PETER McCANN STRAIN

Founder & CTO · AI Engineer

Oxford AI DPhil · Agentic AI · LLMs · Computer Vision · Deep Learning · Data Science — built across medicine, aerospace, recruitment & legal-tech · petermccannstrain.com

AI engineer and founder, nine years and an Oxford DPhil in. I've taken AI from satellite vision at Astroscale and medical imaging at Oxford to enterprise agent systems inside SThree, a FTSE 250, where I was one of two engineers who built the AI function from nothing. A working prototype helped secure **£25M** of board investment, and the team grew from two of us to ten. I'm now founder and CTO of two ventures going to market, Praviar in patent intelligence and Conversico in voice operations for dental practices, both of which I designed and built end to end myself.

The half of AI I care about is the unglamorous one: evaluation, cost, the failure modes nobody demos, and keeping a person in charge of the decision. I write about it at petermccannstrain.com, in an essay series called Architecting the AI Coworker. There's also an AI on the site, trained on my work — ask it whether we'd build well together. I'd rather you interrogate the work than take my word for it.

PROFESSIONAL EXPERIENCE

Founder
2025 – Present

Praviar · AI patent intelligence

Freedom-to-operate is a costly bottleneck in pharma R&D, where one missed patent can sink a programme. Praviar is an AI patent-intelligence platform I conceived, designed, and built end to end. In production, going to market.

- The AI does the reading and the reasoning; deterministic checks and a human reviewer own every conclusion. It reads a patent as evidence, not as text to summarise.
- A computer-vision cascade pulls chemical structures, including the generic Markush forms, straight out of the patent drawings, so the system never relies on text that omits the figure that decides the answer. Validated on a held-out set.
- Built it solo across the stack — web app, service layer, and a checkpointed, resumable analysis engine, so a heavy patent run can be inspected, retried, and continued rather than restarted, with structured validation enforced at every analysis boundary and an internal benchmark tracking recall on known prior art, because here a missed patent is the expensive error.

Co-Founder & CTO
2025 – Present

Conversico · AI voice operations

An AI voice receptionist for UK dental practices, in production and going to market. I designed and built the entire system and wrote all of its code; my co-founders run go-to-market.

- Greeta, the voice agent, answers the calls a busy front desk misses, books and reschedules appointments, and hands back to a human when it should.
- Engineered to fail safely under real conditions: fail-closed tenant isolation so one practice's data can't reach another's, caller-safe degradation when an upstream service breaks, booking safeguards against double-booking, and launch gates that hold release until the evidence clears.
- Built the practice-operator dashboard for non-technical front-desk staff — the live view of recovered calls, bookings, and where Greeta hands back to a human — and own the whole operational spine alone: backend, frontend, end-to-end tests, infrastructure-as-code, CI/CD, and runbooks. A production voice system one engineer can release, observe, and run.

Founder (independent build)
2024 – Present

Personal Assistant AI

A working prototype for an executive-assistant product, which I run every day as my own first user.

- I fire a half-sentence into a messaging app; it works out what I meant, turns it into a structured note, and links it to the people and projects it mentions.
- Eight specialist agents cover diary, memory, research, review, and planning; a layered memory engine moves from keyword to semantic to relationship-graph recall; and an overnight worker fleet runs everything from a morning brief to a nightly reflective phone call.
- A five-role Socratic engine — tutor, devil's advocate, clarifier, connector, and an assessor that keeps it from dominating — for thinking a problem through rather than just filing it, with a quality loop that recalibrates against my own judgement.

Founding AI Engineer
Sept 2024 – Present

SThree plc (FTSE 250)

*One of two founding engineers who built SThree's AI capability from nothing. Helped secure **£25M** of board investment with a working prototype, set the engineering standards, hired and mentored the team from **2 to 10**, set up the company-wide AI education programme, and carried AI delivery into board-level reporting.*

- Designed and built **Aurum** from scratch — SThree's enterprise agentic AI for recruitment: a multi-agent platform that runs controlled hiring simulations end to end, with phased orchestration over reviewable state, synthetic candidate generation, parallel agent dispatch, and durable checkpointing so long runs can be inspected, retried, and resumed.
- Built Aurum's safety and evaluation into the architecture rather than bolting them on: an evaluator layer that catches model disagreement and failure modes before anything reaches a real candidate, statistical fairness diagnostics, adversarial red-team testing, and audit-ready evidence for every decision — all wrapped in a full operator product for configuring runs, inspecting evidence, and reviewing risk, because a hiring decision is high-stakes and hard to undo.
- Built a research layer on top, on self-improving evaluators that rewrite their own scoring instructions (DSPy / GEPA), benchmarked across optimisers and five seeds with frozen data and full evidence governance — the rigour that keeps an evaluation honest.
- Co-designed the production recruitment platform end to end: candidate matching, real-time AI voice interviews, human approval gates, and client reporting.

CONTACT

peter.mccann.strain@gmail.com
+44 7539 882256
petermccannstrain.com
linkedin.com/in/petermccannstrain
United Kingdom

EDUCATION

DPhil (PhD), AI / Clinical Medicine
University of Oxford · 2017–2022
Big Data Institute · MRC **£125K** studentship · Supervisor: Prof. Cecilia Lindgren

MPhys (Integrated Master's), Physics
University of Strathclyde · 2012–2017
Computational Physics, Molecular Dynamics, Complex Systems
Master's project: molecular-dynamics simulation for enhanced oil recovery

Continuing Education
MLOps Engineering (AWS) · Innovation Strategy & Leadership (Oxford)

RESEARCH & WRITING

Author, "Architecting the AI Coworker" — a 22-essay public series on evaluation, cost, and reliability in agent systems · petermccannstrain.com/essays · also on Substack and Medium

Deep-research agent benchmark — 13 orchestration patterns scored by a three-judge reliability panel with equivalence testing and citation/factuality audits, testing whether well-orchestrated small models, fine-tuned with reinforcement learning, can match frontier pipelines.

Self-improving evaluators — reflective prompt evolution (DSPy / GEPA) measured against a held-out key, reported with negative results in full.

AstroGAN — closing the simulation-to-reality gap for satellite imagery.

AWARDS & FUNDING

MRC Studentship
£125,000 · competitive national award (2017–2022)

CivTech 11.5 Exploration Grant
£5,000 · won as a solo applicant against 24+ companies (2025)

NL Education Trust
£2,250 · academic excellence, 3 consecutive years

Graduate of the Term
Peer-nominated, 2× · St Anne's College, Oxford

Gold Saltire Award
250+ hours of community service

- A separate two-week win with outsized payoff: took the language model out of the hot path for production CV skill-tagging, cutting cost roughly **180×** (≈£1,800 → £10 per million CVs) and latency about 100× with retrieval plus a small-model ensemble instead of a bigger model.
- Led SThree's first causal study of recruitment timing across **9,402 placements**: same-day submissions did worse than ones sent around three days later, and the recruiting team mattered far more than speed. A counter-intuitive, commercially material result, delivered to leadership with a recommendation to confirm it in a controlled experiment.
- Built SThree's first data-driven consultant-performance framework and took it to the COO as a decision tool — overturning three inherited management beliefs: revenue and placement success are separate dimensions, several rewarded activities added no real value once controlled for, and top performers come in distinct types rather than one mould.
- Embedded AI across the multinational beyond my own builds — working with the COO, legal, compliance, and product on GDPR-safe data handling and on drawing the lines where human oversight stays non-negotiable.

Co-Founder

March 2023 – Present

Datamise Solutions

GMB Union needed its equal-pay claims digitised, and the vendors it approached quoted heavily for locked-in systems. We built it directly: a two-person company, profitable in its first year.

- Built GMB's equal-pay claims infrastructure: structured intake replacing paper, automated membership verification, contract generation from verified data, and e-signature — one auditable chain from first claim to e-signed contract — the infrastructure behind a **£300 million** equal-pay settlement for women workers across UK local authorities.
- Built **SettleMise** end to end, the engine that turns an equal-pay dataset into a reviewable settlement schedule (Next.js, Clerk, encrypted artifacts, full audit trail). Across 50 runs spanning generated datasets and parameter variations, it matched an independent ground-truth calculator **to the penny**, with zero mismatches — because a number a tribunal might see has to be right, not roughly right.
- Ran Datamise's forensic audit of GMB's gold-standard equal-pay base, the template cloned to every regional campaign, surfacing critical data-loss bugs that would silently delete real claimant submissions on any clone — then handed back the architecture redesign and a costed modernisation roadmap to fix them.
- Traced why the base had gone dark — validation had been failing silently on most submissions for months — and pinned the duplicate-handling and sort faults that would delete the wrong claimant record the moment automations were re-enabled in any regional clone.

Computer Vision Engineer

April 2022 – Sept 2024

Astroscale Ltd (ELSA-M mission)

Hired on a three-month contract to make deep learning work where classical computer vision had stalled. Converted to full-time and stayed two and a half years.

- Built systems that estimate a satellite's exact position and orientation from camera images, for autonomous rendezvous in the ELSA-M debris-removal mission context, with a custom training objective that scores the model against the real geometry of the scene rather than a stand-in for it.
- Designed and built **AstroGAN** to close the gap between simulated and real orbital imagery — the failure mode that breaks space-vision models trained in a simulator. Built on an unpaired-image approach with a custom multi-loss objective, it sharply improved the realism of synthetic orbital imagery on held-out frames.
- Worked the model space for the pose problem, from CNN baselines through Vision Transformers to multi-task designs, and built an adaptive per-image preprocessing pipeline that lifted accuracy on the hardest real-camera lighting conditions.
- Mentored engineers in computer vision and ran the team's journal club on emerging research.

DPhil (PhD) Researcher — AI / Clinical Medicine

2017 – 2022

University of Oxford

Funded by a competitive **£125K** MRC studentship at the Big Data Institute. Thesis: Computational Insights into Mouse Placental Histology

- Built end-to-end deep-learning pipelines for placental histology: a detector that finds individual cells across enormous whole-slide images, an 8-class cell classifier at **89% accuracy**, automatic tissue-boundary tracing, and self-supervised representation learning on a 3-million-image unlabelled set to learn cell structure without labels.
- Turned those model outputs into structured measurements, compared normal against gene-edited placentas, predicted which embryos would not survive, and cross-checked the findings against the biological literature so interpretation stayed tied to evidence.
- Built the annotation and quality-control workflow behind the labelled training data — the foundation the whole pipeline stood on — across 8 placental cell types and gigapixel whole-slide images.
- Developed custom methods to make whole-slide analysis tractable at scale: radial clustering to merge the duplicate cell detections produced at tile boundaries, and concave-hull tracing to pull out tissue regions automatically.

CivTech 11.5 Innovator

Sept – Nov 2025

Scottish Government Programme

- Selected solo against 24+ established companies for the Scottish Government's CivTech exploration phase. In three weeks I designed a human-in-the-loop AI system for Scotland's eight statutory impact-assessment workflows: manual-first, with a supervisor-worker agent layer behind approval gates so a person always decides what enters an assessment.
- Mapped all eight statutory assessment types — BRIA, CRWIA, EQIA, FSDA, ICIA, SEA, Consumer Duty, and DPIA — into a single seven-phase workflow, from screening through to longitudinal review.
- Specified the agent layer as a supervisor-worker LangGraph design of research, analysis, synthesis, and quality agents, sitting behind a human workflow engine that runs with no AI at all — so the system can be adopted manually first and made agent-progressive as trust builds.

LEADERSHIP & GOVERNANCE

Investment Committee Member

St Anne's College, Oxford (2021–2024)

Sole student member among 8 senior professionals overseeing a **£70M** endowment — performance, asset allocation, and ESG policy.

MCR Social Secretary

St Anne's College, Oxford (2018–2022)

4.5-year tenure (4× the usual), **£27K** budget, mentoring incoming graduates.

Earlier roles

Undergraduate Physics Representative (5 years), Physics Ball Chair, College RFC Social Secretary, School Captain / Head Boy.

CORE COMPETENCIES

Agentic AI & Orchestration

Multi-agent orchestration, parallel fan-out dispatch, evaluator architecture, reflective self-improvement loops, long-horizon planning with durable checkpointing and resume · LangGraph, DSPy

Evaluation, Safety & Governance

LLM-as-judge panels with inter-rater reliability, equivalence testing, citation-provenance and factuality audits, bounded fairness diagnostics, red-teaming, human-in-the-loop authority, claim-ledger evidence governance

Cost & Reliability Engineering

Taking the LLM out of the production hot path (≈180× cheaper, ≈100× faster), small/local-model substitution, prompt and semantic caching, checkpoint-and-retry resilience, fail-safe degradation · LangSmith

LLM & Multimodal

Hybrid and agentic RAG, schema-validated structured outputs, LLM-as-judge evaluation, hot-path LLM removal, real-time AI voice, vision-language review · DSPy, QLoRA / GRPO fine-tuning

Generative & Creative AI

An end-to-end generative-media pipeline — AI photo restoration, identity-locked image-to-video, and vision-language quality auditing — that restored and animated a family photo archive into a short film

Deep Learning & Computer Vision

Architectures trained from scratch — Vision Transformers, GANs with custom losses, self-supervised learning — for sim-to-real domain adaptation, satellite pose estimation, whole-slide medical imaging · PyTorch

ML & Statistics

Causal inference (propensity and double-robust methods), equivalence testing, calibration, clustering and dimensionality reduction, logistic-regression heads on frozen embeddings · XGBoost, scikit-learn

Production & Leadership

Sole engineer of entire production AI products, end to end — frontend, backend, and infrastructure (Praviar, Conversico, SettleMise, Personal Assistant AI); scaled a FTSE 250 AI team 2→10 and set its standards · Next.js, TypeScript, Python, AWS, Azure, CI/CD

Platforms & Infrastructure

AWS, Azure, Vercel · Docker, CI/CD · PostgreSQL, Redis, Upstash, FAISS, vector databases · Clerk / NextAuth, RBAC · Playwright, Vitest, MLflow

Domains

Recruitment AI · Legal-tech · Patent intelligence · Medical imaging · Satellite & aerospace vision · Voice operations · Generative media · Public-sector AI