

RAKSHITA DEVURKAR

Austin, TX | (404) 218-7171 | Rakshitadevurkar@gmail.com | [LinkedIn](#) | [GitHub](#)

SUMMARY

Senior Software Engineer with 8+ years building and operating large-scale software systems, from distributed backends and full-stack products to production agentic AI. Owns systems end to end from architecture through launch, with recent focus on LLM evaluation, AI-assisted software development, and developer productivity.

WORK EXPERIENCE

EXPEDIA GROUP, AUSTIN, TX - SOFTWARE DEVELOPMENT ENGINEER III

Aug 2022 - Present

- Designed reusable escalation flows for the agentic Help Center, preserving traveler and booking context across tools and enabling reliable human handoff across **2.5M+ monthly conversations**
- Built a multi-turn evaluation framework with deterministic evaluators & calibrated LLM judges, raising tool-selection accuracy from **0.67 to 0.91** and achieving over **90% agreement with human experts** across **1,400 labeled samples**
- Led engineering delivery of a hotel-disruption servicing flow that drove a 66% increase in successful traveler self-service and an **80% increase in completed relocations** within five days of launch
- Redesigned support handoff around persisted traveler context and explicit escalation reasons, reducing human agent handling time by **52%** and delivering **\$3.5M in annualized savings**
- Scaled full-stack Virtual Agent workflows for changes, cancellations, and refunds across chat and voice, resolving **850K traveler requests** monthly through self-service
- Identified test-booking generation as a recurring bottleneck through engineer interviews; shipped a CLI prototype in 4 weeks. Productized booking APIs as versioned MCP tools; adopted by **700 engineers** org-wide
- Automated incident RCA generation with a Claude Code skill synthesizing tickets, Slack threads, and engineering context, cutting preparation time by **over 96%** from 4 hours to under 10 minutes
- Mentored three engineers and led cross-team architecture reviews, establishing shared patterns for evaluation, rollout, observability, and failure recovery
- Prototyped an evaluation runner targeting **300K monthly trials** with immutable model, prompt, and data snapshots plus booking-state validation for reproducible release checks
- Cut median PR lead time by **35%** across 8 teams by standardizing AI-assisted development workflows and CI checks
- Optimized Claude Code context management using hooks, subagents, and reusable instructions, reducing retained context by **30%** across benchmark workflows

MERIDIANLINK, BOSTON, MA - SOFTWARE ENGINEER

Nov 2018 - Jul 2022

- Architected a shared-UI framework that replaced a third-party dependency and standardized frontend development
- Automated SQL Server ingestion across 50 accounts, reducing manual import effort by **80%**

ONJAX TECHNOLOGIES, APALACHIN, NY – WEB DEVELOPMENT INTERN

Jun 2017 - Aug 2017

- Developed a full-stack form builder for real estate agents, reducing manual workflow configuration by **68%**

METRICSTREAM, BANGALORE, INDIA – MEMBER OF TECHNICAL STAFF

Nov 2015 - Jul 2016

- Built a RabbitMQ/SES notification service processing **50K+ monthly** emails, cutting backend load by **40%**

TECHNICAL SKILLS

Agentic AI	LangGraph, LangChain, RAG, MCP, AWS AgentCore, Prompt Engineering
AI Evaluation	DeepEval, LangSmith, LLM Judges, Multi-turn Evaluation, Guardrails
Languages & Frameworks	Python, Java, Kotlin, TypeScript, JavaScript, C#, Spring Boot, FastAPI, React, Node.js
Backend & Distributed Systems	GraphQL, REST, gRPC, Microservices, Distributed Systems, Kafka, RabbitMQ, SQL
Cloud, Infra & Reliability	AWS Bedrock, Docker, Kubernetes, CI/CD, S3, SQS, SNS, Datadog, Splunk, OpenTelemetry

OPEN-SOURCE PROJECTS

[Evalkeep](#) | Converts production agent failures into regression tests and tracks quality across versions

EDUCATION

STATE UNIVERSITY OF NEW YORK, BINGHAMTON, NY

Aug 2016 - May 2018

Master of Science in Computer Science

RNS INSTITUTE OF TECHNOLOGY, VTU, INDIA

Sep 2011 – Jun 2015

Bachelor of Engineering in Information Science