Europe at a Crossroads: Governance, Business, and Policy Transformation in a Changing Global Economy. Proceedings of the International Conference EU-PAIR 2026

DOI:10.47743/eu-pair.2026-1-22

FROM AUTHORSHIP TO STEWARDSHIP: A GOVERNANCE FRAMEWORK FOR AI INTEGRATION IN ACADEMIC RESEARCH

MIRCEA-ALEXANDRU ROȘU
Alexandru Ioan Cuza University of Iași
Iași, Romania
mircea.roshu@gmail.com

PATRICEA ELENA BERTEA
Alexandru Ioan Cuza University of Iași
Iași, Romania
patricea.bertea@gmail.com

ARMAND FLORIN BERTEA
Alexandru Ioan Cuza University of Iași
Iași, Romania

Abstract
Most researchers already use AI in their work. Recent surveys put the number at around three quarters, with literature reviews, summaries, and writing among the most common uses. Universities and journals have responded with a patchwork of rules: some ban AI tools, some allow them with conditions, most simply ask researchers to declare their use. This paper argues that these rules answer the wrong question. The real issue is not whether researchers use AI. It is whether they can check, explain, and take responsibility for what AI helps them produce. The paper builds on the idea of stewardship: the researcher's value no longer comes from personally writing every word, but from guiding the process, judging the output, and owning the result. This idea is translated into a practical framework for research. It proposes three steps — naming the change openly, building layers of review, and bringing the whole research community along — and a three-level model that separates low-stakes AI help (formatting, grammar), AI-assisted synthesis the researcher must verify, and core arguments the researcher must produce and own. European institutions already require accountability for AI use, through the Commission's Living Guidelines and the European Code of Conduct for Research Integrity. What they do not explain is how to practice it. This framework fills that gap. The paper also extends the governance role of universities from education to research, and gives supervisors, journals, and reviewers a way to evaluate how AI was

used, not just whether it was declared. In short: the question should move from "did you use AI?" to "can you account for what AI contributed?"
Keywords: *AI stewardship; research integrity; academic authorship; research governance; generative AI.*
JEL Classification: I23; O33; I28.

1. INTRODUCTION

Most researchers already work with AI. In a 2024 survey by Oxford University Press covering more than 2,300 academics, 76% reported using some form of AI tool in their research (Oxford University Press, 2024). In *Nature*'s survey of over 1,600 scientists, close to a third said they had used generative AI to help write manuscripts, and more than half expected such tools to become essential to their fields within a decade (Van Noorden and Perkel, 2023). These are not early adopters at the margins of academic life. This is the mainstream of research practice, and it changed in under three years.

The institutional response has not kept pace. Some journals banned generative AI outright, then quietly narrowed the ban. Others allow it under conditions that vary from publisher to publisher. Most universities have produced guidance documents, but these range from prohibition lists to permissive statements of principle, sometimes within the same institution (Hosseini, Resnik and Holmes, 2023). A doctoral student in Iași, Berlin or Lisbon can reasonably ask: what exactly am I allowed to do, and to whom do I answer when I do it? At the moment, the honest answer is that it depends on whom you ask.

This paper argues that the fragmentation is a symptom of a deeper problem: institutions are regulating the wrong object. Almost every existing policy regulates use — whether AI was involved, in which tool, and whether it was declared. What matters for research integrity is accountability: whether the researcher can check, explain and take responsibility for what AI helped produce. A researcher can disclose AI use scrupulously and still publish an unverified, machine-generated literature review. Another can use AI extensively and responsibly, verifying every claim, and have nothing meaningful to disclose under current rules beyond a generic statement. Disclosure tells us what was used. It does not tell us whether anyone checked.

The alternative proposed here is stewardship. The term comes from management research, where it describes actors who protect the long-term welfare of something entrusted to them, rather than maximising their short-term interest (Davis, Schoorman and Donaldson, 1997). Its recent application to AI-augmented work is owed to Brice Challamel, head of AI strategy and adoption at OpenAI, who describes the professional's role as shifting 'from authorship to stewardship': from personally producing every word to guiding an AI-assisted process, judging the result and owning the outcome (Challamel, 2026). This

paper adapts that model to the specific demands of academic research, where the bar for verification sits far higher than in commercial work.

The timing matters for Europe in particular. The European Commission's living guidelines on the responsible use of generative AI in research and the revised European Code of Conduct for Research Integrity already place responsibility for AI-assisted content squarely on the researcher (ALLEA, 2023; European Commission, 2024). What neither instrument specifies is how that responsibility is exercised in daily research practice. The framework developed here is intended as that missing operational layer. The paper proceeds as follows: section 2 clarifies the concepts of authorship and stewardship; section 3 reviews the limits of current policies; section 4 translates stewardship into research practice; section 5 addresses institutional design; section 6 discusses open questions; section 7 concludes.

2. DEFINITIONS AND CONCEPTUAL CLARIFICATIONS

2.1. Academic Authorship: What the Norms Actually Require

Before arguing that authorship norms need a successor, it is worth being precise about what those norms actually require. The most widely used definition comes from the International Committee of Medical Journal Editors, whose criteria have spread well beyond medicine. To qualify as an author, a person must contribute substantially to the work, participate in drafting or revising it, approve the final version and agree to be accountable for all aspects of the work (ICMJE, 2025). The Committee on Publication Ethics takes the same position, for the same reason: an author is someone who can take responsibility for content and respond when it is questioned (COPE, 2023). The European Code of Conduct for Research Integrity, revised in 2023 partly in response to AI, frames the same requirement through its principle of accountability for the research from idea to publication (ALLEA, 2023). Table 1 summarises the three frameworks.

Notice what is absent from these definitions: none of them requires the author to have personally typed the text. Production was always a proxy. We assumed that whoever wrote the words could answer for them, because writing them required understanding. What the frameworks actually protect is the ability to answer for the work. This is precisely why all three bodies concluded, with unusual speed and unanimity, that AI cannot be an author: not because its text is poor, but because it cannot be accountable. It cannot respond to critique, declare a conflict of interest or correct the record (COPE, 2023; ICMJE, 2025).

Table 1. Authorship and accountability across major research integrity frameworks

Framework	What authorship requires	Position on AI
ICMJE (2025)	Substantial contribution; drafting or revising; final approval; accountability for all aspects of the work	AI cannot be an author; its use must be disclosed; authors answer for all AI-assisted content
COPE (2023)	Capacity to take responsibility for the work and to respond to critique	AI tools cannot meet authorship requirements; they cannot declare conflicts of interest or manage copyright
ALLEA (2023)	Honesty, accountability and transparency at every stage of research	Researchers remain fully responsible for content produced with AI assistance

Source: adapted from ICMJE (2025), COPE (2023) and ALLEA (2023)

The key takeaway is that authorship was never about typing. It was about answering. AI did not break this principle; it broke the proxy, by making it possible to produce fluent text without the understanding that production used to guarantee.

2.2. AI Stewardship: The Conceptual Shift

Stewardship has a longer academic history than its recent association with AI suggests. In management theory, stewardship describes a relationship in which an actor manages resources entrusted to them in the interest of a collective, rather than for personal gain (Davis, Schoorman and Donaldson, 1997). Hernandez (2012, p. 174) defines it as 'the extent to which an individual willingly subjugates his or her personal interests to act in protection of others' long-term welfare'. A steward does not own the thing entrusted to them. They answer for it.

Challamel (2026) carries this logic into AI-augmented work. In the older model, which he calls authorship, professional value came from personally producing every fraction of the output, and ownership was easy to explain: you wrote it, so you answer for it. In the new model, value comes from guiding an AI-assisted creation process, judging the result and owning the outcome. Ownership no longer comes from production; it comes from the diligence applied to the process: how the task was framed, how the output was challenged and revised, and how carefully it was reviewed before being released.

Two aspects of this formulation deserve emphasis. The first is that Challamel frames the shift as different, not less: stewardship demands more judgment, more taste and more ownership than authorship, because the steward must hold a line for quality precisely when production becomes easy. The

second is that the model is operational rather than aspirational. It comes with named practices: making the shift explicit and discussable, building layered review systems that focus on outcome quality rather than on who produced the artifact, and building shared AI literacy so that proficiency differences do not harden into a social divide (Challamel, 2026). Section 4 adapts each of these to research.

2.3. Why Research is the Hard Case

Why does research need its own version of this framework? Because research is the setting in which unverified AI output does the most damage. A mediocre AI-generated marketing email costs a company some credibility, and the loss ends there. An unverified AI-generated claim in a published article enters the scholarly record, where other researchers build on it. Scientific knowledge is cumulative: articles cite articles, meta-analyses aggregate findings, doctoral students inherit assumptions. Think of citations as load-bearing walls. Each one supports the claims built on top of it, and a fabricated or misread source weakens everything above.

Peer verification imposes a second constraint. The entire apparatus of scholarly publishing assumes that a reader, a reviewer or a replicating team can trace claims back to evidence. AI-assisted steps that cannot be traced break this assumption silently: the text reads well, the references exist, and the chain of verification has a gap no one can see. Reproducibility imposes a third constraint, because a result that depends on an undocumented AI step cannot be repeated, and a result that cannot be repeated is, by the standards of science, not yet knowledge.

In practical terms, that means a stewardship framework for research cannot simply import enterprise review practices. Commercial work answers to customers, who mostly judge outcomes; research answers to referees and to the record, which judge process. The verification bar must be set where the scholarly record needs it, and that is the design problem the rest of this paper addresses.

3. THE LIMITS OF CURRENT AI POLICIES IN ACADEMIA

How have institutions actually responded so far? Three patterns dominate. The first is prohibition: several journals initially banned generative AI from submissions, a position most have since narrowed to a ban on AI authorship combined with disclosure requirements. The trouble with bans is practical: they are unenforceable, and they push use underground, which is the worst outcome for transparency (Hosseini, Resnik and Holmes, 2023). The second pattern is conditional permission: AI may be used for defined tasks, its use must be declared, and the authors remain responsible for the result. This is now the dominant publisher position, aligned with ICMJE (2025) and COPE (2023).

The third pattern operates at European level and is principle-based. The revised European Code of Conduct for Research Integrity extends honesty, accountability and transparency to AI-assisted work (ALLEA, 2023). The European Commission's living guidelines on generative AI ask researchers to remain accountable for scientific output, to disclose substantial AI use, and to protect unpublished work and personal data from being fed into open systems (European Commission, 2024). The Artificial Intelligence Act completes the picture from the product side, regulating AI systems themselves rather than research conduct (European Union, 2024). Taken together, the European instruments are ahead of most national and institutional policies: they ground AI use in integrity principles rather than in tool lists.

Each layer is reasonable. Together, they still leave the central question unanswered. All three patterns regulate declaration: whether AI was used and whether the use was stated. None of them specifies what responsible use looks like between the prompt and the publication. The living guidelines tell researchers to remain accountable; they do not say what accounting for a machine-generated literature summary involves, how a supervisor checks it, or what a reviewer should ask about it. The rules name the obligation but not the practice.

Academia has solved this problem before. When statistical software spread through the social sciences, nobody seriously proposed banning SPSS, and journals did not merely require authors to disclose that statistics had been used. Methodology sections adapted instead. Authors learned to report which tests they ran, on what data, under which assumptions; reviewers learned to ask. The discipline did not regulate the tool. It regulated the accountability around the tool. AI requires the same adaptation, with two complications: it touches every stage of research rather than analysis alone, and its failures are harder to spot, because fluent prose hides errors better than a regression table does.

The takeaway: current policies are necessary but insufficient. They tell researchers what to declare, not how to work. The next section proposes what the practice underneath the declaration should look like.

4. TRANSLATING STEWARDSHIP INTO RESEARCH PRACTICE

4.1. The Three Stewardship Steps in a Research Context

Challamel's (2026) transition model has three steps: name the pattern, build the review layers, and bring everyone into the transition. Each has a natural research equivalent.

Naming the pattern means making AI use discussable before making it auditable. In organisations, Challamel suggests changing team language, asking 'what have you supervised with AI this week?' instead of treating AI use as a private, slightly shameful habit. Research groups need the same shift: lab meetings where AI-assisted steps are reported as routinely as instrument settings,

and methods discussions that treat prompting strategies as methodological choices. The stigma matters more than it seems. A doctoral student who fears being judged for using AI will not ask the supervisor to check the output, and unchecked output is exactly what the framework exists to prevent.

Building the review layers is the structural step, and the one that demands the most adaptation to research; section 4.2 develops it in detail.

Bringing everyone into the transition addresses the divide between researchers fluent in AI and colleagues who avoid it. Challamel warns that access without shared literacy produces a social divide. In research the divide has a sharper edge, because the AI-avoidant group includes many senior researchers who supervise doctoral students, review manuscripts and sit on evaluation committees. A reviewer who has never verified an AI-generated synthesis cannot evaluate one. Shared literacy, common language and time to compare practice are governance infrastructure, not professional development extras.

4.2. Review Layers for Research: the L1–L3 Model

The centre of the framework is a layered review model. Challamel (2026) describes tagging outputs with human oversight levels, from trusted AI output to high-stakes judgment requiring full human control. Adapted to research, three layers are sufficient (Table 2).

Table 2. Review layers for AI-assisted research

Layer	Typical AI uses	Researcher obligation	Disclosure
L1 – Mechanical assistance	Grammar, formatting, reference formatting, translation polish	Spot-check the output	Generally unnecessary
L2 – Verified synthesis	Literature summaries, data description, code drafting, background sections	Verify every claim against the original sources and document the verification	Stated in methods or acknowledgements, like a statistical procedure
L3 – Owned contribution	Research questions, argument structure, interpretation, conclusions	Produce and own; AI limited to challenge and critique	Author attestation of ownership

Source: author's elaboration, adapted from Challamel (2026)

L1 covers mechanical assistance: grammar, formatting, reference formatting, translation polish. The defining property is that errors at this layer

are visible on inspection and do not touch the substance of claims. A spot-check is adequate review, and disclosure adds little, because knowing the grammar was machine-checked tells a reader nothing about the reliability of the science. Hosseini, Resnik and Holmes (2023) reach a similar conclusion for writing assistance generally.

L2 covers verified synthesis: literature summaries, descriptions of data, drafts of code, first versions of background sections. Here AI produces content with substantive weight, and the defining obligation is verification against original sources: every claim traced back, every citation opened, every number checked. The researcher who signs the manuscript is asserting that this happened. L2 belongs in the methods or acknowledgements, stated the way one states a statistical procedure, because that is what it is: a methodological step whose quality determines the reliability of what follows.

L3 covers the core scholarly contribution: research questions, the structure of the argument, interpretation of findings, conclusions. This layer is not delegated, and the reason is not sentimentality about human creativity. It is accountability arithmetic. L3 is the layer where the author's answerability is the product itself, so producing it is the one task that cannot be outsourced without dissolving authorship altogether. AI can still serve at L3 as a challenger, the sparring partner that attacks the argument and exposes weak joints, but the argument that survives must be the researcher's own.

A note on boundaries. The layers are defined by stakes and verifiability, not by tool or task names. The same coding task can be L1 when it reformats a chart and L2 when it produces the analysis script on which the results depend. The question that assigns a layer is always the same: what happens if this is wrong, and who can tell?

4.3. Domain Expertise as a Prerequisite for Stewardship

There is a precondition the model cannot work around: the researcher must know the domain well enough to judge the output. The best evidence comes from the largest field experiment on AI-assisted knowledge work to date. Dell'Acqua et al. (2025) gave 758 consultants at Boston Consulting Group realistic tasks, with and without access to GPT-4. Inside the frontier of the model's competence, AI users performed markedly better, with quality gains of around 40% for the group using the tool. But on a task deliberately designed to sit outside that frontier, consultants using AI scored 19 percentage points lower than colleagues working without it. The tool did not merely fail to help. It made skilled professionals worse, because they trusted plausible output they could not evaluate.

Translate this into research. A professor reviewing an AI-generated synthesis of her own field will catch the misattributed finding, the inflated effect size, the plausible-sounding citation that does not exist. A first-year doctoral

student reviewing the same synthesis will not. Both can perform L2 review in form; only one can perform it in substance. This is the quiet danger of stewardship language: it can describe rigorous verification or it can dress up rubber-stamping, and from the outside the two look identical.

The concern compounds over time. Gerlich (2025), in a mixed-methods study of 666 participants, found that heavier reliance on AI tools correlated with weaker critical thinking, mediated by cognitive offloading — the habit of delegating reasoning to the tool instead of performing it and checking the result. Higher educational attainment moderated the effect, which is mildly reassuring for academia, but the direction of the relationship is the problem: the people most exposed to offloading are precisely those still building the expertise that stewardship requires. The implication for research training is direct. Stewardship raises the value of domain expertise rather than lowering it, and AI integration in doctoral programmes must be sequenced with competence development, not substituted for it.

5. INSTITUTIONAL DESIGN: WHAT RESEARCH ORGANISATIONS NEED TO BUILD

5.1. From Disclosure to Methodological Accountability

What should institutions build? The first piece is a redefinition of the disclosure statement. Current AI statements typically record which tool was used and for which task. A stewardship-based statement records how the use was supervised: which layer the use belonged to, how verification was performed, and who performed it. The difference mirrors a familiar one. A methods section does not say ‘statistics were used’; it says which test, on what data, under which assumptions, so that a reader can judge whether the conclusions follow. An AI-use statement should do the same.

In a companion paper on AI governance in entrepreneurship education, we argue that declaring AI contribution should be treated like declaring a methodological choice, and that evaluation must address process and authorship, not only outputs (Bertea, Bertea and Roșu, 2026). The same logic, applied to research itself, turns disclosure from a compliance gesture into evidence of stewardship. It also gives the European instruments the operational anchor they currently lack: the living guidelines require accountability, and a layer-based statement is what exercised accountability looks like on paper (European Commission, 2024).

5.2. Peer Review Adaptations

The second piece is peer review. If authors document stewardship, reviewers can assess it, and three questions are enough to start. Are claims traceable, can each substantive assertion be followed back to a verifiable source? Are the human decision points clear, can one tell where the authors exercised

judgment rather than accepted output? Are AI-assisted steps described precisely enough to be repeated? None of these requires new technology. They require updated reviewer checklists and editor expectations, the same way reporting standards for statistics and data availability spread through journals over the past two decades.

Reviewers' own AI use falls under the same layered model, with one addition the European guidelines already make: unpublished manuscripts must not be fed into open AI systems, a confidentiality obligation that precedes any question of review quality (European Commission, 2024). A reviewer may use AI to check reference formatting; a reviewer may not delegate the scientific judgment that the review exists to provide, nor expose an author's unpublished work to a system that may retain it.

5.3. Supervision and AI Literacy as Research Training

The third piece is supervision and training. Powerful, risky tools are normally handled with training rather than prohibition. Cars are dangerous, and the social response is not a ban but drivers' education: supervised practice, with feedback, in increasingly demanding conditions. Handing researchers institutional AI licences without the equivalent of drivers' education is how many universities currently operate, and it inverts the logic. AI literacy, understood as the ability to direct, challenge and verify AI output in one's own discipline, should become an explicit component of research training, on the same footing as research ethics and statistics.

Supervisors carry particular weight here. A supervisor who avoids AI entirely cannot model responsible use, cannot check a student's L2 verification, and cannot tell stewardship from rubber-stamping. As you can expect, the literacy requirement therefore runs in both directions: doctoral students need training in verification, and supervisors need enough fluency to audit it. The institutional point connects to the broader governance argument. Universities sit at the meso level of research systems, translating macro-level norms, such as the European code and the Commission's guidelines, into micro-level daily practice (Bertea, Bertea and Roșu, 2026). If universities do not build this translation, the European instruments remain principles without practice.

6. OPEN QUESTIONS AND TENSIONS

The framework leaves real questions open, and it is more useful to name them than to wave at them.

The first is the long-term effect of AI on researcher formation. Gerlich's (2025) findings on cognitive offloading describe a correlation, not a destiny, but doctoral education is exactly the setting where the struggle is the point: researchers develop judgment by working through problems they do not yet know how to solve. If AI removes the productive difficulty from early research

tasks, the L3 capacities the framework depends on may develop more slowly. Whether structured stewardship training counteracts this is an empirical question, and an urgent one.

The second is attribution. When an AI system, used well at L3 as a sparring partner, suggests the reframing that becomes a paper's central insight, current authorship norms have no category for what happened. The accountability answer is clear: the researcher who recognised, validated and developed the insight owns it, exactly as a researcher owns an idea sparked by a colleague's question at a seminar. The credit answer is less clear, and the history of science suggests that credit disputes do not stay theoretical for long.

The third is the fig-leaf risk. Stewardship vocabulary can legitimise weak oversight: 'human in the loop' can mean rigorous verification or a tired click on approve. The framework reduces this risk by making review layers explicit and reviewable, but no framework survives an institution determined to perform compliance rather than practice it. The honest position is that stewardship raises the ceiling of responsible AI use without guaranteeing the floor.

The fourth is voice. If a generation of researchers drafts with the same handful of models, scholarly prose may converge toward a common register. Early bibliometric evidence suggests this is already measurable: GenAI-assisted publications show statistically significant linguistic convergence toward dominant norms of scientific English (Bao et al., 2025), and repeated reliance on AI-assisted writing negatively affect essential writing skills and critical thinking (Selvi, 2024). Whether that convergence matters scientifically, or only aesthetically, remains unsettled.

Each of these is a research agenda rather than an objection. Pilot studies of layer models in working research groups, comparative analysis of verification behaviour across career stages, and longitudinal designs following AI-trained doctoral cohorts would all test claims this paper can only argue.

7. CONCLUSIONS

This paper began from a mismatch: AI is integrated into research practice at a pace policy has not matched, and the policies that exist regulate declaration rather than practice. The argument developed here is that the academic community needs a stewardship framework, not a thicker compliance checklist. The conceptual move is small but consequential: stop asking whether researchers used AI and start asking whether they can account for what AI contributed. Authorship norms already point this way, because they were always about accountability rather than production. Stewardship supplies the practices that accountability now requires: a named, discussable pattern; review layers calibrated to stakes, from mechanical assistance through verified synthesis to owned argument; and shared AI literacy across career stages.

For Europe, the proposal lands on prepared ground. The revised code of conduct and the Commission's living guidelines already assign researchers responsibility for AI-assisted work; what they lack is the operational layer that specifies what responsible use looks like in a lab meeting on an ordinary Tuesday. Universities, as the institutions that translate European principles into research routines, are where that layer gets built: in doctoral training, in supervision practice, and in how research outputs are assessed. Journals and funders close the loop by asking for stewardship evidence rather than bare disclosure.

The shift is demanding. It asks more judgment of researchers, more attention of supervisors and more sophistication of reviewers, at exactly the moment AI makes production effortless. That is not a flaw in the framework. It is the framework. Holding a line for quality precisely when producing has become easy is what stewardship means, and the scholarly record is worth holding the line for.

References

1) ALLEA (2023). *The European Code of Conduct for Research Integrity – Revised Edition 2023.* Berlin: ALLEA – All European Academies. [online] Available at: https://allea.org/code-of-conduct/ [Accessed 12.06.2026].
2) Bao, T., Zhao, Y., Mao, J., and Zhang, C. (2025). *Examining linguistic shifts in academic writing before and after the launch of ChatGPT: A study on preprint papers*. *Scientometrics*, 130(7), 3597–3627. https://doi.org/10.1007/s11192-025-05341-y, [Accessed 12.06.2026].
3) Bertea, P.E., Bertea, A.F. and Roșu, M. (2026). Reconfiguring entrepreneurial education in the age of AI: a governance perspective. *Paper presented at the EUFIRE 2026 International Conference*, Alexandru Ioan Cuza University of Iași, Romania (in press).
4) Challamel, B. (2026). *From authorship to stewardship.* [online] Available at: https://www.jeremyutley.com/blog/from-authorship-to-stewardship [Accessed 12.06.2026].
5) COPE (2023). *Authorship and AI tools: COPE position statement.* [online] Available at: https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools [Accessed 12.06.2026].
6) Davis, J.H., Schoorman, F.D. and Donaldson, L. (1997). Toward a stewardship theory of management. *Academy of Management Review*, 22(1), pp. 20-47.
7) Dell'Acqua, F., McFowland III, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R. (2025). Navigating the jagged technological frontier: field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Organization Science*, 37(2), pp. 403-793, https://doi.org/10.1287/orsc.2025.21838.
8) European Commission (2024). *Living guidelines on the responsible use of generative AI in research.* Brussels: European Commission, Directorate-General

for Research and Innovation. [online] Available at: https://research-and-innovation.ec.europa.eu/ [Accessed 12.06.2026].
9) European Union (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L 2024/1689.
10) Gerlich, M. (2025). AI tools in society: impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), 6, https://doi.org/10.3390/soc15010006.
11) Hernandez, M. (2012). Toward an understanding of the psychology of stewardship. *Academy of Management Review*, 37(2), pp. 172-193.
12) Hosseini, M., Resnik, D.B. and Holmes, K. (2023). The ethics of disclosing the use of artificial intelligence tools in writing scholarly manuscripts. *Research Ethics*, 19(4), pp. 449-465, https://doi.org/10.1177/17470161231180449.
13) ICMJE (2025). *Recommendations for the conduct, reporting, editing, and publication of scholarly work in medical journals.* [online] Available at: https://www.icmje.org/recommendations/ [Accessed 12.06.2026].
14) Oxford University Press (2024). *Researchers and AI: survey findings.* [online] Available at: https://corp.oup.com/news/how-are-researchers-responding-to-ai/ [Accessed 12.06.2026].
15) Selvi, B. (2024). The Role of Generative Artificial Intelligence Tools in Enhancing Academic Writing. 22-56, https://www.researchgate.net/publication/387857664_The_Role_of_Generative_Artificial_Intelligence_Tools_in_Enhancing_Academic_Writing [Accessed 12.06.2026].
16) Van Noorden, R. and Perkel, J.M. (2023). AI and science: what 1,600 researchers think. *Nature*, 621(7980), pp. 672-675, https://doi.org/10.1038/d41586-023-02980-0.