

JOSH ANGEL

AI RESEARCH & ENGINEERING

Auckland, New Zealand · Hybrid or remote

I build the models, knowledge systems and software behind AI products. I combine hands-on research with deep experience in software engineering, DevOps, and systems and solutions architecture. I can help a business build its AI capability, investigate difficult technical questions, and turn the results into working technology.

WHAT I CAN DO FOR YOUR TEAM

Build a model-development programme. Research architectures, develop reusable foundation models and specialise them for different customers and applications. Work across training models from scratch, continued pretraining and fine-tuning, with data quality, evaluation and deployment considered from the start.

Build agents that use knowledge. Design systems that extract knowledge from data, preserve its evidence and make it available to agentic systems. Connect models, memory, tools and services into useful workflows, with visibility into their decisions and actions.

Turn uncertain questions into evidence. Design controlled experiments, compare approaches and build the tools needed to assess them. Investigate AI model behaviour, quality and efficiency. Recognise when a promising idea does not work, and use that result to choose a better direction.

Engineer beyond the model. Build the applications, data pipelines, inference services and infrastructure that make AI usable. Connect research decisions to practical constraints: compute, memory, latency, reliability and cost.

Provide technical direction while staying hands-on. Help a small team choose worthwhile problems, resolve difficult engineering questions and carry research through to delivery. Contribute directly to the work, not only manage it.

WHERE I WANT TO CONTRIBUTE

A hands-on AI research, model-building, engineering or DevOps role, hybrid or remote from New Zealand, alongside a strong technical team. Individual-contributor or technical-lead work both appeal; the substance of the role matters more than the title.

Contact & work

[LinkedIn](#) · [Hugging Face](#) · [GitHub](#)

SELECTED EVIDENCE

Foundational and specialist models

As Head of AI Research and Custom Models at Arbitr (formerly Straker), my team and I built two multilingual foundation models through continued pretraining of open bases, then used them to design and deliver four customer-specific models and two generic models for the Arbitr platform. Serving environments included AWS and IBM watsonx.ai.

Agents and organisational memory

Implemented research prototypes for grounded knowledge extraction, provenance and agent access. Investigated quality, recall and cost across multi-stage model workflows.

Evaluation research

Developed model-comparison tooling and researched style compliance and explainable quality estimation, including models that identify error spans rather than only returning a score. Style-compliance manuscript prepared; review and publication pending.

Independent model research

Built and trained Rivet-1B from scratch, including its tokenizer. Published the model weights, evaluation artifacts and technical report.

Systems foundation

Spent several years in DevOps rebuilding infrastructure and automating application delivery, skills that remain current and underpin my AI research and development.