Alan Williams

alan.williams@isogonal.io, https://alanswilliams.io

Isogonal **Co-founder** **2017–present**

I lead a small team building software with millions of dollars in revenue impact in industrial and healthcare settings. I am the software lead as well as working with clients on project scoping and solution delivery. Significant projects include

- I worked with New Zealand's national electrical grid operator to scope and develop an application for prioritising maintenance work based on free text data and historical records. I conducted solution discovery, exploratory modeling, and developed the software product, which consisted of a Python API, database, scheduled data pipelines, and a React/Next.js web application. The modelling process involved NLP, learning from human feedback, Bayesian heirarchical modeling and uncertainty quantification. product enables the client to self-mange the training data and model-retraining proces as the data distribution shifts over time. The impact was a $6M cost reduction in the first year, a peer-reviewed paper, and the product is still In daily use many years later. I designed the product and modeling approach and coded much of the software.

- Another project for the grid operator was modeling equipment degradation and failure using100 years of operational data including tens of thousands of natural language documents and terabytes of time-series data, using generative modeling and Tensorflow Probability. This project kick-started modeling individual asset condition based on data within the organisation, rather than applying heuristic rules to asset classes.

- Another project involved developing high-dimensional MPC software for controlling reactive power compensation equipment in the high-voltage network of the national grid, investigating replacing human operators with control software.
- I worked for a US e-commerce firm developing software for hierarchical Bayesian modeling at scale, optimizing daily product pricing for hundreds of thousands of products across hundreds of stores using Tensorflow Probability running on K8s.
- I personally designed, developed and sold a SaaS app for managing and processing geological data to international mining firms. Written in Python (API, data pipelines) with a React/Next.js web application.
- I personally developed a laboratory data management system (LIS) and secure client portal for a regional public health provider in New Zealand, managing PHI data. This saved person-days per week in administration time and is still in daily use five years later. In order to keep maintenance to a minimum, this was built with Javascript, Node, HTML5, CSS and jQuery, with minimal framework use. I also wrote an authentication and user management system from scratch that passed an external penetration test audit.
- I designed and developed an ML-based fraud detection system for a US event ticketing platform using Java and XGBoost. The impact was a $1.2M annual cost reduction in fraud-related costs.
- I developed an Android appliation for outpatient appointment management for the Canterbury District Health Board (New Zealand).
- I developed a system for reducing the cost of urgent blood sample transport from primary practice clinics to a central laboratory. Published in British Medical Journal (BMJ) Simulation and Tech Enhanced Learning

Provisionai.com **Principal Software Engineer** **2019–present**

- I was the first technical hire at ProvisionAI and now lead a team of 4 developers building an enterprise SaaS product with Fortune 500 customers, delivering millions in cost savings per year.
- I personally designed the cloud architecture and delivered the MVP. The product is hosted on GCP and uses Python for the backend (Flask) and data pipelines, C++ for performance-critical libraries, Keycloak for SSO and authentication, and React/Next.JS for web applications.
- I have implemented and significantly extended the founder's shipment linear-program based cost optimization system, and developed heuristics for optimizing 3D bin packing from scratch, using python, CPLEX, and C++. I spent a year researching whether RL and MCTS could improve on these heuristics, but did not arrive at something that had a definitive advantage.

Translucent.io **Senior Data Scientist** **2022–2023**

- Founding member of the Data Science team. Wrote the DS team roadmap for the first six months.
- Designed and built the company's original client-facing backend cloud infrastructure
- Created a transaction network graph from standardizing messy accounting data and models operating on this graph to infer attributes, match entities and predict variables of interest

RetailMeNot **Staff Data Scientist** **2018–2020**

- Designed and built data pipelines and management systems for training thousands of search ranking and user personalization models every 12 hours, processing billions of events daily.
- Designed and built an A/B testing framework enabling measurement of the impact of algorithm changes for the first time
- Micro-optimization in Spark and Scala to reduce cost by 30% in 100+ node EMR clusters

SLI Systems **Research Engineer** **2014–2017**

- Designed a query pre-processing system and a learning-to-rank document ranking system and integrated it into a Lucene-based search engine processing 1B+ searches per month with <100ms 99% latency.
- Design and developed a green-field product recommendation product for the company, processing billions of events

Ströer Labs **Research Engineer** **2013–2014**

- Designed and developed systems for predicting CTR and optimizing second-price auction bids for a European online advertising demand-side platform

Aurecon **Structural Engineer** **2008–2013**

- Achieved PE (Professional/Chartered) engineering certification.
- Thermal-structural finite element modeling and nonlinear time-history analysis for earthquake design

SKILLS

- **Analytical Skills** Experience in bayesian modeling, linear programming, theoretical physics and machine learning
- **Technologies** Full-stack software engineering in python, Typescript, Scala, Java, C++, Rust, and R, using GCP and AWS.

PUBLICATIONS

Nb-doped rutile titanium dioxide nanorods for lithium-ion batteries Solid State Sciences (2018)

Prioritizing Predictive Maintenance Work Using Machine Learning. Electricity Engineers' Association (EEA) Proceedings (2018)

Effective Resource Management using Machine Learning in Medicine: An Applied Example, BMJ Simulation and Technology Enhanced Learning (2018)

EDUCATION

MSc, Mathematics, Distinction, Canterbury
B.E, Hons, First Class, Canterbury University
Winning team, Boston Consulting Group Business Strategy Competition (New Zealand)