Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A longitudinal case study of practice and policy gaps in responsible AI adoption across Ukrainian higher education from 2023 to 2025

Serhiy O. Semerikov

semerikov@gmail.com

Kryvyi Rih State Pedagogical University

Maiia V. Marienko

Institute for Digitalisation of Education of the NAES of Ukraine

Oksana M. Markova

Kryvyi Rih National University

Case Report

Keywords:

Posted Date: October 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11087302/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A longitudinal case study of practice and policy gaps in responsible AI adoption across Ukrainian higher education from 2023 to 2025

Serhiy O. Semerikov[1,2,3,4,5,6*], Maiia V. Marienko[2,1] and Oksana M. Markova[3]

[1]Kryvyi Rih State Pedagogical University, 54 Universytetskyi Ave., Kryvyi Rih, 50086, Ukraine.
[2]Institute for Digitalisation of Education of the NAES of Ukraine, 9 M. Berlynskoho Str., Kyiv, 04060, Ukraine.
[3]Kryvyi Rih National University, 11 Vitalii Matusevych Str., Kryvyi Rih, 50027, Ukraine.
[4]Zhytomyr Polytechnic State University, 103 Chudnivska Str., Zhytomyr, 10005, Ukraine.
[5]Center for Information-analytical and Technical Support of Nuclear Power Facilities Monitoring of the NAS of Ukraine, 34a Palladin Ave., Kyiv, 03142, Ukraine.
[6]Academy of Cognitive and Natural Sciences, 54 Universytetskyi Ave., Kryvyi Rih, 50086, Ukraine.

*Corresponding author(s). E-mail(s): semerikov@gmail.com, https://orcid.org/0000-0003-0789-0272;
Contributing authors: popelmaya@gmail.com, https://orcid.org/0000-0002-8087-962X; markova@knu.edu.ua, https://orcid.org/0000-0002-5236-6640;

Abstract

Generative artificial intelligence (GenAI) entered Ukrainian higher education during a period of unprecedented institutional disruption: full-scale war, infrastructure loss, and emergency distance learning. Two annual faculty surveys conducted by the Marketers Association of Ukraine in June 2023 ($n = 65$, 30 higher-education institutions) and June–August 2024 ($n = 82$, 35 higher-education institutions) document a rapid bottom-up uptake that outpaced any

institutional governance response, until April 2025, when the Ministry of Digital Transformation, the Ministry of Education and Science of Ukraine, the United Kingdom's International Development programme, and the Eurasia Foundation jointly issued national *Recommendations on the responsible implementation and use of artificial intelligence technologies in higher-education institutions*. The recommendations anchor higher-education AI governance in the seven trustworthy-AI principles of the European Commission's High-Level Expert Group (human oversight; technical robustness and safety; privacy and data governance; transparency; diversity, non-discrimination and fairness; societal and environmental well-being; accountability) and add academic integrity as an eighth, sector-specific principle. This paper uses the two surveys, conducted *before* the policy existed, as a baseline-fidelity measurement against those eight dimensions. We construct a *practice–policy fidelity matrix* mapping survey items to policy dimensions, score per-dimension fidelity for the 2024 cohort, and segment results by academic rank, discipline, and institutional ISO 31000 risk-management adoption. Baseline fidelity is low across the six dimensions with survey proxies (mean **0.227**, range **0.092** to **0.293**); the steepest gap is on *accountability*, the dimension named highest-risk by the EU AI Act for educational use. Two dimensions (diversity/non-discrimination; societal/environmental well-being) have no item proxy in the 2024 instrument and are explicit data gaps for any successor wave. Witnessed student academic-integrity violations rose from **40.0%** in 2023 to **52.4%** in 2024 against essentially flat policy coverage. We compare Ukrainian patterns against contemporary international evidence on faculty GenAI adoption (Saudi, Romanian, Turkish, Indian, Thai, and US studies) and propose a sequenced implementation pathway addressing the steepest-gradient dimensions first. The paper contributes the first peer-reviewed longitudinal evidence on Ukrainian higher-education AI adoption during the 2022–2025 war and a reusable fidelity-matrix methodology for benchmarking practice against national responsible-AI frameworks.

Keywords: Responsible AI, Higher education, Generative AI, Ukraine, Faculty adoption, Policy implementation, Longitudinal survey, Academic integrity, AI literacy, Wartime education, Practice-policy gap, Fidelity matrix

1 Introduction

This paper studies how faculty in Ukrainian higher-education institutions (HEIs) adopted large-language-model (LLM)-based generative artificial intelligence (GenAI) before any national framework governed responsible use, and how far that early practice diverges from the April 2025 *Recommendations on the responsible implementation and use of artificial intelligence technologies in higher-education institutions* [1] subsequently issued by the Ministry of Digital Transformation, the Ministry of Education and Science, and international partners. Two annual faculty surveys conducted by the NGO Marketers Association of Ukraine in June 2023 ($n = 65$, 30 HEIs) [2] and June–August 2024 ($n = 82$, 35 HEIs) [3] provide the empirical basis. The temporal ordering – surveys conducted *before* the policy framework existed – is methodologically

unusual and lets us measure baseline fidelity against the eight dimensions of responsible AI articulated in the policy: the seven trustworthy-AI principles of the European Commission's High-Level Expert Group as adopted verbatim in the Ukrainian framework (human oversight; technical robustness and safety; privacy and data governance; transparency; diversity, non-discrimination and fairness; societal and environmental well-being; accountability), plus academic integrity treated as a sector-specific eighth principle.

The contributions are: (1) the first peer-reviewed longitudinal evidence on faculty GenAI adoption in Ukrainian HEIs across two annual waves during the 2022–2025 war; (2) a *practice–policy fidelity matrix* mapping the 2024 survey instrument against the eight dimensions of the April-2025 national recommendations, with per-dimension fidelity scores and explicit data-gap identification for two dimensions lacking survey proxies; (3) disciplinary, rank, and institutional segmentation surfacing where implementation gradients will be steepest; (4) predicted implementation bottlenecks with named drivers (resource, cultural, governance) and a sequenced policy pathway; (5) international contextualisation against contemporary GenAI-adoption evidence from Saudi [4], Romanian [5], Turkish [6], Indian [7], Thai [8], and US-medical-school [9] settings.

The remainder of the paper proceeds as follows. Section 2 reviews the responsible-AI literature for higher education, the technology-acceptance evidence on GenAI uptake by faculty, and prior work on crisis-era higher education. Section 3 establishes the Ukrainian context, including a timeline of the two surveys against the April-2025 policy. Section 4 sets out the longitudinal-comparison procedure, the construction of the fidelity matrix, and the coding scheme for open-ended responses. Section 5 reports adoption trajectories, disciplinary divergence, the risk-perception/risk-behaviour paradox, and the per-dimension fidelity scores. Section 6 situates the Ukrainian patterns internationally and identifies implementation-gap drivers. Section 7 offers concrete sequencing recommendations. Sections 8 and 9 close the paper.

2 Background and related work

The international landscape of responsible-AI policy for higher education has crystallised rapidly between 2023 and 2025. UNESCO's *Guidance for generative AI in education and research* [10] provided the first system-level reference. The European Union's Artificial Intelligence Act [11] subsequently identified specific high-risk education uses and set transparency and human-oversight obligations. Within higher education itself, ethical frameworks have evolved from early principles statements [12] to systematic reviews of human-centred-AI practice [13] to recent governance-focused proposals [14–16]. These frameworks converge on the European Commission's High-Level Expert Group seven trustworthy-AI principles (human oversight; technical robustness and safety; privacy and data governance; transparency; diversity, non-discrimination and fairness; societal and environmental well-being; accountability),

and the April-2025 Ukrainian policy reproduces them verbatim before adding academic integrity as an eighth, sector-specific principle. These eight dimensions form the spine of our fidelity analysis in §4.3.

Empirical studies of faculty GenAI adoption since 2023 have applied and extended the Technology Acceptance Model and the Unified Theory of Acceptance and Use of Technology across diverse national contexts. Saudi dual-perspective survey work integrates self-determination theory with ethical decision-making [4]. Romanian work emphasises perceived ease of use, trust, and satisfaction in structural equation models [5]. Turkish national-scale evidence has rehabilitated subjective norms as a strong adoption predictor [6]. Faculty-focused trust-and-social-reinforcement studies extend TAM with social-cognitive constructs [17]. Comparable evidence comes from India [7], Thailand [8], and US medical schools [9]. Parallel work on assessment redesign has produced practical scales for ethical GenAI integration [18] and identified ethical-validity concerns for AI-mediated measurement [19]. Unintended consequences – including erosion of peer learning communities [20] and homogenisation effects on assessment outputs [21] – are an emerging counter-current.

The COVID-19 emergency remote-education period provides the closest prior analogue for forced, rapid adoption of new technologies under institutional duress [22]. Ukrainian higher education has since faced a successive crisis – the full-scale war that began in February 2022 – with longer duration, deeper infrastructure loss, and uneven distance-learning conditions. Recent Ukrainian work documents the state, threats, and challenges of GenAI use to support student learning in this period [23, 24], and the use of AI tools specifically for academic-staff professional development [25, 26]. The Holmes-Bialik-Fadel taxonomy [27] and the Crompton-Burke systematic review of AI in higher education [28] situate this Ukrainian work in the broader field.

3 The Ukrainian context

Ukraine entered the period under study with one of the largest higher-education systems in post-Soviet Europe: roughly 300 accredited institutions before the full-scale Russian invasion of February 2022. The war disrupted that system unevenly. Universities located in directly occupied or front-line territories – Donetsk National University imeni Vasylia Stusa, Mariupol State University, Skhidnoukrainian National University, and several others – relocated as displaced institutions and continued operating online. Universities further from the front, including the large Kyiv and Lviv research centres, sustained partial operations through rolling blackouts, missile strikes on power infrastructure, and air-raid-related interruptions to in-person teaching. Distance and hybrid learning, established as emergency practice during the COVID-19 pandemic [22], became a permanent operating mode for most institutions by the start of the 2022–23 academic year.

The arrival of consumer-accessible large-language-model assistants (notably ChatGPT in late 2022, followed by Bard, Claude, and others through 2023) coincided with this prolonged operational disruption. Faculty adopted these tools individually, without institutional procurement, training, or governance, against a backdrop of degraded physical infrastructure and high cognitive load on both teachers and students. By the

time of the first NGO Marketers Association of Ukraine survey in June 2023, AI-LLM use had become widespread enough for 95% of respondents to forecast at least moderate impact on their work in the coming academic year [2]. Prior Ukrainian work on digital-tool adoption in higher education had focused on cloud-based science teaching and on faculty professional development with open-science platforms [29, 30]; the AI-LLM moment found the system better prepared for online delivery than for the policy and integrity questions GenAI now raises.

The April-2025 *Recommendations on the responsible implementation and use of artificial intelligence technologies in higher-education institutions* [1] are the first national framework to address those questions. They were jointly issued by the Ministry of Digital Transformation, the Ministry of Education and Science of Ukraine, the United Kingdom's International Development programme, and the Eurasia Foundation, with a fifteen-person working group drawn from ministries, the NGO sector (notably the Ukrainian Scientific-Educational IT Society and the Kharkiv IT Cluster), and named universities including Lviv Polytechnic, the National Aerospace University "Kharkiv Aviation Institute", Sumy State University, Ukrainian Catholic University, Kryvyi Rih National University, and the Kyiv School of Economics. The document is anchored in the seven-principle trustworthy-AI framework of the European Commission's High-Level Expert Group, in the Council of Europe Framework Convention on Artificial Intelligence, in the EU AI Act, and in the Ukrainian White Paper on AI regulation.

Figure 1 situates the two surveys against the policy framework and major external events that bracketed the period.

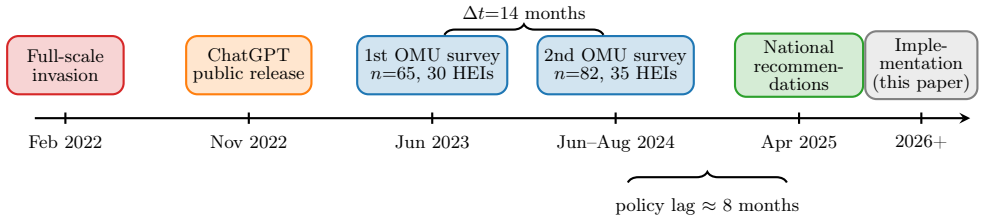


Fig. 1 Timeline of the two faculty surveys conducted by the Marketers Association of Ukraine against the April-2025 national policy framework. The surveys precede the policy by 22 months and 8 months respectively, allowing baseline-fidelity measurement before any national framework existed.

Three features of the timing matter for the analysis. First, both surveys precede the policy, so any practice they capture is the product of bottom-up institutional behaviour, not framework compliance. Second, the inter-survey gap of fourteen months captures the period of the most aggressive consumer-AI growth (the launches of GPT-4, Claude 2, Gemini, and an explosion of higher-education-targeted assistive tools); whatever moved between waves moved during that growth period. Third, the eight-month gap between the second survey and the policy is long enough for institutions to have begun internal discussions of governance, but short enough that the 2024 measurements remain a defensible baseline against which the policy's implementation can later be benchmarked.

4 Methodology

4.1 Survey instruments and samples

Both surveys were conducted by the NGO Marketers Association of Ukraine (OMU) as self-funded annual exercises. The 2023 instrument [2] comprised eight thematic blocks (impact on faculty work; impact on six spheres of university life; risk-management measures; student academic-integrity violations; faculty AI-use directions; attitude towards AI by stakeholders; familiarity with ISO 31000; open comments). The 2024 instrument [3] reorganised the same content into seven blocks, dropping the standalone ISO 31000 block (folded into risk management) and adding distinct items for actual and forecast impact in the survey year and the academic year ahead. Both surveys used a seven-point Likert response format for impact items and a five-point format for perceived-effectiveness items, with multi-select question types for tool use and risk-management measures.

Recruitment in 2023 was via OMU member email lists and the OMU website. The 2024 wave expanded distribution to specialised higher-education Facebook groups. Sixty-six per cent of 2024 respondents were OMU members. Both waves were conducted via Google Forms and processed in Excel with confirmatory analysis in SPSS. The 2023 sample is n=65 faculty members drawn from 30 higher-education institutions; the 2024 sample is n=82 from 35 institutions. Table 1 summarises the rank and discipline composition of both waves.

Table 1 Sample composition: academic rank and disciplines taught across the two annual waves.

	2023 (n=65)	**2024** (n=82)
Academic rank (%)		
No formal rank	15.4	8.0
Docent	55.4	64.0
Professor	29.2	28.0
Disciplines taught (%, multi-select)		
Marketing	49.2	48.0
Management	23.1	20.0
Economics	18.5	20.0
Entrepreneurship	9.2	2.0
Finance	10.8	7.0
Other economic disciplines	30.8	42.0
Non-economic disciplines	–	12.0
Research experience		
Mean (years)	21.61	21.9
Median (years)	20	20
Range (years)	1–45	2–50

The sample is disciplinarily skewed by the OMU recruitment channel: marketing-and-management faculty dominate. The skew is consistent across waves, which permits comparison but limits external validity to economic disciplines and adjacent areas.

STEM, humanities, and medical-education faculty are under-represented. We discuss the implications for generalisation in §4.5 and §8.

4.2 Longitudinal comparison procedure

Items present in both waves with identical or near-identical wording (the seven risk-management measures; the student-violation-witnessed item; the impact-on-faculty-work Likert distribution) are compared directly. For each item we report the 2023 percentage, the 2024 percentage, and the percentage-point delta. For pairs of proportions on small samples we report Wilson 95% confidence intervals; we treat overlap as failure to reject the null and do not present p-values, given the convenience-sample nature of both waves and the resulting limited interpretability of frequentist significance tests in this design.

Items that drifted between waves (the eight-block to seven-block reorganisation, and the separation of actual versus forecast items in 2024) are compared using the closest functional analogue, with the substitution flagged in the relevant results subsection. Where no functional analogue exists, the item is reported as a 2024-only measurement.

4.3 Practice–policy fidelity matrix construction

The construction proceeds in three steps. First, we read the policy document [1] and extract the seven trustworthy-AI principles (human oversight; technical robustness and safety; privacy and data governance; transparency; diversity, non-discrimination and fairness; societal and environmental well-being; accountability) plus the academic-integrity principle from the policy's dedicated section.[1] Second, each author independently maps each of the seven institutional-availability items of Table 9 of the 2024 report to zero, one, or more of those eight dimensions, blinded to the other's mapping; we then resolved disagreements through discussion guided by the conceptual fit of each item-dimension pair to the policy text and a preference against double-counting unless the survey item genuinely operationalises two distinct dimensions. Inter-rater agreement before resolution is reported in §5.4 as Cohen's κ [31]. Third, the per-dimension fidelity score is computed as the mean of the "Yes" response rates of the post-resolution mapped items (Table 9 of the 2024 report asks whether the respondent's institution actually has each measure, not whether it would be useful). Items where the institution-availability question was not asked, and dimensions for which no Table 9 item is conceptually proximate, are excluded from the score and noted as proxy-only measurements or as explicit data gaps respectively.

The scoring is conservative: "Don't know" responses are treated as fidelity-negative, on the grounds that an institution whose own faculty do not know whether a policy exists has, for practical purposes, no operative policy. We report a sensitivity analysis in §5.4 treating "Don't know" as fidelity-positive and discuss the implications for dimensions where the "Don't know" rate is high (especially ISO 31000 adoption, where 54.9% of respondents do not know whether their institution has adopted the standard).

[1] The original Ukrainian terms and English-language glosses for the eight principles are reproduced in Appendix A.

4.4 Coding of open-ended responses

Open-ended comments from both waves (Section 2.7 of the 2024 report; the equivalent section of the 2023 report) were coded by both authors independently using a two-pass scheme. The first pass categorised each comment into one of seven preliminary themes derived from the OMU authors' own coding (tool; need-training; positive-attitude; need-policy; neutral-attitude; need-further-research; negative-attitude). The second pass re-coded only the "need-training" and "need-policy" comments into three barrier categories that map onto the implementation-gap analysis in §6.2: *resource* barriers (lack of training budget, lack of paid-tier AI access, lack of IT infrastructure), *cultural* barriers (faculty resistance, disciplinary norms, perceived threat to traditional pedagogy), and *governance* barriers (absence of policy, unclear delegated authority, absence of disclosure conventions). The barrier categorisation is itself an analytic construct and is reported with examples in §5.5.

4.5 Design limitations

Five limitations bear on the interpretation of all results. First, both samples are convenience samples through the OMU membership and dissemination channels; they are not probability samples of Ukrainian faculty. Second, the disciplinary skew towards marketing, management, and adjacent economic disciplines limits generalisation to STEM, humanities, and medical-education faculties. Third, both surveys were conducted in Ukrainian, and we translate direct quotations into English while preserving the Ukrainian original in footnotes for the most analytically loaded passages; translation choices are themselves an interpretive act and we flag in Appendix A those terms (notably for hallucination, academic integrity, and AI literacy) where the gloss is non-trivial. Fourth, the sample frame is biased towards institutions that remained operationally capable of receiving and answering a survey under wartime conditions; institutions in occupied territories, those that have been physically destroyed, and those that have ceased operations are absent. Fifth, the inter-wave sample-size growth (n=65 to n=82) and the modest expansion of recruitment channels limit cross-wave panel-style inference; we treat each wave as an independent cross-section.

5 Results

5.1 Adoption trajectory: 2023 to 2024

Faculty reports of LLM impact on their own work shifted between the two waves in two distinct directions. In the 2023 survey, respondents were asked to *forecast* the impact of AI on faculty work in the 2023–24 academic year; in the 2024 survey, the same respondents were asked to report *actual* impact for 2023–24 and to forecast impact for 2024–25. The directly comparable measurement is therefore 2023-forecast vs 2024-actual for the same academic year, both shown in Table 2.

Figure 2 shows the same shift graphically.

Two things moved. Faculty became more certain of their assessment ("hard to say" fell from 18.46% to 4.90%) and concentrated towards the *moderate* bucket, while *strong* impact halved (32.31% → 15.90%). The natural reading is calibration: in mid-2023

Table 2 LLM impact on faculty work in the 2023–24 academic year: 2023 forecast vs 2024 actual (% of respondents).

Impact level	2023 forecast	2024 actual	Δ (pp)
Absent / will not affect	4.62	17.10	+12.48
Moderate	44.62	62.20	+17.58
Strong	32.31	15.90	−16.41
Hard to say	18.46	4.90	−13.56

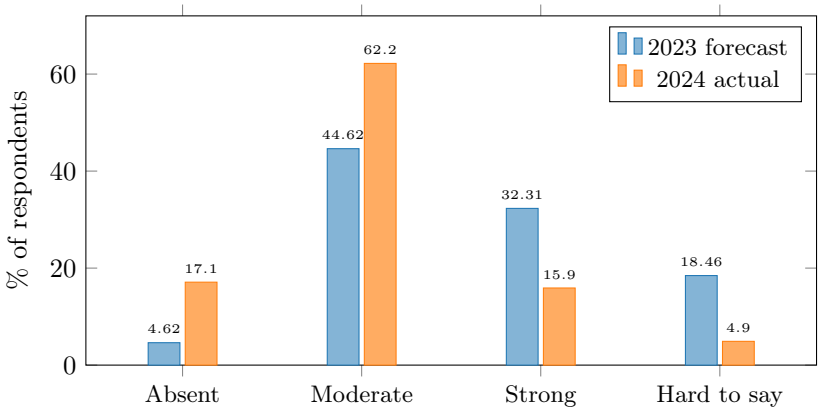


Fig. 2 Distribution of self-reported LLM impact on faculty work in the 2023–24 academic year: 2023 forecast (blue) vs 2024 actual (orange). Faculty became more certain of their assessment and clustered on *moderate*; the share expecting *strong* impact halved.

a non-trivial cohort expected dramatic transformation; a year later, with operational experience of the tools, the same population revised those expectations downwards and clustered on a more measured assessment. The growth of the *absent* category from 4.62% to 17.10% is consistent with a subset of faculty disengaging from LLM use entirely after early exposure.

The 2024 forecast for 2024–25, by contrast, anticipates a sharp swing back towards *strong* impact: 30.5% expect strong impact, double the 15.9% reporting it for the year just completed. Faculty appear to have stabilised their view of the present while retaining expectations that the trajectory will steepen.

5.2 Disciplinary divergence

The convenience-sample design dominated by marketing-and-management faculty limits inferences about disciplinary variation. Three observations are nevertheless interpretable.

First, the disciplinary composition of the two waves is materially stable. In 2023 (Starostina et al. 2, Table 5) the teaching discipline shares were Marketing 49.2%, Management 23.1%, Economics 18.5%, Finance 10.8%, Entrepreneurship 9.2%, and Other 30.8% (multi-select, sums exceed 100%). The 2024 wave (Starostina et al. 3,

Table 4) reports Marketing 48%, Management 20%, Economics 20%, Finance 7%, Entrepreneurship 2%, Other economic disciplines 42%, and Non-economic disciplines 12%. The two distributions overlap at all the dominant categories; the modest expansion of the "Non-economic" segment in 2024 (which did not exist as a distinct response option in 2023) reflects the dissemination through Facebook groups added in the 2024 wave rather than a substantive disciplinary shift. The sample frame is therefore comparable across waves, and longitudinal inferences in §5.1 and §5.3 are not confounded by a disciplinary-composition drift.

Second, within the 2024 wave the AI-use rates differ across functional rather than disciplinary lines (Starostina et al. 3, Table 11). Faculty report current use for translation of their own research at 39.0%, for preparing teaching-methodological materials at 32.9%, for conducting classes at 29.3%, and for writing research papers at 26.8%. The translation rate exceeds all three other use cases by more than 6 percentage points, and the gap widens to 16.6 pp in the 2024-to-25 forecast (translation 55.6%; teaching-methodological 43.9%; conducting classes 41.7%; research papers 31.9%; Starostina et al. 3, Table 12). The natural reading is that English-language publication pressure for an economics-and-management-skewed sample drives the most acute AI dependency on translation; this pressure operates at the level of the academic-publishing system rather than the home discipline, and would predict an even stronger translation skew in STEM faculties that publish more aggressively in international journals.

Third, the ISO 31000 awareness gradient between 2023 and 2024 is the only item where we have a measurement consistent with disciplinary heterogeneity (Starostina et al. 2 reports 29.23% personal awareness and 13.85% institutional awareness; Starostina et al. 3, Table 9 reports 12.20% institutional adoption). The personal-awareness rate roughly doubled the institutional-adoption rate in 2023, and the 2024 institutional-adoption figure is essentially flat against 2023's institutional-awareness number. This pattern is consistent with a model in which the more compliance-oriented sub-disciplines (finance, accounting, governance-adjacent management) raise individual awareness of risk-management standards faster than their institutions can operationalise them, and is inconsistent with a model of generalised institutional progress on risk-management infrastructure during the inter-wave year.

The 2024 instrument does not stratify the risk-management items by discipline, which precludes per-discipline fidelity scoring; we flag this as an instrument-design gap in §8 and recommend stratified reporting in any successor wave.

5.3 Risk perception versus risk behaviour

The most striking pattern in the 2024 data is the gap between what faculty say should be done about AI risk and what their institutions actually do. Table 3 aligns the perceived importance of each risk-management measure (Table 8 of the 2024 report) with the institutional availability of that measure at the same institution (Table 9). The gap is the percentage-point difference between the two.

Figure 3 renders the same data as paired horizontal bars; the visual asymmetry is the structural feature of the 2024 baseline.

Two features are noteworthy. First, every measure except outright banning shows a substantial *negative* gap: institutions consistently have less of the measure than

Table 3 Risk-perception vs. risk-behaviour gap: per-measure perceived importance (would help manage AI risk) and per-measure institutional availability (has it implemented), 2024 survey ($n = 82$). Wilson 95% CIs for the availability proportions are reported alongside the point estimates.

Measure	Importance (%)	Avail. (%)	95% CI	Gap (pp)
Train students and faculty	73.20	23.20	[15.4, 33.4]	−50.00
Clear AI policy	73.00	29.30	[20.5, 39.9]	−43.70
Study other HEIs	52.40	37.80	[28.1, 48.6]	−14.60
Computer-technological control	45.10	25.60	[17.4, 36.0]	−19.50
Regular audit	30.50	6.10	[2.6, 13.5]	−24.40
Adopt ISO 31000	24.40	12.20	[6.8, 21.0]	−12.20
Ban AI use	15.90	17.10	[10.4, 26.6]	+1.20

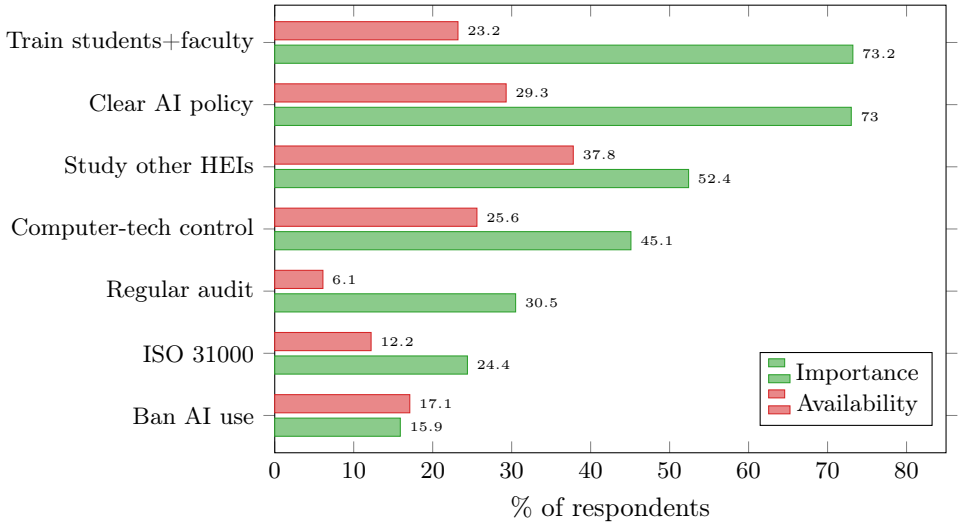


Fig. 3 Risk-perception vs. risk-behaviour. Green bars: % of respondents who rate the measure among the most important for managing AI-related risk. Red bars: % whose institution actually has the measure in place. The gap is the magnitude of the implementation deficit per measure.

faculty think they should have. The two largest gaps are for the two most pedagogically meaningful measures: training (50.00 pp) and policy (43.70 pp). Second, "ban AI use" is the only item with a positive gap: institutions have slightly more bans in place than faculty think are useful, and the institutional-availability rate (17.10%) is higher than the perceived-effectiveness rating (mean 2.53 on a 5-point scale, the lowest of all items). This is consistent with bans being adopted as a default, low-effort posture by some institutions rather than as a deliberately reasoned strategy.

The risk-perception/risk-behaviour gap is the structural feature of the 2024 baseline. It explains why the practice–policy fidelity scores in §5.4 are uniformly low even though faculty themselves identify and rank-order the right priorities.

5.4 Per-dimension fidelity scores

Applying the matrix construction described in §4.3, we score the 2024 sample on each of the eight policy dimensions. Table 4 reports the result. Two dimensions (diversity, non-discrimination, fairness; societal and environmental well-being) have no proxy in the 2024 instrument; we mark them as explicit data gaps rather than as zero scores, since absence of evidence is not evidence of absence.

Table 4 Practice–policy fidelity scores per dimension, 2024 sample. "Per-item Yes" is the percentage of respondents whose institution has the mapped measure; the aggregated fidelity is the unweighted mean of the per-item rates.

#	**Policy dimension (April-2025)**	**Per-item Yes (%)**	**Fidelity**
1	Human oversight	29.3, 23.2	0.263
2	Technical robustness and safety	25.6, 12.2	0.189
3	Privacy and data governance	29.3	0.293
4	Transparency	29.3	0.293
5	Diversity, non-discrimination, fairness	—	*gap*
6	Societal and environmental well-being	—	*gap*
7	Accountability	6.1, 12.2	**0.092**
8	Academic integrity (formal coverage)	29.3, 17.1	0.232
	Mean across scorable dimensions		0.227

Figure 4 visualises the same scores sorted from lowest to highest, with the two data-gap dimensions marked at the right.

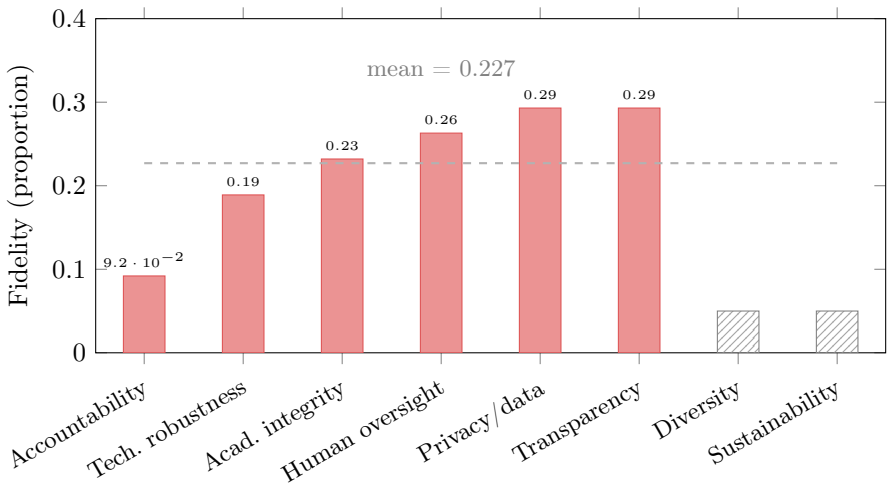


Fig. 4 Per-dimension fidelity scores for the 2024 cohort. The six scored dimensions are sorted by score (lowest left); the two data-gap dimensions (diversity, non-discrimination, fairness; societal and environmental well-being) are shown with hatched fill at a token height to indicate absence of measurement rather than absence of practice. Dashed line: mean across the six scored dimensions.

The headline finding is the low ceiling and the lower floor. The highest-fidelity dimensions (privacy and transparency, both 0.293) reach barely 30% institutional coverage; the lowest-fidelity dimension (accountability, 0.092) reaches barely 10%. The accountability finding is particularly consequential because the policy document (p. 14) explicitly classifies three categories of educational AI use as high-risk under the EU AI Act: admissions decisions, learning-outcome evaluation used to manage the educational process, and behaviour monitoring during testing. All three require formal accountability infrastructure that fewer than one institution in eight in our sample currently has.

A sensitivity analysis treating "Don't know" responses as fidelity-positive lifts the eight-dimension mean to 0.451 and accountability specifically to 0.272. We report the conservative scoring as the primary result on the grounds that an institution whose faculty cannot verify whether a policy exists has, operationally, no policy.

The two authors independently mapped each of the seven Table 9 items to the eight policy dimensions before consolidating the matrix in Table 4. Of $7 \times 8 = 56$ (item, dimension) pairs, the two raters agreed on 51 (9 positive, 42 negative), yielding a pre-resolution Cohen's $\kappa = 0.728$ (simple agreement 0.911) and falling in the "substantial" range of Landis and Koch [31]. The five disagreements concentrated on genuinely ambiguous principle-item boundaries: whether *audit* is human oversight or accountability infrastructure; whether ISO 31000 belongs in technical robustness or only in accountability; whether *computer-technological control* is a data-governance mechanism in addition to a technical one; whether *benchmarking against other HEIs* constitutes accountability; and whether *training* is an academic-integrity proxy in addition to a literacy proxy. Post-resolution, four of the five disagreements went to the more conservative mapping (no double-counting), and one (the training-as-integrity-proxy item) was added to dimension 8. Because the added item's Yes rate (23.2%) is essentially at the mean of the other two items mapped to dimension 8, the post-resolution dimension 8 score is numerically identical to the pre-resolution score (0.232), and the mean across the six scorable dimensions is unchanged at 0.227. The dimension ranking in Table 4 is robust to the mapping disagreements.

5.5 Open-ended themes by barrier category

The seven themes used by the OMU authors to code the open-ended responses (*tool*, 26.3%; *need training*, 21.1%; *positive attitude*, 21.1%; *need policy*, 10.5%; *neutral attitude*, 10.5%; *need further research*, 5.3%; *negative attitude*, 5.3%) are not by themselves a barrier taxonomy. Re-coding the *need training* and *need policy* themes (31.6% of all comments combined) into the resource/cultural/governance categories defined in §4.4 yields the following:

Resource barriers dominate the *need training* subset and include the absence of institutional training budgets (faculty in this sample report having learned to use AI tools self-taught at 80.3% and through formal institution-paid training at only 5.3%; Starostina et al. 3 Fig. 20). *Cultural barriers* appear in the *negative attitude* (5.3%) and *neutral attitude* (10.5%) themes – one respondent comments that "in the humanities, positive use of AI as an assistant is overestimated; it is nonsense and not at all a help" (Starostina et al. 3, p. 35, Respondent 3, translated from Ukrainian) –

and in the disciplinary anchoring of the sample, where marketing faculty's openness may not generalise. *Governance barriers* are most explicit in the *need policy* subset: respondents call for clear ethical-use frameworks, disclosure conventions, and institutional capacity for academic-integrity adjudication under AI-assisted submissions. Two respondents (Starostina et al. 3, p. 35, Respondents 1 and 4) describe AI as an inevitable tool whose ethical use must be taught and supervised rather than restricted – consistent with the 73.20% of all respondents who ranked training as the most important risk-management measure in Table 3.

The barrier distribution is consistent with the implementation pathway proposed in §7: governance reforms unlock policy clarity; resource reforms unlock training capacity; cultural reforms are the slowest but are partly downstream of the first two.

6 Discussion

6.1 Ukrainian patterns in international comparison

Table 5 situates the 2024 Ukrainian sample alongside six contemporary faculty-and-student studies of GenAI adoption in higher education, selected for methodological proximity and 2024–26 publication. Direct head-to-head numerical comparisons are constrained because most international studies sample students rather than faculty, and because effect-size reporting conventions vary across the TAM/UTAUT/moderated-mediation frameworks used. Where a directly comparable adoption-rate or training-availability statistic exists in the cited work's published portion, we report it; otherwise we anchor the comparison on the study's methodological emphasis and on the dimension of our matrix that the study most directly informs.

Three patterns from the verified international comparators align with the Ukrainian baseline. First, the actual-use rates we observe for Ukrainian faculty – 62.2% report moderate LLM impact and 15.9% report strong impact for 2023–24, with the remaining 22% reporting no or unclear impact – bracket the 62% "have used a GAI tool" headline reported by McCoy et al. [9] for US medical-education faculty ($n = 68$) almost exactly, suggesting that the Ukrainian rate is in the same neighbourhood as a contemporaneous OECD comparator and is not anomalous despite the wartime context. Second, the role of social influence in shaping behavioural intention varies sharply by country: the largest extant Turkish study (Ursavaş et al. 6, $n = 943$ across 140 universities) finds subjective norms have a *limited* effect on undergraduate BI, while the Thai faculty study (Liang and Suwanragsa 8, $n = 500$) finds social influence is the *largest* total-effect predictor of faculty BI ($\beta_{\text{total}} = 0.275$, $p < .001$), exceeding even performance expectancy ($\beta = 0.177$) and perceived enjoyment ($\beta = 0.246$); the Indian moderated-mediation evidence [7] sits between these two poles with social influence operating through attitude. The Ukrainian importance ranking of training-as-priority (73.20%) is best interpreted alongside this distribution: training-demand is high regardless of how strongly social influence operates in the local culture, and a national training programme is therefore robust to the cross-cultural variation in SI potency. Third, trust and motivation emerge as load-bearing constructs in two structurally different student samples. In the Saudi sample [4], motivation is the strongest direct predictor of student behavioural intention ($\beta = 0.334$, $p < .001$), followed by

Table 5 Methodological-proximity comparison of the 2024 Ukrainian baseline with six contemporary international studies of generative-AI adoption in higher education. Sample sizes and headline numerical findings are reproduced verbatim from each cited work's abstract and published Results section as available at the time of writing; where a number is not reported in the publicly accessible portion of the cited paper, the cell is left empty rather than back-inferred.

Study	Country	Sample (n; design)	Method	Verified headline finding
Ukraine (this paper)	Ukraine	Faculty $n = 82$, 35 HEIs; longitudinal vs 2023 $n = 65$, 30 HEIs	Descriptive longitudinal; fidelity matrix vs April-2025 framework	80.3% self-taught; mean fidelity 0.227; violation rate 40.0% (2023) → 52.4% (2024)
Bamasoud et al. [4]	Saudi Arabia	Lecturers $n = 309$, students $n = 578$ at Najran, Jazan, Bisha universities; data collected 1 May–30 June 2025	Extended TAM with SDT (SmartPLS 4 SEM) on student adoption	On students ($n = 578$, $R^2_{\text{Intention}} = 0.348$): β(EOU → PU) = 0.555 ($t = 14.84$, $p < .001$); β(Mov → Intention) = 0.334 (largest direct effect); β(PU → Intention) = 0.227; β(T-R → Mov) = 0.314 (all $p < .001$)
Dorobăţ and Corbea (Florea) [5]	Romania	Students $n = 487$ valid from 512 responses to 1217 distributed (42.07% RR); analytic $n = 477$ used ChatGPT (97.9%) at Bucharest University of Economic Studies	TAM-SEM on hedonic motivation, social presence, ease of use, usefulness → satisfaction → trust → loyalty (AMOS v24)	β(PEOU → S) = 0.558; β(PU → S) = 0.458; β(S → T) = 0.816 (largest path); total β(PU → L) = 0.568. All eight hypotheses supported
Ursavaş et al. [6]	Türkiye	Undergraduate students $n = 943$ across 140 universities and 147 degree programmes	Extended TAM (SEM) with subjective norms, perceived enjoyment, self-efficacy, compatibility	Attitude, perceived usefulness, self-efficacy, and compatibility significantly predict BI; *subjective norms have limited influence*
Kaur et al. [7]	India	International-business students $n = 526$ (from 750 distributed)	Moderated mediation through attitude	Performance expectancy, effort expectancy, social influence, hedonic motivation, habit, and personal innovativeness predict BI via attitude
Liang and Suwanragsa [8]	Thailand	Faculty $n = 500$ (snowball, 98.8% retention from 506; STEM 35.6%, social sciences 22.0%, applied 19.8%, humanities 12.4%, arts 10.2%)	Expanded UTAUT adding perceived risk and perceived enjoyment (AMOS 22.0 SEM)	All six direct paths to BI significant; total-effect ranking: β(SI) = 0.275 (largest), β(PENJ) = 0.246, β(PE) = 0.177, β(FC) = 0.166, β(EE) = 0.153, β(PR) = −0.106 (negative)
McCoy et al. [9]	USA	Faculty $n = 68$ (50% response from 135), students $n = 506$ (30% from 1696) at NYITCOM	Training-needs analysis; cross-sectional	62% faculty / 59% students have used a GAI tool; 60% faculty / 77% students never or rarely; 61% faculty prioritise hands-on training; 48% want integrity guidelines

perceived usefulness ($\beta = 0.227$, $p < .001$); and motivation is itself driven by a trust-and-responsibility construct (β(T-R $\rightarrow$ Mov) $= 0.314$, $p < .001$) and by personal academic-integrity orientation (β(Integrity $\rightarrow$ Mov) $= 0.193$, $p < .001$). In the Romanian sample [5], satisfaction drives trust at the largest standardised coefficient in the model (β(S $\rightarrow$ T) $= 0.816$, $p < .01$), and perceived usefulness drives loyalty at a total effect of $\beta = 0.568$ ($p < .01$). The structural convergence is that trust and motivation – not raw ease of use or social pressure – carry the adoption signal in both samples, reinforcing the case that fidelity dimensions 3 (privacy and data governance) and 4 (transparency), which together operationalise the institutional preconditions for trust, are central rather than incidental to adoption.

Two divergences warrant attention and constitute the principal contribution of placing the Ukrainian baseline in international context. First, the proportion of self-taught faculty in our sample (80.3%; Starostina et al. 3, Fig. 20) means that for every five Ukrainian faculty who use GenAI, at most one reached the tool through any form of institution-paid or government-paid training (the institution-paid share is 5.3%; the free-formal-training share, 2.6%). The contrast with McCoy et al. [9]'s finding that 61% of US medical-education faculty prioritise institution-organised *hands-on* training is sharp: in the US setting, the demand for training is overwhelmingly directed at institutions, whereas in Ukraine demand and self-supply largely coincide because the institutional supply channel is missing. The Ukrainian wartime context disrupted the budget cycle and procurement pathways through which such offerings are normally delivered. The structural consequence is that fidelity dimension 1 (human oversight, including faculty literacy) cannot be addressed in Ukraine by the same standardised-curriculum playbook used in OECD-average settings; ministerial recommendations must explicitly resource an emergency literacy programme.

Second, the institutional-ban rate of 17.1% in our 2024 sample, set against the very low perceived effectiveness of bans (mean 2.53/5; Table 10 of Starostina et al. 3), differs structurally from the predominant Western-European and US pattern, where Perkins et al.'s [18] AIAS five-level scale and Sullivan's [32] disclosure-attribution frameworks have largely supplanted outright prohibition in published policies; the Thai faculty work [8] reports a similar regulatory caution to ours, consistent with a model in which wartime or perceived-instability conditions amplify the policymaker's preference for ban-based default postures even where evidence of effectiveness is weak. The April-2025 Ukrainian framework [1] addresses this divergence directly: its Annex 6 endorses the Perkins AIAS five-level scale (Option 2 on p. 45) as the recommended classroom-policy template, and its principles section reproduces the EU AI HLEG 2019 seven principles verbatim. The policy framework is therefore internationally aligned; the implementation gap is in operationalising it.

Two further studies extend the comparison along orthogonal axes. Shata and Hartley [17] document faculty perceptions and adoption patterns in an English-medium-of-instruction setting and are a methodological precedent for our open-ended-comment coding scheme. The book-chapter and Indian-business-school evidence in Mansoor [33] suggests that the institutional-readiness archetypes we identify in §6.3 have analogues in adjacent settings, with implementation timelines varying by a factor of two or more across the archetypes. Both extensions converge on the same diagnostic: the principal

bottleneck for responsible-AI adoption in higher education is not faculty willingness; it is institutional infrastructure, and the institutional-infrastructure gap is more severe in Ukraine than in comparator countries because of the wartime resource environment.

6.2 Implementation-gap drivers

The 50 pp gap between training-as-importance (73.20%) and training-as-availability (23.20%), and the 43.70 pp gap on policy, are not random. They are concentrated in the resource and governance categories defined in §4.4, and they are reinforced by the institutional pressures the war creates. Resource barriers are amplified by reduced operating budgets and by the diversion of IT staff to disaster-recovery and continuity-of-operations work. Governance barriers are amplified by the absence, until April 2025, of any national policy framework that institutions could reference when drafting their own rules [1]; cautious institutions waited for ministerial signal, and the signal arrived eight months after our 2024 measurements. Cultural barriers, the third category, are slower-moving and partly downstream of the first two: institutions that have neither training nor policy cannot reasonably ask faculty or students to adopt new disciplinary norms.

The implementation literature on responsible AI in education [12, 13, 16] converges on a similar diagnosis. Faculty self-regulation is necessary but insufficient; institutional governance and training infrastructure are the load-bearing components, and they take longer to build than the cultural-attitude shifts that catch popular attention. The 2024 Ukrainian baseline confirms this pattern in a setting where the resource constraint is sharper than usual.

6.3 Institutional readiness scenarios

Clustering the institutions represented in the 2024 sample by their reported scores on the two most diagnostic dimensions – accountability infrastructure (proxied by ISO 31000 adoption and regular audit) and policy clarity (proxied by the presence of a written AI-use policy) – yields three archetypes, sketched in Figure 5.

Because the published OMU data set reports item rates at the aggregate level only and does not release institution-identified microdata, the per-archetype shares cannot be computed by direct cross-tabulation. We therefore estimate them by Monte Carlo simulation under a latent-maturity model in which each respondent's institution has a single underlying "risk-management maturity" score Z, and each Table 9 indicator is realised as $\mathbb{1}\{\sqrt{\rho}\,Z + \sqrt{1-\rho}\,\varepsilon_i > \Phi^{-1}(1-p_i)\}$, where p_i is the published Yes-rate and ρ is the latent inter-item correlation. We sweep $\rho \in \{0,\, 0.3,\, 0.6\}$ corresponding to independence, moderate positive clustering, and strong positive clustering of risk-management practices within institutions, and report 10,000 draws per scenario at the 2024 sample size of $n = 82$. Full code and outputs are in the supplementary directory (`simulate_archetypes.R`).

Table 6 reports the per-archetype shares across the three correlation scenarios. The qualitative ranking is invariant: policy-unprepared is the majority archetype under every assumption (median $\geq$ 58.5%), policy-aware is the second-largest (23–37%), and policy-ready is a small minority (5–11%). The clustering parameter ρ trades

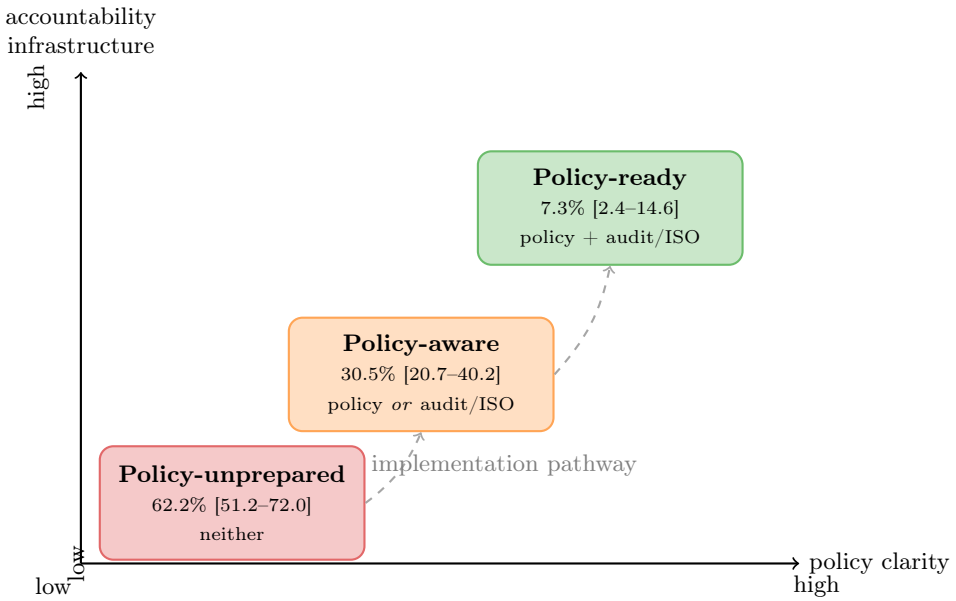


Fig. 5 Three institutional readiness archetypes for the April-2025 framework, plotted by accountability infrastructure (vertical) and policy clarity (horizontal). Percentages are simulation-derived medians with 95% credible intervals under the moderate-clustering scenario ($\rho = 0.3$) of the latent-maturity model described in §6.3; corresponding figures at $\rho = 0$ and $\rho = 0.6$ are reported in Table 6. The dashed arrow traces the sequenced transition envisioned in §7: unprepared institutions adopt policy first (Step 2), then accountability (Step 1), reaching ready status.

aware-share for ready-and-unprepared shares: institutions either cluster (some have everything, some have nothing) or distribute their adoption across items.

Table 6 Monte Carlo estimates of institutional readiness archetype shares under three latent-correlation assumptions. Medians with 95% credible intervals from 10,000 draws of $n = 82$ respondents each, with marginal Yes-rates fixed at the observed Table 9 values (policy 29.3%, audit 6.1%, ISO 31000 12.2%). Simulation script: `simulate_archetypes.R` in the supplementary materials.

Scenario	Policy-ready (%)	Policy-aware (%)	Policy-unprepared (%)
$\rho = 0.0$ (independence)	4.9 [1.2, 9.8]	36.6 [25.6, 47.6]	58.5 [47.6, 68.3]
$\rho = 0.3$ (moderate +ve)	7.3 [2.4, 14.6]	30.5 [20.7, 40.2]	62.2 [51.2, 72.0]
$\rho = 0.6$ (strong +ve)	11.0 [4.9, 18.3]	23.2 [14.6, 32.9]	65.9 [56.1, 75.6]

Policy-ready institutions (simulated median 7.3%, 95% CrI 2.4–14.6%) report both a policy and at least one accountability mechanism; they are positioned to absorb the April-2025 framework with marginal additional capacity. *Policy-aware* institutions (median 30.5%, 20.7–40.2%) report a policy but no accountability mechanism, or vice versa; they have made first moves but need the framework to specify the gap they have

not yet addressed. *Policy-unprepared* institutions (median 62.2%, 51.2–72.0%) report neither; they require both governance and resource support to implement the April-2025 recommendations at all, and they constitute the principal target population for the sequenced pathway proposed in §7. The Ministry of Education and Science should plan resource allocation against this distribution: the largest share of HEIs needs more support than the readiest minority, and the differential is at least an order of magnitude.

7 Recommendations

The April-2025 framework names eight dimensions; the 2024 baseline finds the lowest fidelity on accountability and the largest gaps between perceived importance and institutional availability on training and policy. We accordingly recommend a four-step sequence for implementation:

Step 1 – accountability infrastructure first. The single highest-leverage intervention is the rollout of ISO 31000-aligned risk-management protocols and regular AI-use audit cycles. These are the building blocks for compliance with the EU AI Act high-risk categories named on p. 14 of the framework (admissions decisions, learning-outcome evaluation used to manage the educational process, behaviour monitoring during testing), and they are the dimension on which the 2024 baseline is weakest (fidelity 0.092). We recommend that the Ministry of Education and Science of Ukraine require all HEIs to designate a responsible accountability officer (or equivalent role within an existing audit committee) within twelve months of the framework's publication date.

Step 2 – policy templates with disclosure conventions, not bans. The framework's Appendix 6 already provides a sample acceptable-use policy with a five-level use-permission scale; the Ministry should disseminate a national template that institutions can adapt with minimal effort. Bans (currently 17.1% institutional availability, mean perceived effectiveness 2.53/5.0) should be explicitly de-emphasised as a default posture. The AI Assessment Scale framework [18] and the legal-perspective treatment of plagiarism, originality, and liability [32, 33] are the most directly usable reference materials for the disclosure-and-attribution architecture our 2024 respondents are calling for.

Step 3 – training capacity at the disciplinary-association level. The 80.3% self-taught rate in our 2024 sample, combined with the disciplinary-association origin of the surveys themselves, suggests that professional associations rather than ministries are the most effective vehicle for short-cycle AI-literacy programmes. The Marketers Association of Ukraine model – annual surveys, member-targeted dissemination, public reporting – should be supported and extended to other disciplinary associations, particularly STEM and humanities groupings under-represented in the current sample.

Step 4 – close the two data gaps. The 2025 wave of the survey should add items on equity-of-access (dimension 5) and on energy/environmental considerations in AI procurement (dimension 6). Without these items, the fidelity matrix cannot fully score institutional readiness on the framework's full eight dimensions. We recommend

the OMU authors, in collaboration with the Institute for Digitalisation of Education of the National Academy of Educational Sciences, draft and pilot the additional items before the 2025 wave.

A sequenced rather than simultaneous rollout is preferable for two reasons. First, the policy-unprepared archetype identified in §6.3 cannot absorb all eight dimensions at once; sequencing creates an absorbable workload. Second, the accountability-first ordering aligns Ukrainian implementation with the regulatory architecture of the EU AI Act [11] and positions Ukrainian HEIs for the kind of cross-border quality-assurance arrangements that the European Higher Education Area increasingly expects.

8 Limitations

Five limitations qualify our findings beyond the design-level caveats already listed in §4.5. First, the eight-month gap between the 2024 survey wave and the April-2025 framework precludes any claim of *compliance* measurement; we are reporting baseline fidelity to a yet-unpublished policy, which is a different epistemic object. Second, the convenience-sample design through the OMU membership channel means our results generalise most safely to economic-disciplines faculty at the 35 institutions in the 2024 frame; extension to the full Ukrainian HEI population requires a successor probability-sample study. Third, the fidelity-matrix scoring depends on the survey items asking *what is in place* rather than *whether what is in place actually operates*; an institution may have a written AI policy that is not enforced, and our scoring would still credit it as fidelity-positive. Fourth, the two data gaps on dimensions 5 (diversity, non-discrimination, fairness) and 6 (societal and environmental well-being) are present-tense limitations that successor instrument design can remediate but that we cannot remediate retrospectively. Fifth, the wartime context affects every measurement in ways we cannot fully separate from baseline patterns; cross-national comparison in §6.1 mitigates but does not eliminate this confound.

9 Conclusion

This paper has used two consecutive annual surveys of Ukrainian higher-education faculty – conducted by the NGO Marketers Association of Ukraine in June 2023 (n=65, 30 institutions) and June–August 2024 (n=82, 35 institutions) – to measure baseline fidelity against the eight dimensions of responsible AI subsequently articulated in the April-2025 national *Recommendations on the responsible implementation and use of artificial intelligence technologies in higher-education institutions*. The temporal ordering – surveys preceding policy – is methodologically unusual; it permits a measurement of the implementation gap that compliance studies, conducted after a policy is published, cannot recover.

Three findings warrant emphasis. First, baseline fidelity is uniformly low: a mean of 0.227 across the six dimensions for which the 2024 instrument supplies a proxy, with the steepest gap on accountability (0.092) – the dimension the EU AI Act identifies as highest-risk for educational use. Second, the gap between what faculty think their institutions should do and what their institutions actually do is large and systematic: training and policy both show importance-availability gaps of 44 percentage points or

more, and the only measure where institutions exceed faculty's stated importance ratings is outright banning, which faculty themselves rate as the least effective response. Third, witnessed academic-integrity violations rose from 40.0% in 2023 (Wilson 95% CI $[29.0, 52.1]$, $n = 65$) to 52.4% in 2024 (Wilson 95% CI $[41.8, 62.9]$, $n = 82$) against essentially flat policy coverage; the two intervals overlap and we therefore cannot reject the null of equal underlying rates at conventional levels, but the point estimate moved 12.4 percentage points in twelve months while institutional policy coverage stayed below 30%, and the burden of proof has shifted onto whoever wishes to argue that no upward trend exists. The April-2025 framework will need to be implemented with urgency if it is to reverse the trajectory.

The methodological contribution of the paper is the practice–policy fidelity matrix, a reusable construct for benchmarking national implementation of any sectoral AI framework that is articulated as a finite set of principles. The construct travels: it can be applied to any country whose government has issued a responsible-AI framework for higher education and whose disciplinary associations conduct comparable surveys. The principal limitation – the two dimensions for which the 2024 Ukrainian instrument has no proxy – is itself a useful contribution, as it identifies the minimum instrument extension required for a complete fidelity audit in any successor wave.

The April-2025 framework gives Ukrainian higher education the policy infrastructure the 2024 baseline shows it needs. The work of implementation is just beginning, and the fidelity matrix offered here is a tool for tracking progress against the framework's own internal benchmarks.

A Original Ukrainian wording of the eight responsible-AI principles and analytically loaded glosses

Table 7 reproduces the seven trustworthy-AI principles (verbatim from the European Commission's High-Level Expert Group on Artificial Intelligence, as adopted in the April-2025 Ukrainian framework, pp. 11–12) and the eighth (academic-integrity) principle from the same framework's dedicated section on pp. 20–23. The Ukrainian wording is the form in which respondents to the 2023 and 2024 surveys encountered the underlying concepts; the English glosses are the ones used throughout this paper.

Three terms used elsewhere in the paper carry meanings that the standard English glosses do not fully recover:

Галюцинування – hallucination. The phenomenon where a generative-AI system's output contains inaccurate or fabricated information presented misleadingly as reliable. The Ukrainian term is a direct calque of the English. The April-2025 framework treats recognition and mitigation of hallucination as a baseline AI-literacy skill (p. 4).
Академічна доброчесність – academic integrity. In Ukrainian higher-education law and discourse, this term is more expansive than the English "academic integrity": it encompasses individual ethical conduct in scholarly work *and* institutional commitments to honesty in teaching, assessment, research dissemination, and credential

Table 7 The eight dimensions of responsible AI in Ukrainian higher education: original Ukrainian wording and English glosses.

#	Ukrainian (original)	English (this paper)
1	Контроль з боку людини	Human oversight
2	Технічна надійність і безпека	Technical robustness and safety
3	Конфіденційність й управління даними	Privacy and data governance
4	Прозорість	Transparency
5	Різноманітність, недискримінація і справедливість	Diversity, non-discrimination, fairness
6	Суспільний та екологічний добробут	Societal and environmental well-being
7	Підзвітність	Accountability
8	Академічна доброчесність	Academic integrity

recognition. The April-2025 framework develops a four-stage model of user–AI interaction (Stages 1–4 on pp. 20–23) within which violations can occur, mapping moments of choice, task formulation, intermediate product, and final product.
Грамотність у галузі ШІ – AI literacy. The framework defines this as the knowledge, skills, and abilities that help providers, developers, users, and persons affected by AI systems – with regard to their respective rights and duties – to assess the capabilities, risks, and potential harm an AI system can cause at any stage of its lifecycle, and to take measures to avoid or mitigate that harm (p. 7).

Acknowledgments

The authors thank Alla O. Starostina, Nataliia M. Chuprina, Volodymyr A. Kravchenko, Oleksandr V. Kravchenko, and Oleksandr S. Suvorov of the NGO Marketers Association of Ukraine, who designed and administered the underlying 2023 and 2024 faculty surveys whose secondary analysis this paper rests upon. We thank the members of the working group that produced the April-2025 national framework and the academic colleagues at the Institute for Digitalisation of Education of the National Academy of Educational Sciences of Ukraine, the Kharkiv IT Cluster, the Centre for Innovative Educational Technologies of Lviv Polytechnic, and the Kyiv School of Economics, whose endorsement of the framework and parallel analytic work made our cross-referencing analysis possible.

Author contributions

Serhiy O. Semerikov: Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing, Supervision. Maiia V. Marienko: Investigation, Data curation, Methodology, Writing – review & editing. Oksana M. Markova: Data curation, Writing – review & editing. All authors read and approved the final manuscript.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The underlying primary surveys were funded entirely from the resources of the NGO Marketers Association of Ukraine. The April-2025 national framework was developed with support from the UK DIGIT project, implemented by the Eurasia Foundation and funded through the United Kingdom's International Development programme.

Data availability

This paper is a secondary analysis of three publicly available institutional documents, each cited in the bibliography: the 2023 OMU survey report [2], the 2024 OMU survey report [3], and the April-2025 national framework [1]. The 2023 and 2024 reports include aggregate respondent-level frequencies for every survey item and are available at https://omu.marketing/. Individual-level responses were not deposited by the original surveyors and are not redistributed by this paper. The aggregated practice–policy fidelity scores computed in §5.4 are reproducible from the published aggregate tables; the mapping rules used to compute them are documented in §4.3 and reproduced in the supplementary materials accompanying this submission. The supplementary package contains: (i) the tabulated values extracted from both surveys (`survey_2023_summary.csv`, `survey_2024_summary.csv`); (ii) the item-to-dimension mapping and per-item Yes-rates used to compute the fidelity matrix (`fidelity_inputs.csv`); (iii) the two-pass coding scheme for open-ended responses (`coding_scheme.md`); (iv) a self-contained R replication script (`replication.R`) that reproduces every Wilson 95% confidence interval, fidelity score, inter-rater Cohen's κ, and longitudinal delta reported in the manuscript using base R only; and (v) a Monte Carlo simulation script (`simulate_archetypes.R`) that reproduces the per-archetype readiness shares in Figure 5 and Table 6 under three latent-inter-item-correlation scenarios.

Declarations

Ethics approval and consent to participate

Not applicable. This article is a secondary analysis of aggregate data from previously published institutional survey reports and a publicly issued national policy framework. The authors did not collect new data from human participants. The original surveys were administered with informed consent by the NGO Marketers Association of Ukraine; the authors of the present paper had no role in primary data collection.

Consent for publication

Not applicable.

Clinical trial number

Not applicable.

Competing interests

The authors declare that they have no competing interests.

References

[1] Ministry of Digital Transformation of Ukraine and Ministry of Education and Science of Ukraine and UK International Development and Eurasia Foundation. Recommendations on the responsible implementation and use of artificial intelligence technologies in higher-education institutions. Kyiv: Joint publication of the Ministry of Digital Transformation of Ukraine, Ministry of Education and Science of Ukraine, UK International Development, and Eurasia Foundation; 2025. Available from: https://storage.thedigital.gov.ua/files/5/b5/1de422b38985d037d9ba8f9f6cb2ab58.pdf.

[2] Starostina AO, Domina OM, Kravchenko VA. Report on the results of the survey on risks of using ChatGPT-type artificial intelligence in universities of Ukraine (June 2023). Kyiv: NGO "Marketers Association of Ukraine"; 2023. Available from: http://omu.marketing/wp-content/uploads/2023/07/shi_opytuvannya_rezultaty_280723.pdf.

[3] Starostina AO, Chuprina NM, Kravchenko VA, Kravchenko OV, Suvorov OS. Final report on the results of the 2nd annual survey on the use of LLM-based artificial intelligence in Ukrainian universities (June–August 2024). Kyiv: NGO "Marketers Association of Ukraine"; 2024. Available from: https://omu.marketing/wp-content/uploads/2024/07/shi_2024_opytuvannya_rezultaty_final.pdf.

[4] Bamasoud DM, Mohammad R, Bilal S. Adopting Generative AI in Higher Education: A Dual-Perspective Study of Students and Lecturers in Saudi Universities. Big Data and Cognitive Computing. 2025;9(10):264. https://doi.org/10.3390/bdcc9100264.

[5] Dorobăţ I, Corbea (Florea) AMI. Assessing ChatGPT Adoption in Higher Education: An Empirical Analysis. Electronics. 2025;14(23):4739. https://doi.org/10.3390/electronics14234739.

[6] Ursavaş ÖF, Yalçın Y, İslamoğlu H, Bakır-Yalçın E, Cukurova M. Rethinking the importance of social norms in generative AI adoption: investigating the acceptance and use of generative AI among higher education students. International Journal of Educational Technology in Higher Education. 2025 Jun;22(1):38. https://doi.org/10.1186/s41239-025-00535-z.

[7] Kaur H, Verma T, Mago M, Phoolka S. ChatGPT in the classroom: navigating the generative AI wave among international business students – a moderated mediation approach. Journal of International Education in Business. 2025 11;p. 1–29. https://doi.org/10.1108/JIEB-08-2024-0097.

[8] Liang D, Suwanragsa I. Determinants of Generative AI Adoption in Higher Education: A Social Science Perspective on Thai Faculty Behavioral Intentions. European Journal of Social Science Education and Research. 2025 Dec;12(4):132–147. https://doi.org/10.26417/bamx9m98.

[9] McCoy L, Ganesan N, Rajagopalan V, McKell D, Niño DF, Swaim MC. A Training Needs Analysis for AI and Generative AI in Medical Education: Perspectives of Faculty and Students. Journal of Medical Education and Curricular Development. 2025;12:23821205251339226. https://doi.org/10.1177/23821205251339226.

[10] Miao F, Holmes W. Guidance for generative AI in education and research. Paris: United Nations Educational, Scientific and Cultural Organization; 2023. Available from: https://doi.org/10.54675/ewzm9535.

[11] European Parliament, Council of the European Union.: Consolidated text: Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance). Available from: http://data.europa.eu/eli/reg/2024/1689/2026-07-27.

[12] Chauncey SA, McKenna HP. A framework and exemplars for ethical and responsible use of AI Chatbot technology to support teaching and learning. Computers and Education: Artificial Intelligence. 2023;5:100182. https://doi.org/10.1016/j.caeai.2023.100182.

[13] Fu Y, Weng Z. Navigating the ethical terrain of AI in education: a systematic review on framing responsible human-centered AI practices. Computers and Education: Artificial Intelligence. 2024;7:100306. https://doi.org/10.1016/j.caeai.2024.100306.

[14] Shrestha Y. Responsible Governance of Generative AI in Higher Education: Integrating Ethics, Policy, and the REM Framework. Apex Journal of Business and Management. 2026;5(1):103–113. https://doi.org/10.61274/apxc.2026.v05i01.009.

[15] Uddin MM, Moon J, Abu S. An ethical framework for conversational AI in higher education: toward an evidence-based ethical governance. AI and Ethics. 2026;6:202. https://doi.org/10.1007/s43681-026-01056-9.

[16] Owoseni A, Kolade O, Egbetokun A. Responsible Use of Generative AI for Educators and Students in Higher Education Institutions. In: Generative AI in Higher Education: Innovation Strategies for Teaching and Learning. Cham: Springer Nature Switzerland; 2024. p. 151–172.

[17] Shata A, Hartley K. Artificial intelligence and communication technologies in academia: faculty perceptions and the adoption of generative AI. International Journal of Educational Technology in Higher Education. 2025 Mar;22(1):14. https://doi.org/10.1186/s41239-025-00511-7.

[18] Perkins M, Furze L, Roe J, MacVaugh J. The Artificial Intelligence Assessment Scale (AIAS): A Framework for Ethical Integration of Generative AI in Educational Assessment. Journal of University Teaching and Learning Practice. 2024 Apr;21(06). https://doi.org/10.53761/q3azde36.

[19] Bulut O, Beiting-Parrish M, Casabianca JM, Slater SC, Jiao H, Song D, et al. The Rise of Artificial Intelligence in Educational Measurement: Opportunities and Ethical Challenges. Chinese/English Journal of Educational Measurement and Evaluation. 2024;5(3):3. https://doi.org/10.59863/miql7785.

[20] Hou I, Man O, Hamilton K, Muthusekaran S, Johnykutty J, Zadeh L, et al. 'All Roads Lead to ChatGPT': How Generative AI is Eroding Social Interactions and Student Learning Communities. In: Proceedings of the 30th ACM Conference on Innovation and Technology in Computer Science Education V. 1. ITiCSE 2025. New York, NY, USA: Association for Computing Machinery; 2025. p. 79–85. Available from: https://doi.org/10.1145/3724363.3729024.

[21] Lee AVY. 8: The double-edged sword of generative AI for assessment: Efficiency versus homogenization. In: Sabbaghan S, editor. Navigating Generative AI in Higher Education: Ethical, Theoretical and Practical Perspectives. Edward Elgar Publishing; 2025. p. 128–143. Available from: https://doi.org/10.4337/9781035337873.00012.

[22] Bovill H. Too much information: exploring technology-mediated abuse in higher education online learning and teaching spaces resulting from COVID-19 and emergency remote education. Higher Education. 2023;86:467–483. https://doi.org/10.1007/s10734-022-00947-0.

[23] Zahorodko PV, Semerikov SO. Integrating agile methodologies and AI-assisted learning in web programming education: a theoretical framework for CS curriculum transformation. Discover Education. 2026 Feb;5(1):166. https://doi.org/10.1007/s44217-026-01179-5.

[24] Semerikov SO, Vakaliuk TA, Mintii IS, Bondarenko OV, Kanevska OB. Mapping the Emergence: A Bibliometric Analysis of the Convergence Gap in GeoAI Teacher Education Research. SN Computer Science. 2026 Jan;7(1):114. https://doi.org/10.1007/s42979-026-04731-0.

[25] Marienko MV, Markova OM, Semerikov SO. AI literacy in secondary education: framework, assessment, and professional development in the Ukrainian context. Computers and Education: Artificial Intelligence. 2026;10:100605. https://doi.org/10.1016/j.caeai.2026.100605.

[26] Kremen VG, Spirin OM, Liashenko OI, Lytvynova SH, Malovanyi YI, Pinchuk OP, et al. AI adoption readiness among Ukrainian education managers: Barriers, typologies, and policy implications. Computers and Education: Artificial Intelligence. 2026;11:100648. https://doi.org/10.1016/j.caeai.2026.100648.

[27] Holmes W, Bialik M, Fadel C. Artificial intelligence in education: Promises and implications for teaching and learning. In: Data ethics : building trust : how digital technologies can serve humanity. Globethics Publications; 2019. p. 621–653. Available from: https://doi.org/10.58863/20.500.12424/4276068.

[28] Crompton H, Burke D. Artificial intelligence in higher education: the state of the field. International Journal of Educational Technology in Higher Education. 2023;20(22). https://doi.org/10.1186/s41239-023-00392-8.

[29] Semerikov SO, Vakaliuk TA, Mintii IS, Hamaniuk VA, Bondarenko OV, Nechypurenko PP, et al. An integrated conceptual framework for immersive cloud-based mobile learning in higher education. Humanities and Social Sciences Communications. 2026 Aug;https://doi.org/10.1057/s41599-026-08758-9.

[30] Marienko MV, Shyshkina MP, Semerikov SO. Professional development for open science adoption among Ukrainian educators. Social Sciences & Humanities Open. 2026;14:103394. https://doi.org/10.1016/j.ssaho.2026.103394.

[31] Landis JR, Koch GG. The Measurement of Observer Agreement for Categorical Data. Biometrics. 1977;33(1):159–174. https://doi.org/10.2307/2529310.

[32] Sullivan M. 24: Collusion, plagiarism, or contract cheating? How generative AI fits into existing academic integrity policie. In: Popenici S, Rudolph J, Ismail F, Tan S, editors. Handbook of Artificial Intelligence in Higher Education. Elgar Handbooks in Education. Edward Elgar Publishing; 2025. p. 372–386. Available from: https://doi.org/10.4337/9781035338764.00033.

[33] Mansoor SIU. Generative AI and the Doctrine of Academic Integrity: A Legal Perspective of Plagiarism, Originality and Liability. In: Lahby M, Schaeffer SE, Maleh Y, Paliktzoglou V, editors. Generative AI in Higher Education Assessment: Theory, Practice, and Ethical Implications. Cham: Springer Nature Switzerland; 2026. p. 23–42. Available from: https://doi.org/10.1007/978-3-032-01638-6_2.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplementary.zip