

RESEARCH FINDINGS REPORT

AI-Supported Study Guidance in European Higher Education

Institutional Needs, Adoption Barriers and Solution Directions

Evidence from 45 interviews across European higher education, with 14 institutional cases formally coded

Mario Behrends

Hogeschool Utrecht · Creative Business · 2026

Independent bachelor research

Anonymity statement. This report is shared with the institutions that contributed to the research, in fulfilment of the commitment made during data collection. All participants and their institutions have been fully anonymised. No personal or institutional names appear anywhere in this document. Participants are identified only by an anonymous code and by non-identifying descriptors (role, country and institution type).

Contents

- 1. Executive summary**
- 2. Introduction and context**
- 3. How the research was done**
- 4. Findings**
 - 4.1 Needs and pain points (SQ1, Round 1)
 - 4.2 Adoption barriers (SQ2, Round 1)
 - 4.3 Solution direction and pilot readiness (SQ3, Round 2)
- 5. The wider interview programme (demand-side evidence)**
- 6. What this means for institutions**
- 7. Limitations**

References

Tables. Table 1 Sample composition · Table 2 Evidence pattern across 21 indicators.

Figures. Fig. 1 Cross-case support · Fig. 2 Procurement routes · Fig. 3 Champion archetypes · Fig. 4 Scenario preference.

1. Executive summary

Across Europe, prospective students struggle to navigate fragmented programme information, while the staff who guide them work under structural capacity constraints. This report presents the findings of an independent study into whether higher education institutions (HEIs) would adopt an AI-based study guidance tool, and under what conditions. The project drew on 45 interviews conducted between January and May 2026 with prospective students, university staff and sector representatives. Fourteen institutional interviews, concentrated in the DACH region (Germany and Austria), were formally coded against a structured framework and form the evidence base for the findings below; the wider set, including the student interviews, was analysed to the point of saturation and is reported as supplementary, demand-side evidence in Section 5.

Key takeaways *The seven points that carry the report. Each is expanded in Sections 4 and 5.*

- **The website is the real problem, not missing information.** A high, repetitive enquiry burden is structural across the sector. The institutional website, not a shortage of content, is the root cause of what students bring to staff (see F1, F2).
- **The payoff is reclaimed staff time, not fewer staff.** Guidance teams are structurally under capacity. The realistic return on an AI tool is freeing time for in-depth advising, never headcount reduction (see F3, F8).
- **A marketing-budget pilot beats slow procurement.** Public procurement commonly takes 6-18 months. A pilot scoped to one use case and funded from a department's own budget is a confirmed workaround several institutions already use (see F5, F13).
- **GDPR is a documentation gate, not a wall.** It is cleared, not avoided, with a data-processing agreement, EU data residency and an Article 13 notice, prepared before the first conversation (see F7, F10).
- **An internal champion decides adoption.** An internal champion (a staff member who actively advocates for the tool inside the institution) is the single strongest predictor of adoption. The evidence that convinces one differs by whether they sit in counselling or in marketing and management (see F6, F12).
- **Every solution direction was accepted, and fit follows the bottleneck.** Shown concrete scenarios, institutions rejected none. The right entry point depends on each institution's primary bottleneck, and pilot success is judged on two things at once: capacity relief and factual accuracy, meaning no confident-but-wrong answers on binding content (see F9, F11).
- **Students already use AI to choose, and would trust an institution's own tool more.** The wider interviews show prospective students reaching for general AI assistants but finding them imprecise. They prefer, and trust more, a tool tied to a university's own data, provided it is accurate (Section 5).

These conditions are established for the DACH region. A reader pressed for time can stop here. The sections that follow give the supporting evidence, the visualisations and the methodology behind each point.

2. Introduction and context

The European Higher Education Area offers a very large number of degree programmes across thousands of institutions, with programme information fragmented across many languages and systems (Eurostat, n.d.). For a prospective student, the practical consequence is a discovery problem: the relevant information usually exists, but it is hard to find. University websites remain the most consulted source during the application phase (Heathcote et al., 2020), yet the difficulty lies in how the information is organised on these sites, not in whether it exists (Krejtz et al., 2025; Nagel et al., 2020).

On the institutional side, study guidance and counselling capacity is structurally constrained. Benchmarks place healthy student-to-counsellor ratios far below what most European public institutions can resource (Kearney et al., 2021). AI chatbots are widely discussed as a response, but the existing literature is largely conceptual or focused on student-facing learning applications, rather than on the question of institutional adoption (Labadze et al., 2023). Recent work calls explicitly for country-level empirical evidence on what drives AI adoption in higher education (Abulail et al., 2025). A further structural feature is worth naming: because such a tool is purchased by the institution but used by the student, incentives can diverge, and platforms of this kind tend to optimise for the buyer rather than the end-user (Eisenmann et al., 2006; Segato, 2022).

This study addresses that gap. The DACH region was chosen as the setting because it combines strict data-protection rules and slow, multi-stakeholder public procurement with direct language access for the researcher (Hammerschmid et al., 2013). It is a demanding environment, so conditions that hold here are likely to transfer to lighter-regulated markets.

Research questions

The report answers one main question and three sub-questions. The main question asks: under what conditions would European higher education institutions adopt an AI-based study guidance tool? The three sub-questions break this down, and Section 4 answers them in turn:

- **SQ1** What needs and pain points do institutions experience around student recruitment and guidance? (answered in 4.1)
- **SQ2** What organisational, regulatory and operational barriers shape their readiness to adopt an external AI tool? (answered in 4.2)
- **SQ3** Which solution directions and pilot models would institutions commit to when shown concrete scenarios? (answered in 4.3)

3. How the research was done

In total, 45 interviews were conducted across the project, with prospective students, university staff and sector advisers. The findings below rest on a focused subset: 14 institutional interviews that were formally coded against a structured framework (Table 1). The wider set, including the student interviews, was analysed thematically to the point of saturation and is reported separately in Section 5. The 14 coded interviews used a two-round design. Round 1 (ten interviews) explored institutional needs and adoption barriers; Round 2 (four institutional cases) tested concrete solution scenarios. Participants were selected

purposively for role relevance (Patton, 2002): each had operational visibility into recruitment, counselling or digital strategy at an established HEI. Two supplementary interviews were included for directional comparison: one German-speaking Italian university and one EdTech sector perspective.

Table 1. Composition of the coded interview sample (N = 14, of 45 conducted)

Dimension	Breakdown (number of interviews)
Country	Germany 9 · Austria 3 · Italy* 1 · Netherlands* 1
Institution type	University of applied sciences 6 · Public university 4 · Sector / industry body 2 · Other** 2
Research round	Round 1, needs & barriers 10 · Round 2, solution test 4

*Supplementary interviews included for directional comparison (one German-speaking Italian university; one EdTech sector perspective).

**One university of education and one private medical university.

Interviews lasted 30-90 minutes and were held in German or English between January and May 2026. All participants gave informed consent for recording and analysis, and all data were handled in line with GDPR. Interview design and analysis were anchored in established adoption theory: the diffusion of innovations (Rogers, 2003), the technology-organisation-environment framework (Tornatzky et al., 1990) and the technology acceptance model (Davis, 1989), in the integrated form validated for AI in higher education by Abulail et al. (2025), together with product-market-fit theory from the lean-startup literature (Blank & Dorf, 2012; Ries, 2011).


Each coded transcript was scored against a structured framework of indicators (21 in Round 1, 15 in Round 2), and every indicator marked as confirmed, mixed or a gap. To keep the analysis conservative, only patterns supported by three or more respondents are reported as findings (Eriksson & Kovalainen, 2008), and saturation on the core indicators was reached within the first seven interviews (Saunders et al., 2018). The denominators in the findings differ on purpose. A figure such as “8 of 8” rather than “10 of 10” means that not every topic applied to every interview: the two supplementary, non-institutional interviews are excluded from institution-specific counts, and a few interviews did not touch on a given indicator. “8 of 8” therefore means the pattern held in every case where the topic actually came up.

A note on independence and anonymity

The research was conducted independently as part of a bachelor programme. Questions were framed around institutional problems rather than around any specific product, to reduce bias; as a German-speaking researcher studying the German-speaking region, the author also took deliberate steps to surface rather than assume institutional realities (Dwyer & Buckle, 2009). For this version of the report, every participant and institution has been anonymised: names are replaced by an anonymous code (P1-P14) and by role, country and institution-type descriptors only. German quotes are reproduced verbatim with an English translation in square brackets; interviews held in English are quoted directly. Any third-party names inside quotes have been removed.

4. Findings

The findings are organised around the three sub-questions: SQ1, needs and pain points (4.1); SQ2, adoption barriers (4.2); and SQ3, solution direction and pilot readiness (4.3). Sections 4.1 and 4.2 draw on Round 1; Section 4.3 draws on Round 2. Figure 1 gives an overview of how strongly the Round 1 findings were supported across cases. Table 2 then lists the full evidence pattern across all 21 Round 1 indicators, with the cross-case verdict and signal strength for each.



Bars show the number of institutions in which each pattern was confirmed. Only findings supported by three or more respondents were retained.

Table 2. Evidence pattern across the 21 Round 1 indicators

Indicator	Verdict	Signal
1.1.1 Weekly enquiry volume	Confirmed	Very strong
1.1.2 Response handling time	Mixed	Moderate
1.2.1 Staff-to-enquiry ratio	Confirmed	Very strong
1.2.2 Overflow & delegation	Confirmed	Strong
2.1.1 Programme overview access	Confirmed	Very strong
2.1.2 Student confusion points	Confirmed	Very strong
2.2.1 International complexity	Confirmed	Strong
3.1.1 Procurement pathway	Confirmed	Very strong
3.1.2 Budget authority	Confirmed	Strong
3.2.1 Staff openness to AI	Mixed	Moderate
3.2.2 Internal digital champions	Mixed	Strong
3.3.1 Student adoption likelihood	Confirmed	Strong
3.4.1 Prior AI / chatbot use	Mixed	Strong
4.1.1 GDPR awareness	Mixed	Moderate
4.1.2 Required documentation	Confirmed	Strong
4.2.1 IT security process	Gap	Weak
4.2.2 CMS compatibility	Gap	Weak

5.1.1 Success metrics / KPIs	Mixed	Moderate
5.1.2 ROI logic	Confirmed	Strong
5.2.1 Vendor credibility	Mixed	Moderate
5.2.2 Pilot / proof of concept	Confirmed	Very strong

Verdict colour: green = Confirmed (consistent across cases), amber = Mixed (present but variable), grey = Gap (insufficient evidence).

Signal = strength of the cross-case pattern.

4.1 Institutional needs and pain points (SQ1 · Round 1)

F1. A high, repetitive enquiry burden is a structural feature of the sector

Supported by: 9 of 10 institutions

Counselling and admissions teams absorb a large, recurring volume of prospective-student enquiries. The content is consistent (application status, deadlines, financing, portal navigation) and largely answerable from information that already exists on the website but cannot be found.

“Im letzten Jahr hatten wir 1.600 Telefonate in der Allgemeinen Studienberatung. Wir haben über 600 persönliche Kontakte. Und dann kommen natürlich noch mal 1.700 E-Mail-Anfragen hinzu.” [Last year we had 1.600 phone calls in general student counselling. We had over 600 in-person contacts. And then of course another 1.700 email enquiries on top.]

P2 (general counselling, university of applied sciences, DE)

“um die 20 Anrufe am Tag gehen locker bei mir ein. [...] der Großteil der Fragen, die ich beantworte, wiederholen sich.” [around 20 calls a day easily come in to me. [...] the majority of the questions I answer repeat themselves.]

P10 (admissions, private medical university, DE)

F2. Programme navigation failure on the institutional website is the root cause

Supported by: 8 of 8 Round 1 institutional cases; 4 of 4 Round 2

Every institutional case names its own website as the primary failure point. The problem is information architecture, not a lack of information: programmes are duplicated, scattered across portals, or described in contradictory documents. Structured, consistent content is the technical prerequisite for any AI layer.

“Wir ja 47 Studiengänge haben, die teilweise sehr ähnlich sind. Und wenn man aus der Schule kommt [...], kann das vielleicht auch überfordernd sein.” [We have 47 study programmes, some of which are very similar. And when you come straight from school, that can be overwhelming.]

P3 (marketing & digital, university of applied sciences, AT)

“Wenn auf der Website beispielsweise fünf verschiedene Dokumente sind, zum gleichen Studiengang, die nicht mehr übereinstimmen, dann kann die KI so gut sein, wie sie möchte. Sie wird ihnen nicht das richtige Ergebnis geben.” [If there are, say, five different documents on the website for the same programme that no longer match, then the AI can be as good as it likes. It will not give them the right answer.]

P3 (marketing & digital, university of applied sciences, AT)

F3. Guidance teams are structurally under capacity

Supported by: 10 of 10 contexts

Lean staffing is universal: typically one or two part-time staff serving thousands of prospective-student contacts a year. Capacity will not close through hiring under current public funding. The return on an AI tool is therefore reclaimed staff time, not headcount reduction.

“Ich bin mit 20 Stunden in der Studienberatung tätig. Und meine Kollegin [...] hat auch 20 Stunden [...]. Ansonsten wäre das nicht zu schaffen.” [I work 20 hours a week in student counselling. And my colleague also works 20 hours. Otherwise it would not be manageable.]

P2 (general counselling, university of applied sciences, DE)

“Hier ist, glaube ich, der Schlüssel irgendwie ein paar Tausend pro Studienberaterin und Studienberater. In den USA teilweise ein Berater auf irgendwie 25 Studierende.” [The key figure here is, I think, a few thousand [students] per counsellor. In the US it is sometimes one counsellor per 25 students.]

P5 (managing director, sector coordination body, DE)

F4. International cohorts concentrate the pain and the clearest AI-adoption signal

Supported by: 5 of 10 institutions

International applicants concentrate complexity in language, time zone and qualification recognition. This segment produces the single most direct adoption signal in the dataset and makes a multilingual, out-of-hours tool a functional requirement rather than a nice-to-have.

“Ich glaube, wir liegen bei 25% internationalen Studierendenanteil. Wir haben mittlerweile auch rein englischsprachiges Masterstudienangebot, wo wir sehr viele indische Studierende [haben], die wirklich null Sprachkenntnisse auf Deutsch mitbringen.” [I think we are at a 25% international student share. We now also have an entirely English-language master's offering, with many Indian students who bring zero German-language skills.]

P2 (general counselling, university of applied sciences, DE)

“International students write at 2 a.m. and expect an answer. If you answer them relatively quickly, the next question is often: ‘Are you a bot?’”

P6 (head of recruitment, public university, IT, German-speaking) · interview held in English

4.2 Organisational, regulatory and operational adoption barriers (SQ2 · Round 1)

F5. Public procurement is slow; a marketing-budget pilot is a confirmed workaround

Supported by: Procurement complexity 8 of 8 public cases; marketing-budget pilot 3 of 8, two countries

Public-university procurement is multi-stakeholder and slow (commonly 6-18 months, EU-tender routes longer). A replicable workaround surfaces independently at several institutions: classify the first pilot year as a marketing-campaign tool funded from the marketing budget, which keeps it off the longest approval chain.

“Wir sind eine öffentliche Einrichtung. Und der Albtraum eines jeden, der etwas beschaffen muss, ist eine Beschaffung von Software.” [We are a public institution. And the nightmare of anyone who has to procure something is procuring software.]

P2 (general counselling, university of applied sciences, DE)

“Als Pilotprojekt für ein Jahr finanzieren wir das aus Marketing-Budget-Topf sozusagen.” [As a one-year pilot project, we finance it from the marketing budget, so to speak.]

P4 (digital marketing manager, university of applied sciences, DE)

F6. The internal champion is the strongest single predictor of adoption

Supported by: Typology confirmed 10 of 10; a champion present in 6 of 10

An internal champion is a staff member who actively advocates for the tool inside the institution. Two champion archetypes recur. Counselling champions hold domain knowledge and operational urgency but limited budget authority. Marketing/management champions hold autonomous budget and drive

procurement directly. Institutions without a local champion stall. Champion presence and type are the first qualification criteria.

“Es muss halt immer jemand geben, der von dem überzeugt ist, der dann andere überzeugen kann, dass es ein Vorteil ist.” [There always has to be someone who is convinced of it and who can then convince others that it is an advantage.]

P12 (student counsellor, public university, AT)

“Yes, I even pushed it. [...] I will push it forward.”

P6 (head of recruitment, public university, IT, German-speaking) · interview held in English

F7. GDPR is a universal procurement gate, manageable as documentation

Supported by: 8 of 8 cases

GDPR awareness is universal; operational depth varies. Counselling staff apply a surface heuristic (EU servers plus a DSGVO claim); data protection officers are the formal gate. A vendor that arrives with a Data Processing Agreement, EU data residency and an Article 13 statement moves through the gate faster. GDPR is the cost of entry, not a dealbreaker.

“The provider claims it is GDPR-compliant. The servers are in the EU. That sounds okay to me, but I'm not a legal expert. We have a legal department for that.”

P1 (head of study counselling, public university, DE) · interview held in English

“typische Argumente wie Datenschutz [...] entsprechend auch schon vorbereitet zu haben, um das [...] direkt mitzuschicken an die Stellen, dass die [...] schon verargumentiert ist.” [to have the typical arguments, such as data protection, prepared in advance, so you can send them straight to the relevant offices with the case already made.]

P4 (digital marketing manager, university of applied sciences, DE)

F8. ROI is staff-time reclamation, conditioned on a hard accuracy threshold

Supported by: 10 of 10 respondents

No respondent frames value as headcount reduction. The dominant logic is freeing capacity for the in-depth conversations counsellors regard as their professional identity. A second condition is absolute: hallucination on binding content (admission requirements, deadlines, deposits) is treated as a disqualifying defect because incorrect guidance creates institutional liability.

“With recurring enquiries, I don't have to spend my own time on them. I can invest that time in more in-depth consultations.”

P1 (head of study counselling, public university, DE) · interview held in English

“Wenn so eine KI halluziniert und irgendwelche verschiedenen Hochschulzugangsmöglichkeiten zusammenwirft [...]” [If such an AI hallucinates and conflates different university-entry options [...] [that is the risk].]

P5 (managing director, sector coordination body, DE)

Two of these barrier findings are worth showing directly. Figure 2 contrasts how long the different procurement routes take, and Figure 3 sets out the two champion types and the kind of evidence each one responds to.

Figure 2. Decision-to-deployment time by procurement route




Figure 3. Two internal-champion archetypes and the credibility signal each requires

Counselling champion	Marketing / management champion
• Holds domain knowledge and daily operational urgency • Limited or no budget authority • Convinced by concrete demonstrations on their own content • Not moved by event invitations or formal pitches	• Holds autonomous budget; drives procurement directly • Thinks in integration, compliance and credibility • Convinced by smart integration and trust signals • Wants a documented compliance / integration package

Champion presence was the strongest single predictor of adoption. Matching the credibility asset to the champion type determines whether it lands.

4.3 Solution direction and pilot readiness (SQ3 · Round 2)

F9. All three solution scenarios are accepted; the ranking is institution-specific

Supported by: 4 of 4 institutional cases

Presented with three neutral, unbranded scenarios (a website chatbot, a back-office staff-support tool, and an AI pre-consultation step), no institution structurally rejected any. Rankings diverged by bottleneck: the staff-support tool leads where email volume is high; the chatbot is a universal navigation layer; the pre-consultation step leads where counselling intent is the bottleneck.

“Eigentlich alle drei. Sehr sinnvoll, weil Punkt eins ist einmal der klassische Chatbot auf der Webseite mit Navigation zu den richtigen Informationsseiten.” [Actually all three. Very useful, because point one is the classic website chatbot navigating to the right information pages.]

P14 (counsellor & project lead, university of applied sciences, AT)

“Letzter hätte für mich den größten Mehrwert, dann das erste und dann das zweite. [...] Vorbereitung auf das Gespräch [...], für die Bewerbung [...], der E-Mail-Chatbot.” [The last [pre-consultation] would have the greatest value for me, then the first, then the second: preparing the conversation, then the application chatbot, then the email tool.]

P11 (student counsellor, university of education, DE)

F10. Data protection has hardened from a procurement gate into an infrastructure filter

Supported by: 4 of 4 cases

Data protection no longer blocks only the final purchase. It now blocks practitioner-level testing. One institution operates a mandatory triple-gate (budget + data protection + IT security); at another, a willing champion cannot even trial the email use case because the relevant connectors are disabled institution-wide for compliance reasons. A minimum compliance kit must be ready before the first conversation.

“Ich würde gerne experimentieren [...], aber die Konnektoren zu Outlook sind bei uns [...] noch deaktiviert. Also ich kann dann nicht mal testen [...], weil es bei uns einfach datenschutz[rechtlich] [...]” [I would like to experiment, but the Outlook connectors are still deactivated here. So I can't even test it, simply for data-protection reasons.]

P14 (counsellor & project lead, university of applied sciences, AT)

“Das sind eben Voraussetzungen, die wir erfüllen [müssen]. Also dass das notwendige Budget bereitgestellt wird, dass Datenschutz, IT-Sicherheit befragt werden und ihr Go gegeben haben.” [These are the preconditions we must meet: that the necessary budget is provided, and that data protection and IT security have been consulted and given their go-ahead.]

P13 (management & counsellor lead, university of applied sciences, DE)

F11. Pilot success requires a dual criterion: capacity relief AND accuracy

Supported by: 4 of 4 cases

Every institution defines success in two parts: measurable workload relief, plus output quality high enough to avoid manual correction. Below the accuracy threshold the pilot fails regardless of capacity gained. One respondent also flags an onboarding caveat: quality control initially generates more work before it saves time.

“für mich wäre es ein absoluter Mehrwert, wenn wir dadurch Kapazitäten bei den Mitarbeitern einsparen können. [...] wenn die Qualität der Antworten auch dementsprechend ist, dass man nicht nachbessern muss.” [For me it would be an absolute benefit if it let us save staff capacity, provided the quality of the answers is such that you don't have to correct them afterwards.]

P13 (management & counsellor lead, university of applied sciences, DE)

“Der Chatbot war für uns am Anfang mehr Arbeit, weil wir die ganze Qualitätskontrolle machen mussten. [...] Das war eigentlich am Anfang mehr Arbeit und nicht weniger.” [At the start the chatbot was more work for us, because we had to do all the quality control. At first it was actually more work, not less.]

P12 (student counsellor, public university, AT)

F12. Champion type maps to the credibility signal that lands

Supported by: Mixed, 4 of 4 cases

Counselling champions respond to authentic relevance and concrete demonstrations on their own content, not to formal invitations or events. Marketing/management champions respond to formal credibility signals: smart integration into existing systems, trustworthiness and compliance. Champion type determines which sales asset works.

“Es nutzt nichts, wenn man da schon Einladungen kriegt zu irgendwelchen Events. [...] wenn man Beispiele [...] konkret [...] zeigt, aufzeigt.” [It's no use receiving invitations to events. [It helps] when you show concrete examples and use cases.]

P12 (student counsellor, public university, AT)

“Es müsste smart sein und es müsste integrierbar sein in unser bestehendes System [...] wie ein Baustein smart einfügbar sein. Das muss [...] vertrauenswürdig sein.” [It would have to be smart and integrable into our

existing system, slotting in like a building block. And it must be trustworthy.]

P13 (management & counsellor lead, university of applied sciences, DE)

F13. Pilot readiness is institution-specific and time-sensitive

Supported by: Mixed, 4 of 4 cases

Two structural variables shape readiness beyond Round 1: budget austerity (one public system absorbing spending cuts) and tool-replacement timing (a website relaunch is both a short-term blocker and a clean integration moment). A concrete budget data point recurs: departmental sign-off autonomy below a five-figure threshold keeps a pilot off the longest approval chain.

“Budget ist auch bei uns sowieso immer ein schwieriger Punkt, weil wir jetzt gerade einen Sparkurs haben.”

[Budget is always a difficult point for us anyway, because we are currently on an austerity course.]


P12 (student counsellor, public university, AT)

“Es gibt bei uns schon [...] Richtlinien, dass zum Beispiel ab 10.000 Euro Rechnungen muss die Geschäftsführung mit unterschreiben.” [We do have rules. For example, invoices above 10.000€ must be co-signed by the management board.]

P14 (counsellor & project lead, university of applied sciences, AT)

Figure 4 shows how the four Round 2 institutions ranked the three solution scenarios. The ranking follows each institution's main bottleneck, not the quality of the scenario itself.

Figure 4. Scenario preference by institutional bottleneck (Round 2, N = 4)



No scenario was structurally rejected. The ranking follows each institution's primary bottleneck, not scenario quality.

5. The wider interview programme: demand-side and supplementary evidence

Note. The observations in this section come from interviews that were not coded against the indicator framework, including the prospective-student interviews. They are thematic and exploratory, included to show the direction of the wider evidence, and should be read as supporting context rather than as primary findings.

The 45 interviews conducted across the project went well beyond the 14 institutional cases that were formally coded for the findings above. The wider set included prospective students (the end-users of any tool) and further university staff and sector advisers across several countries. This material was analysed thematically rather than coded against the indicator framework, so it is reported here as supporting observation, not as primary evidence. Two things stand out: it contradicted none of the coded findings, and it added the end-user perspective the institutional interviews could not capture.

The student (end-user) side confirms the problem institutions describe

Prospective students report exactly the discovery and navigation pain that staff describe from the other side of the desk. Choosing a programme is slow and confusing, and poor information leads to wrong-fit choices.

"Sometimes I lost days searching for nothing, because I didn't find anything. With all the time I dedicated to it, I should have come up with at least 20 or 30 master's programmes, but no."

a prospective student

"When I applied, I didn't know any difference between a research university and a university of applied sciences. I didn't know that."

a prospective student

Students already reach for AI, but distrust generic tools and want an institution-specific one

Students are not waiting for institutions to act; they already use general-purpose AI assistants to search for programmes. But they find them imprecise, and they say they would trust an AI tied to a university's own data more, provided it is accurate and shows its reasoning.

"I use a general AI assistant, but it gave me the same ten things, and some were not even related. It was not precise. So I just searched manually, but it was so long."

a prospective student

"I would prefer to use it on the university website, because if it is just on its own, it really needs to be a good one."

a prospective student

"I don't use AI that much, because it gives you the answer without you understanding where the answer came from."

a prospective student

Additional staff and advisers reconfirm the institutional findings

Beyond the coded cases, further staff independently described the same root cause (information that exists on the website but cannot be found), the same value (reclaimed time on repetitive work), and the same pre-consultation use case tested in the coded findings, this time in a different country and system.

"The information is already there on the website. They just can't find it. Whereas with a little AI chatbot you can just ask: am I eligible?"

a university staff member

"This would make my job, especially with the repetitive emails, a lot easier."

a university staff member

"Your software would really help me. I wouldn't have to do that anymore. It would save me a lot of work."

a study adviser at a Dutch university, describing a manual study-matching process (a long survey, a campus day, then a written report per applicant) that mirrors the pre-consultation scenario

In short, the wider data is convergent. End-users feel the same navigation and discovery pain that institutions report, already reach for AI to solve it, and would trust an institution-hosted tool more than a generic one, on the condition that it is accurate. This strengthens, from the demand side, the case the coded findings make from the supply side.

6. What this means for institutions

Read together, the findings describe a sector with a real, structural problem and a realistic path to addressing part of it. Four practical implications follow for any institution weighing an external AI study-guidance tool.

1. **Treat the website as the first fix, not the tool itself.** The navigation problem is one of information architecture. An AI layer only performs as well as the content beneath it; consistent, structured, non-contradictory programme information is the prerequisite for any tool to add value.
2. **Frame value as reclaimed staff time, with accuracy as a hard condition.** The defensible business case is capacity for the in-depth conversations counsellors value, not staff reduction. But factual accuracy on binding content (admission requirements, deadlines, deposits) is a knock-out condition, because incorrect guidance creates institutional liability.
3. **Use a small, reversible pilot to bypass procurement gravity.** A pilot scoped to one use case and financed from a marketing budget keeps a first deployment off the longest approval chain, cutting time to decision from many months to a few. Preparing the minimum compliance kit in advance lets an internal advocate clear the data-protection gate without waiting on the vendor. That kit is a data-processing agreement, EU data residency and an Article 13 notice (Dushi, 2025).
4. **Identify the internal champion early, and match the evidence to them.** Adoption rises and falls on whether someone inside the institution advocates for it. Counselling champions are convinced by concrete demonstrations on their own content; marketing and management champions are convinced by smart integration, trust signals and a ready compliance package.

None of these conditions is exotic. Together they describe a low-risk way for an institution to test whether AI-supported guidance helps its own staff and students, before committing to anything larger.

7. Limitations

Four limitations qualify these findings. First, the Round 2 sample is small (four institutional cases), constrained by a short outreach window and the spring application peak; Round 2 should be read as confirming or complicating the Round 1 conditions, not as quantitative validation. Second, scenario-based interviewing captures stated preference, not actual procurement behaviour, and only a real pilot can test the latter. Third, the sample concentrates on Germany and Austria, with Switzerland unrepresented and the Italian case supplementary; the findings should not be generalised to Southern or Eastern European markets, or to lighter-regulated settings, without further validation. Fourth, the supplementary observations in Section 5 are drawn from interviews that were not coded; they are directional and should not be read with the same weight as the coded findings. The researcher also held an interest in the topic; this was mitigated by problem-framed questioning, role disclosure at the start of every interview, and a theory-derived coding framework applied conservatively (three-or-more-respondent threshold).

References

- Abulail, R. N., Badran, O. N., Shkoukani, M. A., & Omeish, F. (2025). Exploring the factors influencing AI adoption intentions in higher education: An integrated model of DOI, TOE, and TAM. *Computers*, 14(6), Article 230. <https://doi.org/10.3390/computers14060230>
- Blank, S., & Dorf, B. (2012). *The startup owner's manual: The step-by-step guide for building a great company*. K&S Ranch.
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319-340. <https://doi.org/10.2307/249008>
- Dushi, D. (2025, April 3). Ensuring compliance with data protection law when using GenAI: The case of higher education institutions. *Global Campus of Human Rights*. <https://www.gchumanrights.org/preparedness/ensuring-compliance-with-data-protection-law-when-using-genai-the-case-of-higher-education-institutions/>
- Dwyer, S. C., & Buckle, J. L. (2009). The space between: On being an insider-outsider in qualitative research. *International Journal of Qualitative Methods*, 8(1), 54-63. <https://doi.org/10.1177/160940690900800105>
- Eisenmann, T., Parker, G., & Van Alstyne, M. W. (2006). Strategies for two-sided markets. *Harvard Business Review*, 84(10), 92-101.
- Eriksson, P., & Kovalainen, A. (2008). *Qualitative methods in business research*. SAGE Publications.
- Eurostat. (n.d.). Tertiary education statistics. European Commission. Retrieved May 20, 2026, from https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Tertiary_education_statistics
- Hammerschmid, G., Van de Walle, S. G. J., Andrews, R., Bezes, P., Görnitz, A., Oprisor, A., & Štimac, V. (2013). Public administration reform in Europe: Views and experiences from senior executives in 10 countries (COCOPS Work Package 3 Cross-National Report). COCOPS. <http://hdl.handle.net/1765/78038>
- Heathcote, D., Savage, S., & Hosseinian-Far, A. (2020). Factors affecting university choice behaviour in the UK higher education. *Education Sciences*, 10(8), Article 199. <https://doi.org/10.3390/educsci10080199>
- Kearney, C., Akos, P., Domina, T., & Young, Z. (2021). Student-to-school counselor ratios: A meta-analytic review of the evidence. *Journal of Counseling & Development*, 99(4), 418-428. <https://doi.org/10.1002/jcad.12394>
- Krejtz, K., Marcus-Quinn, A., & Duarte, C. (2025). Higher education accessibility information in practice: A report on the accessibility of European universities. *Universal Access in the Information Society*, 24, 2673-2685. <https://doi.org/10.1007/s10209-025-01224-4>
- Labadze, L., Grigolia, M., & Machaidze, L. (2023). Role of AI chatbots in education: Systematic literature review. *International Journal of Educational Technology in Higher Education*, 20, Article 56. <https://doi.org/10.1186/s41239-023-00426-1>
- Nagel, M., Schäfer, S., Zlatkin-Troitschanskaia, O., Schemer, C., Maurer, M., Molerov, D., Schmidt, S., & Brückner, S. (2020). How do university students' web search behavior, website characteristics, and the interaction of both influence students' critical online reasoning? *Frontiers in Education*, 5, Article 565062. <https://doi.org/10.3389/feduc.2020.565062>

- Patton, M. Q. (2002). *Qualitative research and evaluation methods* (3rd ed.). SAGE Publications.
- Ries, E. (2011). *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*. Crown Business.
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- Saunders, B., Sim, J., Kingstone, T., Baker, S., Waterfield, J., Bartlam, B., Burroughs, H., & Jinks, C. (2018). Saturation in qualitative research: Exploring its conceptualization and operationalization. *Quality & Quantity*, 52(4), 1893-1907. <https://doi.org/10.1007/s11135-017-0574-8>
- Segato, G. (2022, November). Why ed-tech startups don't scale. <https://giansegato.com/essays/why-edtech-startups-dont-scale>
- Tornatzky, L. G., Fleischer, M., & Chakrabarti, A. K. (1990). *The processes of technological innovation*. Lexington Books.