

Praneeth Suresh

Singapore | +65 8285 7814 | praneeth.suresh@u.nus.edu
[linkedin.com/in/praneeth-suresh](https://www.linkedin.com/in/praneeth-suresh) | github.com/Praneeth-Suresh

Professional Summary

ML Engineer building trust-sensitive AI systems, agent workflow infrastructure, and production-grade automation. Created Beryl, a repository-owned control layer for AI-assisted development, and currently researching deterministic policy gates for coding agents. Prior work includes React/Node legal automation, NLP pipelines processing 1,000+ documents/day, and sub-200ms retrieval systems. Strong CS & Mathematics foundation (NUS, GPA 5.00/5.00).

Skills

- **ML Frameworks:** PyTorch, scikit-learn, spaCy, NLTK, OpenCV, Keras
- **Infrastructure & DevOps:** Docker, CI/CD, AWS (DynamoDB, S3, Lambda), Google Cloud Platform
- **Databases & Vector Stores:** PostgreSQL, MySQL, SQLite, Chroma DB, AWS DynamoDB
- **Languages:** Python, Java, TypeScript, C
- **Web & APIs:** React, Node.js, Django, Flask, Streamlit
- **Agent Infrastructure:** Repository policy gates, deterministic checks, YAML schemas, benchmark harnesses

Experience

Agile Counsel

AI Legal Automation Intern

Vietnam

May 2026 – Jul 2026

- Built a React and Node.js legal document-automation workflow: validated multi-step intake, deterministic word document generation, and downloadable incorporation packages.
- Helped scale the product from one Delaware incorporation flow to a reusable multi-generator platform, adding token-gated client delivery and automated verification checks.

Kabal Research

Co-Founder & ML Engineer

Singapore

Sep 2023 – Jan 2026

- Built a high-throughput NLP pipeline (spaCy) processing 1,000+ geopolitical news documents per day, delivering structured sentiment signals to downstream risk models with consistent uptime.
- Architected a production graph-based analytics engine (Python, NetworkX) supporting real-time queries across 200+ entities and 4 relationship types for analyst-facing dashboards.
- Engineered a multi-signal scoring pipeline unifying NLP sentiment, game-theoretic payoffs, and graph centrality into a single auditable risk score, reducing analyst review time by ~40%.

ReProse

Co-Founder & ML Engineer

Singapore

Dec 2023 – Jun 2024

- Developed a custom OCR system using DBSCAN clustering (scikit-learn) that outperformed Tesseract by 20% on handwritten input, deployed as a containerised microservice.
- Implemented a RAG pipeline with foundation model integration for automated content generation, handling multi-format document ingestion at scale.
- Delivered end-to-end prototype in 6 months; selected as one of ~20 Social Innovation Fellows (SUTD Create4Good).

Projects

Beryl — Repository Control Layer for AI-Assisted Development

github.com/Praneeth-Suresh/Beryl

- Built an installable control plane for agent-ready repositories, packaging task routing, architecture context, testing policy, security policy, and deterministic check gates into versioned project files.
- Designed Beryl controls for safe, consistent repository governance through scoped changes, protected paths, dependency policy boundaries, forbidden imports, and enforced validation checks.

Beryl Repository Policy Language — CP3106 Independent Research

Jul 2026 – Nov 2026 (projected)

- Designing a 150-run benchmark to measure whether machine-readable policy representation and deterministic enforcement improve coding-agent outcomes over equivalent natural-language repository guidance.
- Building an experiment harness with 12 validated milestones, hidden tests, regression checks, external immutable policies, and tamper detection for condition-blind patch evaluation.

AI Society Corpus — Research Knowledge Graph

aisoc-corpus.aisocietysoc.workers.dev

- Built a static knowledge graph from scratch for the NUS AI Society, mapping studied AI/ML concepts into prerequisite and semantic relationships for onboarding and partner review.
- Implemented build-time semantic linking with Jaccard similarity over shared domains and tags, exposing explorable categories across classical ML, deep learning, generative AI, RL, and world modelling.

Truthflow — Real-time Fact-Checking Engine

github.com/Praneeth-Suresh/Truthflow

- Engineered a RAG-based retrieval pipeline using Chroma DB, achieving **sub-200ms latency** across 10,000+ indexed articles for real-time claim verification.
- Built a multi-format ingestion layer (web scraper + Keras OCR model) supporting diverse input sources with automated preprocessing.
- Reduced LLM inference costs by deploying rule-based linguistic analysis (exaggeration scoring) as a pre-filter, improving throughput without sacrificing accuracy.

Open Climate Fix (OCF)

github.com/openclimatefix

- Contributed production data ingestion pipelines to an open-source solar forecasting model, integrating 3 new dataset sources to improve model training diversity.

Education

National University of Singapore (NUS)	Singapore
<i>Computer Science (AI and Algorithms Focus Areas) with Mathematics (Double Major)</i>	2025 – 2028

- **GPA:** 5.00/5.00
- **Relevant Coursework:** Data Structures & Algorithms, Linear Algebra, Probability and Statistics
- **Awards:** Third Place Manus Hackathon 2025; IDEATE 2025 Semi-Finalist