

Paul Bridger

paul@paulbridger.com, [@paul_bridger](#)

Summary

Technical co-founder with extensive experience taking complex products from conception to production. Successfully created and managed deeply technical teams while establishing a rigorous engineering culture and continuous testing and integration process.

Expertise in modern design and implementation across several domains: machine learning, big-data, full-stack web, iOS, and in multiple environments: Google and Amazon cloud, Linux, and Hadoop/Spark.

Currently focussed on low-latency intelligent video analytics powered by machine learning. Experienced with Pytorch, Gstreamer, Deepstream, TensorRT, CUDA, etc.

Seeking opportunities in entrepreneurial environments where my diverse skillset and leadership capabilities can have a significant impact.

Selected Publications

- [Object Detection from 9 FPS to 650 FPS in 6 Steps](#). (September 2020)
- [Object Detection at 1840 FPS with TorchScript, TensorRT and DeepStream](#). (October 2020)
- [Object Detection at 2530 FPS with TensorRT and 8-Bit Quantization](#). (December 2020)
- [Mastering TorchScript: Tracing vs Scripting, Device Pinning, Direct Graph Modification](#). (October 2020)

Roles and Projects

Machine Learning Consultant

Paul Bridger Consulting Ltd., Bangkok, April 2020—Present

Helped several companies to develop, optimise, and productionise machine learning systems to meet business or product requirements.

Projects spanned several domains: deep-learning visual analytics pipelines, recommendation systems and generative models, and target architectures: datacenter, Nvidia Jetson and iOS.

Senior Machine Learning Engineer

Dive, Madrid, February 2019—January 2020

I led the development of a high-throughput, multi-model intelligent visual analytics system. Besides the core architecture and low-latency frame-processing implementation, I also implemented and productionised state-of-the-art machine-learning models including detection of people and faces, body and face recognition, tracking incorporating a Kalman Filter, demographics and pose extraction. Additional unique capabilities of the system included a novel whole-system calibration process, and the 3D math to fuse all camera inputs into a cohesive tracked space.

Additionally, I was the primary technical salesperson for international business and main contributor of technical prose to proof-of-concept documents and investor/customer pitch decks.

I led the hiring for Machine Learning Engineers, attracting and training a strong team with deep technical ability in performance tuning, research productionisation and model building. I also led our deep learning hardware specification and acquisition, improving productivity and decreasing cost across the group.

Visual analytics models and technologies:

- Gstreamer, Deepstream and Pytorch, CUDA streams
- Body and face detection (SSD300, FasterRCNN, RetinaFace, EXTD_64, MTCNN)
- Body and face identity (OmniScale ReID, ArcFace, FaceNet etc)
- Single-camera tracking based on DeepSORT, extended for multi-camera contiguous spaces
- Demographics (Deep Label Distribution Learning jointly trained for age, gender, ethnicity)
- Pose detection (PifPaf, and Keypoint FasterRCN w/ FPN)

Recommendation system:

- Rebuilt production-ready data pipeline
- Designed and implemented sequence-based user-personalisation deep learning model in Pytorch

- Build highly efficient on-GPU business rules engine

Machine Learning Contractor

Graphyte.ai, Madrid, October 2018—January 2019

- Built customer-next-action LSTM-based sequence model in Pytorch, deployed to AWS Elastic Container Service.
- Implemented collaborative filtering and content models, and a stacked ensemble recommendation model.

Cofounder, Head of Development

Gluent Inc. (gluent.com), London, 2015—August 2018

Gluent successfully attracted many Fortune-500 customers (including several in the Fortune-50), with consistently successful deployments and high stakeholder satisfaction. The Gluent Data Platform is a big-data middleware that connects modern Hadoop-based data sources with established enterprise databases, unlocking value and increasing efficiency.

- Built the first version of the full product, including high-efficiency C++ data pipeline connector, and Python-based CLI data orchestration tools.
- Hired, trained and managed the teams (8 hires) that would a) develop the software to enterprise production grade, and b) implement the solution at many Fortune-500 clients.
- Advised CEO on product development, timelines, business opportunities and strategy.

Founder/Developer

Independent, Singapore and Europe, 2011—2015

- Built and ran numerous small businesses while travelling through Europe.
- SaaS web applications, iOS applications, and independent contracting.
- Competed in several Kaggle competitions, achieving top-10% results using neural networks.

Cofounder, Developer

E2SN Ltd., Singapore, 2009—2011

E2SN is a capacity monitoring and planning middleware for enterprises with under-utilized server estate.

- Developed data-analysis, anomaly-detection, and ETL server components in Scala.
- Developed highly flexible charting and data display front and back-end in Coffeescript and Scala.

Equity Derivatives Risk Developer

Barclays Capital, Singapore, 2006—2009

- Developed new functionality and optimized existing code for the C++ Sophis-based derivatives trading platform.
- Most prolific division-wide C++ interviewer, with over 200 hours of in-person interview experience.
- Developed and deployed production systems in C++, C#, Java, VB.

Contract Developer

New Zealand, 2004—2005

- NZ Government, Te Pūnaha Kōwhiri: developed C# web application with MSSQL back-end.
- Paragon Electronic Design Ltd.: developed C++ audio/fax utility.

Game Developer

New Zealand, 2004—2005

- Developed high-performance 3D game, including gameplay code, artificial intelligence, special effects, physics, audio and graphics. C++ and Python.

Developer

Thomson Reuters Legal and Regulatory, New Zealand, 2001—2003

- Content management system and database development in Python and C++.

Education

Otago University, New Zealand, 1998-2001

- BSc.(Hons) in Neuroscience (1st class honors).
- Neuroscience major with artificial intelligence focus: neural networks, natural language processing, and Bayes theorem.

- Thesis: Bridger, P. and Robbins, A. Spurious Attractors in Hopfield Neural Networks. 2001.

Patents

- *US Patent No.* [US9760604B2](#). (2017). System and method for adaptive filtering of data requests.
- *US Patent No.* [US10069916B2](#). (2018). System and method for transparent context aware filtering of data requests.