

AADIT AGRAWAL

Computer Science Engineer · AI/ML · SWEng · HPC · DevOps

mail@aaditagrawal.com · [+91 99517 74197](tel:+919951774197) · aadit.cc · in/aadit-agrawal · gh/aaditagrawal



EDUCATION

Manipal Institute of Technology, MAHE

Jul 2023 - May 2027

BTech – Computer and Communication Engineering

EXPERIENCE

Head Student Researcher · CASIR, MAHE · ecitis.org

May 2025 - Present

- Architected HPC cluster (96-core Threadripper 7995WX, 48GB A6000, 256GB RAM) – 1,500+ daily inference req, 99.2% uptime.
- Leading the setup of a \$200K+ Compute Cluster using Apple Silicon + NVIDIA Blackwell; optimising for ARM64 + CUDA.
- Real-time CV on Jetson + DJI drone – 30fps edge inference via RTMP.
- LLM semantic translator for Schneider Electric – IEC 61131-3 → IEC 61499 conversion automation for legacy PLC codebases; migration effort reduced by 70%.
- Technical document ingestion pipeline (Schneider Electric) for machine / electronics engineering support – evolved from RAG + BM25 + keyword cluster indexing to trigram indices with custom pi.dev-based search agents. VLM + LLM stack across 2,000+ docs, 94% accuracy, 3× faster than the prior solution.
- Application security analysis platform (Schneider Electric) – source + compiled-binary evaluation in computer-use sandboxes with OS state monitoring (registry / DLL deltas), pattern-matched against a custom-correlated CVE / CWE / MITRE index. Pentest lead time reduced from 3 weeks → 3 days.
- 3 production AI systems; 2 peer-reviewed publications in progress; team learnings published at blog.ecitis.org.

DevOps Engineer, Intern · Pixa AI, Bengaluru

Dec 2025

- Provisioned GPU compute (H100/H200/B200) and CI/CD pipelines for TTS/STT model fine-tuning on Phi-4; achieved 99.5% deployment success across 2 API releases.

Frontend Engineer · Rossmann New Media 2 Inc.

Jan 2026 - Present

- Migrated rossmanngroup.com from WordPress to Next.js; improved page load performance and cross-platform responsiveness.
- Implemented technical SEO overhaul targeting data recovery keywords – moved primary rankings from 100+ to top 5.

Founding President · PRISM

Jan 2026 - Present

- Research club (AI/ML, SWE, HPC, DevOps, Biotech, InfoSec) – 150+ members. Academic Advisory Board for AI Curriculum and Adoption.

Management Committee · MIST

Oct 2023 - Jun 2025

- CTF infra for 300+ users, 100% uptime; cut costs by 80%.
- 12+ workshops for 150+ participants; top 10% national CTFs.

RESEARCH

FL + Digital Twin for Route Optimisation · SCE, MAHE

Jun 2024 - Jul 2025

- GPS telemetry (500K+ datapoints); dashcam CV via OpenCV + YOLO – 12,000+ labelled samples.
- Built FL pipeline for privacy-preserving route optimization across distributed edge nodes.

PROJECTS

Homelab Infrastructure · *Self-hosted cloud*

2022 - Present

- 70+ services on Proxmox/Docker – 98% uptime, 4 yrs, 8TB MergerFS, Blackwell-GPU, 3-2-1 backups.
- Deployed custom object-detection models on security camera feeds achieving 30fps real-time inference; offloaded 95% of compute from CPU to GPU including 4K media transcoding.
- Mitigated live DDoS (50K+ req/min); 95% attack surface reduction.

The TIK Menu API · [GitHub](https://github.com)

2025 - Present

- 5,000+ DAU, 150K+ monthly req, campus adoption in under 7 days. REST API + Cloudflare R2.

Coolstuff.work [Forgejo](https://github.com/Forgejo)

2023 - Present

- Suite of utility tools solving everyday problems for peers – 10,000+ aggregate DAU across sub-apps.
- Built OR-Free, a real-time AI model availability and pricing tracker for OpenRouter, and MIT Courses Search, a course discovery tool with interactive graph visualization (500+ nodes). Next.js, Radix UI.

- Additional tools: transcript-cli (video transcription), Campus Directory (crowdsourced contacts), Mass Contact Maker (bulk VCF generation). See git.aadit.cc.

ophoner · *Mobile agent control loop* 2026 - Present

- Phone-based agent loop for device APIs and filesystem workflows – extensible via tmux and Android shell tooling. Focused on on-the-go file editing/querying without heavy desktop dependencies.

Developer Tooling (OSS) 2025 - Present

- **agentskills**: Published library of Claude Code skills for reusable agentic development workflows.
- **t3code** (fork): customised for parallelised agent-driven development flows.
- Self-hosted CI + sandboxing stack (migrating from Daytona); CodeRabbit integrated for public-repo review.

SKILLS

Languages: Python, C/C++, TypeScript, JavaScript, Go, Zig, Rust, Bash, SQL

Infrastructure & DevOps: Docker, Kubernetes, Proxmox, Linux, Git, GitHub Actions, Terraform, Ansible, Nginx, Caddy, Tailscale

Cloud & Platforms: AWS, GCP, Azure, Hetzner, Cloudflare Workers, Vercel

AI/ML & HPC: PyTorch, JAX, Transformers, CUDA, SGLang, vLLM, llama.cpp, TensorRT-LLM, OpenCV, YOLO, Slurm, NCCL

Web: Next.js, React, Svelte, Astro, Tailwind, Node.js, Bun, FastAPI, tRPC

Databases: PostgreSQL, MongoDB, SQLite, Redis, pgvector, ChromaDB

AI Tooling: Pi (Custom Implementation), Codex, Claude Code, Droid, Cursor, opencode, ampcode, CodeRabbit, MCP, Skills