

Research Article

**Killing the Average, Amplifying the Few: Generative AI, Social Media Patterns,
and the Structural Crisis of Academic Work**

*Ortalamayı Öldürmek, Azınlığı Güçlendirmek: Üretken Yapay Zeka, Sosyal Medya
Örüntüleri ve Akademik Çalışmanın Yapısal Krizi*

Deniz HASANÇEBİ

Araş.Gör. Dr., Ankara Hacı Bayram Veli Üniversitesi

Edebiyat Fakültesi

d.hasancebi@gmail.com

<https://orcid.org/0000-0002-8541-4063>

Makale Geliş Tarihi	Makale Kabul Tarihi
11.12.2025	05.03.2026

Abstract

This article rejects the dominant framing of generative AI in academia as merely a problem of plagiarism, laziness and cheating, and instead reads it as a symptom and accelerator of a structural crisis in the regime of academic labour. It asks two related questions: through which mechanisms does generative AI transform an already metricised and standardised academic production system, and why can this transformation not be adequately captured within the narrow frame of “ethics violations”? Drawing on conceptual analysis and a targeted reading of recent empirical and methodological studies, the article argues that generative AI automates the bulk of “average academic work” – literature summarising, template filling, low-risk repetition – while acting as a powerful multiplier for a small minority of researchers who possess genuine theoretical depth and original insight. At the same time, the spread of synthetic data and synthetic participants raises new questions about what still counts as empirical research when “data” can be generated without encounter with the world. Situating these developments within audit culture, algorithmic infrastructures and the attention economy of social media, the article contends that the real risk of generative AI is not simply individual misconduct but a further algorithmic colonisation of academic subjectivity itself. It concludes by sketching institutional implications for evaluation criteria, the cautious use of synthetic data and the protection of slow, risky and original work.

Key words: Regime of knowledge production, academic capitalism, generative AI, automation of academic labour, original thought.

Öz

Bu makalede, akademiye üretken yapay zekâ tartışmalarını intihal, tembellik ve kopya eksenine sıkıştırılan hâkim çerçeveyi reddederek, meseleyi akademik emek rejimindeki yapısal krizin bir belirtisi ve hızlandırıcısı olarak tespit edilmektedir. İki ilişkili soruya odaklanılmıştır: Üretken yapay zekâ, hâlihazırda metrikleştirilmiş ve şablonlaştırılmış akademik üretim sistemini hangi somut mekanizmalar üzerinden dönüştürmektedir? Ve bu dönüşüm, neden yalnızca “etik ihlaller” çerçevesiyle kavranamayacak kadar derin bir yapısal sorundur? Çalışmada kavramsal analiz ile son dönemdeki ampirik ve yöntemsel çalışmaların hedefli bir okumasına dayanan metin, üretken yapay zekânın “ortalama akademik emek” diye adlandırılan – literatür özetleme, kalıp doldurma, düşük riskli tekrar – işlerin büyük kısmını otomatikleştirirken, gerçek kuramsal derinlik ve özgün kavrayış gücüne

Önerilen Atf /Suggested Citation

Hasançebi, D., 2026, Killing the Average, Amplifying the Few: Generative AI, Social Media Patterns, and the Structural Crisis of Academic Work, *Üçüncü Sektör Sosyal Ekonomi Dergisi*, 61(1), 1159-1176.

sahip az sayıdaki araştırmacı için güçlü bir çarpan etkisi yarattığı savunulmaktadır. Eşzamanlı olarak sentetik veri ve sentetik katılımcı kullanımının, “veri”nin dünya ile karşılaşmadan üretilebildiği bir ortamda ampirik araştırmacının ne anlama geldiğine dair yeni sorular doğurduğu gösterilmiştir. Bu gelişmeleri denetim kültürü, algoritmik altyapılar ve sosyal medyanın dikkat ekonomisi bağlamına yerleştirildiğinde, üretken yapay zekânın asıl riskinin bireysel suistimal değil, akademik öznelliğin daha ileri düzeyde algoritmik kolonizasyonu olduğunu ileri sürülmüştür ve değerlendirme ölçütleri, sentetik veri sınırları ve yavaş, riskli, özgün çalışmanın kurumsal korunması açısından bazı sonuçlar tartışılmıştır.

Anahtar Kelimeler: Bilgi üretim rejimi, akademik kapitalizm, üretken yapay zeka, akademik emeğin otomasyonu, özgün düşünce.

1.Introduction

The current discourse on generative artificial intelligence in academic settings is routinely presented as though a disruptive technological agent had abruptly intervened in a domain presumed to be internally coherent and normatively settled, as if the university had not already reorganised the conditions of its own labour long before such systems appeared. ChatGPT and similar models are described as endangering assessment, authorship and academic integrity, as if the university had only recently started to face a crisis. Yet the fact that a large language model, trained only on textual patterns and with no experiential relation to knowledge, can produce prose that passes through many of our academic filters without resistance already shows that the problem did not begin with AI. It indicates that academic writing in many contexts has been reduced to a reproducible set of moves: recognizable structures, expected transitions, conventional formulations and predictable citation habits. A form that can be mechanically imitated in this way is no longer a transparent expression of intellectual labour but a script that can be filled. Eke’s article tellingly entitled “*ChatGPT and the rise of generative AI: Threat to academic integrity?*” explicitly formulates the concern that such systems may undermine the norms of responsibility, honesty and authorship on which academic work is supposed to rest (Eke, 2023). The question, however, is whether those norms are still effectively embodied in our ordinary practices of writing and evaluation, or whether AI merely exposes how far practice has already drifted from its own ideals.

This becomes clearer when we look at how generative tools are actually used and discussed within academic environments. Raitskaya and Tikhonova’s scoping review shows that AI language tools are now embedded in a wide range of functions: text generation, proofreading, editing, text annotation, paraphrasing and translation, as well as assistance with research writing and feedback (Raitskaya & Tikhonova, 2024, p. 5). The same review notes that the emerging field is increasingly described in terms of “hybrid AI-powered writing of various assignments and genres,” where AI and human input are interwoven rather than neatly separated (Raitskaya & Tikhonova, 2024, pp. 5, 8–9). This does not only mean that new tools have been added to existing practices. It also means that the genres themselves — student essays, theses, articles and reports — have a level of formal rigidity that allows this kind of hybridisation to take place without obvious rupture. If a model can help produce an acceptable discussion section, a plausible literature review or a structurally correct term paper, this suggests that success in these genres often depends more on conformity to a template than on the presence of a unique, conceptually grounded problem. In other words, AI does not so much corrupt a previously “authentic” writing culture as reveal how strongly that culture had already moved toward standardised, easily learnable and therefore easily automatable forms.

The literature on reference “hallucinations” illustrates the same dynamic in a more concrete way. Gravel and colleagues, working with medical questions, show that ChatGPT can generate responses supported by citations that look formally correct yet are fabricated or only partially accurate (2023, pp. 226–230). Alkaissi and McFarlane find similarly high levels of fabricated or erroneous references in citations produced by ChatGPT, concluding that a substantial portion of the bibliographic material generated by the system either does not exist or is distorted in significant ways (Alkaissi & McFarlane, 2023). These findings are usually presented as evidence of a new kind of technological risk: AI appears to “invent” sources and thereby undermines scientific communication. But the situation is more ambivalent. On the one hand, the fabrication of references clearly raises ethical and epistemic problems. On the other hand, the fact that such references can pass through some stages of scholarly review and evaluation shows that

citation has long been treated, at least in part, as a formal marker of legitimacy rather than as a carefully examined network of relations between texts.

Taken together, these strands point to a reversal of the dominant narrative. Generative AI is usually imagined as an external disturbance to a stable academic world. A closer reading of the empirical studies suggests almost the opposite. AI thrives wherever academic writing has already solidified into conventional forms that can be reproduced without understanding. It becomes a problem wherever citation has already become a stylistic performance rather than a traceable link to prior work. It is adopted most intensively where institutional pressures reward productivity, speed and formal correctness over slow, risky and sometimes messy processes of thinking. In this sense, generative AI is not only a technological object to be regulated. It is also a diagnostic instrument. By looking at where AI fits most easily into our academic practices, we learn something about what those practices have become.

On this basis, the article pursues two related questions. First, through which concrete mechanisms does generative AI transform the existing regime of academic labour – in particular the routines that have already been standardised, metricised and bound to institutional performance pressures? Second, why can this transformation not be adequately described within the narrow frame of “ethics violations,” plagiarism or individual dishonesty, but must instead be analysed as a symptom of a structural crisis in how academic work is organised and valued? To address these questions, the article combines conceptual analysis with a targeted, critical reading of empirical and methodological studies on generative AI in academic contexts. The literature considered is not comprehensive in a bibliometric sense; it is selected because it illuminates key points along the argument’s trajectory: the automation of routine writing tasks, the emergence of synthetic data and synthetic participants, the entanglement of universities with audit culture and algorithmic infrastructures, and the consolidation of social media logics in academic life. Studies that do not speak to these structural dimensions are therefore not discussed in detail, even when they are part of the broader AI-in-education debate.

2. Structural Pressures in Contemporary Academia

If generative AI finds such a natural place inside academic writing, it is not only because the technology is powerful. It is also because the existing conditions of academic labour have already normalised a model of production that favours speed, standardisation and measurable output over slow and often unpredictable processes of thinking. The expansion of what Shore and Wright have called “audit culture” in higher education marks a key turning point in this transformation. Audit culture describes the spread of ranking systems, performance indicators and accountability regimes into universities, where they come to define what counts as “quality” and “excellence” in research and teaching (Shore & Wright, 1999; Shore, 2008). In such an environment, academic activity is increasingly organised around what can be counted, compared and ranked, rather than what is conceptually necessary or socially meaningful. The metrics that are fundamental to academia have the duty of measuring existing practices; however, over time, they reshape the entirety of academic labor. They create incentives for certain kinds of work and disincentives for others. Consequently, metrics determine the perceived quality or significance of a research paper. They also change the internal grammar of academic life like a lexicon change, as researchers learn to describe their projects, outputs and careers in ways that align with audit categories.

Again, the metrics make this logic explicit. In academia, how citation counts, journal impact factors and other indicators have become deeply embedded in research assessment and management, even though their methodological limitations and distortive effects are widely recognized (Wilsdon et al., 2015). The quantity/metrics are often used as “proxy measures” for quality almost all cases. Even in the absence of strong evidence that these metrics reliably track intellectual significance, there is a systemic preference for quantitative data. At the same time, this shows us a growing dependence on numerical indicators in hiring, promotion and funding decisions. This combination — skepticism about validity but continued reliance in practice — creates a situation in which academics are structurally encouraged to optimize for what metrics can register. The metricisation of science has generated “new ecologies of misconduct,” where citation manipulation, strategic self-citation, salami-slicing of publications and the use of predatory outlets are rational responses to a system that ties career survival to quantified performance

(Biagioli & Lippman, 2020). This situation has not emerged from the new technologies. In this context, it is unsurprising that a tool like ChatGPT, which can quickly produce formally acceptable text, becomes attractive. It aligns perfectly with an academic regime that measures output units more easily than it measures intellectual depth.

The often-repeated discourse in academia “publish or perish” is therefore not just a slogan. It is the shorthand for a culture in which continuous publication becomes a condition of staying in the profession. This pressure on academics as drive to publish, combined with information overload and the pursuit of journal impact factors, produces a “conflicting tripod of forces” that structures and controls the researchers’ everyday decisions (Grech, 2022). In such a culture, the time available for reading, thinking and rewriting shrinks, while the pressure to generate citable units increases. Some Turkish research shows how this same publish-or-perish dynamic operates in the Turkish context: their bibliometric study of educational administration research reveals a pattern of rising publication counts, concentration in lower-ranking journals, thematic repetition and a tendency to choose “safer” topics that fit current trends (akir, Kuyurtar, & Balyer, 2024). This also highlights that qualitative depth and originality are often sacrificed in favor of producing articles that satisfy promotion criteria in a compressed timeframe. Unfortunately, internal investigations into this field, which is closed to risky and authentic new ideas and directly determined by epistemic communities, remain superficial. Beyond this, the perspective on new factors like artificial intelligence is also based on a return to the paradigm of control and epistemic community. When these structural features are treated as background noise rather than as central determinants of academic practice, the debate on AI is inevitably misdirected.

The integration of generative AI into academic writing is therefore not treated here as an external disruption. On the contrary, AI’s effect on academia is just an intensification of existing tendencies. A system that already favors quantifiable output, formal compliance and rapid turnaround will naturally reward tools that accelerate text production and help writers meet stylistic expectations. In this sense, ChatGPT and similar systems are not foreign bodies. They are functional supplements inside a metric-driven environment. They help resolve the contradiction between ever-increasing demands for publication and the finite time and energy of researchers. This does not make their use unproblematic. But it does mean that focusing exclusively on individual ethics — asking whether a particular student or academic has “cheated” — ignores the structural logic that makes such use rational. Under conditions of audit culture, it is coherent to delegate parts of the writing process to a machine, because the system evaluates the finished text primarily as a formally correct product and only secondarily, if at all, as the trace of a thought process.

At the same time, metricisation reshapes not only how much academics must write, but also how they can write. Shore shows that audit culture encourages universities to pursue “world-class” status through rankings, which in turn pushes scholars to align their topics, methods and writing styles with what appears in high-impact, English-language journals (Shore, 2008). Morrish observes that this environment produces “managed” academic identities, where individuals learn to present themselves as efficient, competitive and endlessly productive, often at the cost of intellectual risk and critical distance (Morrish, 2019). The resulting discursive field is narrow. It privileges certain forms of clarity, certain ways of structuring arguments, certain safe vocabularies. This narrowing of acceptable voice and structure is precisely what makes academic writing so vulnerable to automation. This system of automation doesn’t come from the AI or any other technologies. Academia is itself promoting the “automation” of works before any AI has come for. The more uniform and predictable a genre becomes, the easier it is for a generative model to inhabit it.

When AI is introduced into this already constrained -by self-automation- field, it strengthens some of its most problematic features. It can help scholars tailor their manuscripts to the expectations of specific journals, streamline language to meet reviewer demands, and generate multiple variants of an abstract in minutes. It can be used to rephrase, normalize and smooth out texts that deviate from dominant stylistic norms. In all cases, AI functions, as an instrument of alignment, allows authors to satisfy journals basic demands more efficiently. This alignment did not occur by chance or as a result of the very high skill capabilities of the artificial intelligence. It is the consequence of a long-term convergence between the standardization of academic genres and the pattern-based capacities of generative models.

For this reason, treating generative AI purely as a threat to integrity obscures the deeper issue. Integrity itself has been quietly redefined under the pressures of audit culture and metrisation. It is increasingly equated with formal compliance — correct citation format, adherence to authorship rules, avoidance of plagiarism as detected by software — rather than with a robust relation between the researcher and their object of inquiry. In such a context, it is possible to be “formally honest” while still participating in a system that devalues thought. Generative AI makes this contradiction visible. It fulfils the formal requirements while lacking any epistemic relation to what it writes about. The uncomfortable question for academia is not whether AI can simulate scholarly writing. It clearly can. The question is why so much of what counts as scholarly writing can be simulated in the first place.

3. Synthetic Data and the Displacement of Empirical Labor

A second, and more recent, axis of the AI–academia relationship appears not at the level of writing but at the level of data and “participants.” Social scientists are increasingly experimenting with large language models as substitutes or supplements for human subjects: synthetic survey respondents, synthetic interviewees, synthetic focus groups, and synthetic qualitative datasets. It has been noted that LLM-generated data are now being used to “stand in” for human responses in survey research, experiments and qualitative studies, often with the promise of lower cost, faster turnaround and fewer ethical hurdles (Rossi, Harrison, & Shklovski, 2024). A growing body of methodological work explores how such synthetic respondents can be tuned to mimic particular populations, attitudes or demographic profiles, and many of these studies report surprisingly high correlations between model outputs and real survey data. For example, research on “synthetic survey participants” finds that GPT-4 can generate responses that track human answers across multiple countries, at least on average (Shrestha et al., 2024). At first sight, this looks like a technical innovation. On closer inspection, it raises a fundamental question about what counts as empirical work when “data” can be produced without any encounter with the world.

The same trend appears in experimental and qualitative designs. Studies now describe using LLMs to simulate entire datasets for behavioural interventions, to generate “synthetic participants” for A/B testing, or to pre-run social-science experiments *in silico* before touching the field. Other work explores “synthetic qualitative data,” where models are prompted to role-play interviewees or focus group members with specified characteristics, such as age, gender, political orientation or professional background (Collecting Qualitative Data at Scale with LLMs; MindPort guide). In some cases, these synthetic responses are used only for piloting instruments. In others, they begin to blur into the analytic process itself, shaping codes, themes and interpretations. For a discipline that has long defined itself through empirical engagement with human subjects, this is a significant shift. The labour of going into the field, negotiating access, encountering resistance, misunderstanding and surprise is at least partially displaced by a controllable generator of plausible responses. The notion of “the field” risks being reduced to a parameter in a prompt.

Critical voices have already started to point out the epistemic and ethical stakes of this development. Rossi argues that LLM-based synthetic data rest on “fundamentally different epistemological assumptions” than earlier forms of synthetic data, because here the generator is not a transparent model but an opaque system trained on unknown corpora, with unknown biases (2024). Resnik likewise warns that synthetic data created by generative AI raise “numerous ethical concerns,” including threats to data integrity, scientific reliability and accountability, even when privacy arguments are used to justify them (Resnik, 2024). From a media-ethics perspective, there are serious doubts about whether synthetic data can ever be treated as fully “ethical” when they are used as substitutes for empirical encounter, because they intensify a pre-existing lack of accountability in automated systems: no one has actually said what the data claim to represent. A more defensible position is that such data may be conditionally acceptable in tightly delimited roles – for example, piloting instruments or stress-testing analytical pipelines – but become ethically and epistemically problematic when they are deployed as if they were direct evidence about real populations or events. In each of these critiques, the core issue is not only bias or technical imperfection. It is the substitution of a worldless generator for the messy, asymmetrical and often resistant reality that empirical research is supposed to confront.

At the same time, a number of methodological papers and commentaries present synthetic participants as a pragmatic solution to longstanding problems: cost, access, nonresponse, “WEIRD” sampling, and the slowness of traditional fieldwork. Stanford researchers, for instance, describe how LLMs can emulate “diverse human subjects” to test instruments and assumptions cheaply and at scale (Miller, 2025). Other authors highlight how synthetic respondents can be used to pre-test survey designs, explore hypothetical scenarios or approximate hard-to-reach populations. These promised advantages are not trivial. They speak directly to the structural pressures described in the previous section: the demand for more studies, faster cycles and broader coverage with fewer resources. Under such conditions, synthetic data appear as an almost natural extension of the audit culture. They promise more “data points” without the friction of dealing with actual people.

From the standpoint of the philosophy and sociology of science, however, this convergence is precisely what should concern us. Empirical research has always involved a gap between the concept and the world, between the categories of the researcher and the responses of participants. That gap is where surprise, resistance and learning occur. Synthetic participants, by contrast, are optimised not to resist but to comply. They are trained to produce coherent, contextually plausible answers that satisfy the researcher’s prompt. The more successfully they do this, the more they risk collapsing the very distance that makes empirical work meaningful. If both the questions and the answers are, in effect, produced within a single technical system, then “data” no longer name an encounter with something outside the model, but a variation of the model’s own internal patterning.

In this sense, the turn to synthetic participants and synthetic data is structurally parallel to the earlier shift from thinking to text templates in academic writing. In both cases, generative AI finds a ready-made niche wherever academic practices have already become formalised and predictable (Bin-Nashwan et al., 2023). Just as essay genres can be filled without genuine conceptual engagement, empirical designs can be populated without genuine empirical contact. AI does not invent this drift away from the world. It accelerates it and makes it visible. For a social science committed to understanding real societies, institutions and conflicts, the crucial question is therefore not whether synthetic participants can sometimes be useful. It is whether the long-term normalisation of such tools will further erode the already fragile link between research and the realities it claims to study (Cotton et al., 2024).

The emerging empirical literature no longer treats generative AI as a marginal writing aid but as a system that quietly reorganizes what institutions understand by knowledge, authorship, and epistemic responsibility. Raitskaya and Tikhonova’s scoping review makes this point with unusual clarity: generative tools now participate in “topic development, paraphrasing, summarizing, proofreading, text generation, and reference management,” extending into nearly every stage of academic writing (2024, pp. 5–7). According to this research, academia would be forced “to critically reconsider concepts such as cocreation, ownership, and authorship” (p. 6) in order to work with generative models. So, there is a shift in the academic authorship but this seems, technically incorrect. Because authorship is not only the act of writing but the act of *knowing*, and the question of who wrote a text is inseparable from the question of who stands behind the claims it makes. How to distinguish between the authentic and simulated authorship of an academic product remains a critical question. Academia is based on the journals evaluation system (two-blinded referee etc.) is assuming an essential product-author relation. But, if generative models take over substantial portions of knowledge-bearing work, the chain of accountability that traditionally connects researcher, evidence, and argument will become uncertain.

We can see this uncertainty in the domain of automated citation. Fabrication of the data becomes more visible with the academic references that AI is using. In other words, ChatGPT’s responses to a set of clinically relevant medical questions are “69% of the references were fabricated, although they appeared real” (Gravel, D’Amours-Gravel, & Osmanliu, 2023, p. 226). Their caution that these fabrications were “deceptively real references” which were frequently internally cohesive and arranged well enough to evade cursory examination, is more noteworthy. (p. 229). A identical tendency in a larger textual dataset: ChatGPT frequently generates “factually incorrect or entirely nonexistent references” which they specifically blame on hallucinations. (Alkaissi & McFarlane, 2023). As a result, the traditional appearance of scholarly citations diminishes their epistemic validity. A reference no long ensures that a text exists, let alone that its argument is sound. It turns into an external sign, an aesthetically pleasing indication of intellectual authority.

Bibliometric analyses reveal that the problem is not only individual but systemic. The emergence of “ghost bibliographic references,” showing that fabricated or altered citations can quietly infiltrate citation networks and warp the metrics upon which scholarly evaluation depends. (Orduna-Malea & Cabezas-Clavijo, 2023, pp. 5352–5354). Because of this ghost references emerging, there has been created a new term “sneaked references,” demonstrate that automatically generated citations—whether valid or fabricated—can artificially raise citation counts and distort the patterns through which scholarly visibility is constructed (Besançon, Cabanac, Labbé, & Magazinov, 2024, pp. 1368–1370). Overall, these investigations show that generative AI creates a structurally novel form of epistemic noise that impacts both individual texts and the architectures of scholarly communication, rather than just producing occasional errors.

However, it must be emphasized that the structural degradation at issue here long predates the arrival of generative AI. Fabricated bibliographies and false citations existed well before these systems appeared; there is nothing fundamentally new in this respect. As indicated in the preceding sections, AI merely renders the problem more visible. The contexts criticised in the literature are not novel developments—they are longstanding pathologies that AI now exposes with greater clarity.

Same problems occur in different areas. Clinical fields have already begun to ring alarm bells. Because, many AI-generated references are not just incomplete but “raise basic questions about whether the cited articles exist at all” (Sebo, 2024, p. 247) and because clinical research is an area in which risks must be minimised for the sake of human health, the use of artificial intelligence has once again been positioned as the primary source of danger. Yet to grasp how fragile these domains already were before AI, it is enough to recall the fabricated studies that circulated during the Covid period. The health sector has long been one of the fields most marked by deep structural fractures, where large pharmaceutical companies are effectively able to shape epistemic truth itself. Precisely because the subject matter concerns health, new protocols now prohibit systems like ChatGPT from offering medical guidance. Such restrictions, however, function only as superficial deterrents. There is also some slogans like “beware of references when using ChatGPT as a source of information,” after finding that a substantial portion of its citations could not be verified (Sanchez-Ramos, Lin, & Romero, 2023, p. 356). In the ordinary workflow of an academic journal, is a thorough verification of references and a detailed review process not already required? What, then, does the presence or absence of artificial intelligence actually change? In fact, such warnings function as a muted outcry from scholars who are fully aware of the academy’s deeper structural problems. Knowing how fraught journal evaluation processes already are, these cautions operate as a discreet signal to reviewers who will, inevitably, turn to AI. Overwhelmed by workload and publication pressure, reviewers often lack both the time and the mental clarity to examine submissions with the care they require; they will rely on generative systems, and the proliferation of erroneous references and fabricated citations will continue. Yet even here, the problem does not ultimately lie in AI.

Lechien and colleagues, comparing GPT-3.5 and GPT-4 in otolaryngology, emphasize that even improved models continue to produce erroneous references and argue that automated citation should never be accepted “on trust” (Lechien, Briganti, & Vaira, 2024, pp. 2163–2165). The convergence across disciplines is unmistakable: AI-generated citations *look* authoritative, but their epistemic grounding is often absent. Responsibility becomes ambiguous because the researcher may no longer know whether the references anchoring their own argument actually exist.

At this point, the philosophical stakes become clear. If authorship traditionally binds a scholar to the claims they advance, what happens when a significant portion of those claims—or the sources that justify them—originates from a generative model? Raitskaya and Tikhonova observe that the debate increasingly shifts from *who* writes to the disappearance of a stable “author’s voice” in collaborative human–AI compositions (2024, pp. 21–24). Thorp captures this dilemma succinctly in “*ChatGPT is fun, but not an author*”, writing: “The responsibility for a scientific article cannot be delegated to a machine” (Thorp, 2023, p. 313). What is at stake here is not plagiarism but ontology. When an academic article is assembled largely through algorithmic processes, and when its reference list is generated by tools whose internal logic is inaccessible, the researcher’s epistemic role collapses from that of a knower to that of a signatory. The text may circulate as scholarship, but the chain of reasoning that once connected author and argument becomes increasingly opaque.

This is why treating generative AI simply as a “misuse” problem is analytically insufficient. Ghost references, sneaked citations and hallucinated bibliographies are not merely technical malfunctions. They are symptoms of an emerging regime in which the appearance of scholarly form can be algorithmically produced without the underlying epistemic labor. The deeper concern for philosophy and sociology of science is therefore not whether AI should be banned, restricted or detected. The concern is where to draw the non-negotiable boundary between the act of producing text and the act of knowing. If that boundary is not re-asserted with conceptual clarity, academia may become increasingly productive in quantitative terms while the epistemic content of its outputs grows progressively thinner—an academic sphere rich in citations, metrics and textual abundance, yet poor in responsibility.

4. Algorithmic Infrastructures and The Normalization of AI

The discussion of generative AI in academia often isolates “writing tools” and “synthetic participants” from the wider digital environment in which they have emerged. This separation is misleading. Large language models do not appear in a vacuum. They are the most recent layer of a longer historical transformation in which algorithms have become the basic infrastructure through which information, communication and even culture are organised. Gillespie’s chapter on “public relevance algorithms” already formulated the central point more than a decade ago: “Algorithms are playing an increasingly important role in selecting what information is most relevant to us, a crucial feature of our participation in public life” (Gillespie, 2014, p. 167). Search engines, recommendation systems and ranking mechanisms do not merely filter information. They “produce and certify knowledge” by establishing which items appear, in what order, and with what implied authority (Gillespie, 2014, pp. 167–168). When academic institutions now debate the legitimacy of AI-generated text, they do so within a communication environment that has long delegated the basic work of ordering relevance to opaque computational procedures.

This delegation has a cultural dimension that is not secondary. Striphas argues that over the last three decades we have been “delegating the work of culture – the sorting, classifying and hierarchizing of people, places, objects and ideas – increasingly to computational processes” (Striphas, 2015, p. 395). This is what he calls “algorithmic culture,” a configuration in which cultural judgement and selection are no longer primarily public, contestable and reflexive practices, but background operations embedded in code and platforms (Striphas, 2015, pp. 395–399). In such a setting, the step from ranking search results to generating whole passages of text is not a rupture. It is an intensification. If algorithms already decide what counts as visible, relevant or trending, it is conceptually consistent that similar systems begin to produce the prose that fits into those structures. Generative AI in academia is therefore part of a broader shift in which machine learning models take over tasks that used to be understood as instances of reading, comparing, evaluating and synthesising.

Van Dijck’s analysis of social media platforms as an “ecosystem of connective media with a few large and many small players” helps clarify how this shift is stabilised (Van Dijck, 2013, pp. 14–15). Platforms such as Facebook, YouTube and Twitter encourage users to live inside an infrastructure where connectivity is both an experience and a business model. Early Web 2.0 rhetoric framed this as participatory culture and democratic empowerment. Van Dijck shows, however, that companies “appeared less interested in communities of users than in their data” and that connectivity quickly evolved into a resource which engineers could “code into algorithms that helped brand a particular form of online sociality and make it profitable” (Van Dijck, 2013, p. 14). In other words, platformed sociality is already algorithmically shaped sociality. When students and academics now turn to generative AI, they do so as subjects who have long been habituated to environments where algorithmic mediation is normal, invisible and continuous.

Zuboff’s concept of “surveillance capitalism” pushes the analysis one step further by linking this algorithmic mediation directly to a distinctive regime of accumulation. She writes that “surveillance capitalism unilaterally claims human experience as free raw material for translation into behavioral data” and fabricates these data into “prediction products that anticipate what you will do now, soon, and later” (Zuboff, 2019, p. 13). These prediction products are then traded in “behavioral futures markets,” and competitive pressure drives firms to acquire ever more predictive sources of behavioural surplus (Zuboff, 2019, pp. 13–14). From this perspective, generative AI is not a neutral text engine. It is part of

a larger apparatus in which user interactions, prompts and outputs become further sources of data, refinement and behavioral insight. When AI systems are integrated into academic workflows, the university is not merely adopting a new tool. It is participating in a data regime that treats scholarly interaction, student writing and research communication as streams of extractable behavioral traces.

This is the broader picture of the current problems that are not taken seriously. It can be seen that the normative discourse around AI in academia should not appear more than the academia's structural problems. Because the policy documents always focus on plagiarism, authorship and fairness, almost never acknowledge that universities themselves are now deeply entangled with platform infrastructures and data-driven business models. Learning management systems, citation databases, analytics dashboards and "impact" platforms all operate with similar logics of tracking, ranking and prediction. The same rationality that underpins targeted advertising and behavioural futures also informs dashboards that rank scholars, journals and institutions. When generative AI is added on top of this stack, its outputs are evaluated within an already algorithmically normalised environment. The question is not only whether a student "cheated" by using AI. It is also whether the institution has aligned its own evaluation practices so closely with platform metrics and automated assessment that AI-generated conformity becomes structurally advantageous.

For the philosophy and sociology of science, this raises a sharper and more uncomfortable question than the usual "Is AI good or bad for education?" If algorithms now participate in deciding what is knowable, what is visible and what is writable, then the conditions of possibility for critical knowledge are themselves being reorganised. Generative AI in academic writing and research is only one manifestation of this reorganisation. It sits alongside ranking algorithms, recommendation systems, predictive analytics and surveillance capitalist business models as part of a wider algorithmic infrastructure. To focus solely on individual misconduct or tool misuse is, analytically inadequate. The more appropriate question is how far universities are willing to let their epistemic practices be shaped by infrastructures that optimise for prediction, engagement and profitability rather than for understanding.

5. Automation of Average Academic Labour and the Selection of Original Work

With the arrival of generative AI, the first thing that starts to shake in academia is not "creative genius" or "exceptional intellect," but the everyday labour that fills corridors, sits on committees, completes files and produces two or three "acceptable" articles a year. A very large part of this labour does not consist in genuine theoretical invention but in scanning the literature, extracting summaries, filling out standard article templates and repackaging the same conceptual schemes with small variations. Thousands of theses and tens of thousands of articles repeat the same names and the same arguments in slightly altered configurations, year after year. In such a system the core of the job is already open to automation: the implicit question is simply who can do this technical work faster, cheaper and in a more "properly academic" tone. In most of these tasks the answer is no longer the human academic but the model. What is at stake for the average academic is therefore less "instant unemployment" and more a steady slide into replaceability, invisibility and devaluation.

What is meant by "average academic labour" here is precise. It is the kind of work that produces summaries of canonical thinkers and theories, extends reference lists without really connecting concepts, and confines itself to filling in the "introduction–method–findings–conclusion" template without taking any real conceptual risk. It is used here as an analytical shorthand, not as a moral judgement: an ideal-type that designates forms of work characterised by high volume, low conceptual risk, strong dependence on established templates, and close alignment with metric-driven performance criteria. The only differences between one article and another are often sample size, time frame or country name; the underlying argument chain hardly moves. Generative AI is entering exactly this terrain with almost perfect overlap: it can scan far larger volumes of literature, reconstruct standard argument paths, and package them in clean, grammatically flawless academic prose. In that situation, the gap between what the average academic writes and what a well-prompted model can produce narrows on every axis that institutions currently reward: speed, volume and formal correctness. Under such conditions, insisting on the value of one's own labour reduces to the thin claim "I wrote this myself," a claim that carries almost no weight in systems oriented toward counts and metrics rather than intellectual risk.

The more abstract picture is already backed by concrete data. A large-scale survey of 2,534 researchers at universities in one European country shows that most of those who use generative AI do so in three main modes: as a “work horse” that takes over routine tasks, as a mere “language assistant,” or as a “research accelerator” that streamlines early stages of projects (Andersen et al., 2025). The distribution is not random. Reported use clusters around literature searches, draft generation, summarising, stylistic polishing and basic data handling – exactly the high-volume but low-creativity micro-tasks that make up the bulk of average academic production (Andersen et al., 2025). A global postdoc survey reported in a leading journal paints a similar picture: roughly one third of nearly four thousand respondents say they systematically use chatbots to polish prose, generate or fix code and “tidy up” the literature for everyday research tasks, which places AI squarely in the middle of day-to-day academic work rather than at some exotic periphery (Nordling, 2023, pp. 655–656). In other words, the model is not hovering at the edges of theory; it is being plugged directly into the very layer where routine academic labour lives.

At the macro level, the structure of production makes the consequences even clearer. One highly cited analysis of publication patterns shows that thousands of researchers effectively operate at a pace of “one paper every five days,” revealing an already-industrial rhythm of output in which hyper-productivity has been normalised (Ioannidis, Klavans, & Boyack, 2018). In parallel, work on the “automating of academia” describes generative AI as a technology that meshes seamlessly with existing performance regimes, ranking industries and “excellence” discourses, further hiding the labour behind writing, assessment and student feedback while reinforcing managerial oversight (Watermeyer, Phipps, Lanclos, & Knight, 2024, pp. 446–447). In economic research, detailed case studies show large language models delivering substantial productivity gains precisely in ideation, drafting, background searches, data analysis and coding – that is, in the segments of research most directly tied to putting text and code on the page (Korinek, 2023, pp. 1299–1302). Taken together, these strands indicate that average academic labour has long been bound to an industrial rhythm; AI does not create this factory, it simply makes its products cheaper, faster and even more anonymous.

All this forces a sharper way of drawing the line between the “average academic” and the researcher who operates with genuine depth and originality. For someone with a clear theoretical map, coherent conceptual priorities and a concrete sense of which questions need to be asked, AI functions as a multiplier: it takes over the tiring but non-creative steps and amplifies the reach and tempo of original thinking. For someone whose academic habitus consists mainly in “organising the literature and writing it up nicely,” whose relation to theory and method is essentially nominal and procedural, the same tool is a direct substitute. In this second case the model can match and often surpass the writer in precisely those dimensions that current institutions care about – volume, fluency, formal correctness – while remaining indifferent to the question of whether anything genuinely new is being said. The uncomfortable truth is that a large share of existing academic text is already close enough to the model’s statistical centre that it can be reproduced without much loss, which means the average academic is arguing not for the necessity of their intellect but for the sentimental value of their effort.

From this point on, the question ceases to be one of “efficiency gains” and becomes a question about what will count as “valuable labour” in the future of academia. In an ideal configuration, once AI has taken over the routine tasks, the baseline condition for remaining in academic work would be to truly internalise theoretical and methodological depth, to build connections that a model cannot infer and to take conceptual and empirical risks that no system optimised for statistical plausibility will recommend. In such a scenario the painstaking work that has long been treated as secondary – mastering difficult languages, reading canonical texts in detail, opening archival materials, reconstructing forgotten debates, carrying out critical analyses of complex datasets – would become more valuable, not less. This is precisely where the model’s assistance is weakest: it can support but not replace the historically situated, context-sensitive judgement of a mind that lives inside the text and its world. Read in this way, the current crisis contains not only a threat but also a unique chance to shake off the dead weight that has settled on academic life.

The same holds for ideas that have been labelled “marginal,” “extreme” or simply “outdated” and pushed out of mainstream debate. Much that has been buried under the sediment of paradigm and fashion may acquire new relevance precisely because it does not fit comfortably into the statistical averages on which

the model relies. While AI circulates safely within the centre of its training data, the researcher who works patiently on the periphery – on forgotten authors, local controversies, unfashionable traditions – inhabits a conceptual space in which genuine leaps are still possible. Systems can suggest “similar articles”; they cannot replace the experience of following a line of thought across languages, periods and genres, tracking where it breaks, mutates, or is violently cut off. If academia reads the crisis correctly, the path back to a space where creative and fearless ideas collide may well run through exactly these neglected zones.

What destroys this optimism, however, is the present configuration of market, power and epistemic communities that already shapes academic life. Publication pressure, ranking industries, metric obsession and capital-heavy journal and platform networks have for years systematically marginalised original and disturbing ideas. Work that seriously disrupts paradigms or unsettles entrenched authorities tends to be filtered out in peer review, sidelined by journal policies or starved of funding long before it can reconfigure a field. In such a landscape, the automation of average academic labour does not automatically raise the value of deep, original work; it can just as easily tighten the performance regime that governs it. Original thought, when forced into the mould of quantifiable outputs, risks being transformed into yet another kind of high-speed, high-stress labour that must remain within “acceptable” conceptual boundaries calibrated by algorithms and metrics.

For these reasons, the automation of average academic labour and the selection of original work cannot be narrated as a straightforward story of progress. What we are facing is a crisis with two opposed potentials. On one side lies the prospect that, as routine text production is cheapened, tens of thousands of papers effectively lose their justification, and the work of the average academic becomes indistinguishable from AI output. On the other lies the possibility that a minority of researchers who have truly learned their theories and methods, who have forged a real relation to concepts and texts, can use AI as a multiplier for their own thinking and bring back to life those domains – old texts, marginal ideas, slow and difficult forms of inquiry – that the current regime has written off. Whether the crisis turns into a chance to clear the dead ground or into yet another layer of control does not depend on the technical capacity of AI or the goodwill of individual academics. It depends on whether the academic field continues to be organised under the pressure of markets and power, or whether it can rebuild its own norms in a way that genuinely rewards risk, depth and originality.

6. The Structural Kinship Between Social Media and AI: Where the Real Risk Lies

Most of the debate around generative AI in academia keeps circling the same narrow ring: laziness, cheating, plagiarism, “fake references,” students outsourcing essays to the model, academics relying on “ready-made” text. Guidelines, committee reports and “ethics” papers largely fixate on whether the text is formally honest – who wrote it, who did not – and almost never ask what infrastructure makes that text possible in the first place, what kind of power logic that infrastructure embodies, or how deeply it reshapes the academic subject. In this framing, AI is treated as if it were a late intruder into an otherwise sound academic system, a new cheating device that must be fenced off (Bittle & El-Gayar, 2025, p. 296; Lodge, 2024, pp. 3–6). The claim here is different: the real risk of AI is not “cheating” or “taking the easy way out.” The real risk is that the algorithmic logic that has been operating for years in social media now moves directly into thinking and academic decision-making itself.

For that reason, it is necessary to look at what has already happened in social media before deciding where to place AI on the same line.

6.1. Algorithmic Empire The Social Media Case

The decisive shift in social media is the reversal of the positions of user and customer. On advertisement-based platforms the user is no longer the customer; the user is what is sold. What is traded is not primarily a service but the user’s time, attention and behavioural patterns. Interface elements we have learned to think of as neutral – infinite scroll, like buttons, notification sounds, red badges – are not innocent design choices; they are behaviour-engineering devices built to keep attention locked on the platform for as long as possible (Bhargava & Velasquez, 2020, pp. 321–325).

The psychological side is well documented. Social media operators construct small reward loops using classic variable reinforcement schedules; notifications and micro-interactions repeatedly trigger the

dopaminergic system and generate habit and dependence (Griffiths, 2018, pp. 74–76). Each glance at the screen becomes a pull on a probability machine: perhaps this time there is a new like, a new message, a new argument. The mechanism is essentially the slot-machine logic of gambling, translated into a digital environment and wrapped in the language of “connection” and “community.”

At the economic level this structure has become the backbone of surveillance capitalism. User experience is continuously translated into a stream of data; every search, every pause over a video, every click, every abandoned comment becomes part of what has been called “behavioural surplus.” That surplus is then processed through prediction models and turned into behavioural prediction products sold on futures markets (Zuboff, 2019, pp. 94–100). When a user scrolls through the feed and believes they are “seeing what their friends share,” what they are in fact inhabiting is an environment engineered in line with the economic and political priorities of advertisers and platform owners.

The political consequences of this infrastructure are obvious. The Cambridge Analytica episode forced into public view how social media data can be used to construct psychometric profiles of voters and to target them with finely tuned political messages. Voter data were harvested at scale, psychological profiles were built, and segments were flooded with bespoke fear- and resentment-oriented content. This was not an accident at the margins of the system; it was the political use of a logic already normalised in commercial practice.

The Gaza example shows the same logic laid bare in one of its harshest forms. In the period after 7 October 2023, the systematic removal, down-ranking and restriction of pro-Palestinian content was not just an impression voiced by activists. The over-reliance on automated tools for Palestine-related content and the use of vague “terror” categories in Meta’s enforcement practices have been documented by human rights organisations (Human Rights Watch, 2023, pp. 1–6). Videos of civilian deaths, testimonies from the ground, simple messages of solidarity were repeatedly deleted; users learned by trial and error which slogans, images and words “do not pass” and will get accounts limited. What emerges is not classical censorship alone but a practice of algorithmic suffocation: the feed itself becomes the mechanism that decides which death counts as news and which atrocity is compatible with the “flow.” The algorithm, aligned with corporate interests and state pressure, decides what is visible at all.

The result is clear. Social media entered the scene under slogans of “free expression,” “direct democracy” and “independent news.” It now functions as an algorithmic empire in which blocks of power and capital interlock. This empire not only determines who can speak; it also determines who never appears in the first place.

6.2. Generative AI: the Same Pattern, Now Embedded in Thinking

Placing AI in this picture means recognising that large language models are not merely “programs that generate text.” At the level that matters, they do something structurally similar to social media algorithms: they select, from within a given universe of data, combinations that appear “plausible” and present them to the user as if they were neutral responses. The difference is this: while social media determines which content we meet, AI systems are increasingly shaping which questions we see as reasonable, which conceptual distinctions we treat as “natural,” which argument paths feel intellectually legitimate.

Many academics, often without fully noticing, now begin their research roughly like this: they give the model a broad topic area; ask for “key debates”; let the model’s list of headings steer their literature search; then test possible chapter structures and argument skeletons through the same interface. At that point the model effectively controls the first contact with the literature. The initial frame for what counts as the “main line” of discussion – which authors, which terms, which controversies – is drawn by the system. The researcher often accepts that frame as a natural starting point. The dangerous element here is that the model continuously drags every sentence back towards the statistical middle of its training data. Sharp concepts are rounded off; unsettling examples are dropped; politically risky terms are replaced by safer, more neutral ones. On social media the algorithm quietly removes “too extreme” or “too disturbing” content from the feed and leaves the user with a calmer world. In AI-assisted writing, something similar happens at the level of language: expressions at the statistical edges, radical conceptual pairings, unusual examples are “smoothed out” in favour of phrasing that better fits the learned centre. The user experiences this as moderation, style improvement or “academic tone.” In

reality, the boundaries of their thinking space are being narrowed step by step.

At this point the social media analogy stops being a metaphor and becomes a structural statement: social media algorithms colonise attention; generative AI systems move toward colonising problem-formation and argument-formation. One decides what is worth looking at; the other edges toward deciding what is worth thinking about at all. In this context, “colonise” does not mean a total replacement of human agency, but a gradual structuring of the field of possible options such that certain questions, examples and positions systematically recede from view while others are continuously reinforced.

6.3. Academia’s Contradiction: Internalising Social Media, Panicking Over AI

This brings us to a concrete contradiction in academic practice. On one side, social media use is almost completely unproblematised. Universities and journals now routinely place follower counts, engagement statistics and “impact” graphs alongside citation numbers. Many academics proudly present their Twitter/X, Instagram or LinkedIn visibility in CVs and promotion files. All of this happens in full knowledge of the fact that the attention economy generates dependence, manipulation and surveillance, and that social media platforms have developed systematic practices of suppression and suffocation around politically sensitive content (Bhargava & Velasquez, 2020, pp. 321–329; Human Rights Watch, 2023, pp. 1–6; Zuboff, 2019, pp. 94–100). Yet social media has become an unquestioned background to academic life.

On the other side, the same academic environment reacts to AI with harsh regulations, threats of disciplinary sanctions and detailed prohibition lists. Discussion papers, review articles and policy texts almost without exception code generative AI as a tool that threatens “academic integrity” and must be controlled accordingly (Bittle & El-Gayar, 2025, p. 296; Lodge, 2024, pp. 3–6). The issue is again pulled down to the level of individual morality: students must not cheat, academics must not plagiarise, nobody should hide their use of the model. The algorithmic infrastructure itself, as in social media, slips back out of view.

This is not a minor inconsistency; it is a deeper defence reflex. Social media fits the current market logic of academia: it offers visibility, brand value, “science communication,” rapid circulation of lightweight content, and an easy way to plug universities and researchers into wider markets. Generative AI does the opposite: it exposes how easily the existing system can be automated and how hollow much of its output has become. When standard review articles, interchangeable quantitative studies and endlessly recycled conceptual schemas can be produced by a model in a few prompts, the problem is not the model; the problem is the existing production regime. Panic in this sense is also an attempt to smash the mirror.

From this angle, the real risk of AI for academia is not primarily that it might produce “a large population of unemployed academics.” The deeper risk is that it will allow the thinking space of the academic subject to be colonised algorithmically at a level even deeper than social media. The researcher can slowly turn into a figure whose main work is no longer to wrestle with a real literature, but to decide which prompts to feed into the system. The model decides which sources are read first, which conceptual distinctions count as “the main debate,” which examples are “too political” to be worth pursuing. In such an environment, original, uncomfortable, paradigm-breaking ideas start from a position of structural disadvantage simply because they occupy very little space in the model’s statistical universe.

The conclusion follows from this. Focusing the conversation on plagiarism and laziness puts the risk in the wrong place. Internalising the algorithmic empire of social media while treating AI as merely a cheating device is a way of masking the crisis, not resolving it. The crisis is that the algorithmic logic has already entered the core of academic life. Generative AI has only made that logic more visible, faster and deeper. If discussion remains framed at the level of cheating and effort, academic institutions will, in the coming years, end up legitimising a second wave of algorithmic colonisation after social media – this time not only over how we appear in public, but over what we are able to think.

7. Conclusion

The argument developed here has deliberately refused to treat generative AI as a foreign body that suddenly corrupts an otherwise healthy academic ecosystem. Instead, AI has been read as a technology that makes visible – and accelerates – crises that were already structurally present: the industrialisation

of academic labour, the hollowing out of originality, and the quiet consolidation of algorithmic power over what counts as reasonable, publishable and thinkable. If the debate stays confined to plagiarism, laziness and exam cheating, these crises will remain misdiagnosed as individual moral failures rather than as symptoms of a much deeper transformation.

At the level of labour, the crucial fracture line is not between “those who use AI” and “those who do not,” but between average academic work and genuinely original, theoretically grounded inquiry. The kind of production that consists in summarising familiar literatures, filling standard templates and endlessly rephrasing the same argument is, by its very structure, open to automation. Generative models can already perform much of this work faster, cheaper and in a more formally acceptable register. This means that large segments of academic output are de facto replaceable, whether institutions admit it or not. At the same time, precisely the same tools can act as powerful multipliers for a minority of researchers who have actually internalised their theories and methods, who can use AI to clear away routine tasks and concentrate on the sort of conceptual and empirical risks that no statistical engine will take on their behalf. In principle, this could become a selection mechanism in favour of depth and originality; in practice, it collides with publication regimes, markets and epistemic communities that remain structurally hostile to work that genuinely disturbs the status quo.

The second fracture line lies in the kinship between AI and social media. Both are built on the same algorithmic logic: selecting, ranking and shaping human behaviour in accordance with metrics and optimised outcomes. Social media colonises attention by deciding what users see; generative AI moves toward colonising problem formation and judgement by steering what researchers ask, which distinctions feel “natural” and which lines of argument appear “balanced.” The striking contradiction is that academia has largely normalised social media – with all its documented manipulations and dependencies – while reacting to AI with alarmist regulations framed almost entirely in terms of academic integrity. That contradiction is not accidental. Social media supports the existing market logic of visibility and branding; AI threatens to expose how much of academic production can be mass-produced without real thought, and how far the field has already drifted into algorithmic governance.

At the level of institutional design, this diagnosis points toward at least three concrete directions. First, appointment, promotion and funding criteria would have to loosen their dependence on sheer publication counts and journal metrics, and give greater weight to slow, risky and conceptually demanding work – for instance through strengthened peer-review panels that read a small number of substantial outputs in depth rather than scanning long lists of articles. Second, the use of synthetic data and synthetic participants requires clear epistemic boundaries: their role could be restricted to instrument piloting, scenario exploration and methodological testing, while empirical claims about real populations would remain tied to encounters with actual subjects and contexts, with full disclosure of any synthetic components. Third, universities would need to review their own reliance on platform metrics and commercial analytics systems, so that AI-generated conformity does not become structurally advantageous in evaluation processes that ostensibly reward originality.

Under these conditions, the real question is whether academic institutions will allow this second wave of algorithmic power to settle quietly into the heart of their practices, or whether they will take the opportunity to rebuild their own norms. If AI is treated simply as a faster way to manufacture “acceptable” texts, the outcome will be a further devaluation of average academic labour and a tighter coupling of even the best work to metrics and models. If, instead, the crisis is read clearly, it can force a reorientation toward forms of labour that are not easily automated: slow reading, archival excavation, risky conceptual innovation, engagement with marginal and forgotten traditions, genuine theoretical and methodological learning. The choice is stark. Either academia becomes a willing province of the algorithmic empire it helped to normalise, or it uses the shock of AI to recover its own capacity for judgement, conflict and originality.

References

- Alkaissi, H., & McFarlane, S. I. (2023). Artificial hallucinations in ChatGPT: Implications in scientific writing. *Cureus*, 15(2), e35179. <https://doi.org/10.7759/cureus.35179>
- Andersen, J. P., Degn, L., Fishberg, R., Gravensen, E. K., Horbach, S. P. J. M., Kalpazidou Schmidt, E., ... Sørensen, M. P. (2025). Generative artificial intelligence (GenAI) in the research process: A

- survey of researchers' practices and perceptions. *Technology in Society*, 81, 102813.
- Besançon, L., Cabanac, G., Labbé, C., & Magazinov, A. (2024). A great opportunity or a Pandora's box? A survey on the use of large language models in academic writing. *Research Integrity and Peer Review*, 9, 4.
- Bhargava, V., & Velasquez, M. (2020). Ethics of the attention economy: The problem of social media addiction. *Business Ethics Quarterly*, 31(3), 321–359.
- Biagioli, M., & Lippman, A. (Eds.). (2020). *Gaming the metrics: Misconduct and manipulation in academic research*. Cambridge, MA: MIT Press.
- Bin-Nashwan, S. A., Sadallah, M., & Bouteraa, M. (2023). Use of ChatGPT in academia: Academic integrity hangs in the balance. *Technology in Society*, 75, 102370. <https://doi.org/10.1016/j.techsoc.2023.102370>
- Bittle, K., & El-Gayar, O. (2025). Generative AI and academic integrity in higher education: A systematic review. *Information*, 16(4), 296.
- Çakir, A., Kuyurtar, D., & Balyer, A. (2024). The effects of the publish or perish culture on publications in the field of educational administration in Türkiye. *Social Sciences & Humanities Open*, 9, 100817. <https://doi.org/10.1016/j.ssaho.2024.100817>
- Cotton, D. R. E., Shipway, J. R., & Van der Velden, L. (2024). “Chatting and cheating”: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2), 227–238.
- Eke, D. O. (2023). ChatGPT and the rise of generative AI: Threat to academic integrity? *Journal of Responsible Technology*, 13, 100060.
- Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 167–194). Cambridge, MA: MIT Press.
- Gravel, J., D'Amours-Gravel, M., & Osmanlliu, E. (2023). Learning to fake it: Limited responses and fabricated references provided by ChatGPT for medical questions. *Mayo Clinic Proceedings: Digital Health*, 1(3), 226–234. <https://doi.org/10.1016/j.mcpdig.2023.05.004>
- Grech, V. (2022). Publish or perish, information overload, and journal impact factors – A conflicting tripod of forces. *Saudi Journal of Anaesthesia*, 16(2), 204–207. https://doi.org/10.4103/sja.sja_632_21
- Griffiths, M. D. (2018). The evolution of the “components model of addiction” and the need for a confirmatory approach in conceptualizing behavioral addictions. *Dusunen Adam: The Journal of Psychiatry and Neurological Sciences*, 31(3), 179–184.
- Human Rights Watch. (2023). “Meta's broken promises”: Systemic censorship of Palestine content on Instagram and Facebook. New York, NY: Human Rights Watch.
- Ioannidis, J. P. A., Klavans, R., & Boyack, K. W. (2018). Thousands of scientists publish a paper every five days. *Nature*, 561(7722), 167–169.
- Korinek, A. (2023). Generative AI for economic research: Use cases and implications for economists. *Journal of Economic Literature*, 61(4), 1275–1315.
- Lechien, J. R., Briganti, G., & Vaira, L. A. (2024). Artificial intelligence and ChatGPT in otolaryngology research and clinical practice. *American Journal of Otolaryngology*, 45(1), 103892.
- Lodge, J. M. (2024). *The evolving risk to academic integrity posed by generative artificial intelligence: Options for immediate action*. Canberra, Australia: Tertiary Education Quality and Standards Agency.
- Miller, K. (2025, July 29). Social science researchers use AI to simulate human subjects. *Stanford*

- Report*. <https://news.stanford.edu/stories/2025/07/ai-social-science-research-simulated-human-subjects>
- Morrish, L. (2019). *Pressure vessels: The epidemic of poor mental health among higher education staff* (HEPI Occasional Paper 20). Oxford, UK: Higher Education Policy Institute. <https://www.hepi.ac.uk/reports/pressure-vessels-the-epidemic-of-poor-mental-health-among-higher-education-staff/>
- Nordling, L. (2023). How ChatGPT is transforming the postdoc experience. *Nature*, 622(7983), 655–657.
- Orduna-Malea, E., & Cabezas-Clavijo, Á. (2023). ChatGPT and bibliometrics: Reflections on artificial intelligence and scientific communication. *Profesional de la Información*, 32(2), e320204.
- Raitskaya, L., & Tikhonova, E. (2024). Appliances of generative AI-powered language tools in academic writing: A scoping review. *Journal of Language and Education*, 10(1), 1–27.
- Resnik, D. B. (2024). Vulnerable subjects. In *The ethics of research with human subjects* (pp. 293–329). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-031-82757-0_8
- Rossi, L., Harrison, K., & Shklovski, I. (2024). Synthetic data and the problems of LLM-generated evidence in social research. *Sociologica*, 18(2), 145–168.
- Sanchez-Ramos, L., Lin, L., & Romero, R. (2023). Beware of references when using ChatGPT as a source of information to write scientific articles. *American Journal of Obstetrics and Gynecology*, 229(3), 356–357. <https://doi.org/10.1016/j.ajog.2023.04.004>
- Sebo, P. (2024). ChatGPT in medical research: The risk of fabricated data and references. *Family Practice*, 41(3), 494–496.
- Shore, C., & Wright, S. (1999). Audit culture and anthropology: Neo-liberalism in British higher education. *Journal of the Royal Anthropological Institute*, 5(4), 557–575.
- Shore, C. (2008). Audit culture and illiberal governance: Universities and the politics of accountability. *Anthropological Theory*, 8(3), 278–298.
- Shrestha, P., Krpan, D., Koaik, F., Schnider, R., Sayess, D., & Binbaz, M. S. (2024). Beyond WEIRD: Can synthetic survey participants substitute for humans in global policy research? *Behavioral Science & Policy*, 10(2), 26–45. <https://doi.org/10.1177/23794607241311793>
- Striphas, T. (2015). Algorithmic culture. *European Journal of Cultural Studies*, 18(4–5), 395–412.
- Thorp, H. H. (2023). ChatGPT is fun, but not an author. *Science*, 379(6630), 313.
- van Dijck, J. (2013). *The culture of connectivity: A critical history of social media*. Oxford, UK: Oxford University Press.
- Watermeyer, R., Phipps, L., Lanclos, D., & Knight, C. (2024). Generative AI and the automating of academia. *Postdigital Science and Education*, 6(2), 446–466.
- Wilsdon, J., Allen, L., Belfiore, E., Campbell, P., Curry, S., Hill, S., ... Thelwall, M. (2015). *The metric tide: Report of the independent review of the role of metrics in research assessment and management*. Bristol, UK: Higher Education Funding Council for England.
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. New York, NY: PublicAffairs.

Araştırma Makalesi

Killing the Average, Amplifying the Few: Generative AI, Social Media Patterns, and the Structural Crisis of Academic Work

Ortalamayı Öldürmek, Azınlığı Güçlendirmek: Üretken Yapay Zeka, Sosyal Medya Örüntüleri ve Akademik Çalışmanın Yapısal Krizi

Deniz HASANÇEBİ

Araş.Gör. Dr., Ankara Hacı Bayram Veli Üniversitesi

Edebiyat Fakültesi

d.hasancebi@gmail.com

<https://orcid.org/0000-0002-8541-4063>

Genişletilmiş Özet

Bu makale, üretken yapay zekâ tartışmalarında baskın hale gelen etik merkezli çerçeveyi sorgulayarak başlar. Son yıllarda yayımlanan yönergeler, makaleler ve raporların büyük bir kısmı, ChatGPT ve benzeri modelleri tembellik, kopya, intihal, sahte kaynak kullanımı ve “ödevini yapay zekâyâ yaptıran öğrenci” panikleri etrafında tartışmaktadır. Bu söylemde yapay zekâ, dışarıdan gelmiş bir “hile aracı” olarak kodlanır; sanki iyi işleyen, özgün ve adil bir akademik sistem, bu araçların gelişyle birlikte bozulmuş gibi konuşulur. Makalenin temel iddiası tam tersidir: Yapay zekâ, zaten çürümüş ve büyük ölçüde otomasyona açık hale gelmiş bir akademik üretim rejiminin içindeki krizi görünür kılmaktadır. Dolayısıyla sorun bireysel ahlâk veya niyet problemi değil, yapısal bir emek ve iktidar meselesidir.

Bu çerçevede ilk odak noktası, “ortalama akademik emek” olarak adlandırılabilir üretim biçimidir. Birçok alanda akademik üretimin büyük kısmı, belirli düşünürlerin veya teorilerin özetlenmesinden, standardize edilmiş makale iskeletlerinin doldurulmasından ve benzer argümanların küçük varyasyonlarla tekrarlanmasından ibarettir. On binlerce tez ve makale, her yıl aynı kaynaklara atıf yaparak, aynı kavramsal çiftler etrafında dönerek, aynı metodolojik prosedürleri yeniden sahneye koyarak çoğalmaktadır. Makale, bu üretim tarzının en başından beri otomasyona açık olduğunu vurgular: Eğer işlev, esasen derleme, özet ve kalıp doldurma ise, bu işi insandan daha hızlı ve ucuz biçimde yapabilen bir algoritmanın devreye girmesi şaşırtıcı değildir. Bu noktada yapay zekâ, var olan bir boşluğu doldurmakla kalmaz; boşluğun kendisini, yani ortalama akademik emeğin ne kadar düşük özgünlük düzeyine sahip olduğunu da ortaya çıkarır.

Metin, bu tespiti soyut bir teşhis olarak bırakmak yerine, üretken yapay zekânın fiilî kullanım biçimlerine eğilir. Farklı ülkelerde yürütülen anketler ve vaka çalışmalarının ortak bulgusu, araştırmacıların büyük çoğunluğunun bu araçları literatür taraması, metin taslağı üretme, özet çıkarma, dilsel cilalama ve temel veri işleme gibi mikro işlerde kullandığıdır. Bu işlerin ortak özelliği, yüksek hacimli fakat düşük yaratıcılık gerektiren, “görünmeyen emek” kategorisine giren görevler olmasıdır. Makale, bu noktanın altını kalın biçimde çizer: Bugün YZ’nin en yoğun yerleştiği alan, teori kurma veya paradigma kırma değil, tam anlamıyla ortalama akademik emeğin kalbidir. Bu da ortalama akademisyenin yaptığı iş ile iyi yönlendirilmiş bir modelin çıktısı arasındaki farkı hız, maliyet ve dilsel özgünlük boyutlarına indirgemektedir.

İkinci önemli boyut, üretimin makro ölçekte zaten endüstriyel ritme bağlanmış olmasıdır. “Her beş günde bir makale” çıkaran binlerce araştırmacının bulunduğu, yayın sayısının temel performans göstergesi hâline geldiği ve bilimsel üretimin “yayın fabrikası” mantığı ile çalıştığı bir ortamda, yapay zekânın devreye girmesi yeni bir rejim kurmaktan ziyade mevcut rejimi hızlandırır. Makale, bu noktada üretken YZ’nin yalnızca tekil araştırmacıların verimliliğini artıran bir araç olmadığını; aynı zamanda metriklere ve sıralama endüstrisine dayalı akademik alanın yapısına yerleşen bir otomasyon vektörü olduğunu savunur. Yazma, değerlendirme, öğrenci geribildirim ve bürokratik raporlama gibi alanlarda emeği görünmezleştiren ve yönetsel denetimi artıran bu süreç, ortalama akademik emeğin giderek daha ucuz, daha kolay ikame edilebilir ve daha az saygı gören bir forma bürünmesine yol açar.

Makalenin dördüncü bölümünde geliştirilen tez, bu çerçevede keskinleştirilir: Yapay zekâ, bir yandan ortalama akademisyenin yaptığı işin büyük bölümünü otomatikleştirerek, onu görünmez ve değersiz hâle getirirken; diğer yandan derin kuramsal donanıma ve özgün kavrayış gücüne sahip az sayıdaki araştırmacı için güçlü bir çarpan etkisi üretmektedir. Burada kritik fark, teoriyi ve yöntemi gerçekten içselleştirmiş, kendi sorularını kendisi kuran araştırmacı ile, sadece literatürü toparlayıp düzgün yazmayı “işin özü” zanneden profil arasındadır. Birincisi için YZ, yorucu ama yaratıcı olmayan adımları devralan bir asistanlar ordusuna dönüşür; ikincisi için ise doğrudan ikame edici bir teknolojidir. Makale, bu ikiliği bir “doğal seçim” retoriği ile basitleştirmek yerine, bunun yapısal bir ayrışma olduğunu vurgular: YZ, ortalama akademik emeğin zaten var olan boşluğunu görünür kılarken, özgün düşünceye dayalı emeği daha da zorunlu hâle getirir; ancak bunun gerçekleşmesi, akademik alanın kendi normlarını nasıl yeniden kuracağına bağlıdır.

Çalışmanın altıncı bölümünde, bu emek tartışması yapay zekâ ile sosyal medya arasındaki yapısal benzerlik üzerinden daha derin bir düzeye taşınır. Sosyal medya, dikkat ekonomisine dayalı algoritmalarla, hangi içeriğin görünür olacağına, hangi görüntülerin ve cümlelerin akışta kalacağına karar veren bir imparatorluk kurmuştur. İnsanların neye öfkeleneyeceği, neye alışacağı, hangi acıya karşı duyarsızlaşacağı büyük ölçüde bu algoritmik mantık tarafından belirlenir. Makale, üretken yapay zekânın bu mantığı düşünmenin içine taşıdığını savunur: Sosyal medya dikkati kolonize ederken, YZ problem kurma ve karar alma alanını kolonize etmeye yönelir. Artık hangi sorunun “makul”, hangi kavramsal ayrımın “doğal”, hangi tonun “akademik” olduğuna dair sezgi, modelin istatistiksel ortalamasına göre şekillenmektedir. Bu, sadece teknik bir kolaylık değil, epistemik özerkliğin kaybı anlamına gelir.

Burada açığa çıkan çelişki özellikle önemlidir: Akademik alan, sosyal medyanın bağımlılık, manipülasyon ve sansür üreten algoritmik rejimini büyük ölçüde içselleştirmiş; pek çok kurum sosyal medya görünürlüğünü bir performans göstergesi hâline getirmiştir. Aynı alan, yapay zekâ söz konusu olduğunda ise birdenbire sert yönetmeliklere, ağır disiplin cezalarına ve “her türlü kullanım suçtur” tonuna yönelmektedir. Makale, bu çelişkinin basit bir tutarsızlık değil, piyasa mantığıyla uyumlu olan algoritmalara gösterilen hoşgörü ile, sistemi çıplaklaştıran algoritmalara yönelen savunma refleksi arasındaki fark olduğunu öne sürer. Sosyal medya, görünürlük ve marka değeri ürettiği için tolere edilirken; yapay zekâ, akademik üretimin büyük bölümünün “makine tarafından da yapılabildiğini” gösterdiği için tehdittir.

Sonuç olarak makale, yapay zekânın akademi için gerçek riskinin intihal, kopya veya tembellik olmadığını; asıl tehlikenin algoritmik mantığın akademik alanın çekirdeğine yerleşmesi olduğunu savunur. Eğer tartışma kopya paniği ve sınav güvenliği çerçevesine sıkıştırılırsa, hem ortalama akademik emeğin sessiz tasfiyesi hızlanacak, hem de özgün düşüncenin açtığı imkân penceresi kapanacaktır. Buna karşılık akademi, krizi doğru okursa, yapay zekâ şoku, özgün ve riskli çalışmaların gerçekten değer kazandığı; eski metinlerin, marjinal fikirlerin ve yavaş düşünmenin yeniden merkezî hâle geldiği bir dönüşüm için tetikleyici olabilir. Makalenin vardığı nokta, bu ikili imkânın teknik kapasitelerden çok, akademik alanın piyasa ve iktidar baskısına karşı kendi normlarını yeniden kurma cesaretine bağlı olduğudur.