

— HIGHER EDUCATION · VERTICAL 03

For institutions whose product is learning, and whose risk is misinformation at scale.

Cohorte's programs for universities, business schools, and accredited online providers. Built on the OPIT AI tutoring system: -60% professor support, +40% student progression. Faculty governance, academic integrity, accreditation defensibility.

FOR

Provosts, Deans, Heads of Digital Learning, CIOs

REGIMES COVERED

Faculty governance · AACSB · EQUIS · AI Act · ISO 42001

CLOSEST REFERENCE

OPIT · Riccardo Ocleppo, CEO

DIRECT LINE

charafeddine@cohorte.co

01 WHY HE NEEDS DIFFERENT TRAINING

AI shows up in three places at once.

The student is using it (often invisibly). The faculty is being asked to teach with it. The institution is being asked to assess learning that has been AI-mediated. Each of those is a governance problem. Generic AI training answers none of them.

PEDAGOGY AT SCALE

An AI tutor is only as good as its verification layer.

A confident wrong answer to a foundational question teaches a generation to be confidently wrong. Install the verification discipline before the tooling. OPIT is the existence proof (-60% professor support, +40% student progression, -80% time-to-feedback).

ACCREDITATION

Accreditors are asking questions you don't yet have answers to.

AACSB, EQUIS, ACBSP, AMBA, HCERES, QAA, ANVUR: all increasingly asking how AI fits into assurance-of-learning, integrity, and teaching quality. Documented response is required before the next review cycle.

FACULTY TRUST

Earned through evidence, not top-down rollouts.

An academic faculty adopts AI when they see colleagues use it well in the same field. The programme has to be peer-credible, not vendor-credible. This is where published research and named institutional references matter most.

02 WHY COHORTE

The methodology that built the OPIT tutoring system.

Cohorte's founder architected the AI tutoring system at OPIT, the LLM-native accredited online university. Open-source. Inspectable. Defensible to accreditors.

CONFORMAL PREDICTION

Provably right within stated bounds.

The mathematical primitive that converts a model output into a statistically rigorous confidence interval: a bound on every student-facing answer, so a confident wrong one is caught before it reaches a student.

Source: The conformal-calibration method, from our reliability-certification paper. teams.cohorte.co/research.

SELF-CONSISTENCY

Detecting confabulation.

If the same prompt run five times produces five different answers, the model is not reasoning. The verification gate that catches this before output reaches a student-facing channel or an audit log.

Source: The self-consistency method, from the same paper. teams.cohorte.co/research.

EXPLOITATION SURFACE

The red team your CISO asked for.

Prompt injection. Indirect injection through retrieval corpora. Tool-call hijacking. The systematic taxonomy your security team uses to scope penetration testing of an LLM application.

Source: Our 10,000-trial exploitation-surface taxonomy. teams.cohorte.co/research.

REFERENCE STACK

Open-source. Inspectable. Yours.

TrustGate (verification gates), Guardrails (policy enforcement), Agent-Auth (authorization), Agent-Monitor (observability). Six public repositories. Your engineers audit the layer before deploying it.

Where: github.com/Cohorte-ai · MIT / Apache 2.0.

03 USE CASES WE TRAIN AGAINST

Where AI touches academic outcomes.

The AI workflows higher-education institutions are actually shipping. Where each sits. Pedagogical risk. Verification primitive.

USE CASE	WHERE IT SITS	PEDAGOGICAL RISK	VERIFICATION PRIMITIVE
Student tutoring agent	Teaching & learning	Misinformation at scale	Citation-grounded answers, curriculum-bounded retrieval, faculty review queue
Faculty preparation copilot	Teaching & learning	Pedagogical drift	Source attribution, style alignment to syllabus
Research literature synthesis	Research	Citation accuracy	Paper-level retrieval, DOI verification, publication-date check
Grant application drafting	Research administration	Funder compliance	Funder-specific template lookup, eligibility check
Admissions essay review	Admissions	Fairness, AI-content detection	Decision logs, human-in-the-loop, bias monitoring
Student support / advising	Student services	Hallucinated policy	Policy-grounded answers, escalation gates, audit trail
Assessment generation	Teaching & learning	Validity, integrity	Bloom-level mapping, variant generation, validation testing
Plagiarism / AI-content review	Academic integrity	False positives	Confidence intervals, human review, documented threshold
Institutional knowledge agent	Institution-wide	Hallucinated procedure	Permission-aware retrieval, DLP gates, audit logs

04 SAMPLE CURRICULUM · 12-WEEK HE BOOTCAMP

Module by module. Tuned to your stack in scoping.

Six modules across the twelve weeks. The actual curriculum is tuned to your stack and use cases in the pre-bootcamp scoping call with Charafeddine.

MODULE	TOPIC	WHAT YOUR TEAM PRODUCES
01	Scoping & operating discipline. Process First. The Thinking Stack. LUMEN scoping brief drafted live for one institutional workflow (tutoring, advising, research synthesis, assessment).	Mission charter and LUMEN scoping brief for one institutional workflow, signed off by the faculty sponsor.
02	AI Engineering Foundations. The Three V's (Vibes, Variance, Vendor). Accountable Development Lifecycle. Prompt architecture, multi-model patterns, MCP tool use.	Working system v0.1 against your LMS, curriculum content, and faculty knowledge base.
03	Trust & Verification. The Confidence Problem. Curriculum-bounded retrieval. Citation-grounded answers. Conformal prediction for confidence intervals on student-facing output.	Verification gates and citation-grounded retrieval wired in. Faculty review queue for borderline outputs.
04	Accountable Agents · 4-Layer Architecture. Platform Protocol. Agent-Auth (least privilege). Guardrails as architectural constraints. Agent-Monitor.	Agent-Auth and Guardrails layered in. Escalation rules to faculty review documented.
05	Academic Governance. Faculty review queue. Academic integrity controls. Accreditation defensibility (AACSB, EQUIS, AMBA, HCERES, QAA, ANVUR). AI Act Article 12 logging.	Accreditation-defensible governance log. Faculty review queue. Immutable audit trail.
06	Capstone & Sponsor Brief. Red-team workshop. Founder critique. Sponsor demo to leadership. Operating brief written, reviewed, signed.	Production-grade HE system. Sponsor brief defensible to your accreditors.

05 THE HE PORTFOLIO

Same operating layer. Academic-rhythm tuning.

Four programs. Sequenced to the academic rhythm. Programmes fit around teaching loads, accreditation cycles, and faculty governance.

ENTRY · DE-RISKER

HE Pilot

€8K-€12K · 4 weeks · 2-3 people

One institutional workflow, taken to proof: a working, verified prototype (tutoring, advising, research synthesis, or assessment). Half-fee refund clause. 100% credit toward follow-on within 90 days.

SINGLE-TEAM DEPTH

HE Team Bootcamp

€4,200 / seat · 12 weeks · 6-15 seats

Private cohort. Your faculty, your LMS. Pre-bootcamp scoping with Charafeddine. 12 months Engine Room per graduate. Capstone defensible to your accreditors.

FIRM-WIDE BREADTH

HE Curriculum License

€12,000 / year · up to 25 seats

Five role-based paths (Faculty, Researchers, Admin, Students, IT). Completion dashboards. SSO. Quarterly Q&A with Charafeddine. 12 months Engine Room access for all seats.

TRANSFORMATIONAL

HE AI Readiness Program

from €35,000 · 3 to 6 months

Two-day on-site assessment. Faculty training. Monthly 2h founder advisory. Written AI Operating Model playbook at close. The institution-wide engagement.

Funding (honest). Cohorte is not Qualiopi-certified today (2026 roadmap, in progress). Direct CPF and direct OPCO not available right now. Standard path: direct L&D budget. For OPCO routing: via Qualiopi-certified partner. Email teams@cohorte.co.

06 HE OBJECTIONS WE ALWAYS HEAR

The honest answers. Before the call.

The four objections an enterprise HE buyer raises in every conversation. Answered here so the call can spend its time on the work.

"Our faculty senate has not approved any AI tool yet."

Good. The curriculum starts from there. The Trust & Verification module installs citation-grounded retrieval and conformal-bounded output specifically so a faculty member can defend an AI-mediated lesson plan. **The accreditation-defensible governance log in module 5** is the artifact that lets a senate approve institution-wide deployment.

"OPIT is your only HE reference."

Honest. OPIT is the existence proof: -60% professor support, +40% student progression, presented at Microsoft Milan May 2025. Riccardo Ocleppo, CEO, takes reference calls arranged after a mutual NDA. **Additional named HE references extended quarterly** as engagements close.

"How do we explain this to accreditors?"

Module 5 produces the governance log mapped to your accreditation framework by section number (AACSB, EQUIS, AMBA, HCERES, QAA, ANVUR). The mapping is appended as a regulatory annex to every bootcamp deliverable. **The artifact is designed to be appendix-ready** for the next assurance-of-learning review.

"Budget is committed; AI is not on this year's plan."

Cohorte is not Qualiopi-certified today (2026 roadmap, in progress), so direct CPF and direct OPCO are not available right now. Standard path is institutional faculty-development or executive-education budget. **The Pilot at €8K-€12K** typically fits inside discretionary spend mid-year.

DIRECT LINE · NO SALES GATE

Talk to Charafeddine directly.

Sixty-minute discovery. No deck. We help you scope the right entry point: Pilot to de-risk, Bootcamp to install, License to scale, AI Readiness to transform.

charafeddine@cohorte.co

Cohorte SAS · Paris & Rabat · Founded 2022
teams.cohorte.co · cohorte.co/teams