Oliver Thomas, PhD

Senior Engineer, Machine Learning & AI

Brighton, UK · oliverthomas@outlook.com · github.com/olliethomas · linkedin.com/in/oliver-t-thomas

PROFILE

I design and build production machine learning and AI systems end to end: from whiteboarding the problem, through training models and writing the services that run them, to automating the monitoring that keeps them improving. I've built the systems and processes that let me and the wider team ship models with confidence.
Selected strengths: Production ML · Evaluation and Monitoring · Model Design · Python · IaC · Agent Tooling (MCP)

EXPERIENCE

Senior Engineer (AI & Optimisation)

Metasonixx · *Oct 2026 – present · Fixed term · Brighton, hybrid*

Building an ML and optimisation pipeline that designs metamaterials to cancel noise in target frequency ranges.

- Turned research code into a reusable Python system.
- Rewrote the pipeline in JAX, adding gradient-based optimisation alongside the existing genetic search.
- Made every run reproducible, validating the models against published results and independent libraries.

Senior Engineer (AI & ML)

Scoreline AI · *Aug 2024 – Sep 2026 · Remote*

Owned the technical direction of the AI behind Scoreline's fantasy football (£2.4m ARR, c.60,000 paying subscribers) and cricket products, serving predictions and constrained optimisation against fixed deadlines. Led the rebuild of the football product's legacy prediction system, using proven, modern architecture (normalised data, event-driven services, serverless infrastructure) to make it faster, cheaper to run and safer to change.

Delivery

- Moved the prediction service from a 6-hour batch cycle to event-driven updates, so new data and admin-panel overrides reach served predictions in seconds.
- Redesigned the database, cutting the data fetch per player from 27 to 0.1 seconds, making event-driven updates across all 600+ Premier League players possible.
- Moved the core MILP optimiser out of the main app into its own auto-scaling microservice, modernising it with tests and a formal interface that made user-reported issues reproducible and the optimiser easy to extend.
- Wrote and secured an MCP server that lets members use the product through AI assistants, with Auth0 login and features gated by membership tier.
- Kept the prediction services running unattended with thorough CI, monitoring, alerting and proactive log reviews, removing the need for matchday cover.
- Co-designed the serverless rebuild of the cricket product's agency-built backend, which handled a 10-minute traffic spike before each match and cut its cloud bill by 99%, from circa £15k to £120 a month.
- Built a culture of fearless deployment, moving from infrequent releases to shipping multiple times a day.

Modelling and evaluation

- Built the platform behind every predicted number on both products, from team strength to expected points, where each target's model could be replaced independently once a better one proved itself.
- Took a betting-market-led model from research code to production at team and player level, with tests, registry versioning and fallbacks for fixtures without odds.
- Built neural networks in PyTorch, from simple baselines to temporal attention over a player's recent matches.
- Researched beyond the shipped models: graph neural networks learning from the links between players, teams and matches, and networks predicting a distribution over in-match statistics, not a single number.
- Defined the promotion criteria for models, including time-based backtesting and domain validation metrics.
- Built model health checks that catch degenerate predictions before a model is promoted and while it serves, with Datadog alerts when accuracy drops or predictions collapse.
- Owned the constrained-optimisation model that turns predictions into squad selections.

Architecture and standards

- Authored the reference architecture for a multi-sport prediction platform: training/serving pipelines, a model registry with lifecycle semantics, versioned data contracts, event-driven feature freshness and non-functional targets (e.g. graceful degradation).
- Defined the testing standard for the prediction platform's backend services: a layered suite spanning unit, integration, contract, property and BDD tests, enforced in CI with a dedicated QA suite before release.
- Broke platform work into scoped increments as requirements evolved, documenting judgement calls as Architecture Decision Records that engineers across technical disciplines could build against.

Stack. Python · PyTorch · Lightning · Weights & Biases · Optuna · Hydra · Pandas/NumPy · scikit-learn/XGBoost · Postgres · Pulumi (IaC) · GCP Cloud Run · GCP Pub/Sub · pytest · MILP · FastAPI · CI/CD

Postdoctoral Research Fellow

Predictive Analytics Lab, University of Sussex · *Jan – Sep 2022 and Sep 2023 – Aug 2024*

Re-appointed to the same role after a career break; the second appointment was held through the Basque Center for Applied Mathematics (BCAM), Bilbao.

Lab funded by a European Research Council grant ("Bayesian GDPR"). Maintained the lab's shared engineering practice and supported other researchers, alongside contributing to projects as a technical generalist.

- Promoted peer-review habits in the lab to develop researchers: reading groups, forums for presenting work in progress, and feedback on drafts.
- **Carbon capture intervention modelling:** satellite imagery and public data, built with social scientists for policymakers assessing Sustainable Development Goals, part of a £12.5m Defra/BEIS-funded landscape-modelling programme with Kew. Contributed code, wrote up work as tickets and groomed the backlog.
- Built a satellite-imagery pipeline: label a feature such as vineyards, and a U-Net finds them across the wider map, linking each to local public data, so users could quickly test hypotheses like "vineyards improve biodiversity".

Career break

Sep 2022 – Sep 2023 · Travel across Central and South America after my PhD submission.

Consultancy Project Lead

Predictive Analytics Lab, University of Sussex · *Feb 2020 – Jan 2021*

- Led five researchers to deliver a documented Python library that learns causal structure (Bayesian networks, via MCMC) from questionnaire data and estimates intervention effects, building much of its pipeline myself.
- Agreed the deliverable directly with the client and introduced backlog grooming, standups, retrospectives and client demonstrations.

Technology Graduate

American Express · *Sep 2017 – Sep 2018 · Burgess Hill*

- **Data Architecture:** built a tool that generated Spring Boot data access APIs from ER/Studio entity-relationship models, and planned database migrations for teams as part of their move to a microservice architecture.
- **Enterprise Cloud Platform:** built reference applications for teams adopting cloud systems, including distributed tracing for microservices with Zipkin.

EDUCATION

PhD, Machine Learning

University of Sussex · *Sep 2017 – Dec 2022*

First year part-time, concurrent with a graduate role at American Express.

Thesis: *Fair Representations in the Data Domain* (185pp) · Supervisor: Prof. Novi Quadrianto. Built adversarial, generative (VAEs, GANs, normalising flows) and self-supervised models in PyTorch that learn fair representations which stakeholders can inspect and explain. Wrote a Bayesian active-learning framework for benchmarking the group's proposed methods against established baselines.

Papers at CVPR 2019, ECCV 2020 and ACM EAAMO 2021.

BSc (Hons) Computer Science, First Class

University of Sussex · *2013 – 2017*

Dissertation: *Learning Using Privileged Information: A Unified Approach*. Two faculty prizes: Outstanding Computer Science Student, and Best AI Final Year Project.