Euan Goodbrand

London

Github: github.com/euangoodbrand
Email: euangoodbrand@gmail.com
LinkedIn: linkedin.com/in/euangoodbrand
Mobile: +44 7956 028983
Website: euangoodbrand.com
Stack Overflow: EuanG

Education

Imperial College London — London, UK
***MSc Computing (Artificial Intelligence and Machine Learning)**, Distinction* — *Oct. 2023 - Oct. 2024*

- **Thesis:** Deep Learning on Noisy and Imbalanced Systems Data
- Corporate Partnership Programme Individual Project Prize in Computing Science (best thesis/ research prize)
- Average coursework mark 95%

University of Sheffield — Sheffield, UK
***BSc Computer Science**, First Class Honours* — *Sep. 2020 - Jul. 2023*

- **Thesis:** Procedural Content Generation of 3D Graph-Based Structures
- Douglas Lewin Memorial Prize for best exam performance
- Software Hut Prize for client software in a team

Publications

Deep Learning for Imperfectly Labeled Malware Data — ACM CCS 2025
F. Alotaibi, **E. Goodbrand**, S. Maffeis.

- SLB framework for robust deep learning under severe label noise via dynamic clean/noisy set partitioning and pseudo-labelling, boosting macro F1 from 74.5% to 96.0% on noisy malware benchmarks.
- Demonstrates scalable training under weak and imperfect supervision, with techniques suited to domains where labels are partially automated, inconsistent, or expensive to obtain.

Experience

Mercedes-AMG Petronas Formula One Team — Brackley, UK
Machine Learning Engineer — *Sep. 2024 - Present*

- Adapted transformer-based 3D architectures for high-fidelity aerodynamic mesh/point-cloud data (millions of points per sample) to infer unobserved aerodynamic fields (surface/volume pressure, velocity components, surface friction), trained across thousands of geometries with continual weekly updates under distribution shift, with uncertainty estimation to quantify confidence in predicted fields.
- Designed transformer fusion of sparse on-car pressure taps to dense 3D geometry via cross-attention from tap tokens to mesh tokens, using masked sensor tokens to handle missing or unreliable taps and varying sensor setups across race weekends and weather conditions.
- Developed spatio-temporal formulations for high-frequency (200 Hz) streaming sensor data with millisecond-level latency constraints, handling low-SNR turbulence-dominated noise, time-synchronisation issues, and non-stationary regimes (speed, yaw, ride height) without access to future context.
- Implemented RoPE-based geometric encodings aligned to geodesic and mesh-adjacency ordering, improving stability under extreme geometric variability and supporting reference-frame and geometry invariances required by vehicle physics.
- Scaled training and ablations on an on-site NVIDIA H200 cluster (24 GPUs) and benchmarked training throughput and efficiency across 8× GH200, 8× A200, and 8× AMD MI300X clusters to identify best-performing configurations for high-resolution 3D workloads.
- Optimised model inference for track-relevant latency, reducing end-to-end inference time by 5× while preserving aerodynamicist-facing quality on Cp-field predictions.
- Reproduced and stress-tested state-of-the-art 3D vision and geometric learning methods for internal use, delivering stronger baselines and up to 10× reductions in memory and training time through architecture and training modifications suited to proprietary high-resolution 3D data.
- Contributed to the Mercedes–Meta AI aero 3D vision collaboration (with Oxford VGG), applying early-access tooling (e.g., SAM 3) for trackside vision analysis and delivering robust camera calibration for a rigid fisheye setup operating at near 300 km/h to support consistent cross-run comparisons.

Imperial College London London, UK
Machine Learning Research Assistant *Feb. 2024 - Present*

- Lead co-author on deep graph-based learning for brain graph super-resolution with Dr Islem Rekik; coordinated ∼20 co-authors, designed experiments and rigorous ablations, and ensured reproducible pipelines.
- Implemented graph-learning baselines and message-passing variants in PyTorch/NetworkX; evaluated with calibration/error metrics and task outcomes, with visualisations for interpretability.
- Explored graph fusion and scaling strategies; maintained experiment tracking and code quality for reliable research artefacts.

HotelMap and RoomCard London, UK
Machine Learning Engineer *Jun. 2024 - Sep. 2024*

- Built PDF floorplan extraction classifiers and vector retrieval to accelerate large-scale venue analysis.
- Automated data pipelines with crawlers and RPA on AWS; improved NLP-based parsing for structure and search.
- Prototypes demonstrated to Google, Adobe, The Pokemon Company and enterprise clients.

PROJECTS

Thesis: Deep Learning from Label and Feature Noise in Cybersecurity Systems

- Developed MIMICRY (synthetic label noise), CLEAN-STOP and SENTINEL (ensemble cleaning); improved robustness under label/feature noise and class imbalance across malware and NIDS datasets.

Procedural Modelling and Generation Software Research SDC 2023 highlights

- Algorithms for 3D procedural content; produced playable levels with geometric/topological correctness; presented to 1,000+ developers at Sumo Digital Developer Conference 2023.

Optuna v4.6.0 (open-source contribution) (hyperparameter optimization framework) Optuna v4.6.0 release link

- Contributed to a widely used open-source HPO library by modernising type-checking and import patterns across core modules (`optuna.importance`, `FanovaImportanceEvaluator`, `study/_optimize.py`).
- Refactored typing-only imports under `TYPE_CHECKING` to reduce unnecessary runtime dependencies and improve static analysis hygiene; changes were merged and highlighted in the v4.6.0 release (e.g., #6278, #6279, #6280).

ML Software for Hospital Kidney Disease Detection GitHub

- Kubernetes-based real-time AKI detection from HL7 lab streams with pager alerts and fault-tolerant inference services.
- Reliability proven via chaos-monkey style failure injection (network partitions, pod/node evictions, DB restarts, message backlogs), sustaining 100% uptime over a two-week run with on-call runbooks.
- Ops: Prometheus/Grafana monitoring and SLO alerts, autoscaling and graceful rollouts; exceeded NHS baseline with F_3-score of 99.9% .

AWARDS & CERTIFICATIONS

Corporate Partnership Programme Individual Project Prize (2025): best MSc Computing thesis, Imperial College London.

Douglas Lewin Memorial Prize (2023): best examination performance in Computer Science, University of Sheffield.

Software Hut Prize (2023): most effective software for a real client, University of Sheffield.

AWS Certified Cloud Practitioner (2024-2026).

Global Engineering Challenge - Best Communicated Solution (2022).

SKILLS

Vision & Geometry: PyTorch3D, Open3D, OpenCV; camera models and SE(3) transforms; point-cloud/mesh ops; instance/semantic segmentation; keypoint/pose tracking; basics of COLMAP workflow.

Learning: PyTorch, JAX, TensorFlow, PyTorch Lightning; multi-GPU training (DDP), mixed precision; label-efficient/weak supervision, ensembling, calibration, rigorous ablations.

Data & Visualisation: NumPy, Pandas, Polars, Matplotlib; Blender, Houdini for 3D data generation/visualisation.

Systems/MLOps: CUDA, Linux, Docker, Kubernetes, Git, CI/CD, AWS, Azure.

Programming: Python, C++, C#, Java.