouver, BC

- Deployed and maintained production web applications on AWS (S3, EC2, Lambda, CloudFront) serving a live AI healthcare product, with zero downtime across the 5-month internship.
- Built a React.js + Tailwind landing page for an AI clinical assistant and wired a SendGrid email workflow through AWS Lambda, converting form submissions into automated lead-capture emails without a backend server.
- Cut admin console load times by 70% by replacing full-table fetches with API-based pagination and restructuring third-party API calls, going from roughly 8s to under 2.5s on the main dashboard.
- Tightened the AI chatbot UX for healthcare professionals by reworking chat history storage, response state flow, and error states, which reduced support tickets about conversation drop-offs during the final sprint.

TECHNICAL PROJECTS

[JobLens AI](#) | [Demo](#) | Python, Pandas, scikit-learn, FastAPI, PostgreSQL, Docker, AWS

Apr – Jun, 2026

- Built a Canadian job-market intelligence platform that pulls first-party postings from Greenhouse, Lever, and Ashby, runs Groq-powered skill extraction on each posting, and refreshes the dataset automatically every week.
- Designed a role-fit scoring engine using TF-IDF similarity and role-specific weighting across various technical categories (backend, data, DevOps, etc.) so candidates see a numeric fit score and ranked list of missing skills, not a black-box match.
- Deployed a Streamlit dashboard backed by FastAPI and PostgreSQL with free-text search, CSV upload, saved analyses, and one-click Markdown/PDF report export.
- Containerized the full stack with Docker, deployed to AWS using ECR, ECS Fargate, ALB, private RDS, and Secrets Manager, and kept 195 Pytest tests green in CI throughout.

[Adventure of the Ages](#) | C++, OpenGL, GLFW, GLM, SDL, FreeType, GLSL

Jan – Apr, 2026

- Built a 2D side-scrolling platformer in a 4-person team, shipping multiple playable levels and themes within a single semester with a fully modular rendering and gameplay architecture.
- Owned camera control, shader-driven effects, real-time input handling, and state-based UI feedback, covering roughly a third of the engine's core systems.
- Tracked down a major FPS drop in tutorial and UI-heavy scenes using Valgrind memory analysis and caching optimizations in the text-rendering path, resolved before the final demo.
- Designed reusable content pipelines so new levels and themes could be added without touching the core engine, cutting integration time for new assets across the team.

[Graider \(Hackathon\)](#) | Python, Reflex, Generative AI, Groq APIs, Llama 3.2, OOP

Oct 18 – 20, 2024

- Built a full-stack AI grading tool in 48 hours with a Reflex front end and an OOP backend managing students, rubrics, question parts, and reference answers.
- Integrated Groq + Llama 3.2 to auto-generate question breakdowns, reference answers, and per-criterion rubrics, all editable through the UI before finalizing.
- Implemented per-part answer segregation and rubric-driven evaluation so AI feedback was granular enough for graders to approve or override each criterion individually, not just a single score.