



Capability catalogue

Machine learning for work where being wrong is expensive.

Reinforcement learning · Medical informatics
Biomaterial research · Data science
Explainable AI · AI governance · Software engineering

Nodal Savvy Limited · RC 1802674
Hanwa, Zaria, Kaduna State, Nigeria
info@nodalsavvy.com · +234 703 555 1484 · nodalsavvy.com

What we do

Nodal Savvy applies research-grade machine learning to problems in healthcare, materials science and regulated industry. Eight practices, set out below. Most engagements draw on several of them at once.

Reinforcement learning	Sequential decisions under uncertainty: treatment schedules, process control, adaptive experiments.
Medical informatics	Clinical and health data that is messy, sensitive and governed.
Biomaterial research	Formulation, process and characterisation data connected to properties you can design for.
Data science	The groundwork that decides whether anything built on top of it will hold.
General-purpose AI	Foundation models pointed at real workflows, with an honest test of fit.
Explainable AI	Model reasoning made legible to the person who acts on it and the person who signs it off.
AI governance and assurance	Controls, documentation and evidence to deploy, audit and defend an AI system.
Software and web engineering	The systems that turn a model into something people can actually use.

Who we work with

Hospitals and health systems. Medical device, pharmaceutical and materials manufacturers. Research institutes and universities. Industrial process operators. Any organisation that has to deploy an AI system under regulatory scrutiny, or explain one to an auditor, a clinician or a board.

Delivery is remote across Nigeria and internationally, and on site where the work requires it — laboratory instruments, or clinical systems that cannot be reached from outside.

Who we are

The person behind the work.




Dr Emmanuel Okafor

Founder, Executive Director

A PhD in Artificial Intelligence from the University of Groningen, and Assistant Professor at the Center of Artificial Intelligence for Security at Naif Arab University for Security Sciences in Riyadh. Previously a postdoctoral fellow at the SDAIA-KFUPM Joint Research Center for Artificial Intelligence, and a decade on the faculty at Ahmadu Bello University in Zaria.

The practices in this catalogue are not adjacent to that research. They are that research. Reinforcement learning for industrial control and robotics. Explainable diagnostic systems for lung disease and skin lesions. Machine learning for hydroxyapatite scaffolds and perovskite property prediction. Each was put through peer review before it was ever offered to a client.

Peer-reviewed publications by year



Twenty-one journal articles, ten conference papers and a book chapter. Principal Investigator on funded research into AI-driven threat detection and explainable anomaly detection; contributor to MANTIS, a €30 million EU consortium of 47 partners.

Qualifications

- PhD, Artificial Intelligence — University of Groningen
- MSc, Control Engineering, with distinction — Ahmadu Bello University
- BEng, Electrical Engineering — Ahmadu Bello University
- Registered Engineer — COREN, R54313
- MIT Empowering the Teachers Fellow, 2022

Selected publications

- Control optimization of steam boilers via reinforcement learning. IEEE Access, 2025.
- Explainable deep learning diagnostic system for prediction of lung disease from medical images. Computers in Biology and Medicine, 2024.
- Deep reinforcement learning with light-weight vision model for sequential robotic object sorting. Journal of King Saud University - CIS, 2024.
- Explainable machine learning for predicting the band gaps of ABX3 perovskites. Materials Science in Semiconductor Processing, 2023.
- Ensemble learning prediction of transmittance at different wavenumbers in natural hydroxyapatite. Scientific African, 2020.

Research and domain practices

Four disciplines where the subject matter, not the algorithm, decides whether a project succeeds.

Reinforcement learning

Decisions taken in sequence, where each one changes what comes next.

- Offline and batch RL from historical logs
- Off-policy evaluation with confidence bounds
- Contextual bandits and Bayesian optimisation
- Simulator and reward design, with reward hacking review
- Constrained and safe policies for regulated settings

Medical informatics

Clinical data, rarely collected for the question you are asking.

- EHR extraction, cohorts and computable phenotypes
- Mapping to OMOP CDM and FHIR
- De-identification and information governance
- Models reported to TRIPOD+AI
- Calibration, subgroup performance, survival analysis

Biomaterial research

Small, expensive, unbalanced datasets that defeat standard methods.

- Structure-property and composition-property modelling
- Design of experiments and active learning
- Image analysis for micrographs and scaffold morphology
- Surrogate models for expensive simulation
- Degradation and shelf-life modelling from few samples

Data science

The measurement and pipeline work that everything else depends on.

- Data audits: what exists, what is trustworthy, what is missing
- Pipelines, warehousing, reproducible environments
- Causal inference and experiment design
- Forecasting and demand modelling
- Monitoring, drift detection and retraining policy

Applied across every engagement

Four practices that run through the work regardless of subject, and which increasingly decide whether a system is allowed into production at all.

General-purpose AI

Language and foundation models, evaluated before they are adopted.

- Retrieval over your documents, cited back to source
- Evaluation sets built on your material, before model choice
- Structured extraction from notes, reports and literature
- Fine-tuning and adapters where prompting has run out
- Guardrails, abstention and human review points
- Cost, latency and lock-in modelled before you commit

Explainable AI

Explanations tested with the people who must rely on them.

- Intrinsically interpretable models where accuracy allows
- Feature attribution, with its limits stated plainly
- Counterfactuals: what would have had to differ
- Saliency maps checked against expert judgement
- Calibrated uncertainty, and knowing when to abstain

AI governance and assurance

The evidence that lets a system be deployed and defended.

- EU AI Act risk classification and conformity readiness
- ISO/IEC 42001 design and gap analysis
- NIST AI Risk Management Framework mapping
- Model cards, datasheets, intended-use statements, audit trails
- Impact assessments, bias evaluation, post-market monitoring
- Independent review of systems built by others

Software and web engineering

Delivery, not a notebook and good wishes.

- APIs and services with versioning and audit trails
- Internal tools, research portals, annotation interfaces
- Containerisation, CI/CD, cloud or on-premise deployment
- Accessibility and performance as delivery requirements
- Handover documentation written for your engineers

How we work

Three engagement shapes, chosen by how well defined the problem is. Most clients begin with discovery.

Discovery sprint

Two to three weeks, fixed price

We assess your data, constraints and prior attempts, then tell you what is achievable and what it costs. Governance and approvals are reviewed too, since they set the timeline more often than the technical work.

You end up with A written recommendation, a costed plan, and a clear statement of what would make the project fail.

Scoped build

Six weeks to six months

Delivery against an agreed specification, alongside your team, with fortnightly demonstrations of something that runs.

You end up with Working software, the evaluation that justifies it, and documentation written for whoever maintains it.

Embedded advisor

A few days a month, rolling

For teams with capability in house who want a second opinion: design and evaluation review, technical due diligence, and support hiring.

You end up with Fewer expensive detours, and a team measurably stronger than last quarter.

Fees

Discovery is a fixed fee. Builds are quoted against a written scope, with the rate for any addition stated upfront, so the cost of a change is known before you request it. Every individual's rate appears in the proposal. Proposals are unbilled; third-party costs are passed through at cost.

What you can expect from us

- A negative finding delivered in week two, not month six, if the data cannot support the question.
- Code, models and documentation handed over with tests and training, so you can remove us without disruption.
- Success defined and agreed in writing before work starts, against the outcome that matters to you.

Terms and company details

Stated here rather than discovered at contract stage.

Who does the work	The people named in the proposal. Any substitution is agreed with you in advance, never reported afterwards.
Your data	Processed inside your environment wherever possible. Where it must move, it moves under a written agreement, encrypted, and is deleted on request with written confirmation. Never used to train models for anyone else; never sent to a third-party AI service without your specific written consent.
Intellectual property	You own the code, models, documentation and derived data from the moment it is written, not on final payment. We retain general know-how only, and publish nothing about an engagement without your written approval.
Our use of AI tools	We use AI coding assistants in delivery. Output is reviewed by the named engineer, and your confidential material never enters a third-party tool without the consent above.
What we decline	We are not a law firm, notified body or trial sponsor, and work alongside your advisers on regulatory matters. We do not act for directly competing clients concurrently. We decline covert surveillance work, and systems making consequential decisions about people with no route to human review.

The company

Registered name	NODAL SAVVY LIMITED, a private company limited by shares.
Registration	RC 1802674, registered with the Corporate Affairs Commission of Nigeria on 6 June 2021. Verify at icrp.cac.gov.ng/public-search .
Location	Hanwa, Zaria, Kaduna State, Nigeria.
References	We publish engagement summaries only with written client agreement. Once a conversation is serious we arrange a reference call and share a redacted example deliverable.

Start a conversation

Describe the problem, not the technique you think you need. The first call is thirty minutes at no cost, and ends with a clear view of whether this is worth pursuing.

info@nodalsavvy.com • +234 703 555 1484 • nodalsavvy.com