Carrow Morris-Wiltshire – Master CV

Core Identity & Contact

- **Location:** Newcastle upon Tyne, UK (Relocating to London - Available for hybrid/on-site)
- **Clearance/Citizenship:** British Citizen
- **Portfolio/Web:** carrow.me.uk
- **GitHub:** github.com/carrowmw
- **LinkedIn:** linkedin.com/in/carrow-morris-wiltshire-0492b6188/

Professional Summary

[TARGET ROLE: Full-Stack AI Engineer / Machine Learning Engineer / Data Architect] specialising in the deployment of resilient, decoupled AI pipelines for complex systems. Combines PhD-level expertise in spatial deep learning (GNNs, Transformers) with rigorous enterprise-grade software engineering (Docker, CI/CD, strict data contracts). Proven track record of architecting secure infrastructure, from city-wide mobility monitoring requiring enterprise-level data architecture to integrations with national digital twin pilot projects. Specialises in designing disconnected, scalable systems where individual components can be seamlessly upgraded, with domain applications spanning urban systems, defence, and enterprise FinTech.

Core Technical Skills

Technologies & Frameworks

- **Languages:** Python, TypeScript, SQL, Bash
- **Machine Learning & MLOps:** PyTorch, PyTorch Geometric, Scikit-learn, Weights & Biases (W&B), Optuna, MLflow
- **Data Engineering & Cloud:** GCP, Supabase, Docker, PostgreSQL, PostGIS, Redis, Parquet, Pydantic, Pandera
- **Backend & APIs:** FastAPI, NestJS, Node.js/Express, Prisma ORM, AsyncIO
- **Frontend & WebGL:** React, Next.js, Vite, Zustand, Deck.gl, MapLibre GL
- **Geospatial & Analysis:** GeoPandas, OSMnx, Turf.js, GEOSAM

Technical Concepts & Methodologies

- **AI Architectures:** Graph Neural Networks (GNNs), Transformers, Neural Processes, Masked Auto-Encoders (MAE)
- **Applied Mathematics:** Zero-Inflated Negative Binomial (ZINB) modelling, Uncertainty Quantification
- **Systems Architecture & Security:** CI/CD, Container Isolation, Decoupled Architecture, WebAuthn (Passkey), JWT Security, Vulnerability

Management
- **Operating Systems:** Linux (CompTIA Linux+ in progress)

Professional Experience & Core Engineering Projects

PhD Researcher in Resilient AI & Spatial Deep Learning

Geospatial Systems CDT, Newcastle University (Sept 2023 - Present)

- **Hybrid GAT-Transformer Architecture:** Designed hybrid Graph Attention Networks (GATs) and Transformer encoders to dynamically ingest data from irregular IoT traffic sensors, enabling accurate nowcasting and real-time uncertainty estimation even during severe sensor outages.
- **Geospatial Data Engineering:** Engineered custom preprocessing pipelines to ingest hundreds of millions of highly irregular time-series data points, augmenting dynamic data with static spatial context via Google Places APIs.
- **Advanced Statistical Engineering:** Solved "Null-Zero Conflation" in ambiguous sensor data via custom Zero-Inflated Negative Binomial (ZINB) loss functions, allowing architecture to differentiate hardware failure from zero-activity without prior imputation.
- **Uncertainty Quantification (UQ):** Developed XAI frameworks separating aleatoric and epistemic uncertainty, implementing "classification with rejection" workflows to build operational trust for engineering deployment.
- **HPC & MLOps Infrastructure:** Built production-grade, containerised Python 3.10 codebases. Enforced strict data validation (Pydantic/Pandera) and automated distributed hyperparameter studies (Optuna, W&B, 400+ tracked experiments) across private 4x A100 GPU servers via secure SSH.
- **Decoupled Architecture:** Designed "Artifact Twinning" data bridge decoupling a concurrent Python/FastAPI backend from a Next.js/Deck.gl frontend, integrating Redis and PostGIS for spatial interpolation at 60fps.

Turing Enrichment Student

The Alan Turing Institute (Oct 2025 - Present)

- **Advanced Model Architecture:** Developed "TransformerNP," a new Neural Process paradigm integrating with the DeepSensor framework for environmental forecasting. The model delivered faster training and inference times while significantly improving accuracy and the handling of extreme weather events.
- **HPC & Distributed Training:** Orchestrated model training on massive HPC clusters, utilizing SLURM workload managers to target specific compute nodes for optimized resource allocation.
- **Enterprise MLOps:** Centralized training infrastructure utilizing a Weights & Biases (W&B) enterprise license, enabling a 4-person research team to run, track, and collaborate on distributed experiments in a shared workspace.

- **Agile Engineering & CI/CD:** Managed project timelines and version control workflows via GitHub Kanban boards, enforcing rigorous code reviews, extensive branch management, and multi-reviewer pull request approvals.

Full-Stack Data Architect

Solo Architect – Open Banking Financial Dashboard (June 2025 - Present)

- **Zero-to-One Infrastructure:** Architected and deployed a decoupled TypeScript monorepo containing a NestJS API, React/Vite frontend (MUI X Charts), and strict shared data contracts via PostgreSQL/Prisma. Designed the system with scalable architecture natively capable of supporting thousands of concurrent users.
- **Live Data Ingestion & Compliance:** Navigated stringent security audits to secure a restricted Plaid Production License, enabling the live ingestion of 10+ years of historical transaction data across 10+ connected banking institutions.
- **Identity & Access Management:** Implemented strict WebAuthn (Passkey) 2FA protocols, JWT guards, network rate-limiting, and encryption-at-rest for sensitive financial access tokens.
- **LLM Integration (In Progress):** Currently architecting an LLM integration layer to analyze secure transaction history and deliver automated, personalized financial advice.

Technical Lead (NOVA Digital Twin)

Department for Business and Trade Sprint Challenge (March 2026)

- **Geospatial Engine & MCE Model:** Architected a scenario-based Multi-Criteria Evaluation (MCE) model within a Node.js backend. Implemented complex weighting systems allowing stakeholders (DBT, councils, developers) to define policy priorities and dynamically generate polygon suitability scores.
- **Interactive Digital Twin Features:** Engineered frontend tools allowing users to drop wind turbines and solar panels onto the map, utilising backend data models to provide real-time estimated capacity calculations for site development.
- **Dynamic Presentation Reporting:** Built a fully automated reporting system integrating React components inside Reveal.js presentations. Utilised Turf.js to analyse selected sites and generate printable slides detailing area, distance to grid/roads/protected sites, and solar/wind generation potential.
- **Agile Architecture & Leadership:** Managed project workflows via Kanban, acting as the technical liaison between data, policy, and research teams to translate stakeholder requirements into actionable architectural decisions. Enforced strict workflows with daily PR reviews and testing.

Project Manager & Technical Lead (Industry Collaboration)

Knight Frank Challenge Week (March 2025)

- **Cloud Infrastructure:** Led technical infrastructure for a 20-person PhD cohort, managing cloud environments (GCP/Workspace) for a national housing submarket classification project.
- **Computer Vision Pipeline:** Engineered a semantic segmentation pipeline using GEOSAM to process high-resolution satellite imagery (SPOT/Pléiades) for granular Grey Belt land identification.
- **Stakeholder Delivery:** Orchestrated and delivered the final technical presentation directly to Knight Frank stakeholders.

Team Lead (Industry Collaboration)

Airbus / Local Resilience Forum Sprint (April 2023)

- **Pilot Demonstrator Delivery:** Directed a cross-functional team of 12 researchers to deliver a situational awareness pilot for emergency responders within a strict one-month sprint, supporting a multi-million-pound grant application.
- **Multi-Modal Threat Integration:** Architected a dashboard integrating spatial risk-ratings for waste/fire hazards with ML-driven change-detection capabilities.
- **Real-Time NLP Pipeline:** Engineered and integrated a real-time social media scraping pipeline to ingest, parse, and rank the severity of localised hazard reports (e.g., snow/power grid failures) during crisis events.

Data Scientist

Verco Global (June 2022 - Sept 2022)

- **Predictive Modelling (PowerQuery):** Developed statistical models using PowerQuery to quantify climate impact on energy consumption across a massive portfolio of hundreds of enterprise sites spanning all continents.
- **Process Automation:** Automated legacy spreadsheet workflows, integrating siloed global sources into a single truth and reducing manual workload by 80 hours annually.

Academic Teaching & Mentorship

- **Complex Systems & Sustainability:** Mentored over 48 engineering students (8 groups) across Civil Engineering and Public Policy modules, providing feedback on analysing and communicating "wicked" global problems through the lens of ethics and sustainability.
- **Product-Driven Data Science:** Mentored Masters-level Computer Science students on the "Data Science in the Wild" module (in partnership

with Google), advising on how to structure real-world data pipelines to build "sellable" data products.

- **Advanced Machine Learning:** Delivered highly commended guest lectures on Time-Series Prediction and Graph Neural Networks (GNNs). Assessor for Masters-level Python Machine Learning coursework.
- **Core Programming & Data Analysis:** Taught, demonstrated, and marked foundational programming modules, including Java (Comp Sci), Python (Engineering/CS), and an R-based curriculum focused on structuring data science projects for a Data Science Masters.
- **Practical Field Engineering:** Instructed advanced practical surveying modules (Total Stations, Leica Pegasus Backpack mobile mapping, UAVs, GPR).

Academic Foundation & Publications

- **Education:** PhD in Urban Digital Twins (In Progress), MRes Geospatial Data Science (Distinction), MEng Civil and Structural Engineering (First Class Honours)
- **Selected Publications:** IJCNN 2026 (Spatiotemporal DL), EGU 2026 (Neural Processes), OFEME/CUPUM 2025 & GISRUK 2024 (Hybrid GNN-Transformer Models for Urban Mobility)