



AI Agents for Higher Education

Agents · Automations · Real-World Workflows

A Companion Extension to the AI-Driven Leadership Workshop

The European Conference on Education 2026

"The next frontier isn't AI that answers questions — it's AI that gets things done."

What We'll Cover in This Session

01

5 min

Chatbots vs. Agents vs. Automations

Clearing up the terminology

02

10 min

How AI Agents Work

The perceive–reason–act–learn loop

03

15 min

Automation Use Case Gallery

10 HE workflows ready to automate

04

15 min

No-Code Tools Landscape

Make, Zapier, Power Automate, n8n, Copilot Studio

05

10 min

Building Your First Agent

Live demo walkthrough + prompt engineering

06

5 min

Risks & Human-in-the-Loop Design

When to keep humans accountable

Chatbots vs. Agents vs. Automations — What's the Difference?

Chatbot

Reactive

Responds to a single message with a single reply

Examples:

- › FAQ bot on university website
- › Teams/Slack Q&A helper
- › Course information bot

⚠ *Stateless. No memory. No action. Can't do anything.*

AI Automation

Rule-Based

A pre-defined workflow triggered by an event — AI may generate content within it

Examples:

- › New application → AI email draft → send
- › Form submit → data extract → log to sheet
- › Weekly → AI report draft → email to dean

⚠ *Follows fixed paths. No reasoning or decisions.*

AI Agent

Autonomous


An autonomous system that perceives context, reasons over goals, picks its own tools, and acts across multiple steps

Examples:

- › Monitor applications, chase docs, update CRM, alert advisor — no human trigger
- › Research 5 data sources, synthesise report, flag anomalies, email VP

⚠ *Needs guardrails. Can make mistakes. Requires oversight.*

The Agent Loop: Perceive → Reason → Act → Learn



Tools the agent can use:

Web search · Email / Calendar · Database read/write · Document generation · API calls · Slack / Teams

Anatomy of an AI Agent — The Building Blocks



LLM Brain

The reasoning core (GPT-4o, Claude, Gemini). Interprets goals, decides next steps, generates outputs.



Memory

Short-term (conversation context) and long-term (vector DB, structured store). Enables continuity across tasks.



Tool Library

A set of callable functions: search, email, calendar, API, file system. The agent picks which to invoke.



System Prompt

You define the agent's persona, goals, constraints and tone. This is how you program behaviour without code.



Guardrails

Rules about what the agent CAN'T do: spend money, delete data, contact externals without approval.



Feedback / Logs

Every action is logged. Dashboards show what the agent did, when, and why — for audit and improvement.

10 HE Workflows Ready to Automate Right Now



**Enrolment
Pipeline**



**Applicant
Chaser Bot**



**GDPR
Compliance Audit**



**Staff
Onboarding**



**Accreditation
Reporting**



**Budget
Anomaly Alerts**



**Data Request
Fulfillment**



**Student
Welfare Triage**




**Research
Grant Matching**



**Facilities
Scheduling**

Each of these can be built with no-code tools in under a day. We'll explore how on the next slides.

Walkthrough: The Enrolment Pipeline Agent



X Manual version

Admin staff manually open each email, copy data to spreadsheet, email applicants one by one. Takes 2–4 hrs/day at peak.

✓ Agent version

Runs 24/7 with zero manual effort. Processes 200+ applications/day. Consistent, personalised, audit-logged.

Walkthrough: Student Welfare Triage Agent



Data Inputs

LMS logins, assignment submissions, library access, counselling appointments, grades



AI Risk Scoring

Agent scores each student daily against a risk model trained on historical dropout data



Tiered Routing

Low risk → no action. Medium → nudge email. High → alert personal tutor. Critical → welfare team



Personalised Reach-out

Agent drafts a warm, contextual check-in email using student's name, course and recent pattern



Outcome Logging

Records every alert, action and student response for PSRB reporting and continuous model improvement

Real-World Impact

14 days

earlier intervention
vs. manual monitoring

+25%

improvement in
retention rates

90%

accuracy in
identifying at-risk students

0 hrs

manual data collation
required per week

⚠ Always keep a human in the loop for high-risk escalations. Agent should flag, not decide.

The No-Code Automation Landscape for HE

Make.com

Best for beginners

Free tier available

- ✓ Visual drag-and-drop builder
- ✓ 1,500+ app integrations
- ✓ HTTP modules for any API
- ✓ Excellent for HE workflows

Best for: Admissions, comms, reporting pipelines

Power Automate + Copilot Studio

Best for Microsoft HE

Included in M365

- ✓ Integrated with Teams, SharePoint, Outlook
- ✓ Copilot Studio for AI chatbot/agent layer
- ✓ Policy-compliant for most UK/EU institutions
- ✓ No extra vendor procurement needed

Best for: Any Microsoft 365 campus environment

Zapier

Easiest to start

Free → paid tiers

- ✓ Quickest setup (minutes)
- ✓ 7,000+ app connections
- ✓ AI steps built in
- ✓ Good for simple trigger→action flows

Best for: Quick wins: form→email, submit→notify

n8n

Best for technical teams

Open source / self-hosted

- ✓ Full self-hosting (data stays on campus)
- ✓ Code nodes for custom logic
- ✓ LangChain / AI agent nodes built in
- ✓ Ideal for IT-led projects

Best for: Data-sensitive or custom enterprise workflows

System Prompts: How You Program Agent Behaviour

✗ Vague system prompt (bad agent)

*"You are a helpful assistant.
Answer questions about the university."*

Problems: No role, no scope, no guardrails, no tone, no tools, no escalation path → unpredictable behaviour

✓ Structured system prompt (good agent)

ROLE: You are an admissions support agent for Greenfield University.

GOAL: Process incoming applications, identify missing documents, and send chase emails.

TOOLS: You have access to: email_send, crm_update, document_checker.

TONE: Professional, warm, and concise. Always use the applicant's first name.

LIMITS: Do NOT make admissions decisions. Do NOT contact applicants more than twice. Escalate to human if documents remain missing after 5 days.

FORMAT: Log every action with timestamp and reason.

Live Demo: Build a Student Query Agent in Make.com

1

Create a free Make.com account

make.com → sign up with work email

2

Create a new Scenario

Click '+' → name it 'Student Query Agent'

3

Add a Webhook trigger

This is where queries arrive (from a web form or email)

4

Add an OpenAI / Claude module

Connect your API key. Paste your system prompt.

5

Add an Email or Slack Send module

Route the AI response to the right person or channel.

6

Add a Google Sheets / Airtable Log module

Record: query, response, timestamp, category.

7

Test with a real query, review the output

Iterate on your system prompt until the tone is right.

8

Activate the Scenario

Toggle ON. Your agent is live — monitoring continuously.

Estimated build time: 25–35 minutes for a first-time user with no prior automation experience.

Map Your Own Agent

In pairs or small groups, sketch your own agent:

- 1. Choose a workflow:** Pick one admin or leadership task at your institution that currently takes too much manual time.
- 2. Define the trigger:** What event starts the process? (Email arrives, form submitted, date reached, data changes...)
- 3. List the steps:** What does the agent do — in order? What tools does it need at each step?
- 4. Write the system prompt:** Use the 6-part framework: Role · Goal · Tools · Tone · Limits · Format
- 5. Identify the risks:** Where could this go wrong? What's the human checkpoint?

Agent Sketch Template

Workflow name:

Trigger:

Steps (1–5):

Tools needed:

System prompt sketch:

Human checkpoint:

Groups share back: What workflow did you choose? What's the hardest step to automate?

Risks of Agentic AI — And How to Mitigate Them



Hallucination & Errors

→ Never let agents make final decisions without a human review gate. Log all outputs.



Data Privacy Breaches

→ Agents must process only the minimum necessary data. Audit all API connections for GDPR compliance.



Bias & Unfairness

→ Regularly audit agent decisions for demographic disparities — especially in admissions and HR workflows.



Runaway Automation

→ Build hard stop limits: max emails per day, no deletes, no financial transactions without 2-step human approval.



Staff Trust & Buy-In


→ Involve staff in designing the agent. Communicate clearly what it does and doesn't do.



Single Point of Failure

→ Agents can break. Always maintain a manual fallback process. Test failure modes in staging.

When to Keep Humans in the Loop — A Decision Framework



Your AI Agent Starter Pack — Leave Here Ready to Build



Start with ONE workflow

Pick the most painful manual process. Build one agent. Learn. Then expand.



Use the 6-part system prompt

Role · Goal · Tools · Tone · Limits · Format. Print it, use it every time.



Pick a no-code tool this week

Make.com (free), Power Automate (you have it), or n8n (open source).



Design for humans first

Every high-stakes decision needs a human gate. Build it in from day one.



Log everything

An agent you can't audit is an agent you can't trust. Always enable logging.

Useful Resources

Tools

- › [make.com](#)
- › [n8n.io](#)
- › [zapier.com](#)
- › [powerautomate.microsoft.com](#)

Learning

- › [learnprompting.org](#)
- › [deeplearning.ai](#) (free courses)
- › [EDUCAUSE AI resources portal](#)

Community

- › [Make.com Community Forum](#)
- › [n8n Community](#)



Go Build Something.

Questions · Discussion · Next Steps

"The institutions that will lead the next decade are already experimenting with AI agents today — in small, responsible, human-centred steps."

rubina@aua.am · #ECE2026 · Workshop resources