

Julian CHAN

✉ julianchan8897@gmail.com | ☎ +447767766646 | 🐙 github.com/julian-8897 | 💼 [linkedin.com/in/julian-chan-a86b2921b](https://www.linkedin.com/in/julian-chan-a86b2921b)

SUMMARY

AI/ML research engineer and final-year PhD researcher building learned dynamical models, differentiable ML pipelines, LLM agents and GPU software. Experience spans JAX/PyTorch training, model evaluation and numerical validation.

SELECTED AI PUBLICATION

Prior or Feedback? What an LLM Uses When Adapting Neural Operators

Accepted to the NeurIPS AI4Science Workshop.

RESEARCH AND ENGINEERING EXPERIENCE

Differentiable ML Systems with Tesseract (Turbulence) (Inverse methods) 2026

Collaboration with Pasteur Labs; Honourable Mention, Tesseract Hackathon 2025.

- Trained a PyTorch correction model through a JAX spectral solver in-process, reducing 500-step rollout MSE by **71.7% versus the solver alone** and 15.7% versus a matched offline-trained baseline, with lower error on all 32 test trajectories. Separately demonstrated training through two containerised services.
- Built a shared gradient interface for a differentiable solver, interchangeable JAX/PyTorch PINNs and a flow-matching posterior. Delivered an installable CLI and Streamlit app with atomic checkpointing and CI across both backends.

Neural Surrogates and Symbolic Regression 2024 – 2026

PhD research, University of Surrey

- Developed parameter-conditioned Neural ODE surrogates for N-body simulations, reaching approximately **1% median fractional error on held-out trajectories**; also evaluated an unseen higher-resolution simulation. [First-author paper accepted to MNRAS](#).
- Combined the pretrained surrogate with dimensionally constrained symbolic regression to infer compact equations for black-hole-binary evolution. *First-author manuscript submitted.*

Erebus: GPU-accelerated N-body Simulation Software 2026

- Developing Python/CuPy simulation software for fast multipole gravity and regularised black-hole dynamics, using AI-assisted development for custom GPU kernels.
- Implemented CPU-reference comparisons, conservation checks and restart regressions; added configurable simulation runs, HDF5 checkpoints and SLURM workflows for HPC execution.

Research Brief Agent: LLM Service for Cited Technical Briefs (GitHub) 2026

- Built a streaming FastAPI service in which an LLM agent searches a persistent Qdrant index, retrieves arXiv papers and reads full-text and web evidence to produce **source-grounded technical briefs**.
- Controlled context length and execution cost through evidence compaction, depth presets and limits on agent turns and tool calls; recorded model and tool traces, token use and cost with Langfuse.
- Added citation validation and scenario evaluations in GitHub Actions, API-key authentication, rate limiting and a Docker Compose stack for the API and vector database.

SKILLS

AI/ML: Python, PyTorch, JAX, LLM agents, model evaluation

Systems: CuPy, FastAPI, Qdrant, Docker

EDUCATION

PhD Astrophysics & Machine Learning, University of Surrey 2023 – Present

MSc Computer Science *Distinction*, Swansea University 2021 – 2022

MSci Physics with Particle Physics & Cosmology, University of Birmingham 2016 – 2020