Oliver Adams

Melbourne, Australia | oliver.adams@gmail.com | oadams.github.io | GitHub | Google Scholar
Australian & British citizen

I'm an ML engineer and researcher with 6+ years of industry experience building and deploying machine learning systems across computer vision, NLP, speech, and conversational AI. I have a PhD in NLP from the University of Melbourne and did a postdoc at Johns Hopkins. I'm comfortable working at the architecture level of transformers and LLMs (I read and implement from current research) and have production experience spanning vision transformers, multimodal document AI, speech recognition, and RAG systems. I have a track record of leading small teams through end-to-end ML development, from problem scoping and data curation through model training, rigorous evaluation, and deployment. I place a strong emphasis on data quality, systematic error analysis, and knowing when a simpler approach will outperform a complex one. I have a deep interest in AI safety research and am actively building my expertise in this area.

Technical Skills

ML/DL Frameworks: PyTorch, Hugging Face Transformers, scikit-learn

Model Architectures: Transformer-based language models (BERT, GPT-family, open-source LLMs), multimodal models (LayoutLMv2), vision transformers (DINOv3, ViT, RT-DETR) for custom classification and object detection, hierarchical Bayesian models

Domains: Computer vision (classification, object detection, synthetic data generation), NLP (information extraction, NER, document AI, semantic similarity), speech (ASR, wake word detection, turn-taking, escalation detection), conversational AI (RAG, LLM fine-tuning, prompt engineering, agentic tool use)

Tools & Infrastructure: Python, Rust, Docker, Weights & Biases, DVC, Git/Git-LFS, GCP (Vertex AI, AutoML)

Languages: English (native), German (intermediate)

Professional Experience

Senior Data Scientist

zData (acquired by Atos) — Remote (Australia) — September 2019 – Present

Technical lead on ML consulting engagements for enterprise clients across computer vision, NLP, speech, and conversational AI. Led or co-led teams of 2–5 through full project lifecycles on 20+ engagements.

- **Weapon Detection for School Security (2025–2026):** Co-led development of a weapon detection system (guns and knives) for CCTV surveillance in school environments. Fine-tuned Meta's DINOv3 vision transformer (300M and 7B parameter backbones) with custom classification and object detection heads, achieving strong performance on held-out real-world evaluation data across diverse scenes. Implemented a custom object detection architecture following the DINOv3 paper with modifications for single-pass inference. Explored RT-DETR as an alternative detection backbone. Built a synthetic data generation pipeline using generative image models to address training data scarcity. Established zero-shot VLM baselines (DINOtxt, SigLIP) to validate the fine-tuning approach. Drove iterative data quality improvement through systematic manual error analysis, addressing annotation inconsistencies and demographic gaps in training data.
- **Document AI & NLP (2019–2023):** Built and deployed information extraction and text classification systems across clinical, legal, and industrial domains using BERT and LayoutLMv2. Handled documents of 100+ pages with custom inference pipelines for 2D PDF chunking and cross-page context. Developed reusable error analysis tooling and evaluation frameworks. Repeatedly diagnosed underperforming models by identifying training data quality issues — inconsistent annotations, label extraction bugs, conflicting labels — and drove substantial performance improvements through data remediation. Identified and mitigated model bias through targeted evaluation and data collection.
- **Speech & Audio ML (2024–2025):** Built a turn-yielding model using a small LLM on ASR transcripts to detect conversational turn completion. Developed a custom wake word detection model for a voice-activated bot. Created an escalation detection system to identify emotional escalation from speech in security contexts. Benchmarked multilingual ASR models (3B–7B parameters) on Gulf Arabic dialects for bodyworn camera applications.
- **Conversational AI & RAG Systems (2023–2025):** Drove key technical decisions across multiple RAG-based chatbots for enterprise clients including a university information system and domain-specific

assistants. Evolved the architecture from naive retrieve-always embedding search to a tool-oriented agentic system that invoked retrieval selectively. Shifted team approach from fine-tuning-heavy to prompt-engineering-first, simplifying pipelines and reducing synthetic data dependency. Evaluated emerging LLM orchestration frameworks for production use.

• **Additional projects:** Built an LLM evaluation harness in Rust implementing LLM-as-a-judge scoring. Person detection benchmarking (Google AutoML vs. Amazon Rekognition) for security camera footage. Predictive maintenance using hierarchical Bayesian models across 150K+ IoT devices. Retail change detection. Semantic text similarity for information retrieval. Prepared internal training materials on ML ethics.

Postdoctoral Fellow

Center for Language and Speech Processing, Johns Hopkins University — May 2018 – May 2019

Researched multilingual pretraining and adaptation of speech recognition models. Published work on massively multilingual adversarial speech recognition at NAACL 2019. Collaborated with teams at the Center for Language and Speech Processing under Dan Yarowsky and Shinji Watanabe.

Research Assistant

International Computer Science Institute (UC Berkeley) — January – March 2018

Short-term research appointment in speech and language processing.

Software Engineer

Aikuma Language Preservation Project (University of Melbourne) — June 2012 – June 2014

Primary developer of Aikuma, an Android app for collaborative language documentation. Contributed to UI/UX design and functionality decisions. Won the $8,000 Grand Prize at the 2013 Open Source Software World Challenge.

Education

PhD in Computer Science

University of Melbourne — 2012 – 2017

Thesis: *Automatic Understanding of Unwritten Languages.* Supervisors: Steven Bird, Trevor Cohn, Graham Neubig. Research on speech recognition and machine translation for low-resource and endangered languages. Included research internships at IBM Research Australia, Nara Institute of Science and Technology (Japan), and Xerox Research Centre India.

BSc (Honours) in Computer Science

University of Melbourne — 2008 – 2011

Selected Publications

O Adams, M Wiesner, S Watanabe, D Yarowsky (2019). Massively Multilingual Adversarial Speech Recognition. *Proc. NAACL-HLT 2019.*

O Adams, T Cohn, G Neubig, H Cruz, S Bird, A Michaud (2018). Evaluating Phonemic Transcription of Low-Resource Tonal Languages for Language Documentation. *Proc. LREC 2018.*

O Adams, G Neubig, T Cohn, S Bird, Q T Do, S Nakamura (2016). Learning a Lexicon and Translation Model from Phoneme Lattices. *Proc. EMNLP 2016.* **Best Short Paper Award.**

O Adams, A Makarucha, G Neubig, S Bird, T Cohn (2017). Cross-Lingual Word Embeddings for Low-Resource Language Modeling. *Proc. EACL 2017.*

Full publication list at Google Scholar.

Selected Awards

EMNLP 2016 Best Short Paper Award — Learning a Lexicon and Translation Model from Phoneme Lattices

Open Source Software World Challenge Grand Prize ($8,000 USD) — Aikuma language documentation app, 2013

Australian Postgraduate Award — PhD scholarship, University of Melbourne