

Sam Armstrong

Website: sam-armstrong.github.io
Email: samuel.armstrong2014@gmail.com
GitHub: github.com/Sam-Armstrong
LinkedIn: linkedin.com/in/sam-a-a5557112a

PERSONAL PROFILE

I am currently completing my Master's degree in Computer Science, after 3+ years of industry experience building software that is maintainable, scalable, and genuinely useful for customers. I am driven by developing, shipping and iterating rapidly, and am seeking a company culture that shares this same philosophy.

EDUCATION

2025–2026

University of Warwick

MSc Computer Science

Estimated Grade: *Distinction*

Focusing on advanced topics in high performance computing, agentic artificial intelligence, computer security, networks, and computer science theory. All modules have been completed with an average grade of 79%, with the thesis currently in progress.

2018–2021

University of Sheffield

BSc Artificial Intelligence
& Computer Science

Grade: *First Class Honours*

Studied for a Bachelor's degree in AI and Computer Science, achieving a final year grade average of 79.5%. The course focused on the theory and applications of deep learning, software engineering, and system design.

MISCELLANEOUS

Anthropic London Hackathon (2023)

Created a GitHub bot to automatically add docstrings to functions using the Claude API.

Warwick Hackathon: Whack (2025)

Created an AI learning tool that lectures over pdf slides and asks questions to the student.

Publication in MABS 2022

Ivy's integration into Kornia

My Sheffield graduate profile

NOTABLE PROJECTS

CozyCasts

Created a platform for streaming AI-generated podcasts to help you learn, or for entertainment. The backend is designed for scalability, while maintaining low latency and cost-conscience, and uses an API/worker pattern to decouple the compute-heavy TTS from the API.

- API service is deployed serverless on Google Cloud Run
- Worker is deployed on GKE with KEDA for scale-to-zero
- API service schedules worker jobs using Google Cloud Pub/Sub
- Redis Pub/Sub is used for streaming audio back to the API as it's generated in real-time, then on to the user via a WebSocket
- We use cache-aside to prevent a race condition on the WebSocket
- Metadata is stored in PostgreSQL (Cloud SQL) and audio/images in GCS
- Frontend webapp is built using React/Next.js

ENTREPRENEURSHIP

2024–2025

Co-Founder — Ivy

Co-founded Ivy — which allows any ML models or code to be converted between frameworks such as PyTorch and JAX — as a spin-off from our previous company, Unify. Primarily working with:

- DevOps with cloud infrastructure (GCP) and CI/CD
- Advanced Python programming; AST; ML frameworks

I continue to maintain Ivy now as an open-source project, which is:

- Available via PyPI ('pip install ivy') and has ~14,000 GH stars
- Integrated into Kornia, the popular computer vision library

EMPLOYMENT

2023–2024

Machine Learning Engineer — Unify AI (YC W23)

- Led a team of five developing a computational graph tracer which can extract a minimal functional graph from arbitrary Python code — enabling any Ivy model to be compiled into an efficient representation.
- Developed the multi-backend framework Ivy: creating backend functions, forming the CI, and reviewing PRs from contributors.

2022–2023

ML Software Engineer — Intradiem

- Orchestrated the creation of Java data pipelines for transforming raw data from an AWS S3 datalake into bucketed time-series data. These were used for both ML applications in both training and inference, and required performant algorithm design for scalability.
- Led development of a Python micro-service using Flask/MariaDB, containerised with Docker, to document customer/product relationships and make this information easily accessible to services across the business.