

# Justin Kong

201-681-7507 | [jhk174@scarletmail.rutgers.edu](mailto:jhk174@scarletmail.rutgers.edu) | [linkedin.com/in/justin-hkong](https://www.linkedin.com/in/justin-hkong) | [github.com/JustinK33](https://github.com/JustinK33) | [justinkong.app](https://justinkong.app)

## EDUCATION

---

### Rutgers University

Bachelor of Science in Computer Science

**Relevant Coursework:** Data Structures & Algorithms, Object-Oriented Programming, Discrete Mathematics, Linear Algebra, Data Management

New Brunswick, NJ

Expected May 2028

## EXPERIENCE

---

### Software Engineer Intern

Deepiri | Agile, Scrum (Plaky)

Nov 2025 – Present

Pittsburgh, PA (Remote)

- Scaled backend systems to **110+ msg/sec** at **10 ms p95** by tuning Kafka partition layout and async consumers.
- Reduced coordination overhead by **60%** through automating **30+** Python onboarding and PR workflows.
- Prevented downstream failures under traffic spikes by implementing **Redis rate limiting** across **microservices**.
- Reduced producer errors to **zero** in production by building a **Prometheus** and **Grafana** observability stack.
- Cut retry storms by **97%** in distributed systems using **exponential backoff** and **circuit breakers**.

### AI/ML Engineer Extern

Outamation

Oct 2025 – Dec 2025

Southlake, TX (Remote)

- Parsed **200+** page mortgage PDFs reliably by combining **PyMuPDF** extraction with **Tesseract OCR** fallback.
- Indexed **2,000+ chunks per file** by building a chunked retrieval pipeline for large-scale PDF ingestion.
- Boosted **LlamaIndex** retrieval by integrating **SentenceTransformer** embeddings with FAISS vector indexing

## PROJECTS

---

**NoteTube** | Python, Django, Redis, AWS, Nginx, Docker, LangChain, gRPC

Dec 2025 – Present

- Achieved **19%** p95 improvement by deploying **Django** to **AWS EC2** behind an **Nginx** reverse proxy with **TLS**.
- Decoupled note generation via **gRPC unary** calls, enabling independent scaling without monolithic redeployment.
- Cut API latency by **86%** (**313ms** → **44ms**) by caching high-traffic user and session queries in **Redis**.
- Accelerated semantic search by **90%** by building a **LangChain RAG pipeline** with **PGVector**.

**LinkNest** | Ruby, Rails, SQL, MySQL, AWS, Active Storage, Docker

Feb 2026 – Present

- Secured multi-tenant Rails profiles with slug routing and **AWS RDS MySQL** scoping for isolated URLs.
- Cut profile load times by **50%** resolving N+1 query chains across tenant-scoped Rails associations.
- Boosted analytics job performance by **90%** by replacing per-row queries with bulk database operations.

**DailyNode** | JavaScript, Next.js, Discord.js, Google Cloud, GitHub Actions, Docker

Feb 2026 – Mar 2026

- Deployed **Dockerized GCP** service with **CI/CD** and resilient persistence for reliable zero-downtime production.
- Scaled to **1,000+** users by building an event-driven Discord bot with per-guild isolation and job scheduling.

**RUMyValentine** | Python, FastAPI, SQLAlchemy, PostgreSQL, Pydantic

Oct 2025 – Feb 2026

- Delivered sub-**400ms** latency on REST APIs for a 70 dimensional weighted-vector matcher under production load.
- Maintained **99.9%** uptime for **2,000+** users and **10,000+** matches through scalable backend architecture.
- Secured auth with Google OAuth 2.0 SSO using OIDC, PKCE, ID token verification, and JWT protection.

## TECHNICAL SKILLS

---

**Languages:** Python, Java, TypeScript, JavaScript, SQL, C, Ruby, Bash

**Frameworks & Libraries:** FastAPI, Spring Boot, Django, Node.js, Express, Next.js, Flask, Ruby on Rails, gRPC

**Databases:** PostgreSQL, MySQL, Redis, MongoDB

**Cloud & Infrastructure:** AWS (EC2, RDS, VPC), Docker, Kafka, GCP, Nginx, Git, GitHub Actions, Linux

**Observability & Testing:** Prometheus, Grafana, pytest, JUnit, k6, Postman