BLOG 01 — LAUNCH POST · AI GOVERNANCE

I Built My Own AI Degree. It Led Me Somewhere I Didn't Expect.

Twelve months, three universities' worth of curriculum, and one conclusion I couldn't shake.

Robin Beetge Co-founder, laibrary 12 min read

Twelve months, a study plan, and a very patient collaborator.

Last September I did what quite a few people in senior technology roles were quietly doing at the time. I went shopping for an AI education.

MIT. Johns Hopkins. Stanford. Serious institutions running serious programmes. Twelve to eighteen months, real mathematical rigour, and price tags to match. I read the prospectuses properly, breaking them down into the sum of their parts while realistically contemplating how I was ever going to fit a new education into my schedule.

But I kept snagging on the same thought.

A university curriculum takes roughly two years to design, review and approve. The field it describes reinvents itself about every eleven weeks. I looked at module lists written before anyone outside a research lab had said the words "tool use" or "context window" out loud, and I did the arithmetic. Enrol now, graduate in eighteen months, and a meaningful slice of your syllabus has quietly become archaeology.

To be fair to the universities, that criticism only lands on part of the programme. Linear algebra does not rot. Probability theory is not going anywhere. The foundations are the foundations precisely because they outlast the applications sitting on top of them.

But the applications are where the work is. And the applications were changing faster than any accreditation committee on earth could move.

So I did the obvious, slightly mad thing. I pulled the published curricula from all three, laid them side by side, kept the skeleton, and rebuilt the rest myself.

Well. Claude and I rebuilt it. Twelve months, at a kitchen table, with a study plan and a very patient collaborator.

Periodisation, but for your brain

I coach cyclists. The first thing you learn as a coach is that you cannot skip base. Everybody wants to do intervals in January because intervals feel like progress. The athletes who win in August are the ones who spent the winter doing dull, unglamorous aerobic work while everyone else was collecting Strava segments.

Learning AI is precisely the same, and almost nobody does it. The temptation is to start at the shiny end — prompt engineering, agent frameworks, whatever launched last Tuesday — because that end has immediate utility. It also has a shelf life measured in months.

So I periodised it. Base, build, peak. And base meant going back further than I'd have liked.

PHASE ONE

Vectors and calculus

Matrices, vector spaces, eigendecomposition, then multivariable calculus and the chain rule — makes the entire field possible (although I can't claim more than a few hours of Khan Academy and 3Blue1Brown videos).

PHASE TWO

How the machine actually works

Embeddings, attention, the transformer architecture. Why context windows behave the way they do. Why models hallucinate.

PHASE THREE

The modern stack

MCP layers. Tools. Skills. Orchestration. Agents. The week it stopped feeling like a chatbot. Things started to get interesting. I started to hear topics I was reading about come up in daily conversation, and started weaving them into our product strategy and roadmap.

PHASE FOUR

Ethics, governance and deployment

Meant to be a month. I've now been in phase four for the better part of a year. I found gaps. Disconnects intensified by the speed of change and the ability to harness and drive that change in the intended direction.

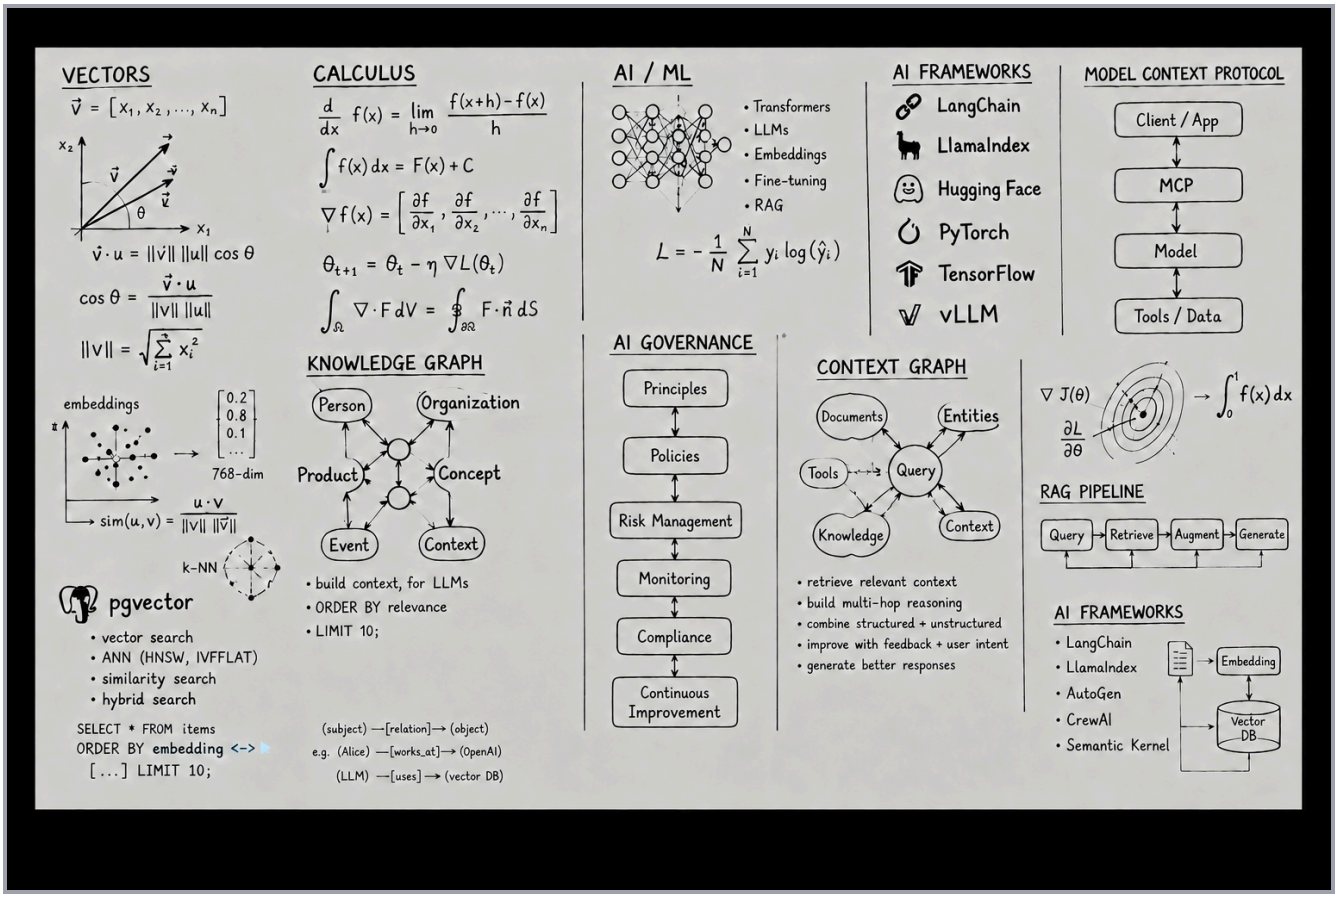

A DIY AI curriculum

Phase one was vectors and calculus. Genuinely humbling. I first got Claude to draft this blog, and he stated I had not thought seriously about eigenvectors since university - facts is, I'd never thought about them. Matrices, vector spaces, eigendecomposition, then multivariable calculus - undecided whether right of passage or self flagellation.

But there is payoff for all that suffering. Once you actually understand gradient descent, the mystique evaporates. You are standing on a hillside in thick fog, feeling the slope under your feet, and stepping downhill. That is it. That is the engine of the revolution. Do it a few trillion times across a few billion parameters and you get something that can write a sonnet and debug your Python.

Understanding that did more for my judgement about AI than any number of think-pieces. You stop being mystified, and mystification is the enemy of good decisions.

Phase two was how the machine actually works. Embeddings, attention, the transformer architecture. Why context windows behave the way they do. Why models hallucinate, which stops being a scandal and starts being an entirely predictable property of the thing once you see the mechanism.

Phase three was the modern stack, and this is where the course I'd designed started outrunning the courses I'd been shopping for. MCP layers. Tools. Skills. Orchestration. Agents.

There was a specific week where the thing stopped feeling like a chatbot and started feeling like a colleague with API credentials. It could read my files, call my systems, run my code, remember what we'd agreed last Tuesday and act on it without being asked twice.

That week changed the trajectory of everything that followed.

Phase four was meant to be a month on ethics, governance and deployment. The tidy little capstone module every AI curriculum bolts on at the end, usually taught by someone from the law faculty, usually the module everyone skims.

I've now been in phase four for the better part of a year.

Every road led to the same door

Here is what kept happening. I'd go deep on something technical — retrieval architectures, agent orchestration, evaluation methodology — and I'd surface with the same question every time.

> ### Who is accountable when this thing acts?
>
> Not in a philosophical sense. In an operational one. Which human. Which control. Which record. Which evidence, produced how, retained where, shown to whom.

So I followed that question into the frameworks. And the frameworks are where things got genuinely strange.

People will tell you there are hundreds of AI frameworks. That's true and useless, in the way "there are thousands of laws" is true and useless. The number depends entirely on what you're counting, and until you sort them by what they actually do to you, the count tells you nothing.

There are eight kinds.

Once you see the eight, the landscape stops being a fog bank and becomes a map.

01	**Comprehensive AI statutes**	Laws that cover AI horizontally, across every sector. There are three in force worldwide. The EU AI Act, Korea's AI Framework Act, and Japan's Promotion Act, which is a statute with no teeth by deliberate design. Brazil is close. That's the whole list. The thing everyone talks about is the rarest thing on the board.
02	**International treaties**	Exactly one: the Council of Europe's Framework Convention on AI. The first binding international AI treaty in history, signed by the UK, the US, Canada, Japan and the EU, ratified by almost nobody yet, and watched by almost nobody outside the profession. Ratification status is one of the best leading indicators we have of where binding law arrives next.
03	**Sub-national statutes**	This is where the actual obligations live if you're American. Texas, California, Colorado, Illinois, New York. Over two thousand bills introduced. And here's the part that catches people: reach attaches through where your users are, not where your company is. You cross into a state's jurisdiction silently, through ordinary growth, usually discovering it via a customer's security questionnaire.
04	**Sectoral regulators**	The rules that already bound you before anyone said "AI Act." Banking supervisors replaced fifteen years of model risk doctrine in April. Insurance regulators in more than twenty US states now require a written AI programme, examined through market conduct. The FDA reversed its position on clinical decision support in March. If you're in a regulated industry, these bind you today, and most AI governance conversations ignore them entirely.
05	**Certifiable management standards**	ISO/IEC 42001 and its family. The first AI management system standard you can actually be audited and certified against — and the certification scheme only became real in 2025, when the standard governing the auditors themselves was published. This is the closest thing we have to a globally portable governance credential. SOC 2 sits just outside this tier and is worth the distinction: it is an attestation your customer reads about controls an auditor tested, where 42001 is a certification you hold.
06	**Risk management frameworks**	Chiefly the NIST AI Risk Management Framework. Entirely voluntary, and simultaneously a statutory safe harbour in Texas, the baseline for US federal agencies, and the most common questionnaire in American enterprise procurement. A voluntary framework that became load-bearing by accident.

07	**Security and threat taxonomies**	OWASP's Top 10 for LLMs and for agentic applications, MITRE ATLAS, the Cloud Security Alliance's controls matrix. These are where governance stops being paperwork and becomes test cases. The 2026 OWASP list moved Excessive Agency up to number three, built on nearly eight thousand real-world incidents. That single ranking change tells you more about where this is heading than most published research.
08	**Principles and soft law**	OECD, UNESCO, the UN, the G7. Non-binding, routinely dismissed, and quietly the most influential tier of all. When the EU needed a definition of "AI system" for a law carrying fines of up to seven percent of global turnover, it didn't invent one. It reached for the OECD's. Principles don't bind you. They supply the vocabulary that the things which do bind you get written in.

Eight tiers. Roughly nine hundred technical standards, ninety-odd national frameworks, and a few dozen genuinely binding laws.

And underneath all of it, once you strip the jurisdictional packaging away, the same twelve requirements:

01 **Know what you have.**	02 **Classify it.**	03 **Manage the risk.**	04 **Assess the impact.**
05 **Govern the data.**	06 **Document it.**	07 **Log it.**	08 **Disclose it.**
09 **Keep a human in the loop.**	10 **Make it accurate and secure.**	11 **Watch it after launch.**	12 **Govern your suppliers.**

Twelve common steps. Endless permutations.

The conclusion I couldn't get away from

Two things became true for me at roughly the same time, and they are both true at once.

The first: this is a genuine industrial revolution. Not a hype cycle, not a bubble with a good story. I have spent thirty years in enterprise software watching things get called transformational that were, on inspection, features. This isn't that. Steam, electricity, the microprocessor, the internet — AI belongs in that sequence, and I'd argue it arrives faster and reaches further than any of them.

The second: the management layer is completely, blatantly missing.

And this is the bit that made me put down the textbooks.

We spent a hundred and fifty years building the discipline of managing human workers. Think about what that actually involves.

150 YEARS OF MANAGING PEOPLE	HOW WE DEPLOY DIGITAL WORKERS
You write a job description before you hire.	No job description.
You check references.	No access review.
You onboard.	No performance record.
You provision access deliberately — this system yes, that one no.	No escalation path.
You set objectives.	No offboarding.
You review performance.	It reads our documents, calls our APIs, talks to our customers and moves our money.
You keep a record.	
You have an escalation path when something goes wrong.	
And on the last day, you take the badge back, revoke the credentials, and document the exit.	

None of that is bureaucracy for its own sake. All of it is hard-won institutional learning about what happens when capable agents act on your behalf, with your authority, in your name.

Now look at how we're deploying AI.

We hand a digital worker credentials, tool access, persistent memory and increasing autonomy. It reads our documents, calls our APIs, talks to our customers and moves our money. And we give it no job description. No access review. No performance record. No escalation path. No offboarding.

Over half of organisations cannot produce a list of their AI systems.

Read that again with the HR analogy in your head. It means over half of organisations cannot list their employees.

Nobody would run a company that way. We are all currently running our AI that way.

So we designed laibrary

What have you got to prove?

That's the whole origin story. There's no dramatic pivot in it. I went looking for an education, built one, and the education kept depositing me at the same unattended door.

The thesis is deliberately narrow. AI governance is a workflow and evidence problem, not a policy-document problem.

Most of what's sold as AI governance today produces documents. Policies, registers, assessments, attestations. All necessary. None sufficient. Because when a regulator, an auditor or a notified body eventually turns up, they will not ask whether you have a policy. They will ask you to demonstrate that a specific control operated, on a specific date, for a specific system, and to show them the artefact.

That's not a document problem. That's a workflow problem with an evidence chain attached. Which is, as it happens, exactly the kind of problem enterprise software has been solving for thirty years, just never pointed at this.

So laibrary treats every AI system and every agent as something you inventory, classify, govern, evaluate continuously and offboard deliberately. It maps obligations across frameworks so that one fix closes findings in several places at once, because those twelve requirements really are the same twelve requirements wearing different regulatory hats. And it produces the evidence as a by-product of the work rather than as a scramble three weeks before an audit.

I'll be writing a lot more about all of it here. The eight tiers get a piece each. The frameworks get pulled apart properly. The deadlines get tracked, because there are more of them arriving sooner than most people realise.

But this is the argument underneath everything else, so I'll say it plainly and then get out of the way.

Governance is not the brake pedal. It's the instrument panel.

You can't go fast in the dark.
