

ISHAN AGARWAL

Software Engineer | Backend, Platform & Applied AI | NUS Computer Science 2027
Singapore | ishan_agarwal@u.nus.edu | linkedin.com/in/ishan-agarwal-nus | ishan-agarwal.com

Final-year NUS Computer Science student graduating May 2027. Four software internships across robotics simulation, backend microservices and applied AI, including six months building simulation tooling at Hyundai's Singapore R&D centre.

EDUCATION

National University of Singapore Aug 2023 – May 2027 (Expected)
Bachelor of Computing (Honours), Computer Science; Minors in Mathematics and Quantitative Finance Singapore

- Focus areas: Artificial Intelligence; Computer Security. Coursework: Data Structures & Algorithms, Operating Systems, Computer Networks, Information Security, Cryptography, Software Engineering, Introduction to AI & Machine Learning, Natural Language Processing.

TECHNICAL SKILLS

Languages: Python, Java, JavaScript/TypeScript, SQL, Apex, LaTeX
Backend, Web & Infrastructure: Flask, Node.js, React, Flutter, REST APIs, Docker, Git, CI/CD, MySQL
AI, Data & Simulation: scikit-learn, XGBoost, Hugging Face Transformers, LangChain, LangSmith, NumPy, Pandas, NVIDIA Isaac Sim, Omniverse Kit, USD

EXPERIENCE

Hyundai Motor Group Innovation Center Singapore (HMGICS) Jan 2026 – Jun 2026
Digital Twin & Simulation Engineering Intern | Python, NVIDIA Isaac Sim, OmniGraph, USD Singapore

- Built a code generator that turns Kawasaki AS robot programs into fully wired Isaac Sim control graphs, cutting **roughly one engineer-week of manual wiring per subprogram to about four hours**, review included. A single factory cell has 40+ subprograms.
- Scoped the project alone from a one-line brief, then benchmarked **seven generation approaches** before choosing deterministic Python for auditability inside an air-gapped network.
- Reverse-engineered **four undocumented OmniGraph and USD API behaviours** by diffing generated scene output against a hand-built baseline, then centralised the fixes into a shared module.
- Shipped the simulation reporting system and three Omniverse Kit UI extensions on a config-driven architecture every later feature reused. Reports generate per run with joint-angle charts, tool paths and a parts inventory.
- Audited asset traceability against an 80% target by cross-referencing **36 physical parts** against external inventory systems, lifting coverage **from 31% to 44%** and showing the remaining gaps were missing data rather than code defects.

TechFour Engineering Solutions May 2025 – Jul 2025
Software Engineering Intern | Python, Flask, Flutter, MySQL India

- Built and deployed **three backend microservices** for user management (authentication, OTP, password reset), document handling and internal communications, built as reusable components so the company's later projects could drop them in instead of rewriting them.
- Integrated WhatsApp and email APIs and shipped the Flutter client.
- Added SonarQube quality gates and documented every endpoint in Swagger.

PwC India Dec 2024 – Jan 2025
AI Engineering Intern | Python, LangChain, LangSmith, Streamlit India

- Prototyped a retrieval-augmented generation assistant over **several hundred internal policy and research documents**, for policy lookup and for surfacing prior work so teams stopped repeating research. Built in LangChain, with LangSmith for tracing and offline evaluation.

PROJECTS

LectureAI | AI Lecture Notes Platform (Co-Founder) Feb 2025 – Dec 2025

- Co-founded an EdTech startup turning lecture recordings into AI-generated study notes. Led full-stack development and ran customer discovery across NUS cohorts. Wound down after failing to reach product-market fit.

Extractive Question Answering | Margin-Triggered Reranking Aug – Dec 2025

- Fine-tuned RoBERTa to an 84.3 exact-match baseline on SQuAD v1.1, found the gold span already sat in the model's top five candidates for 95% of questions, and added a bi-encoder reranker that fires only on those low-margin cases, gaining accuracy at near-zero compute cost.

Fake News Detection | Metadata and Fairness Trade-offs Jan – May 2025

- Benchmarked 18 model and feature configurations on the LIAR political-claims benchmark, from TF-IDF and Word2Vec to a fine-tuned DistilBERT, and found that adding speaker metadata lifts macro-F1 by **2 to 4 points** partly by learning who is speaking, a measurable fairness cost.

ADDITIONAL

NTSE Scholar, All India Rank 21 (National Talent Search Examination) | Salesforce Administrator & Developer certified | Cambridge English C2 Proficiency (Grade A) | Volunteer teacher, Teach SG