UniAgent·Lex — Conference Script

English · ~60 minutes · accompanies the UniAgent·Lex deck · cue [L] to switch slides to Spanish · [P] for presenter mode

0 · Opening (3 min)

■ SLIDE — Hero: AI agents in higher education

Good morning, colleagues. Thank you for being here. Before we start, a small experiment: raise your hand if, in the last seven days, you have used a generative AI tool — ChatGPT, Claude, Gemini, Copilot, anything. *[pause]* Now keep your hand up if you have used what we call an *agent* — a system that didn't just answer you, but actually *did something* on your behalf: booked, searched the web, edited a file, wrote and sent an email. *[pause]*

That gap between the first hand and the second hand — that is what this hour is about. We have spent two years debating chatbots. Meanwhile, a different technology has quietly arrived on our campuses, and most of our regulations, our syllabi, and our academic-integrity policies were written for the previous generation.

My goal today is threefold. First, I want to give you a precise, non-mystical definition of what an AI agent is, and what it is not. Second, I want to walk you through the European legal framework that already governs how we may, and may not, deploy these systems in a university. And third — and this is where it gets uncomfortable — I want to open a frontier most faculty have never heard of: **neurorights**, and why they will be the next big conversation in higher education.

We will keep the language bilingual on screen — press **L** at any moment if you want to switch — and every claim I make is backed by a primary source you can pull up in the side panel on the right.

1 · What is an agent, really? (7 min)

■ SLIDE — Capabilities pillars

Let's start with vocabulary, because the word "agent" is doing a lot of heavy lifting right now, and most of it is marketing.

A *language model* — GPT-4, Claude, Gemini — is a function. You give it text, it returns text. It has no memory beyond the prompt, no tools, no ability to act in the world. It is, in the strict sense, a very expensive autocomplete.

An *assistant* is a language model with a conversational wrapper, perhaps some retrieval, perhaps a system prompt. ChatGPT in its 2023 form was an assistant.

An *agent* is something categorically different. An agent is a language model embedded in a *loop*, with access to *tools*, the ability to *plan* multiple steps, and — critically — the autonomy to *decide* which tool to call, in what order, and when to stop. When I say "book me a flight to Madrid under 300 euros, departing Friday morning," an agent will search, compare, fill forms, and in the most advanced versions, actually complete the purchase.

Three properties define an agent, and I want you to remember these three words because they will come back when we discuss the law:

Autonomy. The system takes decisions without a human in every step of the loop.

Tool use. It interacts with external systems — browsers, APIs, databases, your learning management system, your student records.

Persistence. It maintains state across interactions. It remembers. It can run for minutes, hours, sometimes days.

Now, why does this matter for a university? Because every single one of those three properties — autonomy, tool use, persistence — collides with categories that European law treats as high-risk when the affected person is a student, an applicant, or an employee.

■ SLIDE — Live Lab

Let me show you, very briefly, what this looks like in practice. *[demo of LiveLab, ~90 seconds]* Notice three things: I never told it which sources to consult; it chose. I never told it to stop; it decided when the answer was complete. And if I had given it write-access to a calendar or an email account, the same loop would have produced an action in the real world, not just a paragraph on a screen.

That is the shift. We are moving from systems that *advise* to systems that *act*.

2 · University use cases (8 min)

■ SLIDE — Uses pillar

Where is this already happening on campuses? Let me give you a tour of six categories, from the most benign to the most legally sensitive.

One: research assistance. Literature reviews, summarisation of long PDFs, code generation for data analysis. This is the use case faculty know. It is also the least regulated, because the human researcher remains in the loop and bears full authorship responsibility.

Two: administrative automation. Drafting emails, scheduling, processing repetitive forms, triaging student inquiries. Many of our administrative offices are quietly piloting this right now, often without formal authorisation from the data protection officer. We will come back to that.

Three: personalised tutoring. Khan Academy's Khanmigo, OpenAI's study mode, custom GPTs built by individual faculty. Promising pedagogically, but the moment the agent stores conversational history linked to a student identifier, you are processing personal data under GDPR.

Four: assessment and feedback. Automated grading of essays, code, even oral exams. This is where we cross a legal threshold I will explain in detail in the next section.

Five: admissions and student selection. Resume screening, application scoring, interview analysis. We have a landmark European case on exactly this — the *Schufa* judgment of December 2023 — and I will walk you through it.

Six: proctoring and behavioural monitoring. Remote exam supervision, attention tracking, gaze analysis, keystroke biometrics. This is where neurorights enter the conversation, and where, frankly, some current vendor practices are already non-compliant with the AI Act that will be fully enforceable next year.

Each of these categories sits at a different point on a risk gradient. The mistake most institutions make is treating them as a single phenomenon — "AI in education" — and writing a single policy. The European legislator has done the opposite: it has built a tiered system, and we need to

understand that tier system before we deploy anything.

3 · The EU legal framework (12 min)

■ SLIDE — Timeline

So let's talk about the law. And let me reassure the non-lawyers in the room: I will not read articles to you. I will tell you what changes for your daily work.

The European AI Act — Regulation 2024/1689 — was published in the Official Journal on the 12th of July 2024 and entered into force on the 1st of August 2024. It is the first horizontal regulation of artificial intelligence anywhere in the world. It applies extraterritorially: if your university uses a model trained in the United States or hosted in Singapore, but the output affects a person in the European Union, the regulation applies.

The Act follows a phased calendar, and I want you to memorise four dates.

February 2025. The prohibitions of Article 5 became enforceable. Certain AI practices are simply banned in the European Union — among them, systems that infer emotions in educational or workplace contexts, except for medical or safety reasons. Read that again: **emotion inference in education is, as of February 2025, prohibited.** Every proctoring vendor that advertises "stress detection" or "engagement scoring" is selling you a product that is illegal to deploy on a European student.

August 2025. Obligations on general-purpose AI models — the GPT-4s, the Claudes, the Geminis — entered into force. This concerns the providers more than us, but it means that by now every foundation model offered in the EU must publish a summary of training data and document copyright compliance.

August 2026 — that is in two months. Annex III obligations kick in. Annex III is the list of *high-risk* AI systems, and three categories on that list are directly relevant to us: AI used to determine access to education, AI used to evaluate learning outcomes, and AI used to monitor and detect prohibited behaviour during tests. If your institution uses an automated admissions filter, an automated grader, or a proctoring system that flags suspicious behaviour, you will, from August 2026, need: a conformity assessment, a risk management system, human oversight, technical documentation, logging, and registration in the EU database.

August 2027. The remaining obligations on Annex I high-risk systems become enforceable. Full deadline.

■ SLIDE — Legal tiers

The Act classifies systems into four tiers.

Unacceptable risk — prohibited. Social scoring by public authorities, manipulative subliminal techniques, real-time biometric identification in public spaces with narrow exceptions, and, as I said, emotion recognition in education and the workplace.

High risk — permitted but heavily regulated. Education, employment, access to essential services, law enforcement, migration. This is where most of our use cases land.

Limited risk — transparency obligations. Chatbots must declare they are chatbots. Synthetic content must be labelled.

Minimal risk — no specific obligations. Spam filters, video game AI.

Now, the AI Act does not replace the GDPR — it stacks on top of it. The GDPR — Regulation 2016/679 — still governs the personal data dimension. Article 22 of the GDPR remains the operative rule for *automated individual decision-making*: a student has the right not to be subject to a decision based solely on automated processing if it produces legal or similarly significant effects on them. Admission. Grading. Expulsion. Scholarship allocation. These are "similarly significant" decisions, and they require meaningful human involvement.

4 · Case law (10 min)

■ SLIDE — Case Law

Let me ground this in three cases that should be on every university lawyer's desk.

Case one: Schufa, Court of Justice of the European Union, 7th December 2023, C-634/21.

Schufa is a German credit-scoring agency. A consumer was denied a loan because of her Schufa score. She invoked Article 22 GDPR. Schufa argued: "We don't make the decision — the bank does. We only provide a score." The Court rejected this. The Court held that the *automated generation of the score itself*, when the score *plays a determining role* in the subsequent human decision, qualifies as automated decision-making under Article 22.

Translate this to a university. If you deploy an AI system that pre-scores applications, and the admissions committee then "merely confirms" the algorithm's ranking — you are, in the eyes of European law, making automated decisions. Calling it "decision support" does not save you. The standard is whether the human has *meaningful* discretion, with the capacity and the time to override.

Case two: Ofqual, United Kingdom, August 2020.

When COVID cancelled A-Level exams, the UK exam regulator Ofqual deployed an algorithm to predict grades. The algorithm systematically downgraded students from disadvantaged schools — because it incorporated the historical performance of the school as a variable. Forty percent of grades were lowered relative to teacher estimates. After massive public protest — including the now-famous chant *"fuck the algorithm"* in front of the Department for Education — the system was withdrawn within four days. The ICO opened an investigation. The political cost was enormous.

The lesson is not that algorithms are biased — we knew that. The lesson is institutional: the system was deployed without an adequate impact assessment, without meaningful student consultation, and without a viable rollback plan. Every one of those failures is now, under the AI Act, a documented legal obligation, not a best practice.

Case three: New York City Local Law 144, in force since 5th July 2023.

Not European, but instructive. NYC requires that any *Automated Employment Decision Tool* used in hiring within the city undergo an independent bias audit, published annually, with candidates notified that AI is being used. Several universities with US campuses are already subject to it. The European equivalent — for high-risk systems under Annex III — arrives in August 2026 and is, if anything, stricter.

I'll add a fourth, briefly, because it is the cautionary tale: **SyRI**, Netherlands, 2020. The Hague District Court struck down the Dutch government's algorithmic welfare-fraud detection system as a violation of Article 8 of the European Convention on Human Rights. The reasoning — opacity, disproportionality, discriminatory impact — applies *mutatis mutandis* to any opaque student-monitoring system.

5 · Neurorights — the emerging frontier (12 min)

■ SLIDE — Neurorights

Now we come to the part of this lecture that, two years ago, I would not have included. Today I consider it the most important section.

Neuroscience has reached a point where consumer-grade devices — EEG headbands, eye-tracking webcams, even ordinary smartphone cameras combined with machine learning — can infer cognitive and emotional states with non-trivial accuracy. Attention. Confusion. Fatigue. Stress. In experimental settings, even rough categories of intent.

The legal community has responded with a new generation of rights, collectively called *neurorights*. Five are now widely recognised:

One. The right to cognitive liberty — the freedom to determine one's own mental states without external interference.

Two. The right to mental privacy — protection of neural and cognitive data from unauthorised access.

Three. The right to mental integrity — protection against harmful manipulation of brain activity.

Four. The right to psychological continuity — protection of personal identity from technologies that could alter it.

Five. The right to equitable access to cognitive enhancement, where such technologies exist.

These are not academic abstractions. **Chile** amended its constitution in October 2021 — Law 21.383 — to explicitly protect brain activity and the information derived from it. It is the first country in the world to do so. **Spain** included a right to neurorights in its 2021 Digital Rights Charter. The **Declaration of León**, signed in October 2022 under the Spanish EU presidency, committed European states to a coordinated approach. The **Council of Europe**, through the DH-BIO committee, has been drafting standards since 2020. And in November 2025 — three weeks ago — **UNESCO** adopted a global Recommendation on the Ethics of Neurotechnology.

Why does this matter for a *university*?

Because remote proctoring with gaze-tracking is a neurotechnology. Because adaptive learning platforms that profile a student's confusion patterns are processing cognitive data. Because the next generation of EEG-based "focus enhancement" devices — already on sale to consumers — will arrive in our classrooms within this academic cycle, and we have no policy.

Let me be concrete. If a proctoring vendor offers your institution a system that flags "loss of attention," three legal triggers fire simultaneously:

— It is **emotion inference in education**, prohibited under Article 5 of the AI Act since February 2025.
— It is **special category data** under Article 9 GDPR if it touches on health or psychological state.
— It potentially infringes the **right to mental privacy** as codified in Spanish and Chilean law, and as recommended by UNESCO.

The compliant answer is not "let's add a consent checkbox." Consent under GDPR must be freely given. A student cannot freely consent to be monitored if the alternative is failing the exam. The European Data Protection Board has been explicit on this point in the employment context, and the same logic applies to mandatory assessment.

6 · What to do on Monday morning (5 min)

■ SLIDE — Resources

I want to leave you with practical guidance, because critique without action is useless.

First: inventory. Within the next thirty days, every faculty and every administrative unit should produce a list of every AI system currently in use — including the informal ones, the ChatGPT-Plus subscription a professor pays for personally, the custom GPT a department built last semester. You cannot govern what you have not mapped.

Second: classify. Apply the four-tier model. Mark each system as prohibited, high-risk, limited-risk, or minimal-risk. Anything in the first two tiers needs a documented decision: discontinue, redesign, or formalise with a conformity assessment.

Third: human-in-the-loop, meaningfully. Audit every process where an AI output reaches a student. Is the human reviewer given the time, the training, and the authority to disagree with the system? If not, you are running an Article 22 violation regardless of how many humans signed off on the workflow.

Fourth: transparency by default. Students have a right to know when AI is involved in a decision that affects them, what data was used, and how to contest it. Build this into your academic regulations now, not in 2027 when the regulator asks.

Fifth: train your faculty. Not on prompt engineering — on legal literacy. The single highest-leverage investment a university can make this year is two hours of structured training for every academic staff member on the AI Act, GDPR Article 22, and neurorights.

Sixth and last: protect mental privacy as a constitutional value. Even where current law does not yet require it, treat student cognitive data with the same seriousness you treat health data. The regulation will catch up; the reputational damage of being on the wrong side of it will not be reversible.

7 · Closing (3 min)

■ SLIDE — Hero again

I opened with a show of hands. Let me close with a question you do not need to answer out loud.

What kind of university do we want to be in 2030?

One option: an institution that treats AI as a productivity tool, deploys it wherever it saves money, and discovers — case by case, lawsuit by lawsuit, scandal by scandal — what the limits were.

The other option: an institution that takes seriously the fact that we are the place where Europe trains the people who will design, regulate, and live with these systems. An institution that models, in its own internal practice, the legal and ethical standards it teaches in its lecture halls.

The AI Act, the GDPR, the neurorights framework — these are not obstacles to innovation. They are the *grammar* of a kind of innovation that is compatible with the European project. They are what allows us to compete on something other than surveillance.

The agents are already here. The law is largely written. The neurorights conversation is just beginning. Our job — yours and mine, as faculty and as administrators — is to make sure that the university remains a place where human cognition is cultivated, not harvested.

Thank you. *[pause]* I'll take your questions now.

Speaker notes

• Total reading time at conversational pace: **58–62 minutes**.

• Built-in flex: the LiveLab demo (Section 1), the show of hands (Opening), and the Q&A handoff (Closing) absorb ±5 min.

• Press **L** on stage to flip the slides to Spanish without losing your place; press **P** for presenter mode on the sources panel.

• All cited cases and dates match the deck's *Case law* and *Timeline* sections — citations are copy-ready in the side panel.