PALASH THAKUR

+1 343 553 3441 ◇ Ottawa, ON, Canada ◇ Available Sep – Dec 2026 (Co-op)
palasht75@gmail.com ◇ LinkedIn Profile ◇ Personal Website

EDUCATION

University of Ottawa, Ottawa, ON, Canada *Sep 2025 – Apr 2027*
MEng, Electrical and Computer Engineering (ECE); CGPA: **9.60/10.00**
Design of Scalable Intelligent Systems.

Shri Ramdeobaba College of Engineering and Management, Nagpur, India *Aug 2018 – May 2022*
Bachelor of Engineering in Electronics Engineering; Minor in Computer Science
CGPA: **9.30/10.00**; Minor SGPA: **9.00/10.00**

TECHNICAL SKILLS

Languages: Python, SQL, Bash
Backend & APIs: Splunk, FastAPI, SQLAlchemy, REST APIs, Microservices concepts
Cloud & Containers: AWS (EC2, RDS, S3, Lambda), Docker, Kubernetes, Helm, Terraform, Nginx
Linux & Tooling: Linux, Git, CI/CD (GitHub Actions, GitLab CI, Jenkins)
Testing: Pytest, **Unittest**, Integration testing, Load testing
Datastores: PostgreSQL, MongoDB, Redis
AI & ML: LLMs, RAG, FAISS, Agentic workflows, NLP, Prompt engineering
Security Concepts: IAM, authentication vs authorization, secure secrets handling, TLS, hashing, network security fundamentals
Other: React, Streamlit, Graphviz

PROFESSIONAL EXPERIENCE

Nokia, Ottawa, ON, Canada *May 2026 – Present*

Security Software Developer Co-op

- Developed a lightweight In-Service Upgrade (ISU) and rollback implementation to replace a legacy heavyweight upgrade mechanism, reducing upgrade-time resource usage by **35%**, improving upgrade execution speed by **40%**, and cutting rollback recovery time by **50%**.
- Built AI-based workflow automations for security software teams, reducing repetitive triage, investigation, and engineering support effort by **30–40%** across recurring internal workflows.
- Developed an AI assistant that connects knowledge across JIRA, Confluence, internal codebases, and documentation, reducing legacy-code investigation time by **25%+** for debugging and onboarding use cases.
- Assessed vulnerabilities in existing software by reviewing security findings, affected code paths, exploitability, and remediation options across **multiple product components**.
- Created reusable agentic AI skills for workflow automation, knowledge retrieval, and developer assistance, accelerating reuse across **5+ security engineering workflows**.

Persistent Systems Ltd., Maharashtra, India *Dec 2021 – Aug 2025*

Senior Software Engineer *Clients: Intuit, Connxus HIE, Harvard Medical School*

- Led a team of **3 engineers** delivering backend services for **QuickBooks Finance Agent** used by **1M+ customers**.
- Built Python REST APIs, FastAPI services, and AWS data workflows supporting finance, healthcare, and AI-assist use cases across **3 client domains**.
- Optimized high-traffic REST endpoints, reducing average response time by **66%** for **1M+ QuickBooks Pro users**.
- Delivered scalable FastAPI services connecting React UI to LLM endpoints for an AI-assist chatbot handling about **10k calls/hr**.
- Increased automated test coverage to **99%** with Pytest and unit tests, reducing post-release defects by **70%**.

- Improved PostgreSQL performance by **40%** through query optimization, indexing, and execution plan analysis.
- Served as primary DBA point of contact and contributed to AWS HealthLake workflows for cloud-based health data processing.
- Built NLP, UMLS, and ML components for EMR extraction and clinical research workflows with validation and error handling.
- Developed image enhancement and survival analysis tools supporting reproducible oncology research workflows.

SELECTED PROJECTS

LeaseCheck: Multi-Shot RAG Clause Auditor for Canadian Leases *2025*

- Built a GPT-4o mini + FAISS RAG system to flag illegal or unfair residential lease clauses with cited statute snippets, focusing on explainability and traceability.
- Implemented function-calling logic to decide when to invoke a legality checker combining regex shortcuts with semantic search.
- Shipped a Streamlit chat UI and a scheduled GitHub Action to rebuild the vector index from updated statutes.
- **URL:** https://github.com/palasht75/LeaseCheck-Multi-Shot-RAG

Explainable LLM Evidence Summarizer *2025*

- Created an audit-friendly summarization toolkit that outputs bullet points with sentence-level evidence IDs for traceability.
- Supported dual backends, cloud and local, exposing CLI, Python API, and Streamlit demo for integration across environments.
- **URL:** https://github.com/palasht75/llm-evidence-summarizer

OPEN-SOURCE CONTRIBUTIONS

Keycloak Identity and Access Management Platform *2026*

- Contributed upstream to Keycloak IAM through **3 merged or submitted PRs**, improving OIDC token introspection client-authentication error handling, Admin Console timestamp display, and HMAC terminology.
- **Links:** PR #48945 — PR #49031 — PR #48988

UMLS Python Client Published: *Sep 2024*

- Developed and published umls-python-client, a Python package that simplifies integration with UMLS REST APIs through a modular, developer-friendly interface.
- Featured on the National Library of Medicine UMLS community resources page: NLM UMLS Community.
- **Links:** GitHub Repository — Documentation

HONORS, PUBLICATIONS, AND LEADERSHIP

- **Winner, Google Cloud Hackathon**: Won first place in the Google Cloud Hackathon; invited to record a winning video with **7,800+ views** on YouTube.
- **Google Cloud APAC Speaker**: Presented as a young developer at HumBanayenge, hosted by Google Cloud India.
- **Best Paper Award, ICESC-2020**: Recognized for research on traffic sign detection for driverless cars.
- **Academic Publications**: Published **2 peer-reviewed papers** on traffic sign detection and emotion recognition, including work in *BioScience Biotechnology Research Communications* and Springer *ICMEET-2023*. Paper 1 — Paper 2.
- **Mentorship and Community Leadership**: Mentored **8K+ students** through Google Cloud Developer Community outreach, **200+ undergraduates** through RCOEM technical training and mock interviews, and new joiners at Persistent Systems on backend engineering, testing, and delivery workflows.