AI agents inside the university

What they can read, what they can change, and who decides

IN BRIEF

01 Universities and school systems have put AI assistants in front of applicants, students and staff. In a published trial at an American university, admitted students sent a chatbot's text messages were 3.3 percentage points more likely to start in the autumn. Some universities now run their own: one in California hosts open models on its own infrastructure.

02 The record is mixed. The Los Angeles school district took its chatbot offline when the company that built it had no staff left to supervise it. The suppliers are moving in different directions: one says nothing its assistant drafts "goes out until an educator reviews it", and another says its agent "doesn't just suggest or brainstorm. It takes action." None of the vendor pages we read offers on-premises hosting.

03 This paper is about the institution's own work: the reading, checking and drafting done by staff and researchers in admissions, student services, the registry, teaching support, the research office, the library and the quality office. It is not about tutors or anything aimed at children. It sets out what an agent can read, how it reaches an institution's systems, what it may change without a person, and what the rules already say about admissions, marks, student data and research funding. It ends with seven questions to put to any supplier, and with a place to start.

04 The position throughout is simple. An agent prepares. A named person decides. Nothing reaches an applicant or a student unseen, and no agent admits, marks or judges anyone. The agent runs on infrastructure the institution controls.

IN THIS PAPER

01 Where universities and school systems are today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 6
06 A worked example: one university 8
07 Seven questions before an agent touches a student record 9
08 Where to start 9

PART 01

Where universities and school systems are today

3.3 points
more likely to begin the autumn term, among students sent a university chatbot's texts

Three months
after a school district announced its chatbot, the builder had furloughed most of its workers

Two
of the supplier pages we read state a step where an educator reviews first

None
of the vendor pages we read offers on-premises hosting

Chat with applicants came first. For Georgia State University, a start-up "built an artificially intelligent, text-messaging chatbot" for admitted students. The researchers "randomly assigned half of all the accepted GSU students to receive Pounce outreach". Among students who intended to enrol, "those who received the text outreach were 3.3 percentage points more likely to begin their fall semester at GSU" [1].

Some universities now run their own. UC San Diego says: "Open-weight models used by TritonAI are hosted on UC San Diego infrastructure at the San Diego Supercomputer Center, and in some cases through enterprise cloud providers where there is a UC contract" [2]. The University of Amsterdam built its own chat tool; its programme manager says it "wants to reduce its dependence on major suppliers and maintain more control over its own infrastructure" [3]. Both are assistants more than agents, and both show that institutions want this kept close.

Two warnings. In 2024 the Los Angeles school district announced an AI chatbot called "Ed". Three months later the company that built it had furloughed most of its workers, and the district "immediately pulled the signature Ed chatbot offline, district officials said, because there was no AllHere staff available to supervise it" [4]. Vanderbilt University disabled its AI writing detector in 2023 and concluded: "we do not believe that AI detection software is an effective tool that should be used" [5]. Nothing in this paper is pupil-facing or detects anything; the lesson is about depending on a supplier's service.

The suppliers disagree on how far an assistant acts. D2L says: "Nothing D2L Lumi drafts goes out until an educator reviews it" [6]. Microsoft says its education feedback features "produce suggestions that an educator reviews before returning them to a learner" [7]. Instructure says: "IgniteAI Agent doesn't just suggest or brainstorm. It takes action" [8]. Salesforce lists agents for recruitment, financial aid and transfer credit; its page does not say whether a person reads a reply before it reaches an applicant [8].

Where they run is rarely said. None of the vendor pages we read offers on-premises hosting. D2L says its assistant "operates within controlled environments" [6]; the Instructure and Salesforce pages say nothing about hosting [8]. So the buyer has to ask: where it runs, which tools it may call, as whom, with whose approval, and what record is kept.

One we can speak to. KAUST ran the full platform in-country, in the university's own environment.

What the published accounts leave out. Only two of the supplier pages we read state a step where an educator reviews before anything goes out.

None describes who approves a change to a student record.

PART 02

The work nobody has packaged

Both

admissions templates we opened send as soon as the workflow decides

0 to 100

the score one of them gives each applicant, and routes on

We listed the jobs done around a student and a research project and looked for each in the template libraries and prebuilt agents of the agent and workflow platforms. They cover admissions, enquiries and feedback. They thin out wherever a named person or a board signs something.

WHERE TEMPLATES OR PREBUILT AGENTS EXIST	WHERE NONE WAS FOUND
Applications scored, with the outcome emailed	**An admissions file checked for completeness, with the officer deciding**
Admission letters generated and emailed	**A transcript summarised for the registrar to verify**
Outreach to prospective students	**A research ethics application checked against the committee's checklist**
Answers to financial aid questions	**A grant draft checked against the funder's rules**
Transfer credit estimates	**A data management plan checked against the funder's requirements**
Feedback suggestions on a student's work	**Accreditation evidence mapped to each standard**
	A purchase request checked against policy
	A reading list checked against the library's holdings

Two findings matter more than the list. First, both admissions templates we opened send as soon as the workflow decides. One scores each applicant from 0 to 100 and routes on the score: "Auto-Accept (95-100): Acceptance letter with scholarship details", and "Reject (<70): Respectful rejection with improvement suggestions". A person appears only as a "Committee review option for borderline cases". The other emails an admission letter once the applicant's email address passes a check, and lists its upkeep as "Minimal (Set & forget)" [9].

Second, both are built from hosted services: a hosted form, a hosted model, a hosted mailbox, a hosted spreadsheet. That settles where applicants' data goes before anyone has asked [10].

The gaps sit where the reading is slow and the decision already belongs to a named person or a board. That makes them the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In a university or a school system that means three kinds of thing.

Documents.

Regulations, policies, programme handbooks, admissions criteria, funder rules, the ethics committee's checklist. These are indexed once as knowledge sources, with each reader's permissions kept, so an agent answers with the clause it relied on and says when the documents are silent.

Systems of record.

Student records, courses and enrolments, the library catalogue, research outputs, the service desk. Most are reached through an interface the institution already runs. The institution gives the platform the description of that interface, chooses which operations become tools, and decides which of them need approval.

Institutional data for analysis.

Enrolment, progression, survey results and research spend, held in the data warehouse or a reporting database, and read only. Examination and proctoring systems are not one of the doors.

All three pass one gate: every tool call is checked against the institution's policy, which can allow it, record it, hold it for a person, or block it.

Where the institution's interface checks the caller, the agent acts as the person using it and cannot claim another.

SIX WAYS INTO AN INSTITUTION'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	An application, a transcript, a module outline, a grant call, an ethics form	Nothing is connected. The safest start
02	**A knowledge source**	Regulations, handbooks, admissions criteria, funder rules, research procedures	Permissions are kept per reader
03	**An imported description of an interface**	Student records, enrolments, the library catalogue, research outputs, the service desk	You choose the operations. Changes wait for approval
04	**A reporting database**	Enrolment, progression, survey results, research spend	Read only
05	**A server from the tool catalogue, run in your cluster**	Databases, web search, document conversion	It runs inside; check what each tool does
06	**Public sources**	Funding calls, research indexes, regulators' guidance	Reads only, from the publisher

PART 03 · CONTINUED

32
public interface descriptions tried, each interface counted once

30
read by the platform's importer

2,892
operations in the 30 it read

1,186
of those operations are sent as changes

What is published. We tried 32 public education and research interface descriptions, counting each interface once, among them learning platforms, school data and roster standards, credential standards, research repositories and research indexes. The platform's importer read 30 and refused 2, the files of one learning platform published in an older format. That is a limit of the importer today, not of the platform those files describe. The 30 carry 2,892 operations [10].

What that exercise taught us. The open education interfaces are built to do things, not only to read. By method, 1,706 of the 2,892 operations read and 1,186 are sent as changes. One school data standard alone carries 1,222 operations, and 459 of them change something.

What is not published. Most student information, admissions and research information systems publish no open description. Where nothing is published, the institution supplies its own description, or the agent reads from the data warehouse.

Tests and classrooms. The safe design keeps the agent on the office side. It takes no part in a test, and no pupil talks to it.

So choose the operations at import, start with the ones that read, and add nothing that enrols, records a mark, deposits or submits until the approval step has been seen working.

PART 04

What an agent may do

One rule runs through every agent in this paper: it prepares, and a named person decides.

01	**It reads** what the person using it could read, and no more.
02	**It drafts** the reply to an enquiry, the completeness check, the progression list, the compliance note, with each statement tied to its source.
03	**It lists** what is missing, what does not match, what falls outside the regulations.
04	**It stops** before anything that changes a record or reaches an applicant, a student, a funder or a regulator. A tool that changes something waits for approval.
05	**It decides nothing about a person.** It does not rank applicants, does not score a student or a member of staff, and does not infer anyone's emotions.

Four lines are drawn tighter than the rest.

Admissions.

An agent may check an application against the programme's stated requirements and summarise a transcript with the page each grade came from. It admits no one, rejects no one and ranks no one. The admissions officer decides and sends.

Students and what they are told.

An agent may draft the reply to a student's enquiry from the handbook, the regulations and the student's record, with the clause cited. It sends nothing to a student unseen. A member of staff reads and sends every reply, and the adviser gives the advice.

Marks and progression.

An agent may draft written feedback against the published marking criteria for the marker to edit, and set results against the progression rules for the exam board. It sets no mark and proposes none. It decides no progression, no award, no appeal and no misconduct case, and it proctors no test. The marker, the board and the panel decide.

Research and grants.

An agent may check a draft application against the funder's rules, and check an ethics application against the committee's checklist. It writes no part of the case for support and no peer review, and it submits nothing. The researcher writes and submits; the committee decides.

PART 05

What the rules already ask

This part needs a lawyer's read before it is relied on. It reports what the texts say; it is not advice.

Admissions and assessment

Deciding admission or evaluating learning is high-risk. The European Union's AI Act lists as high-risk "AI systems intended to be used to determine access or admission or to assign natural persons to educational and vocational training institutions at all levels", and "AI systems intended to be used to evaluate learning outcomes" [11]. An agent that lists what an application is missing admits no one and evaluates nothing. One that scored applicants and sent the outcome, as the library template in section 2 does, would be read against this text.

So is software that polices a test. The same list includes systems "intended to be used for monitoring and detecting prohibited behaviour of students during tests" [11]. No agent here takes part in a test.

The dates have moved. Regulation (EU) 2026/1744 set the high-risk duties to apply from "2 December 2027 as regards AI systems classified as high-risk pursuant to Article 6(2) and Annex III" and from "2 August 2028" for those tied to Annex I [12]. Education sits in Annex III, so its date is 2 December 2027. A high-risk system must be built so that it "can be effectively overseen by natural persons during the period in which they are in use" [11].

PART 05 · CONTINUED

Inferring emotion is banned outright. The AI Act prohibits "the use of AI systems to infer emotions of a natural person in the areas of workplace and education institutions", except "for medical or safety reasons" [11].

In school and college qualifications, AI may not mark alone. The Joint Council for Qualifications tells exam centres in the United Kingdom: "Where centres do permit AI tools to be used to mark student work, an AI tool cannot be the sole marker." [13] It does not cover university degrees.

Students and their data

Decisions about a person stay with a person. Under the GDPR, a person "shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her", with stated exceptions [14].

People must be told. The AI Act asks that systems "intended to interact directly with natural persons" be built so that "the natural persons concerned are informed that they are interacting with an AI system", unless that is obvious [11].

A student's work is the student's. England's education department says: "Schools and colleges must not allow or cause students' original work to be used to train generative AI models unless they have permission, or an exception to copyright applies" [15]. An agent that runs on a model the institution serves sends no student's work to a supplier's model.

Research and funders

A funder may limit AI in an application and in its review. UK Research and Innovation says: "Applicants must not use generative AI tools to generate an entire application, or sections of an application, without human involvement". Its assessors must "not use generative AI tools as part of their assessment activities except for language refinement purposes" [16]. Here the agent checks a draft against the funder's rules and writes none of it, and no agent drafts a peer review.

Confidential material must not leave. UKRI says: "The release of confidential material into generative AI tools constitutes a breach of confidentiality and integrity" [16]. Whether a tool on the institution's own infrastructure counts as such a release is for the funder and counsel. Keeping an unpublished application inside is the cautious design.

Access and security

MARKET	WHAT WE READ	WHAT IT ASKS
European Union	NIS2 Directive, Annex II point 7	Lists "Research organisations" among other critical sectors; essential and important entities must take "appropriate and proportionate technical, operational and organisational measures" [17]
European Union	NIS2 Directive, Article 2(5)	Member States may apply the directive to "education institutions, in particular where they carry out critical research activities" [17]
European Union	GDPR, Article 32(1)	Controllers and processors "shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk" [14]
England	Department for Education guidance	"It is recommended that personal data is not used in generative AI tools" [15]

We did not open the law on student records in the United States, the rules of any other funder, or the rules in the Gulf, India or Australia. An institution there should ask its own counsel.

PART 05 · CONTINUED

What the rules have in common

01	**A person decides** admission, marks, progression and anything else that judges a student.
02	**Students are told** when they deal with AI, and nobody infers how they feel.
03	**The researcher's work stays the researcher's,** and a funder's confidential material stays inside.
04	**Student records and unpublished research need protecting,** which is easier where the institution controls the system.

PART 06

A worked example: one university

Take a made-up university: about 18,000 students, 2,400 staff, six faculties and a research office. Nothing here is a customer's story, and every figure is illustrative; it is section 4 put to work.

01

In admissions, an officer uploads an application and its documents. The agent lists two gaps: no language certificate, and a transcript whose dates do not match the form. The officer writes to the applicant and, later, decides.

02

In student services, a student asks about interrupting her studies. The agent drafts a reply from the regulations and her record. An adviser edits it and sends it.

03

In the registry, before the exam board meets, the agent sets 1,150 students' results against the progression rules and lists 64 cases that need a decision, each with the rule. The board decides every one.

04

In the research office, a researcher uploads a draft grant application. The agent lists three breaches of the funder's rules: a budget line over the cap, a missing data management plan and a CV over the page limit. The researcher fixes them and submits.

05

For the ethics committee, the agent checks an application against the checklist and finds no consent form for one group of participants. The secretary returns it.

06

In the quality office, the agent maps 310 documents to the 42 standards of an accreditation and lists five standards with no evidence. The quality lead writes the submission.

Every one of these reads documents or records the university already has, changes no record by itself, sends nothing to an applicant, a student or a funder, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent touches a student record

01 **Where does it run?**
On our infrastructure, in our private cloud, or in the supplier's service? Where is the model? Can it run with no connection to the internet?

02 **What can it read?**
Does it act as the person using it, with their permissions, or as one all-seeing account?

03 **What can it change?**
Which tools change an application, a student record, a mark, a research record or a payment, and does each wait for a person?

04 **Can it speak for us?**
Show us that no reply, offer or decision reaches an applicant or a student before a member of staff has read and sent it.

05 **Who decides?**
Show us the step where a named person or board approves, and what the agent does when nobody has.

06 **What is recorded?**
Who asked, which tool ran, what was decided, and for how long is that kept?

07 **Can we stop it?**
One agent, one tool, or all of them, at once, by us.

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and no student waits on the agent.

Questions on your own regulations

Staff ask what the institution's regulations and policies say and get the answer with the clause cited. No student record is touched, so it can run on the first day.

Grant application compliance check

The agent lists where a draft breaks the funder's rules. The researcher fixes it and submits.

Accreditation evidence

The agent maps documents to standards and lists the gaps. The quality lead writes the submission.

One department, one process, invented data first.
Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/EDUCATION

SOURCES

Read on 22 November 2026. Figures are the publishers' own and were not audited by us.

01 Hunter Gehlbach and Lindsay C. Page, "Freezing 'summer melt' in its tracks: Increasing college enrollment with AI", Brookings Institution, 11 September 2018 (the trial's own researchers). brookings.edu

02 UC San Diego, TritonAI, "Responsible AI, Privacy and Hosting". tritonai.ucsd.edu

03 University of Amsterdam, "UvA AI Chat: Why was it created, and what is it exactly?", 17 December 2025. uva.nl

04 The 74, "LA Unified Faces Criticism After Collapse of Splashy AI tool 'Ed'", 28 July 2024. the74million.org

05 Vanderbilt University, "Guidance on AI Detection and Why We're Disabling Turnitin's AI Detector", 16 August 2023. vanderbilt.edu

06 D2L, "Our AI story" (D2L Lumi). d2l.com

07 Microsoft, "Application card: AI in Microsoft 365 Education", 21 October 2026. learn.microsoft.com

08 Instructure, IgniteAI page, instructure.com; and Salesforce, Agentforce for education, recruitment and admissions page, salesforce.com

09 n8n workflow templates 9574 (university application evaluation and scholarships) and 9504 (university admission process). n8n.io

10 Template libraries and prebuilt agents of agent and workflow platforms, searched 22 November 2026; and 32 interface descriptions fetched from their publishers and standards groups the same day and run through the platform's importer (Katonic research).

11 Regulation (EU) 2024/1689 (AI Act), Annex III point 3(a), (b) and (d); Article 5(1)(f); Article 14(1); Article 50(1).

12 Regulation (EU) 2026/1744, 8 July 2026, the amended Article 113.

13 Joint Council for Qualifications, "AI Use in Assessments: Your role in protecting the integrity of qualifications", revised 20 November 2026. jcq.org.uk

14 Regulation (EU) 2016/679 (GDPR), Articles 22(1) and 32(1).

15 Department for Education (England), "Generative artificial intelligence (AI) in education", updated 12 August 2025. gov.uk

16 UK Research and Innovation, "Use of generative artificial intelligence in application preparation and assessment", 23 September 2024. ukri.org

17 Directive (EU) 2022/2555 (NIS2), Articles 2(5)(b) and 21(1), and Annex II point 7.