The Third Knowledge Regime: Artificial Intelligence as an Emergent Response to the Limits of Classical Academia and Open Repositories

El tercer régimen del conocimiento: la inteligencia artificial como respuesta emergente a los límites de la academia clásica y los repositorios abiertos

Carlos Eduardo Ravello Joo
ORCID: 0009-0007-5631-7436
carlosravello.com
Trujillo — Lima, Peru

DOI: 10.5281/zenodo.24164493

November 2026 · Version 2 (v1: May 2026)

Version note

Version 2 (November 2026). This version corrects and qualifies the May 2026 version, which remains available on Zenodo.

Corrected:

- The SWE-bench figure, now sourced to the AI Index 2025.
- The share of notable models produced by industry: nearly 90%, in 2024.
- Time to doctorate: a median of 7.2 years from graduate school entry and 5.7 from doctoral program start (NSF/NCSES).
- The *Nature* survey: 52% "significant crisis", not 83%.
- Richard Smith's years as editor of the *BMJ*: 13, not 25.
- The DOI and journal of Carneiro et al. (2020).
- The title of Topaz et al. (2026), the description of its corpus and the causes it proposes.
- The date and authorship of the Wikidata notability reform proposal.
- The attribution of the crawler-blocking figures: BuiltWith, Tollbit and Cloudflare.

Removed: figures without a verifiable source (PyPI, npm, Wikidata use in Wikipedia, commercial integrations) and the claim that Wikidata is the primary source of Google's Knowledge Graph.

Qualified: the use of "failure", several manifesto-style sentences, and the reading of Carneiro et al. and Gartenberg et al.

Added:

- A section on the domains where slow review must remain, with the distinction between reviewing out of usefulness, out of obligation or by default.
- Evidence on AI use by students and on how AI systems crawl, summarize and cite.
- The author's own case as a critique of the "circular lock" argument.
- A note on the preprint status of the text.
- Expanded limitations and the missing references.

- An appendix with a reproducible recalculation of the Reproducibility Project: Psychology figures.

Author note

This work is part of a series of independent preprints by the author. It connects with conceptual frameworks developed in Ravello Joo (2026a, 2026b, 2026c, 2026d), deposited in Zenodo. No external funding was received. The author declares no conflict of interest.

Abstract

Artificial intelligence does not surpass humans through intrinsic cognitive capacity. The observable operational advantage emerges from both computational capacity and the elimination of accumulated institutional frictions that limit collective human cognition. This article examines that diagnosis from three simultaneous fronts: the structural slowness of the traditional academic system, the epistemic noise produced by open repositories without adequate filtering, and the anonymous gatekeeping that operates over the knowledge graphs that AI systems ingest. The concept of epistemological coordination costs is introduced as a framework to understand why both classical academia and open repositories fall short, each in its own way, of the current scenario, and how AI enters that historical gap. Each regime reviews for a different reason: the open one out of usefulness, the classical one out of obligation, while AI does not review at all; it summarizes by default. That is why the rigor demanded of knowledge cannot depend on the regime it travels through. The article also marks where the thesis does not apply: for knowledge whose errors cannot be reversed, such as clinical knowledge, slow review must remain. It is argued that knowledge no longer travels only through human institutions, and that documenting rigorously before automated systems fill that vacuum with imagination without method is a task that falls on those of us who do research, not an academic ornament. Information, like any physical system, tends toward disorder without organized energy flow. Knowledge institutions were the mechanisms humanity built to locally reduce that entropy. When those mechanisms fall short, disorder advances faster than correction.

Keywords: institutional friction, academic validation, open repositories, Wikidata, knowledge graphs, epistemological coordination costs, artificial intelligence, peer review, informational entropy

Table 1. Three epistemological regimes

System	Speed	Validation	Result
Classical academia	Low	High on paper, uneven in practice	Rigidity + invisibility
Open repositories	High	Minimal or none	Noise at scale
Artificial intelligence	High	Signal compression	Operational advantage

1. Introduction

As of March 24, 2025, Hugging Face hosted 1,533,973 artificial intelligence models (González et al., 2025). Not as spectacle. As symptom.

The symptom points to a question that knowledge institutions are not answering with sufficient honesty: why does an architecture without metabolism, without professional career, without reputational fear, without biographical time, operate cognitively faster than systems specifically designed to produce and validate knowledge?

The usual answer—that AI has more computational power—is true but incomplete. Computational power explains processing speed. It does not explain why a five-year clinical study can remain invisible to the systems that today build knowledge at scale, while a poorly structured GitHub repository appears indexed within hours.

This article proposes a more uncomfortable answer: AI's operational advantage emerges both from computational capacity and from the elimination of accumulated institutional frictions that limit collective human cognition (Gartenberg et al., 2026; Maslej et al., 2025). This is not an intelligence problem. It is an architecture problem.

Information, like any physical system, tends toward disorder without organized energy flow (Shannon, 1948; Schrödinger, 1944). Knowledge institutions were the mechanisms humanity built to locally reduce that entropy. The formal development of that framework connects with the author's prior work (Ravello Joo, 2026a, 2026b, 2026c). What this article examines is what happens when those mechanisms fall short at the same time, from two opposite directions.

2. Two Simultaneous Failures

I use "failure" in a precise sense: a system that worked for its era and cannot keep up with the speed and volume of the current one.

2.1 The Slow Failure: Academia as Scarcity Architecture

Peer review was born in a world where the problem was informational scarcity. Little knowledge was produced, and the challenge was protecting it from contamination. Slowness was an adaptive feature, not a defect: careful validation was the reasonable price of reliability.

That world no longer exists. In 2024, nearly 90% of notable AI models came from industry (Maslej et al., 2025). On SWE-bench, a benchmark of real programming problems introduced in 2023, AI systems solved 4.4% of the problems in 2023 and 71.7% in 2024 (Maslej et al., 2025). In the United States, the median time to a research doctorate in 2022 was 7.2 years from graduate school entry and 5.7 years from the start of the doctoral program (NSF/NCSES, 2024).

The risk is that relevant research becomes obsolete during its own validation process. Gartenberg et al. (2026), using the full corpus of 6,957 initial submissions and 10,389 reviews handled by *Organization Science* since 2021, documented a 42% rise in submissions since the release of ChatGPT, with lower writing quality, and conclude that the system is shifting toward more research rather than better research.

More seriously: the system does not consistently produce what it promises. The Open Science Collaboration (2015) found that only 36% of replications of psychological studies produced statistically significant results, and that 39% were judged by the teams themselves to have replicated; replication effects were, on average, half the size of the originals (recalculation in the Appendix). Camerer et al. (2018) replicated 21 social science experiments published in *Nature* and *Science*: 13 (62%) replicated, with effects about half the original size. The Retraction Watch database exceeded 58,000 retractions in April 2025 (Retraction Watch, 2025). In the first seven weeks of 2026, in the open-access subset of PubMed Central, one in every 277 papers had at least one nonexistent reference, compared to one in 458 in 2025 and one in 2,828 in 2023. The authors point to paper mills, deliberate misconduct and the uncritical use of AI writing tools as possible origins (Topaz et al., 2026).

Richard Smith (2006) wrote from inside the system: peer review is a flawed process at the heart of science and journals. And he said it as someone who ran it, as editor of the *BMJ*, for thirteen years.

2.2 The Fast Failure: Repositories as Noise Architecture

The response to the slow failure was speed without filtering. Open repositories—GitHub, Zenodo, Hugging Face, arXiv, OSF—solved the problems of technical invisibility and bureaucratic slowness: a DOI in hours, global visibility, versions and the possibility of being cited.

But speed without methodology produces its own failure. With more than one and a half million models on Hugging Face alone as early as March 2025, no in-depth human filter can keep up.

The most revealing data comes from Carneiro et al. (2020): the difference in reporting quality between preprints and peer-reviewed articles is about 5%. After months of editorial process, the improvement is small. What that means depends on how it is read: either peer review does not add as much value as it claims, or open repositories are not as chaotic as their critics argue. If the difference in reporting is 5%, the question of how much review contributes, against what it costs, is a legitimate one. The authors themselves conclude that the improvement exists, even if it is small.

Noise is not just a quality problem. It is a training problem. The language models that hallucinate most are those operating on scarce, contradictory, or low-quality data. When open repositories produce noise at scale, that noise becomes signal for systems building automated knowledge. Chaos scales.

2.3 The Third Regime: AI as Abundance Compression

Classical academia and open repositories appear opposed. Functionally they are distinct mechanisms for solving the same problem: how to filter information under human limits of attention and coordination.

Academia manages uncertainty through hierarchy, scarcity, credentials, and exclusion. Open repositories, through abundance, iteration, massive exposure, and memetic selection. Both are responses to different historical conditions. Neither was designed for the current scenario.

AI enters as a third epistemological regime. It does not validate like academia. It does not explore chaotically like the open internet. It synthesizes, recombines, compresses, and accelerates cognitive navigation. But it has a critical structural limit: its quality depends on the quality of what it ingests.

A Stanford team showed that, by pairing a language model with verifiable queries over Wikidata, useful answers could be given to 96% of the questions in a benchmark's development set; on that same set strict accuracy was 76%, and 65% on the test set (Xu et al., 2023). AI is not epistemically self-sufficient. It is a compression machine that amplifies what it receives, good or bad.

This introduces the central concept of this article: epistemological coordination costs. The contemporary knowledge crisis does not emerge simply from an excess or absence of filters, but from the institutional inability to reconcile epistemic validation and cognitive speed under conditions of informational overproduction. AI radically reduces those coordination costs. Human institutions, by contrast, kept accumulating them until they became very hard to sustain.

3. The Circular Lock

3.1 Wikidata as Operational Filter

In 2020, Google said its Knowledge Graph held more than 500 billion facts about five billion entities, drawn from hundreds of sources, with Wikipedia among the most commonly cited (Sullivan, 2020). Wikidata is one of the open databases that feed graphs of this kind.

Wikidata is not an encyclopedia repository. It is part of the knowledge graph architecture on which the systems that build the semantic world that LLMs ingest operate. Without a node in Wikidata, it is harder for Google to resolve the entity. Without entity resolution, dispersed mentions on the web do not accumulate toward any center. The work exists, but for the system it is nobody.

That filter is administered in part by anonymous editors with 19th-century notability criteria, applied with 20th-century tools, over 21st-century phenomena. An editor without name, without verifiable expertise in the domain, can delete in minutes the technical work of an independent researcher invoking notability criteria that do not distinguish between research without institutional affiliation and content without value, with no real avenue of reply and no one answerable for the impact.

3.2 The Mechanism and Its Technical Consequences

The proposal to reform Wikidata's notability policy, opened in December 2025, acknowledges it: certain knowledge has been and is being structurally marginalized, leading to less coverage in reliable sources and therefore an increased barrier to demonstrating notability (Wikidata, 2025).

That is the circular lock in its most precise formulation: the system demands notability to enter, but notability is built with the visibility the system denies by not entering. The problem lies in the structure, not in an isolated governance bug.

The technical consequences are concrete and scalable. Without a node in Wikidata, entering Google's Knowledge Graph becomes harder, and so does everything that follows: the knowledge panel, mentions accumulating on a single node, models recognizing the entity. The work tends not to appear in AI-generated responses, and the virtuous cycle of semantic authority does not start.

My own case complicates this argument. My Wikidata item was deleted and, even so, Google keeps an entity under my name in its Knowledge Graph (checked on November 12, 2026). I have no record of whether that entity appeared while the item existed or afterwards, so I cannot know whether Wikidata was the seed and the entity survived the deletion, or whether Google assembled it from other signals. What it does show is that Wikidata is not the only door. That is why I say here that the absence of a node makes entry harder, not that it prevents it.

The same period that saw this anonymous gatekeeping expand saw ClaudeBot blocked on some 5.8 million websites according to BuiltWith (Claburn, 2025), and an estimated 73,000 pages crawled for every visit referred back in June 2025 (Cloudflare, 2025). Machines are hungry for verified knowledge. The systems that should feed them block them or leave them without nodes to anchor what they find.

3.3 Active Technical Invisibility

Publishing in Scopus, *Nature*, or Elsevier is not just slow. In 2025, it is also an act of progressive technical invisibility. Major academic publishers operate with hard paywalls that serve Googlebot only metadata or initial snippets. According to Tollbit, as cited by *The Register*, sites blocking AI crawlers increased by 336% in the year to the second quarter of 2025 (Claburn, 2025). A Rutgers and Wharton study found that large news publishers who blocked language model crawlers lost about 7% of weekly visits in the six weeks that followed (Zhao & Berman, 2026).

University repositories complete the picture: unstructured URLs, orphaned pages without interlinks, no schema markup, unoptimized sitemaps, pre-modern loading speeds. The Googlebot

has a crawl budget. It is a time economy. It will not get lost in a 2003 technical architecture. It goes where the food is served: Zenodo, arXiv, OSF, GitHub Pages.

The researcher who publishes only in traditional channels is not being more rigorous. They are being invisible. And in the current ecosystem, invisible to machines is invisible to the knowledge already being built at scale.

4. The Validators Before the Threshold

This article does not ask to abolish peer review. It asks that reviewers be at the level of the moment they guard.

The day the first AI said hello changed the minimum threshold of what it means to do science seriously. That threshold rose for everyone. For those who research and for those who validate. The problem is that only some are being required to notice that the floor moved.

A peer reviewer who does not master the tools with which the knowledge they review is produced is not a guardian of rigor. They are a bottleneck with a title. Not out of bad faith—most operate ad honorem, with their own academic workload and criteria inherited from a system that never updated its standards. That is why the demand is structural, not individual: validation systems must incorporate criteria that include whether the work can confront the reality it describes, whether it can be found, whether it can be falsified by something more than three colleagues in a closed circuit.

Popper did not ask for anonymous peers or institutional titles. He asked for real exposure to refutation. The system that today claims that heritage produces—according to *Nature*'s survey of 1,576 researchers (Baker, 2016)—that more than 70% have attempted and failed to reproduce another scientist's experiment, and more than 50% have failed to reproduce their own. 52% consider that there is a significant reproducibility crisis. Exposure to refutation had already weakened before AI arrived, quietly, within the same system that claimed it. It did not disappear: it shifted.

Open repositories, public benchmarks, code that either works or does not, apply more genuine Popperian pressure on a claim than eight months of peer review followed by a paywall that blocks Googlebot.

4.1 Where Slow Review Has to Stay

There is a domain where this argument has to stop. What has been said so far applies to knowledge that can be reversed: a badly measured benchmark is corrected in the next version, a broken package is withdrawn, a mistaken preprint is versioned. That is what speed and downloads are for: often it is tested by whoever finds a use for it, discarded just as fast, and a comment is left along the way. That is a kind of exposure that science, in certain domains, cannot ignore. But a diagnostic criterion, a psychological intervention or a drug do not work the same way. Neither does a badly built bridge or a new safety measure. There, review is not a luxury: it requires that people with the experience and the credentials examine it, and tear it apart if necessary. Here credentials do count, because whoever signs with them answers for the harm; what failed in section 4 was the title without command of the tools, not the title. The error is not paid by the system but by a person, and for that person there is often no version 2.

It is worth separating candidate knowledge, which can circulate quickly and be exposed to refutation in plain sight, from evidence, which is what is used to make decisions about someone. The question that orders that boundary is who absorbs the error when the error cannot be undone.

Psychology is in an uncomfortable position here. It is one of the disciplines hit hardest by the replication crisis (Open Science Collaboration, 2015) and, at the same time, one of those most in need of careful review, because it works with people. For psychology, the demand is a different one: that the review it gets be real.

In both cases the rigor required is the same: studies and statistics have to be equally solid in the slow regime and in the fast one. What changes in this era is who pays for the error: in the irreversible domain it is paid by whoever receives the intervention; in the reversible one, increasingly, by whoever is gathering that information right now. From my point of view, denying information to artificial intelligence systems is no longer a realistic option; that debate is being left behind. AI systems seek us out even when we try to shut the door on them. In June 2025, Anthropic crawled some 73,000 pages for every visit it referred back (Cloudflare, 2025), and according to Tollbit, as cited by *The Register*, 13.26% of AI bot requests ignored robots.txt directives in the second quarter of that year (Claburn, 2025). They summarize us and cite us, but almost no one goes to the source: in an analysis of 900 U.S. users, in only 1% of visits to a page with an AI summary did anyone click on a link in the summary itself (Pew Research Center, 2025). And they guarantee neither self-regulation nor prior verification of the data. So the burden, less visibly, falls on students who are learning to cite and working on their first research projects, and who already use these tools massively (Digital Education Council, 2024; Freeman, 2025). The available surveys come from the United Kingdom and from an international study; I do not have an equivalent figure for Peru. They are the ones who will inherit the problem.

This is not about saying which regime is better. It is about why each one reviews: the open one reviews out of usefulness, the classical one out of obligation, and AI does not review: it summarizes by default, with no reason of its own. That is why rigor cannot depend on the regime.

5. The Obligation Before It Is Too Late

Algorithmic weight does not distinguish between truth and opinion. It distinguishes between engagement and silence. An influencer with 10 million followers opining on vaccines, economics, or mental health generates more signal in the graph than a researcher with 40 years of clinical work publishing in a repository that Googlebot never visits. Not because the algorithm is malicious. Because the algorithm optimizes what it was asked to optimize: attention, screen time, interaction.

If rigorously documented knowledge cannot compete in that graph, the problem is no longer only epistemological. We are not discussing whether academia is slow. We are discussing who is going to train the systems that will build the reality of the next generations. If that training is done with the algorithmic weight of opinions without method, without falsifiability, without verifiable record, the result is not a faster world. It is a world where the difference between truth and narrative disappears not through ignorance but through architecture.

Knowledge no longer travels only through human institutions. More and more it passes through algorithms, and what those algorithms find well documented carries more weight. AI has a real ceiling of comprehension and too much imagination. LLMs hallucinate more when verified sources are scarce or blocked. If humans do not leave a well-documented record, machines will fill that vacuum with the imagination they have to spare. And they will do it with the speed that institutions never had.

Documenting well is not academic vanity: it is the surest way we have for the knowledge we produce today to survive the algorithm that decides tomorrow what is real.

Limitations

This work constitutes an analytical essay with critical narrative review of verifiable literature. It is not a formal empirical study. Source selection introduces selection bias inherent to all non-systematic reviews. Some quantitative data cited come from industry reports whose exact methodology is not public. The technological landscape described evolves at a speed that may render specific aspects of the analysis outdated in the short term.

The article is a position paper, not a neutral review, and its tone reflects that. The examples come mostly from technology and software, where speed favors the thesis; in other fields the balance may differ. The concept of epistemological coordination costs is not yet operationalized: I do not propose here how to measure it, and as long as it cannot be measured it cannot be refuted either.

This text is a preprint without peer review that criticizes peer review. The contradiction exists and I do not hide it. Version 2 corrects figures and references that were wrong in the first; version 1 remains available on Zenodo for anyone who wants to compare them. It is the kind of correction in plain sight that the article defends.

Appendix. Reproducible Verification of the Reproducibility Project: Psychology Figures

What it is and what it is not. The body of the article cites the replication rates of the Open Science Collaboration (2015). This appendix recalculates them from the project's public data. It is not a new finding: the RPP has been reanalyzed many times. It is a check done in plain sight: what the article cites, the author recalculated, and the code and data allow anyone to repeat it.

Data. `rpp_data.csv`, "Analysis" component of the RPP project on OSF (https://osf.io/download/fgjvw/), CC0 1.0 license. File SHA-256: `4d87a2636181d8e31230422854cf0fabeaecd2f1b47a59627482c05a2b56d54d`, which matches the hash declared by OSF. The file has 168 rows, of which 100 are completed replications (`Completion (R) = 1`). It is encoded in latin-1, not UTF-8.

Procedure. Two Python programs that use only the standard library: `rpp_reanalisis.py` and `diag_p.py`. The first stops if the input SHA-256 does not match the declared one, and lists blank cells and out-of-category values separately, without silently discarding them. It was run twice, in two different Miniconda environments: the base environment, with Python 3.13.5, and a new one with Python 3.15.2. The output summary was byte-for-byte identical in both (SHA-256 `07de851a5821c5ef798323c6d1211784c9e28e1866eff978a6070a700c36e8c6`).

Table A1. Recalculation results against the published values

Indicator	Recalculation	Published (OSC, 2015)
Replications with $p < .05$, among originals with recalculated $p \leq .05$	34 of 92 (37.0%)	36%
Same, with OSC's denominator (originals reported as significant)	34 of 97 (35.1%)	36% (35 of 97)
Replication teams' subjective judgment: replicated	39 of 100 (39.0%)	39%
Original effect within the replication's 95% CI	44 of 94 (46.8%)	47%
Mean r, original / replication (97 pairs)	0.396 / 0.197	0.403 (SD 0.188) / 0.197 (SD 0.257)

Median of the replication r / original r ratio, per study	0.43	—

By discipline (completed replications): cognitive psychology (n = 43) replicated between 1.7 and 2.1 times as often as social psychology (n = 57), depending on the criterion:

- **Significance:** 52.5% vs. 25.0%.
- **Teams' judgment:** 53.5% vs. 28.1%.
- **Original effect within the CI:** 61.0% vs. 35.8%.

What does not match, explained.

The significance denominator. The article counts 97 significant originals; the strict criterion (recalculated $p \leq .05$) counts 92. The diagnostic identifies the 8 cases left out:

- **Three originals with a non-significant reported p:** $p > .1$, $p = .48$ and p_rep = .18. These are the ones the article also excludes: $100 - 3 = 97$.
- **Four originals reported as significant or marginal whose recalculated p falls just above .05:** between .0503 and .0516.
- **One original reported as p < .001** with no recalculated p in the file.

None of those five replications has $p < .05$ in the file. With OSC's denominator, the recalculation gives 34 of 97 (35.1%), against the published 36% (35 of 97). One replication remains unaccounted for. There is a likely, unverified explanation: study 59, reported as $p < .001$ in the original, has its replication judged successful by its team but no p in the file; if the article counted it as a significant replication, its numerator would be 35. This is reported as a hypothesis.

Study 135 shows a recalculated original p of .91 and a replication with $p = 3 \times 10^{-10}$. It does not affect the counts, since the original was not significant; the cause was not investigated.

The "half" effect. Comparing means, replication effects are 50% of the original ones, as the article states. Per study, the median ratio is 0.43, so "half" is the more favorable reading.

The original mean r. The recalculation gives 0.396 over 97 pairs, and the article publishes 0.403 (verified in the article's structured abstract in *Science*, accessed November 12, 2026). The difference probably comes from not using the same set of pairs; this was not verified.

Availability. Both programs, the output table (`tabla_rpp.csv`) and the summary (`resumen_rpp.txt`) accompany this version in the Zenodo record.

Dataset reference: Open Science Collaboration. (2015). *Reproducibility Project: Psychology* [Data set]. OSF. https://osf.io/ezcuj/

References

Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature*, *533*(7604), 452–454. https://doi.org/10.1038/533452a

Camerer, C. F., et al. (2018). Evaluating the replicability of social science experiments in *Nature* and *Science* between 2010 and 2015. *Nature Human Behaviour*, *2*(9), 637–644. https://doi.org/10.1038/s41562-018-0399-z

Carneiro, C. F. D., et al. (2020). Comparing quality of reporting between preprints and peer-reviewed articles in the biomedical literature. *Research Integrity and Peer Review*, *5*, 16. https://doi.org/10.1186/s41073-020-00101-3

Claburn, T. (2025, December 8). Publishers say no to AI scrapers, block bots at server level. *The Register*. https://www.theregister.com/software/2025/12/08/publishers-say-no-to-ai-scrapers-block-bots-at-server-level/2064025

Cloudflare (Lee, J.-H., Gupta, D., Mitchell, B., Tatoris, R., & Clausen, H.). (2025, July 1). *Control content use for AI training with Cloudflare's managed robots.txt and blocking for monetized content*. https://blog.cloudflare.com/control-content-use-for-ai-training/

Digital Education Council. (2024, August 7). *What students want: Key results from DEC Global AI Student Survey 2024*. https://www.digitaleducationcouncil.com/newsroom/what-students-want-key-results-from-dec-global-ai-student-survey-2024

Freeman, J. (2025). *Student generative AI survey 2025*. Higher Education Policy Institute & Kortext. https://www.hepi.ac.uk/reports/student-generative-ai-survey-2025/

Gartenberg, C., Hasan, S., Murray, A., & Pierce, L. (2026). More versus better: Artificial intelligence, incentives, and the emerging crisis in peer review. *Organization Science*, *37*(3), 795–812. https://doi.org/10.1287/orsc.2026.ed.v37.n3

González, A., et al. (2025). *How do pre-trained models support software engineering? An empirical study in Hugging Face*. arXiv:2506.03013. https://arxiv.org/abs/2506.03013

Maslej, N., et al. (2025). *AI Index Report 2025*. Stanford Institute for Human-Centered Artificial Intelligence. https://hai.stanford.edu/ai-index/2025-ai-index-report

NSF/NCSES. (2024). *Doctorate recipients from U.S. universities: 2022* (NSF 24-300). National Science Foundation. https://ncses.nsf.gov/pubs/nsf24300

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, *349*(6251), aac4716. https://doi.org/10.1126/science.aac4716

Pew Research Center. (2025, July 22). *Google users are less likely to click on links when an AI summary appears in the results*. https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/

Ravello Joo, C. E. (2026a). Metacognición 2.0: Diseño deliberado de identidad digital ante sistemas predictivos de inteligencia artificial. Zenodo. https://doi.org/10.5281/zenodo.20092009

Ravello Joo, C. E. (2026b). El perfil que no escribiste: Identidad digital, biometría ocular y autoridad semántica en sistemas algorítmicos predictivos. Zenodo. https://doi.org/10.5281/zenodo.20112409

Ravello Joo, C. E. (2026c). La década del influencer como colapso termodinámico de la información: Cierre operacional, alostasis del conocimiento y bifurcación del ecosistema digital (2012–2025). Zenodo. https://doi.org/10.5281/zenodo.20260383

Ravello Joo, C. E. (2026d). La inteligencia artificial como navaja suiza epistemológica: cuando la autoridad percibida reemplaza a la verificación [Artificial intelligence as an epistemological Swiss Army knife: when perceived authority replaces verification]. Zenodo. https://doi.org/10.5281/zenodo.24116983

Retraction Watch. (2025, April 26). [Blog entry with the database count]. https://retractionwatch.com/2025/04/26/

Schrödinger, E. (1944). *What is life?* Cambridge University Press.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*(3), 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Smith, R. (2006). Peer review: A flawed process at the heart of science and journals. *Journal of the Royal Society of Medicine*, *99*(4), 178–182. https://doi.org/10.1258/jrsm.99.4.178

Sullivan, D. (2020, May 20). *A reintroduction to our Knowledge Graph and knowledge panels*. Google. https://blog.google/products/search/about-knowledge-graph-and-knowledge-panels/

Topaz, M., Roguin, N., Gupta, P., Zhang, Z., & Peltonen, L.-M. (2026). Fabricated citations: An audit across $2 \cdot 5$ million biomedical papers [Correspondence]. *The Lancet*, *407*(10541), 1779–1781. https://doi.org/10.1016/S0140-6736(26)00603-3 (Erratum: *The Lancet*, *408*(10551), 218.)

Wikidata. (2025, December). *Requests for comment: Notability policy reform*. Wikimedia. https://www.wikidata.org/wiki/Wikidata:Requests_for_comment/Notability_policy_reform (accessed November 12, 2026)

Xu, S., Liu, S., Culhane, T., Pertseva, E., Wu, M.-H., Semnani, S. J., & Lam, M. S. (2023). Fine-tuned LLMs know more, hallucinate less with few-shot sequence-to-sequence semantic parsing over Wikidata. *Proceedings of EMNLP 2023*. arXiv:2305.14202. https://arxiv.org/abs/2305.14202

Zhao, H., & Berman, R. (2026). *Strategic response of news publishers to generative AI* (v4). arXiv:2512.24968. https://arxiv.org/abs/2512.24968