

ABHISHEK SHIVAKUMAR

Polymath Engineer · Audio, ML & Low-Level Systems

abhishek.shivakumar@gmail.com · Cambridge, UK · github.com/godofecht · abhishek-shivakumar.com

SUMMARY

Polymath engineer with a decade of daily C++, deep across real-time audio, machine-learning systems, low-level C++ and embedded. 25 plugins to 500,000+ users, a sub-2ms neural inference engine, my own compiled language, 20+ published apps, and an independent research programme.

TECHNICAL SKILLS

C++ (RAII / MISRA), Python, C#, PyTorch, TensorFlow, JUCE, Qt, Rust, Zig, LLVM, Vulkan / CUDA, Unity, React / Vue / Node

EXPERIENCE

Founder & Lead Engineer, Quilio 2022 - Present

- ▶ Shipped 25 cross-platform audio plugins to a base of 500,000+ users, owning DSP, architecture and release.
- ▶ Built a from-scratch C++ neural inference engine running speech synthesis in under 2ms single-core, roughly 3x faster than TFLite and ONNX at half the memory.
- ▶ Built the full web stack and ran the product end to end while leading a team of 10+ engineers.

Co-Founder & CTO, Neuraspin 2022 - 2024

- ▶ Built a neural decoder and language model for real-time imagined-speech decoding, owning the full sensing-to-inference stack.
- ▶ Took the company to the TechStars final round.

Lead Engineer, Casperaki 2022 - 2024

- ▶ Built the cross-platform C++ music-generation engine behind several AI music products.
- ▶ Implemented production DSP and neural vocal synthesis with a full CI/CD pipeline.

Engineering Lead, Ogilvy / Mondelez 2024 - 2025

- ▶ Led a 10-person team across 4 time zones to ship a production AI generation pipeline for a global brand.
- ▶ Owned architecture, cloud infrastructure (AWS) and delivery.

Software Consultant, Affluent 2024 - Present

- ▶ Delivered 8+ end-to-end engagements across C++, Python and TypeScript.
- ▶ Built a low-latency audio-over-IP library with custom buffering and jitter correction.

Research Consultant, Newcastle University 2021 - 2022

- ▶ Built the software powering a brain-computer interface for rest-state neuromodulation experiments.

SELECTED PROJECTS

Flow. A statically-typed compiled language with codegen for C, LLVM/MLIR, SPIR-V, Metal and WebAssembly, with algebraic effects and autodiff.

Neural inference engine. A from-scratch C++ runtime running neural audio in under 2ms single-core, beating TFLite and ONNX.

Neural speech codec (QLAC). A neural speech codec and a learned audio codec on quantised latents, for very low bitrates while holding quality.

Piano to Flute GAN (undergraduate thesis). A GAN converting piano into flute, replicating the airy plosive attack and dropping the piano hammer transient.

Real-time depth NN library (masters thesis). A lightweight real-time neural-network library for stereo depth via Slow Feature Networks and a Disparity Network.

RayTherion. A real-time GPU ray tracer on Vulkan/SPIR-V and CUDA with custom acceleration structures.

EDUCATION

MPhil Basic & Translational Neuroscience, University of Cambridge · BSc (Hons) Music Technology, Royal Birmingham Conservatoire