



NOIRIS AI

INTELLIGENCE · FOR THOSE WHO SEE FIRST

PROGRAM SYLLABUS · COHORT 7

The *Quantum AI* Mentorship AI operator training

Eight weeks. A real operation of yours running in production at the end – plus the offer, the price and the contract to sell it again.

8

WEEKS

16

LIVE SESSIONS

41h

OF CLASS

16

ARTIFACTS

2×

PER WEEK

DANIEL ROCHA · FOUNDER & CEO, NOIRIS AI · [NOIRISAI.COM/QUANTUMAI/EN](https://noirisai.com/quantumai/en)

Tool courses expire. A craft does not.

Most AI programs teach a tool. Tools change every six months and the student is back at the start. Quantum AI teaches an operation: looking at an expensive, slow human process, deciding what the machine does, building it, getting it standing, keeping it standing — and charging for that.

- **No make-believe exercises.** In Week 1 the student picks a real operation — from their own company, their own job or a pilot client — and it runs through all 16 sessions. Every artifact is a piece of that operation.
- **Measure before you touch.** Baseline in Week 2, measured again in Week 8. The result presented is the student's number, not a promise of ours.
- **Technical and commercial in the same program.** Whoever builds and doesn't sell becomes a hobbyist. Whoever sells and doesn't build becomes a middleman.

What this program is not. It is not a programming course — nobody leaves a software engineer. It is not a machine learning course — nobody trains a model from scratch. It is not a course in “magic prompts” — here a prompt is an engineering specification, with tests and an acceptance criterion. And no revenue is promised: what is promised is method, artifacts and a working operation.

What you build, week by week

Four phases. You start by measuring what the process costs today and finish by presenting the new number, the offer and the contract.

PHASE I · WEEKS 1–2

Observation

See the system before touching it. Map the real process, measure the baseline and specify the task with a test and an acceptance criterion.

You leave with: map, baseline and a tested prompt

PHASE II · WEEKS 3–4

Superposition

Build the parts: orchestration, real integrations, an agent with tools and a knowledge base that cites its source.

You leave with: automation, agent and knowledge base running

PHASE III · WEEKS 5–6

Entanglement

Tie it all into a living operation: front line in front of real people, guardrails, dashboard, alerts and a documented handoff.

You leave with: an operation in production

PHASE IV · WEEKS 7–8

Collapse

The operation becomes a business: priced offer, pipeline in motion, model contract and Demo Day with a panel.

You leave with: offer, pipeline and contract

Weekly ritual. Session A: concept with guided building, opening with 15 minutes of “what broke this week”. Session B: open workbench — the student builds, the mentor circulates and corrects, and nobody leaves without delivering the artifact. Between sessions, 3 to 5 hours of directed work on their own project.

The 16 sessions

PHASE I · OBSERVATION Weeks 1-2 · See the system before touching it

01 The AI Operator: the craft, the arsenal, the target 2H

Understand the craft and leave with the environment set up and the target operation chosen.

CONTENT

- What an operator does that a chatbot user does not
- The market's four roles and what each one charges
- Input → decision → action → record → exception
- Process selection criteria; credential hygiene

LABORATORY

- Set up the environment, first API call in the orchestrator, validate the spend ceiling

ARTIFACT 1

Working environment + justified target operation

02 Cartography of the operation: measure before you touch 3H · LAB

Turn a messy process into a map with numbers.

CONTENT

- Notation: step, actor, input, output, decision, exception
- The four numbers: volume, time, cost, rework
- How to gather this when nobody has measured anything
- The human line: what the machine does not touch

LABORATORY

- Draw the map live and defend the scope
- Calculate what the process costs per month today

ARTIFACT 2

Operation map + signed baseline

03 How the model actually works (and what it costs) 2H

Choose a model by task, cost and risk — not by fashion.

CONTENT

- Token, context window, input vs output
- Hallucination: when it shows up and the three real defences
- Extended reasoning vs fast answer; latency
- Routing, fallback, context caching

LABORATORY

- The same task across three models, measuring cost, latency and quality

ARTIFACT 3

Model matrix with cost per unit

04 Context engineering: the prompt as a specification 3H · LAB

Write production instructions — versioned, tested, with an acceptance criterion.

CONTENT

- Role, context, task, constraints, format, criterion
- Structured output: the data that feeds a system
- Evaluation: a test bench with the ugly cases included
- Versioning and the anti-pattern of the unreadable prompt

LABORATORY

- Run against 10 cases: typical, edge, cases that must be refused, and historical ones where the human got it wrong

ARTIFACT 4

Production prompt + measured suite of 10 tests

05 The nervous system: orchestration with n8n

3H · LAB

Master the layer that ties together AI, data and the client's systems.

CONTENT

- Why an orchestrator and not a script: who maintains it later
- Nodes, executions, the four trigger types
- Credentials and variables: holding a client's secrets
- Local vs cloud; the flow is code

LABORATORY

- End-to-end flow: real input → AI → validated output → real destination

ARTIFACT 5

Flow v1 running on a real trigger

06 Integrations that pay the bills

3H · LAB

Connect the operation to the systems where the business actually happens.

CONTENT

- WhatsApp: routes, cost, risk and platform rules
- Two-way email, spreadsheets, CRM and calendar
- Dirty data: phone numbers, ID documents, duplicates
- Rate limits and error 429 at the first real volume

LABORATORY

- Plug in two real systems and run a batch of dirty data

ARTIFACT 6

Integrated automation with real data

07 An agent is not an automation

2H · LAB

Know when an agent is the right answer — and build one that works.

CONTENT

- Goal, tools, loop, memory, stopping criterion
- Tool use: describing a tool and validating what it asks for
- Maximum steps, cost ceiling, human approval
- When *not* to use an agent

LABORATORY

- An agent with three real tools and a readable decision trail

ARTIFACT 7

Agent with 3 tools and configured ceilings

08 A brain with sources: knowledge base and RAG

2H

Make the system answer about the client's business without making things up.

CONTENT

- Long context, RAG or fine-tuning: what each one costs
- Ingestion and chunking; why the wrong cut destroys everything
- Hybrid search and source citation as a requirement
- Maintaining the base — a line item in the retainer contract

LABORATORY

- Ingest the real base and answer 10 questions with the source checked
- Adversarial test: does the system admit it doesn't know?

ARTIFACT 8

Base answering with a verifiable source

09 The front line: service and sales in front of real people

3H · LAB

Put the operation in contact with real humans and survive it.

CONTENT

- The six moments of the conversation
- Human handoff: trigger, context, out of hours
- Logging: a conversation that didn't become data didn't happen
- Transparency about talking to a machine

LABORATORY

- Round-robin with drawn profiles: rushed, sceptical, confused, angry, ready to buy

ARTIFACT 9

Front line tested with 5 real people

10 Voice and multimodal

2H

Choose and implement the right channel — text is not always the answer.

CONTENT

- When voice beats text, and when it loses
- Transcription, reasoning, synthesis: accumulated latency
- Voice cloning: documented consent
- Structured extraction from documents

LABORATORY

- Implement an additional channel with the criterion declared beforehand

ARTIFACT 10

Second channel with a measured metric

11 Nothing breaks in silence: reliability, cost and guardrails

3H · LAB

Turn a pretty prototype into a system that survives Monday morning.

CONTENT

- Failure taxonomy; retry, queue, idempotency
- The degradation path — never vanish in silence
- Guardrails validated on the output, not just asked for in the prompt
- Cost ceiling, alerts, dashboard and data protection in practice

LABORATORY

- **Live fault injection:** the mentor breaks the operation on purpose

ARTIFACT 11

Operation that passed the break test

12 Production and handoff: delivering without becoming a hostage

2H

Deliver in a way that lets the operation live without you — and that keeps you paid.

CONTENT

- Development vs production; backup and recovery
- The client's credentials stay with the client
- Delivery dossier and training the team
- The retainer conversation and your price

LABORATORY

- Handoff test: another student operates the system using only the dossier

ARTIFACT 12

Operation in production + approved dossier

13 Offer and price: sell the result, not the hour

2H

Build an offer that holds up and a price you can defend without stammering.

CONTENT

- The ladder: paid diagnostic → setup → retainer
- API cost passed through, built in, or capped
- Anchoring on the Week 2 baseline
- An honest guarantee: what is guaranteed and what never is

LABORATORY

- Present the offer in 3 minutes; the cohort attacks the price and the student defends it

ARTIFACT 13

Priced offer with calculated margin

14 Pipeline: from zero to a booked diagnostic

3H · LAB

Leave the session with real conversations under way.

CONTENT

- ICP: who has the expensive process you just solved
- The unfair advantage — you have the number
- The 45-minute diagnostic
- The one-page proposal and follow-up

LABORATORY

- Live prospecting: messages written and sent inside the session

ARTIFACT 14

Active pipeline with 30 contacts approached

15 Contract, risk and the difficult conversation

2H

Close business without taking on risk you cannot afford.

CONTENT

- Ownership of the flow, the prompt, the base and the data
- A realistic service level and what depends on third parties
- AI clauses, human review and data protection
- Protocol: communicate, contain, correct, compensate

LABORATORY

- Negotiation simulation with real objections, recorded and dissected

ARTIFACT 15

Model contract + incident protocol

16 Demo Day and the operator's routine

4H · LAB

Prove what was built and leave with the next step defined.

CONTENT

- 12 minutes per student: 7 presenting, 5 under questioning
- Live demonstration — slides don't count
- The weekly routine that stops the operation from rotting
- 90-day plan and the panel's verdict

LABORATORY

- Questioning: what breaks if the volume multiplies by ten?

ARTIFACT 16

Operation presented + 90-day plan

Sixteen pieces, not a certificate

Every week produces real artifacts from the student's own operation, corrected publicly against acceptance criteria published from Week 0. At the end, the student has an operational portfolio and a running system.

#	ARTIFACT	SESSION	PHASE
01	Working environment + target operation defined	1	I
02	Operation map + signed baseline	2	I
03	Model selection matrix with cost per unit	3	I
04	Production prompt + suite of 10 test cases	4	I
05	Orchestration flow v1, end to end	5	II
06	Integrated automation running on real data	6	II
07	Agent with 3 tools and cost limits	7	II
08	Knowledge base with source citation	8	II
09	Service front line tested with 5 people	9	III
10	Second channel implemented (voice or document)	10	III
11	Operation with guardrails, dashboard and alerts	11	III
12	Operation in production + delivery dossier	12	III
13	Complete priced offer	13	IV
14	Active pipeline with diagnostic script	14	IV
15	Model contract + incident protocol	15	IV
16	Operation presentation + 90-day plan	16	IV

The Quantum Operator Seal. There is no exam. The seal is granted to whoever meets three conditions: delivering the 16 artifacts on time, having the operation running in production by Session 12, and presenting at Demo Day with the Scoreboard closed — baseline against result. Whoever falls short receives a record of participation, without the seal. The seal is worth something because it can be withheld.

Daniel Rocha



Founder and CEO of **NOIRIS AI** — an applied intelligence house that builds agents and automation systems for companies. Not an agency, and not generic software.

An AI strategist based in Canada, combining **executive education from Harvard, MIT, Stanford and Yale** with hands-on implementation: agents and automations in production, not in slides.

What is taught in Quantum AI is not theory assembled for a classroom — it is the method NOIRIS itself uses to get an operation standing, including the parts nobody publishes: dirty data, rate limits, human handoff, cost at volume and the day the model fails in front of the client.

STANFORD UNIVERSITY

Artificial Intelligence Professional Program

MIT PROFESSIONAL EDUCATION

Machine Learning for Big Data

HARVARD UNIVERSITY

Data Science & Computer Science for Business

HARVARD BUSINESS SCHOOL

Leadership & Management

YALE UNIVERSITY

Negotiation & Business Strategy

Who gets in, and who should not

Entry is by application and a selection conversation, not by checkout. The promise depends on the student having a real operation and real time.

Prerequisites

- Comfort with technology at an advanced-user level
- **No programming required** — what is required is not being afraid of configuration screens
- About 9h per week: 5h live on average and 3 to 5h of directed work
- A real target operation, your own or a pilot client's with guaranteed access

It is not for you if

- You want to learn AI out of curiosity — the program is far too heavy for that
- You have no real process you can access to work on
- You expect guaranteed income at the end of eight weeks
- You don't actually have the 9 weekly hours available

The program is the same. What changes is how much attention you get.

UP TO 12 SEATS

Quantum Cohort

- 16 live sessions · 41h
- Public correction of artifacts
- Cohort channel and recordings
- Full template arsenal

UP TO 5 SEATS ·

RECOMMENDED

Quantum Core

- Everything in Cohort
- Individual written review of every artifact
- 4 one-to-one sessions of 45 minutes
- Direct access to the mentor between sessions

1 SEAT

Quantum Private

- Curriculum adapted to your operation
- Full individual support
- NOIRIS builds alongside you
- Schedule set with the student

Exit Guarantee. Up to the end of Session 4, if you conclude the program is not for you, we refund you in full — and you keep the artifacts already produced.

Operation Guarantee. If you deliver the 16 artifacts on time and your operation is not running in production by the end of Week 8, we stay with you, at no extra cost, until it runs.

What is never guaranteed: revenue, number of clients, perfect accuracy of any AI system, or any result that depends on third parties.

The seat is for someone with an operation waiting.

If you have the process, the access and the hours, we build the rest together over the next eight weeks.

Full program: noirisai.com/quantumai/en

Selection conversation: calendly.com/contato-noirisai/45min