

Sidharth Rajmohan

India | +91-8921-628177 | dumko.raj@gmail.com
github.com/dumko2001 | linkedin.com/in/sidharth-rajmohan

EXPERIENCE

Polymath

Jun 2026 – Present

Software Engineer (Contract) | Remote

- The take-home assignment was my first experience building RL environments. Built 15 environments in less than 2 months, across cybersecurity, distributed systems and other specialized domains.

Displays and Graphics Ltd

Apr 2026 – Present

Software Engineer (Contract) | Remote / UK

- Through Safe CLI, I was introduced to Prof. Rafal K. Mantiuk at Cambridge. With little graphics experience, I learned Vulkan and built Android code to control variable-rate shading and frame rate in games.
- Wrote repeatable test scripts and used AGI/Perfetto traces to compare rendering modes, verify 30/60 FPS pacing and measure GPU workload.

Safe CLI

Apr 2026 – Present

Software Engineer (Contract)

- I started contributing to open source for exposure. My autoresearch fix led Rushaan, an engineer at Meta, to reach out. With his overall direction, I built Safe CLI end to end: package-security tooling for npm and PyPI with command-line, SDK and agent integrations.
- Implemented the API and analysis workers for recursive dependency scanning and AI package inspection, with cached results for unchanged artifacts, incremental version scans and saved audit trails.

Independent

2025 – Present

Automation Engineer (Subcontract)

- Delivered subcontracted automation for Meteoric Boost and Isotta Homes, cutting onboarding manual work by 93% and removing a recurring 2-hour billing cycle by turning Google Sheets data into Xero invoices.

Celer AI

Mar 2025 – Sep 2025

Founder

- Founded Celer AI and built a voice-to-text workflow for clinicians using Next.js and Supabase, cutting medical-report generation from 15 minutes to 40 seconds.
- This was where I realized I wanted to help other people with technical work, which led me into independent automation.

OPEN SOURCE CONTRIBUTIONS

I work with AI agents to investigate code, review their changes and test fixes.

PyTorch

- Fixed reported CPU matrix-multiplication errors with fewer OpenMP threads, reworking the first failing patch into separate compute/reduction loops and isolating the regression test. [#169128](#)
- Reduced peak temporary memory for GPU norm/max reductions on mixed-size tensors by 96% (26.2 MB to 1.06 MB) in an RTX 4090 benchmark. Reworked data transfer after the first version slowed uniform-size inputs. [#190549](#)
- Combined 400 DDP gradient copy-out calls into one batch in a test model, keeping individual copies for higher-order gradients after review caught a correctness issue; fixed compilation of batch-normalization default arguments. [#190524](#) [#169126](#)

Harbor | Agent evaluation

- Prevented one judge answer from being reused across criteria by rejecting duplicate names; the maintainer moved validation before model calls. Also contributed root-level judge discovery to a broader maintainer update. [#2331](#) [#2334](#) [#2821](#)

Andrej Karpathy's autoresearch

- Exposed Python tracebacks on failed runs by making the agent read the last 50 log lines before attempting a fix; this contribution led to the Safe CLI collaboration. [#17](#)

NVIDIA NemoClaw

- Restricted authentication-file permissions. Review caught cases with no API key or symlinked agent folders; the merged fix added unconditional, symlink-aware checks. [#188](#)
- Restricted permissions on deployment secrets and migrated config files, and made onboarding wait for dashboard readiness; four fixes merged. Separate shell-injection analysis informed a maintainer's credited TypeScript rewrite. [#174](#) [#183](#) [#186](#) [#187](#) [#171](#) [#1886](#)

Google Workspace CLI

- Made token-storage failures visible and replaced blocking file operations after automated review. Created secret files with private permissions and kept Meet-link request IDs stable during retries. [#542](#) [#500](#) [#506](#)

LangChain & LangGraph

- Fixed a tool-runtime failure caused by checking parameterized types as ordinary classes. Added generic-origin checks confirmed by the maintainer; corrected Mistral model names, passed Groq structured-output settings through and removed an unsupported Ollama argument. [#6509](#) [#34046](#) [#34053](#) [#34114](#)

Prime Intellect verifiers

- Prevented response-parsing crashes when compatible providers omitted optional token-usage details. Replaced strict SDK types with local parsing models. [#2019](#)

ClawMetry

- Reported and patched login-token XSS and cron-log path traversal; a maintainer reapplied both fixes after routes moved, with credit. Restored foreground sync. [#1072](#) [#1074](#) [#472](#)

TECHNICAL SKILLS

Languages: Python, Rust, TypeScript, C++, Go, SQL

Systems: PyTorch, Vulkan, Docker, Kubernetes, AWS, GCP

Applications: LangChain, LangGraph, Next.js, Supabase, Inngest, n8n

ACHIEVEMENTS

- 3rd of 1,200+ entries, Splunk Global Hackathon 2025: built an OPA security add-on for real-time threat detection and privilege-escalation auditing.

EDUCATION

B.Tech, Electronics & Instrumentation

Manipal Academy of Higher Education, India | 2019 – 2023