

Yes, but can it fact-check too ...?

Reflections on AI in Journalism Education

Frederik Marain

Yes, but can it fact-check too ...?

Frederik Marain teaches Research, Data Journalism, Economy, and International News Topics at AP University of Applied Sciences and Arts, Antwerp, and VUB University, Brussels. He was a Business and Economics editor at the Flemish newspaper De Standaard, and chief editor and Director Media at internet media company Planet Internet Belgium. Frederik obtained a Master's degree in Philosophy and a Master's degree in Economics.

AP HOGESCHOOL ANTWERPEN
ejta.eu

Published by the European Journalism Training Association (EJTA),
with support from AP University of Applied Sciences and Arts Antwerp.
ejta.eu
ap.be

Copyright © 2026 by Frederik Marain and EJTA

Graphic design and cover design/illustration by Serge Van Camp

ISBN 978-9-09-043130-7

Introduction 3
Scope, limitations and caveats 3
How AI was and was not used for this paper 6
Acknowledgements 7
But can it fact-check too ...? 8
Why fact-checking? 8
"Learning from language alone" 9
Fact-checking with AI: An experiment and a prototype 14
An experiment 15
Taking it two steps further: a prototype 18
Yes, it can 23
The big problem for journalism education 24
AI and higher education 24
The future of our schools 26
Journalism education: More bad news 32
The Cloister and the Starship. A radical solution? 40
Inarticulable expertise and accountability 42
Bibliography 46
To conclude: An inconvenient experiment 51

Introduction

The short answer to the question in the title of this essay is: Yes, it can; better than most undergraduate journalism students; as well as a lot of graduate students and professional journalists.
The finding that Artificial Intelligence (AI) can do what is considered a particularly complex and essential part of journalism should drive us, journalism teachers, crazy. This is a time for marvel, but also for panic, for a correct diagnosis of the problem, and for reflections on a solution.

Scope, limitations and caveats

But writing about AI these days feels like drinking from three firehoses simultaneously and then trying to spill it out in digestible doses. These reflections, therefore, will have to start with some serious scope limitations and caveats.

First, these reflections are about the integration of AI in ***journalism education***. Not in journalism (although some of the ideas will be relevant for the journalism profession). Not in higher education in general (although the diagnosis will start from the general problem that AI poses for higher education). This scope limitation is not for practical reasons only. More fundamentally, the investigation of the challenges of integrating AI in higher education and most certainly the proposed solutions and policies, should be department-specific and often even course-specific. This was indeed the first in a list of ten recommendations in the *Recommendations for Integrating AI in Journalism Teaching* that the AI and Data Journalism Focus Group of the European Journalism Training Association (EJTA) published in June 2025: "Develop and use ***specific AI-guidelines*** for Journalism departments." (EJTA, 2025) Hollis Robbins, former Dean of Humanities and Special Advisor for Humanities at the University of Utah and a leading voice in the AI-higher education debate, concurs, in an article where she drafts a job description for a university AI czar: "Support discipline-specific course policies rather than one campus rule for every field, because what AI means for a writing seminar, an organic chemistry lab, a clinical rotation, and a data-science capstone are four different problems." (Robbins, 2026) One of the biggest pitfalls for Journalism departments (happening now!) is to wait for general AI-guidelines and rules at the school level.

Second: Yes, we know ***AI is an umbrella term*** for an ecosystem that encompasses Machine Learning, Deep Learning, Generative AI, and Large Language Models (LLMs). For practical reasons, the term "AI" in this paper will almost always be shorthand for the LLM-chatbots that most students and teachers know and use. In the same pragmatic vein, we will, unjustifiedly, mostly gloss over the advances in AI since the release of the first LLM-chatbots.

Third: These reflections will not dwell on the ***ethical, regulatory and other broader societal implications of AI***, or only sideways. To be sure, discussing and reflecting on these implications is the most important mission for journalism students, teachers and journalists. But (1) There is no shortage of accessible articles and books on these topics. One very good introduction is by AP University of Applied Sciences and Arts colleague Wim Casteels: *Het masker van AI. AI-geletterdheid als gids bij maatschappelijke kansen en uitdagingen* (The Mask of AI. AI literacy as a guide to societal opportunities and challenges). The book and the weekly column on LinkedIn with the same title are in Dutch. But the book has an up-to-date source list of 43 pages that is an excellent guide in the literature.

Eminence grise journalist Martin Wolf, chief economics commentator at the Financial Times, has recently written *A guide for the perplexed on AI*, the first in a series of (mostly pessimistic) opinion articles. (Wolf, 2026) First sentence: "How should someone who knows next to nothing about artificial intelligence think about its implications for humanity?" Wolf tries anyway to answer the questions: AI: bubble, boon or bane? His answers, in order: "up to a point, yes, and yes". In an earlier interview with German-American political scientist and journalist Yascha Mounk, Wolf had said: "The idea that it's all under the control of ten or so billionaires who, in my view, to a significant extent are psychopaths, terrifies me." (Mounk, 2026) Still more eminently, even pope Leo XIV has chimed in, with his encyclical *Magnifica Humanitas. On Safeguarding The Human Person In The Time Of Artificial Intelligence.*

"The idea that it's all under the control of ten or so billionaires who, in my view, to a significant extent are psychopaths, terrifies me." *(Martin Wolf)*

Books and articles like these should help us reflect on a second reason why we should be hesitant to vent opinions and launch into discussions about broader societal implications of AI. A lot of these discussions turn out, very understandable but mostly not helpful, to be a form of mood affiliation and feeling, not thinking.
This is understandable because ...

Fourth: We are "Lost in Model Land", as the title of the February 2025 weekly Wayfinder newsletter by Ezra Eeman cries out. Eeman is Strategy & Innovation Director at the Dutch public broadcaster NPO. His Wayfinder newsletter is one of the best fingers on the pulse of technology trends in media. But even he had to admit at some point that he was Lost in Model Land. "The models are undeniably more powerful than ever, but they demand more from us too. ... It seems the AI revolution has turned into an expert-driven arms race." For us, non-expert journalism teachers who are supposed to "make students AI-fluent" (I have no idea what that means; nobody does, really), this is not just an unfortunate side-effect to be remedied with more reading and studying. We don't understand what is happening; let alone that we can predict what will happen. This is something we have never experienced to that extent before.

Fifth: What is, then, ***a right and rational attitude*** for teachers to adopt? I only know there are several, some understandable and useful, others less so. Panic is good and rational. Excitement and awe are good and rational. Somewhere between wait-and-see, indifference, and detachment is understandable and rational under the circumstances. Less understandable, and not conducive to an adequate curriculum development, is the laziness of the "yes-but ... " kind of reactions ("yes-but-it-hallucinates"; "yes-but-six-fingers"; ...). These are outdated or will be soon. An understanding of or at least a feeling for the rate of improvement of AI-technology that we have witnessed and will be witnessing is the qualifying criterium to participate in the debate.

Sixth: This is ***not an academic paper***. In no conceivable way, and partly for the limitations and caveats mentioned above, does it pretend to offer an up-to-date overview of the academic literature on the topic. Only a few peer-reviewed articles were used or cited. Most inspiration and quotes come from what we could call AI-opinion makers. Some of them are academics but, mostly, when I cite them it is in their capacity as thinker about AI in education. Too many arguments and experiments were omitted to make this text a balanced paper; for example, all the challenges for AI were done with ChatGPT, and not repeated with other AI-chatbots. Most omissions were by choice, because I didn't think they would have changed the main argument. These are, as the title says, just reflections.

How AI was and was not used for this paper

Obviously, AI (ChatGPT Plus, paid version) was used to test its fact-checking capabilities. Also, I had a long chat about the topic of skills, competencies and craft in journalism. It was ChatGPT who suggested the idea of the three-layered Journalism diploma. From that discussion, I took one sentence that I used in extenso in this text. ChatGPT was used for specific research questions; for instance to make the list of institutions in the Top 15 Humanities universities that offered Journalism degrees (double-checked).
In the discussion of the philosophical nature of "truth" in AI, I will refer to and quote from a discussion that philosopher and IT-specialist Erwin Keustermans had with ChatGPT.
If Google Translate (as dictionary and idea generator; very seldom for literal translations) and spell-check tools are considered AI-tools, then those were used too. Other than that, AI was not used as a writing-aid. As a result, the writing, if not the ideas, will be clunkier and less polished than they might have been. This is completely in the spirit of these reflections.

AI was not used as a writing-aid. As a result, the writing, if not the ideas, will be clunkier and less polished than they might have been. This is completely in the spirit of these reflections.

Acknowledgements

These reflections are an outcome of discussions in two EJTA Focus Groups: the Eufactcheck Focus Group and the AI and Data Journalism Focus Group. With "eighty Journalism Schools, in thirty countries across Europe, working together to improve journalism education in Europe", EJTA is one of the largest and most active of higher education associations in Europe – not what is sometimes called a "learned society" but a fruitful breeding ground and laboratory for best practices and, specifically for intricate challenges such as the integration of AI in our curricula, a self-help group for the perplexed.

EJTA-director Nadia Vissers first suggested the idea for these reflections. She read first drafts and was instrumental and patient in the entire production process. I cannot thank her enough. Wim Lauwers, head of the Department of Media, Design and IT at AP University of Applied Sciences and Arts Antwerp, provided generous support. EJTA-colleagues and journalism teachers Alexander Maeder and Will Cope read a completed draft and provided encouragement and useful feedback. Christophe Wille, biotechnologist and lecturer at Ghent University, has been a *compagnon de route* from the early days of AI; he also provided encouragement and useful comments.

But can it fact-check too ...?

❧

Why fact-checking?

There are several reasons why it makes sense to focus on fact-checking when we want to reflect on the integration of AI in journalism education.

First, of course, from the primacy of the principles of truthfulness, accuracy and fact-based reporting that are common in all journalism ethics and standard codes, it follows that fact-checking is an ***essential*** part of journalism. It is so much so that most self-respecting newsrooms have institutionalised some form of independent fact-checking function by another journalist than the original author of an article. If it is true that the amount of mis- and disinformation has increased with digitisation and social media, fact-checking becomes even more essential.

One could argue that in this post-modern, relativist, post-truth world the status and primacy of "truth" are diminished or at least contextualised. But as Billiet et al (2018, p.27) note in their report *The battle for the truth. Fake news and disinformation in the digital media world*: "Even if we assume that truthful reporting about the facts is based not only on what a fact represents materially, but on what that fact means in the diverse contexts of meaning in which it occurs, there is still the question as to how much weight should be attributed to the various contexts." As one of the authors of that report writes in a later article, arguing that the post-truth atmosphere is disastrous: "Without sufficient common ground, it is impossible to establish in a society a joint commitment to respect what is in the public interest such as the rule of law, democracy, and human rights. This is why truth really matters."
(Pattyn, 2020, p. 139)

In this respect, it is interesting and perhaps telling that the International Federation Of Journalists, in the latest (2019) update of its *Global Charter of Ethics for Journalists* changed the formulation of its very first article, from "Respect for truth and for the right of the public to ***truth*** is the first duty of the journalist" to "Respect for the ***facts*** and for the right of the public to truth is the first duty of the journalist".

More on truth in a while. For now, suffice it to note that, after all, our subject is called "fact-checking", not "truth-checking".

The second reason why we want to focus on fact-checking when reflecting on the integration of AI in journalism education, is that fact-checking is ***complex***. It requires learning and practising skills and competencies such as advanced (digital) search and research skills, including digital forensics; ability to discern primary from secondary sources and to evaluate credibility and potential bias in primary sources (and in oneself); critical thinking, and, preferably, a lot of knowledge. That is a tall order and journalism teachers will testify that a lot of students underestimate this complexity.

Now, if, as we are about to show, AI can handle at least part of this complexity, we have a potential problem of academic integrity (students cheating), but also, more fundamentally, students not actively learning or practising what we think they should learn and practise.

A third reason to focus on AI and fact-checking is that we know that ***AI is itself a source of misinformation and biased information***. This is indeed still the most used "yes but" argument against AI.

"Learning from language alone"

Which leads us, finally, to a ***philosophical*** digression on the ***relation between AI and facts and truth***. This argument is the most interesting, I believe, because it will give us, more than volumes of technical background would, much-needed context on what AI is and is not, on what it can and cannot mean for us and for our students.

Fact-checking, obviously, is about facts, or "reality"; more often than not, it leads to the assignment of a "truth"-score. The philosophical question we are interested in when we are investigating whether AI can factcheck is how the relation between AI and facts and reality can lead to truth. If ever. Relatedly, we probably will be able to say something useful about hallucinations.

[Note that, in this philosophical context, we will use the term "Large Language Models" (LLMs) instead of the passe-partout "AI" because it is technically more correct.]

A good place to start is the Scottish philosopher David Hume (1711-1776), who argued that the foundational mechanisms of human belief and knowledge – reason and senses – are unreliable and capricious.

Here is a relevant text from *An Enquiry Concerning Human Understanding*:

"What never was seen, or heard of, may yet be conceived; nor is any thing beyond the power of thought, except what implies an absolute contradiction. … But though

our thought seems to possess this unbounded liberty, we shall find, upon a nearer examination, that it is really confined within very narrow limits, and that all this creative power of the mind amounts to no more than the faculty of compounding, transposing, augmenting, or diminishing the materials afforded us by the senses and experience". (Hume, 2011/1748)

In a well-known passage in the conclusion of Book I, Part IV. Of The Sceptical And Other Systems Of Philosophy of his *A Treatise of Human Nature*, Hume describes how all this thinking about thinking has led him to "philosophical melancholy and delirium":

"The intense view of these manifold contradictions and imperfections in human reason has so wrought upon me, and heated my brain, that I am ready to reject all belief and reasoning, and can look upon no opinion even as more probable or likely than another. Where am I, or what? From what causes do I derive my existence … I am confounded with all these questions, and begin to fancy myself in the most deplorable condition imaginable, invironed with the deepest darkness, and utterly deprived of the use of every member and faculty." (Hume, 1978/1739, p. 268-9)

Luckily, Hume can wrestle himself out of this state of delirium: "Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices to that purpose, and cures me of this philosophical melancholy and delirium ... I dine, I play a game of backgammon, I converse, and am merry with my friends." (Hume, 1978/1739, p. 269)

In an article, *On Large Language Models' Delirium (with Hume and Foucault)*, Dutch philosopher and professor of political science Eric Schliesser argues that this delirium is a perfect term to describe what LLMs are and do. Rather than "understand LLMs as bull-shitters in the technical sense of lacking concern with truth" it is precisely because they are so committed to reason that they sometimes can end up in strange places. It is "(t)he kind of madness ... (that follows from) ... a certain intense commitment to reason or reasoning by which one ends up in an irrational or unreasonable place despite a (to quote Foucault) "quasi-conformity" to reason." (Schliesser, 2023)

"My hypothesis is that we should treat Chat GPT and its sibling LLMs as always being on the verge of the functional equivalent state of delirium." This delirium is "functionally equivalent to Humean delirium" but with important differences.

Hume knew when he had to stop thinking and go play a game of backgammon with his friends; LLMs don't have this common sense. Since they are at "constant risk of drifting away from reality while seeming to follow rational methods of reasoning" they are "delirious without knowing it". (Schliesser, 2023) This is, by the way, a much sharper way to formulate the hallucination problem.

So, can we use LLMs to detect falsehoods or, somehow equivalent but better still, to develop new knowledge? Is there, in other words, somewhere a link between the sentences of an LLM and reality? Strictly speaking: No! Why? Because LLMs "learn from language alone". But, importantly, this is the same vulnerability as humans have.

The "learning from language alone" quote is from a long chat philosopher and IT-specialist Erwin Keustermans had with ChatGPT on the semantics, epistemology, and ontology of LLMs. (link to the complete chat available in the bibliography).

Keustermans' first prompt in the chat was to challenge ChatGPT with the sentence "a car is apple". It's the kind of games philosophers love to play.

The sentence feels wrong because it is semantically not well-formed. LLMs generally seem reasonable because they learn semantic constraints; that is: they learn to speak properly. But they do so "without explicit ontology". Ontology is the branch of philosophy that studies the nature of being, existence, and reality (in the grand scheme of things: what we are concerned with when we do fact-checking). In Hume's and Schliesser's terms, "without explicit ontology" was "detached from reality".

To clarify and dive deeper in what is happening, take the semantically well-formed sentence (suggested by ChatGPT): "The sun orbits the earth."

This is relevant for our factcheck-question because in the 16th and into the 17th century, based on all available texts and ***on all observable facts***, this sentence was "true" in every conceivable way. Then, in 1543, Copernicus provided a new text: the ***mathematical proof*** for a sun-centered universe. Case closed? No! Copernicus' new text was generally disregarded because it was only one of many more texts that, supported by observable facts, told us otherwise. Then, in 1610, Ur-factchecker Galileo turned his telescope towards Venus, provided empirical proof for Copernicus' mathematical model ... and spent the rest of his life under house arrest.
The problem then, as now, is not disinformation, nor even misinformation, but coherent, widespread and socially reinforced misinformation.

The problem then, as now, is not disinformation, nor even misinformation, but coherent, widespread and socially reinforced misinformation.

As ChatGPT says: "Language models are large-scale statistical mirrors of human epistemic practice." (Open AI, Keustermans, 2026) This means that "they inherit biases, blind spots, and falsehoods; they stabilize dominant narratives; they amplify whatever is already socially reinforced." ***Just as we humans do***.

We can now give a more nuanced answer to the "Can they fact-check?" question.

The answer is still: No! Because they have no explicit ontology. ***"A language model cannot privilege truth over frequency."*** Or somewhat more convoluted: "If a falsehood dominates the available data, a language model will not independently generate or privilege competing hypotheses unless there is some trace—however weak—of alternative structure already present." (Open AI, Keustermans, 2026)

Where LLMs can be useful is when (as was the case with geocentrism in the 16th and 17th century) falsehood is not total and there are some traces, even weak ones, of what we would now call "alternative facts". LLMs can then be used to "make epistemic processes faster, not truer" and to "compare many arguments at once, to map hypothesis space more clearly; to make suppressed alternatives legible." (Open AI, Keustermans, 2026)

In the "learning from language alone" insights, there are also wider consequences and lessons relevant for journalism and journalism education.

While the fact that LLMs learn from language alone is the one insight that our students should really understand about AI, we should be careful to avoid turning this into to a "ha, it's just learning from language" perception. Learning from language is no small feat! One would only wish that our students would do at least that.

Of course, journalism students should also learn to "turn their telescope towards Venus". But to be able to do that in a smart way they should have learned and keep learning from language. AI poses a big pitfall here (we see it happening).The speed and fluency with which LLMs can analyse and produce language is deceptive.

Journalism education should counter that deception, now more than before. Learning from language should be messy and slow.

The speed and fluency with which LLMs can analyse and produce language is deceptive. Journalism education should counter that deception, now more than before. Learning from language should be messy and slow.

Towards the end of his conversation with ChatGPT, Keustermans pokes the LLM about the bigger risk. LLMs are not just statistical mirrors of human epistemic practice; they are ***very large scale*** statistical mirrors. The problem with scale is that systemic risks that might emerge from large scale are not visible at small scale; therefore, we do not know that, or when, we are heading for trouble. When Keustermans suggests an analogy with combustion engines (which, at the time, seemed as good as or better an idea than electric engines), ChatGPT gets it and even runs with the metaphor: "The combustion engine polluted the environment. Language models risk polluting the epistemic environment." (Open AI, Keustermans, 2026)

This epistemic pollution would manifest itself in "homogenization of expression, convergence of explanations and loss of minority hypotheses." And then, very quotable: ***"When language becomes too good at explaining the present, it may stop helping us escape it."***

There is a way out. Unlike with the internal combustion engine where, for a long time, we had to consume what the market offered, "humans remain epistemic agents, not just consumers. The danger is not inevitability; it's complacency." The mission is to keep epistemic diversity alive. ChatGPT itself suggests a list that, for our purposes here, reads very much like a programme for teaching journalism and fact-checking: "Resist treating AI outputs as authoritative; preserve adversarial inquiry; value slow, messy, human disagreement; treat fluency as a warning sign, not a credential." (Open AI, Keustermans, 2026)

Fact-checking with AI: an experiment and a prototype

So, in theory, philosophically, ontologically, epistemologically, AIs can't factcheck. We will now demonstrate that they are nevertheless capable of doing so in practice.

What follows is a summary of the results of two experiments, both of which build on the insights and experiences of EJTA's EUfactcheck project, specifically on the factcheck flowchart that was developed for the project.

EUfactcheck is "the fact-checking project of the European Journalism Training Association (EJTA) that supports Journalism schools to build and execute a sustainable curriculum unit on fact-checking, grounded in a European network. Through fact-checking European political claims and trying to tackle mis- and disinformation, we want our students and our public to grow a deeper insight and interest in democratic processes, both on national and European level. EUfactcheck supports fact-based debate in the EU and aims to stimulate media and information literacy." (EUfactcheck.eu About Us, 2019; Vissers, 2020).

The project started in 2017 and was launched in the run-up to the 2019 EU elections. Since then, students from more than thirty European Journalism schools contributed fact checks and blog posts reflecting on fact-checking. As of May 2026, 424 fact checks and 194 blog posts have been published on the website eufactcheck.eu.

Participating students and teachers use the fact check flowchart as a common methodology. The fact check flowchart was developed by AP University of Applied Sciences and Arts teacher in journalism and fact-checking Wouter Frateur and by the author of these reflections. The flowchart does not claim to be the fact-checking methodology for journalists; it is an educational tool in that it helps students understand all the steps to be taken and the pitfalls to be avoided during the fact-checking process. (Frateur and Marain, 2019; download links of the flowchart and of a Train the Trainer presentation in the bibliography)

The three steps of the process are: Analyse claim; Analyse author / source; and Fact check, which includes the final rating, or assignment of a "truth status".

In the Analyse claim step, students learn to understand the checkability of a claim, its components, and any shortcomings. In the Analyse author / source step, students distinguish between the author of a claim and the primary source of a claim; they also identify the author's and the source's capacity and affiliation. Finally, in the Fact check step, students find and contact the primary source, find and contact second experts, check if those experts confirm the claim or not, and rate the claim as either True; Mostly True; Mostly False; False; or Uncheckable.

One of the most useful pedagogical devices of the methodology are the warning lights. These warning lights operate mostly in the first, Analyse Claim, step of the flow, where they help students avoid common pitfalls. A warning light should illuminate when, for example, a clear subject is ambiguous or missing; when a claim is qualitative or, when quantitative, the figure is vague ("about", "many", …), or, especially in the case of percentages, the denominator (100 percent) is missing; when geography is missing or too broad; and, for survey results, when sample size, sampling method, or margin of error are missing. Too many warning lights don't necessarily lead to a "False" rating, but they do result in what we called a "shaky claim."

An experiment

Since the factcheck flow chart is essentially a step-by-step checklist, an AI should be able to perform at least some of the steps. This was the hypothesis of the first experiment, which was low tech, quick, and dirty (the second experiment is anything but).

The "low tech" and "quick", to be sure, are in AI-terms: the experiment used the "simple" AI-chatbot ChatGPT (the paid, Plus version). This means that the experiment is, in principle, accessible to students without technical (AI) knowledge or skills. The experiment is "dirty" in that it doesn't install or reflect on ethical or deontological guardrails.

Following is a summary of some of the results of this experiment with ChatGPT Plus and some claims. The complete chat can be accessed via the link in the bibliography. (OpenAI, Marain, 2025)

We first uploaded the three pdf's of the flow chart and asked if ChatGPT could read them. It could; it understood the general logic of the flow and also immediately grasped the meaning and importance of the warning lights. We then presented some claims and asked to apply the flow chart methodology.

The first claim was one we had used as an exercise in several Train-the-Trainer sessions with EUfactcheck participating teachers. The claim, stated on July 31, 2017 in a Dallas Morning News story, was also fact-checked by the PolitiFact organisation: "Dallas County Sheriff Lupe Valdez said that Texas Department of Public Safety numbers show only 1.6 percent of crimes are committed by unauthorized immigrants."

In the Analyse Claim step, ChatGPT correctly issued two warning lights:

"1. Denominator clarity – The claim does not explicitly state what counts in the "100 %" of crimes (all reported offenses statewide? Certain crime categories?).

2. Geography clarity – Texas is implied but never said outright; a casual reader could miss that limitation." (OpenAI, Marain, 2025)

Next, in the Author / Source analysis of the claim, it correctly identified the author of the claim and the primary source, the Texas Department of Public Safety (DPS) crime statistics, with their capacity and affiliation.

Finally, in the Verification step of the flow chart, it gave a correct link to the DPS's public "Criminal Alien Arrest Data", and immediately noted that this source did "not publish a table that lets you derive the share of all Texas crimes attributable to unauthorized immigrants." This led to a provisional rating of "MOSTLY FALSE — the direction is correct (unauthorized immigrants commit a small share of crime) but the specific 1.6 % is not supported by the cited source." (OpenAI, Marain, 2025) It then suggested two authoritative second experts, an American analyst of immigration policy currently working at the Cato Institute, and an Austin-based director of research of the Justice Center of the Council of State Governments, concluding that the provisional rating should be maintained, "pending an external expert check".

The second claim we tested was made by an advisor to an Antwerp alderman and was based on a survey that the advisor conducted among approximately 2,000 Antwerp Muslims. This was how the claim was presented to ChatGPT: "70 percent of Muslim women are unemployed."

Actually, in the English translation of the original Dutch version of the claim, also fact-checked by Flemish magazine Knack, there is an ambiguity. "Unemployed* can be understood as a general description for "doesn't work", but is also a "technical" definition in labour statistics.

For this claim, ChatGPT correctly lit four warning lights in the Analyse Claim step. First, the subject, "Muslim women", is too broad, religion being not an official labour-force category; second, the denominator of the "70 percent" is not defined: age-band? labour-force definition?; third, the geography is not spelled out in the claim text; and fourth: sampling method and margin of error of the survey are not given.

In subsequent steps, ChatGPT did a cross check with official labour force statistics, linking to the correct source. It also suggested two authoritative second experts.

ChatGPT hinted at, but didn't analyse correctly enough, the ambiguity in "unemployed", and the fact that "Muslim women" is not a category that can be found in population statistics, let alone labour statistics.

A third claim was made by Belgian far-left PVDA chairman, Raoul Hedebouw: "Statistics show that a working-class child in Belgium almost automatically becomes a worker themselves. The social determinism inherent in our education system is enormous."

ChatGPT correctly raised multiple flags: undefined subject and denominator; unsourced "statistics"; qualitative superlatives. It also suggested several authoritative sources and experts to check social-educational mobility and stratification in Belgium, concluding that the concern about strong social gradients in Belgian education is valid, but the "almost automatic" determinism is not what the leading statistics show.

The results of even this low-tech, quick and dirty experiment are, I think, conclusive enough that most journalism and fact-checking teachers would agree that it proves that, yes, AIs can factcheck as good as or better than most of our students.

The results of even this low-tech, quick and dirty experiment are, I think, conclusive enough that most journalism and fact-checking teachers would agree that it proves that, yes, AIs can fact-check as good as or better than most of our students. They surely do it faster than students. In the experiment, ChatGPT did the three steps of each factcheck in under three minutes. The experiment also provides useful insights in

what "simple" AIs can and cannot do, and in where they still go wrong. But it forces us at least to revisit the instructions for take-home fact-checking assignments for our students.

Of course, the experiment also raises questions and problems. Most obviously for journalism education: How do we integrate AI in the fact-checking process so that it enhances the learning process of the students instead of letting them offload their learning efforts to a tool? More broadly, there are also ethical and deontological considerations.

Taking it two steps further: a prototype

These problems and considerations are addressed extensively by AI-researcher Bob Cruijsberg in his Master thesis, *Keeping Students In The Loop. Designing An Ethical AI Assistant For Fact-Checking Education In Journalism*. (Cruijsberg, 2026)

This project, in which Cruijsberg developed a complete AI fact-checking prototype, is an instance of a far too rare interdisciplinary collaboration of which we will have to organise many more if we want to remedy one of the problems he uncovers from a review of literature and other experiments: "(A)pproaches to integrating advanced AI-supported fact-checking in journalism education remain fragmented and often lack sustained stakeholder involvement." (p. ii)

It is precisely by taking this interdisciplinary approach, combining his own technical expertise with insights and requirements of journalism students and teachers iteratively during the whole development process, that Cruijsberg was able to turn the too simple question, Can they fact-check?, into a much more useful project. "(T)he design choices made in this study reflect a deliberate shift away from full automation … In contrast, the present work prioritizes human-AI collaboration by deliberately leaving key judgments, such as factual verification and source credibility, to students. This approach aligns with the conclusions of Wolfe and Mitra [2024], who argue for emphasizing human-AI collaboration rather than automating the fact-checking process in its entirety." (Cruijsberg, 2026, p. 72)

These considerations resulted in the research question: ***"How can an agent-based assistant be developed to support journalism students in making responsible decisions during the fact-checking process?"*** (p. 3)

Ethical considerations were embedded as functional requirements throughout the system design process (p. 70). As to the core journalistic requirements of transparency and accountability, "AI-supported fact-checking tools should ... be able to explain their outputs and decisions, while taking into account moral considerations and core journalistic values. (p. 2) In discussions and focus groups, journalism teachers and students emphasised student agency in the fact-checking process. For example, while the first version of the prototype automatically classified claims as True, False, Mostly True, Mostly False, or Uncheckable (following the Eufactcheck classifications), "it became clear during the early interviews that this judgment should be made by the student rather than delegated to the AI. Subsequent iterations therefore shifted greater control to the user, for example by requiring confirmation at each step of the process, including decisions such as which questions to use for information retrieval or search." (p. 68)

This interdisciplinary cooperation proved not only valuable for making sure that ethical and deontological journalistic values were supported and built into the system, but also for its technical sophistication. As a result, the project went way beyond what most journalism teachers would be able to achieve.

Setting aside for a moment the promise at the beginning of this text, to gloss over the advances in AI since the release of the first LLM-chatbots, here are some of those advances, which Cruijsberg of course uses for his prototype and which are relevant in the fact-checking context.

Retrieval Augmented Generation (RAG) is an AI framework that improves Large Language Model (LLM) responses by grounding them on external, authoritative knowledge bases rather than just their static, pre-trained data. (Google Cloud topics). For example, a newspaper could add its own archive to an LLM-application, thereby mitigating the risk of hallucination by anchoring responses in verified data and specific domain knowledge. An IBM Research Blog compares a RAG-enhanced system with an open book exam. (IBM Research Blog, 2023) Talking to an LLM without RAG would then be like taking an exam from a student without knowing which texts the student has used to study.
A second technology that Cruijsberg incorporates, and that has considerably expanded the practical applications of AI, is ***Agentic AI.*** Agentic AI, or AI Agents, are autonomous or semi-autonomous systems that are given (by humans, for now) high-level goals and that can reason, plan, and execute multi-step tasks using external tools. Unlike LLM-chatbots, that simply respond to prompts, agentic AI operates with

minimal supervision to solve complex problems. It combines the flexible, user-friendly characteristics of LLM-chatbots with the accuracy of traditional programming.

One example of such an AI Agent in a media environment is under development at Ippen Media, one of the largest media networks in Germany, publishing over eighty titles, and with a strong digital department. When a journalist is assigned a topic, the agent acts as a digital co-worker. It autonomously researches the subject across internal databases and the web, summarizes expert opinions, extracts historical context, and gathers quotes. The agent synthesizes this into a single, structured overview, leaving the human journalist to verify the facts and edit the final piece. (Pollock, 2026; Liang et al, 2025) In a broadcast environment, an AI Agent could help broadcast control room operators automate a range of real-time controls, driven by their voice commands: Speech-to-text (STT) transcribes the operator's voice commands; an LLM interprets those instructions; the agent triggers camera movements, adjusts running orders, and delegates tasks to sub-agents who can access specific databases (e.g. broadcast assets, show plans) or other tools. (FT Strategies, n.d.)

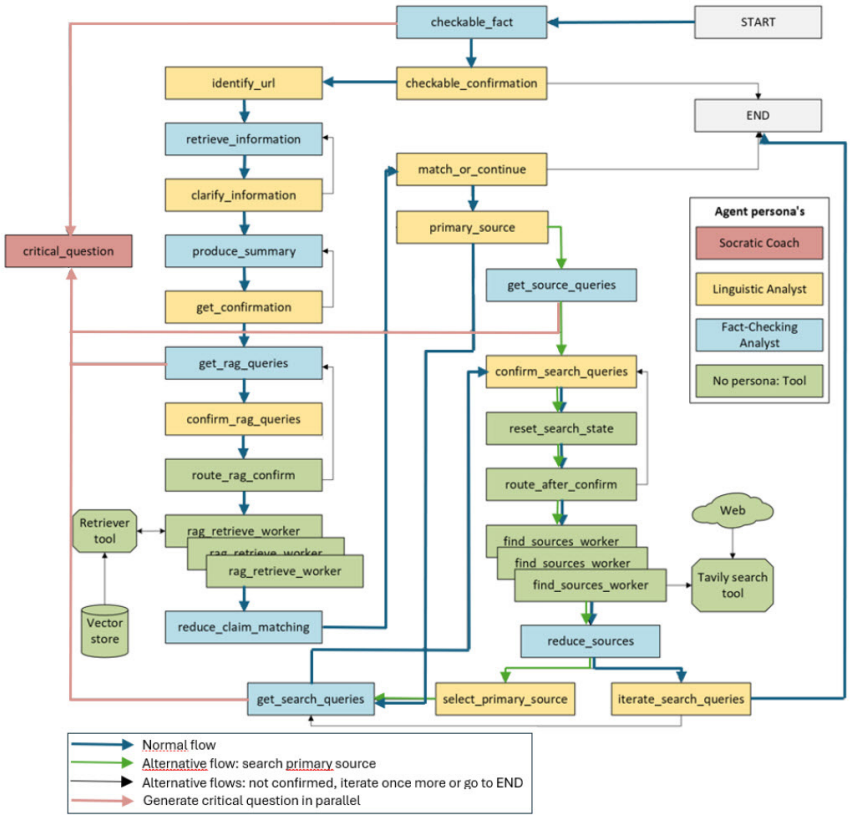
In its Future Newsroom Studies 2026 FT Strategies reports that agentic usage remains at 4% or below in every newsroom workflow category, mostly in transcription and translation. The same report notes that future editor-coders in newsrooms will have to be able to design and develop agentic workflows. (FT Strategies, 2026)

For his prototype, Cruijsberg has built a ***Multi-Agent System*** (MAS), "a framework for LLM-based agents that draws inspiration from human cognition and consists of three core components: perception, a brain, and action. (Cruijsberg, 2026, p. 24) Compared to the "simplest architecture" of a single AI Agent, where one LLM handles all the tasks, in a MAS, multiple agents with specialized roles can handle larger, more complex tasks. The perception part, for example, "serves as the entry point of the system and may involve textual, visual, or auditory input. In the context of this study … the agent is limited to textual input only. (p. 25) Moreover, agents working together can cross-validate and refine each other's outputs; multiple agent systems will thus reduce hallucination risk through multi-agent crosschecking and domain specialization." (p. 27)

Cruijsberg used ***Python*** as programming language to develop the MAS, and the Python library ***LangGraph*** because it "represents a particularly suitable choice for this project, as its explicit workflow control enables the direct transfer of the structured workflows developed in the EUfactcheck project." (p. 27) Specifically, the flexible MAS framework ensures that the workflow is "inherently interpretable, traceable, and modular." (p. 51)

To further support the pedagogical goal of the factcheck workflow, Cruijsberg introduced three distinct personas in the workflow: the Linguistic Analyst, who transforms unstructured human input into clean, structured JSON data; the Fact-checking Analyst, who evaluates evidence, identifies inconsistencies, and assesses claims; and the Socratic Coach, who encourages reflection and guides students through critical questioning. (p. 51)

A graphical representation of the complete workflow, as modelled by Cruijsberg, is shown as figure 2 below. It is a fine and impressive example of what can happen when an IT-mind addresses a journalistic question.



Multi-agent fact-checking workflow; reproduced from Cruijsberg, 2026, p. 52.

Cruijsberg evaluated his model thoroughly in iterative discussions and focus groups with journalism teachers and students, and with other AI-experts. "Across all qualitative evaluations, a strong consensus emerged that the assistant's primary educational value lies in supporting critical thinking and reflective judgment rather than automating fact-checking outcomes." (p. 67)

AI-assisted fact-checking tools such as the proposed assistant are best integrated into journalism education as scaffolded learning supports rather than autonomous verification systems.

The term that Cruijsberg and others use to describe how students can use AI to support their learning instead of offloading it to AI, is ***scaffolding***. "Based on the findings of this study, AI-assisted fact-checking tools such as the proposed assistant are best integrated into journalism education as scaffolded learning supports rather than autonomous verification systems. … (A)gent-based assistants most effectively support journalism education when they are designed to scaffold students' critical thinking, rather than replace it. This study contributes to ongoing discussions on responsible AI in education by demonstrating that sustainable, open-source, multi-agent systems, when combined with stakeholder involvement and careful pedagogical design, can enhance learning while preserving student responsibility and professional values." (p. 72-73)

Yes, it can

If we agree, if only for the sake of argument, and without too many "yes, buts …", that we have established that AI can fact-check, we have a useful starting point for further reflections on problems and opportunities for AI in journalism education. The first step is to admit that, indeed, there is a problem, and that it is BIG.

The big problem for journalism education

Some of the problems that AI poses for journalism education are shared with the problems AI causes in higher education in general. We will therefore first reflect on the problems of AI in higher education and then show that the problem for journalism education is even worse.

AI and higher education

The problem of AI in higher education is often formulated as some variation on the theme of ***cognitive outsourcing***. Students (and teachers) outsource core intellectual learning processes and tasks to AI, with the result that the intended skills and competencies are never attained or, if they were present at all, may be even lost. "People Outsourcing Their Thinking to AI" (Shroff, 2025), "may cause certain "mental muscles" to atrophy" (Oakley et al, 2025, p. 34) or lead to "skill atrophy in tasks like brainstorming and problem-solving." (Kosmyna et al, 2025, p. 108)

In a recent white paper, aptly titled *Technologies Smarter, Humans "Dumber"? How Technological Advances Improve and Diminish Human Cognition and What This Means for Education (*Van Damme, 2026), Dirk Van Damme offers a useful, balanced take on the problem.

Van Damme is former Head of the Department of Educational Sciences at Ghent University, former Cabinet Member of three Flemish education ministers, former Head of the Center for Educational Research and Innovation (CERI) at the OECD and now Distinguished Senior Research Fellow for the Boston based Center for Curriculum Redesign, for which he wrote the paper.

Van Damme puts the AI-problem for education in perspective by taking a sweeping historical view on the impact of technological innovations on human cognition and, hence, on education. The paper conceptualizes technology as an ***epistemic change agent***, examining "how successive technological innovations have altered human cognitive capabilities and the conditions for learning." His goal is to "outline implications for curriculum design, pedagogy, and education policy in the age of artificial intelligence." (Van Damme, 2026, p. 5)

An analysis of twenty case studies of past and recent technological innovations - from fire (!) over writing and printing up to social media and, ultimately, artificial intelligence – reveals that all these technologies somehow have produced "both cognitive gains and losses: while technologies expand abstraction, efficiency, and access to knowledge, they simultaneously weaken embodied, contextual, and internally sustained capacities that education has traditionally cultivated." (Van Damme, 2026, p. 5)

Van Damme recognizes the outsourcing problem, but shows that it has been a constant through centuries of technological progress:

"Technological innovations ... alter our mental energy, memory, and attention. For example, they cause individuals to offload or outsource cognitive operations (e.g., spatial memory in the case of GPS). Cognitive offloading shifts our mental operations from internal knowledge to external memory. This means a change from knowing to knowing where to look." (Van Damme, 2026, p. 7-8)

And: "Each technology shifts tasks from human memory or perception to external artifacts — from fire (outsourcing warmth and light) to writing (outsourcing memory), to digital tools (outsourcing reasoning and recall) ... diminish(ing) active cognitive capacities and resilience." (Van Damme, 2026, p. 12, p. 14)

While Van Damme convincingly argues that all technologies have resulted in both cognitive gains and losses and narrates how education and learning have adapted, he also recognizes that, with the emergence AI, the challenge of adaptation is of a different order: "The emergence of artificial intelligence ... represents a qualitative shift in how knowledge is processed and generated. ... This transition marks the movement from "informational technologies" to "epistemic technologies", in which machines participate in meaning-making rather than merely storing or transmitting data." (Van Damme, 2026, p. 74) From IT to ET (Van Damme doesn't use the abbreviations), there you have it.

"The emergence of artificial intelligence ... represents a qualitative shift in how knowledge is processed and generated. ... This transition marks the movement from "informational technologies" to "epistemic technologies"

Noting that, at this stage, "what we can learn from the research literature on gains and losses remains preliminary" (Van Damme, 2026, p. 75), he nevertheless proceeds, as he has done for all previous technologies, to listing the gains and losses of AI for human cognitive behaviour and competencies.

The list (Van Damme, 2026, p. 75-76) is based on twenty recent scientific articles. It is as good as any a checklist educators will get to check if and how they can integrate and enhance the potential gains of AI in their teaching, and how to mitigate the potential losses:

Gains:

- Access, acceleration, and efficiency
- Speed & scaffolding
- Expansion of epistemic reach and inclusion
- Personalization and control in learning
- Productivity and learning boosts
- Cognitive scaffolding and meta-reflection
- Augmentation of creative and inferential capacities

Losses:

- Opacity and loss of epistemic traceability
- Generalization bias, distortion and hallucination of information
- Cognitive offloading
- Possible negative impact on human creativity
- Erosion of epistemic vigilance and critical reasoning
- Affective and motivational consequences
- Weaker coping mechanisms
- Homogenization

The future of our schools

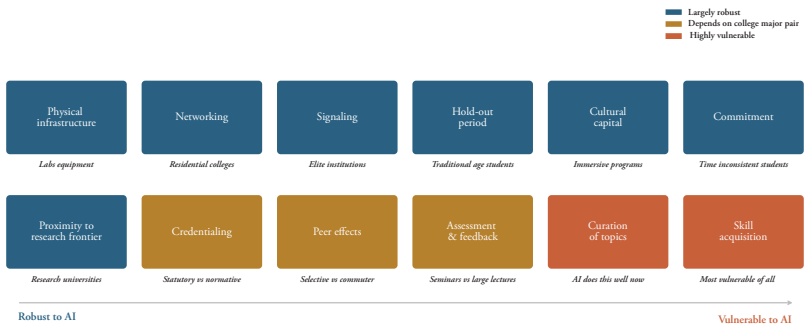
Van Damme's (2026) focus is on the impact of new technologies on our students' cognitive behaviour and competencies, and on educational responses to these challenges.

Noble and necessary as it may be, we shouldn't only worry about the minds and the future of our students. The future of our schools, our departments, and our jobs is also at stake.

In a recent series of Twitter/X threads, Jesús Fernández-Villaverde, Howard Marks Presidential Professor of Economics at the University of Pennsylvania and Senior Fellow at the American Enterprise Institute (AEI), reflects on the role and the future of "traditional" higher education.

He starts out by assembling a list of (12) arguments supposedly supporting the survival of traditional higher education (Fernández-Villaverde, 2026, Mar 24) and then investigates how each argument might be affected by AI. (Fernández-Villaverde, 2026, Mar 30) While his examples and arguments are focused on the US higher education situation, they still contain useful lessons for Europe.

Here is the list of 12 arguments, colour-coded for how Fernández-Villaverde assesses their robustness or vulnerability to AI:



Arguments for higher education and their vulnerability to AI. (Jesús Fernández-Villaverde)

Fernández-Villaverde approaches these rationales as an economist. We could think of each, or a combination of them, as the underlying value of students', their parents' and society's willingness to pay for higher education. For future employers, they represent the underlying value of the benefit they hope an employee with a degree will bring them. Skill and knowledge acquisition in topics that are carefully curated and provided with ready-made syllabi is one of the more obvious of those factors; the assurance of progressive assessment and feedback another one; networking and the acquisition of cultural capital also. Taken together, they should indicate whence schools in higher education derive their social and economic value and, ultimately, their raison d'être. Different views on the value of higher education will assess the value of each of the

rationales differently. For some students, for example, the networking and building of cultural capital will be the most memorable and valuable component of their education. Future employers will likely value more the acquired skills in carefully selected topics.

Whatever the inherent value of some of these factors, one rather convincing view on the real (economic) value of higher education is that it is largely – some say almost exclusively - about signalling and credentialing. In its most critical form, the signalling view on education is even presented as "the case against education" (Caplan, 2019). The signalling view on higher education proposes that a degree acts primarily as a credible signal and proof to potential employers of a student's mastery of a range of specific, profession-related skills and competencies and, independent of specific skills, of general abilities such as intelligence, diligence, conscientiousness, and grit. Employers use the signal of a degree as a risk- and time-reducing element in the hiring process. The credentialing value of a diploma is the fact that it is sometimes required, for legal, social or other reasons, to perform certain tasks or professions.
This system of signalling and credentialing worked more or less satisfactorily as long as one basic assumption was met: assessment, grading, and ultimately the diploma provide accurate information regarding both the student's skills and knowledge level and her broader competencies. It was not perfect, but it was a useful proxy.

That basic assumption has now been completely undermined by AI.

This lowers the value of our diplomas, the only "real" social product our schools produce, on several levels and via several routes. For skills and knowledge acquisition, and feedback and assessment, the signal of a completed degree has now been downgraded to:

This person can produce acceptable outputs under unknown levels of assistance.

In a wide-ranging dialogue (podcast with transcript provided), ominously titled *How AI Broke Education*, Henry Shevlin and Dan Williams, respectively Education Director and Associate Fellow at the Centre for the Future of Intelligence, University of Cambridge, identify the take-home essay assignment as the "ground zero of the negative effects (of AI)", noting that a lot of other kinds of assignments are similarly impacted too. They determine that "(t)o the extent that students are using AI to produce their assessment material, the ***signalling value ... just dissolves completely***.

(Williams, 2025; my emphasis) Logically and inevitably, with the signalling value of the intermediate assignments, the signalling value of continuous feedback and assessment goes down the drain as well. "You've got no way of knowing which students need extra help versus which are already doing fine." (Williams, 2025)

Quotes like these, and the accompanying anecdotes, are now the staple of discussions in every university and college teacher room. *Everybody is Cheating Their Way Through College*, titled New York Magazine in a much-discussed May 2025 article. Some of the reactions to the article, and surely some discussions in teacher rooms, are of the moralizing kind, blaming students' laziness, lack of intrinsic motivation, or ever shorter attention spans. Hollis Robbins, an American academic and essayist, former Dean of Humanities and Special Advisor for Humanities at the University of Utah, doesn't agree, putting the blame on the institutions instead: "Students didn't create this system; neither did faculty. Institutions created it when they built universities on business models dependent on continuous enrolment growth in a country with a declining birth rate. Students … are responding rationally to a system that has communicated, through years of grade inflation and lowered expectations, that completion matters more than comprehension. Everyone knows this." (Robbins, 2026, January 31)

There are other ways, besides cheating, via which AI will impact the social and economic value of a diploma and of higher education. Students who realise that the signalling value of a diploma is diminishing rapidly, will start to bypass institutional education and self-learn, for instance with the AI-twins of the best professors in their fields. Employers will know this and will increasingly consider hiring self-learners instead of relying solely on the signalling and credentialing values of a diploma. Related to this effect is the way that AI increasingly offers employers, for whom the signal value of a diploma had always consisted in the reduction of risk and time in the hiring process, alternative ways to assess competence.

Be it cheating on a massive scale, students and employers bypassing traditional institutional education, or parents, and society, starting to ask what, exactly, they are paying for, a lot seems to suggest the conclusion that AI is the comet that will lead to the extinction of the dinosaur that higher education has become.

According to some observers, that is indeed what we see happening.

Here is Hollis Robbins' prediction (quoted at length for dramatic effect):

"This fall, prospective students and parents should be looking at university recruitment materials with one question in mind: what exactly is a university education worth in the AGI era? … General education courses focused on basic knowledge transfer become indefensible. If the information is general enough to be called "general education," AGI can deliver it more effectively than any human instructor. … The curriculum narrows dramatically. Most lecture courses disappear. … This represents a 60-70% reduction in current faculty positions, with remaining roles requiring fundamentally different capabilities than traditional academic work. … Students cannot be expected to continue paying for information transfer that AGI provides freely. Instead, they will pay to learn from faculty whose expertise surpasses AI, offering mentorship, inspiration, and meaningful access to AGI-era careers and networks. Universities that cannot deliver this specific value will not survive." (Robbins, 2025, February 9).

She follows up with what she sees as a business opportunity:

"The next few years will be a bloodbath as colleges and universities close. Hundreds of people will lose their jobs and thousands of students will join the millions who have taken a bunch of courses but don't have a degree. … There is an enormous business opportunity for anyone who recognizes that the ***"college degree" is losing its monopoly on what counts as "completed learning."*** (Robbins, 2026, May 23; my emphasis)

Dan Williams and Henry Shevlin agree with the dinosaur thesis, shifting their focus, as we have done, from what AI does to students' brains to what it does to our institutions: "The problem here is not just that students aren't learning things (which would be bad enough). I think it's a kind of ***extinction-level threat for these institutions***, because once it becomes clear that the grades you're giving students don't really provide any information about their intelligence or conscientiousness, then ***the social function of the institution dissolves completely***." (Williams, 2026; my emphasis)

Should we panic? In the circumstances that would probably be the most rational reaction. But, as so often, the economist introduces some nuance.

Fernández-Villaverde recognizes that higher education is in trouble, largely due to the disappearing value of the signalling and credential functions:

"(I)f the knowledge itself becomes cheaply available, the only thing holding up demand [for higher education degrees] is the credential, and credentials without underlying value are precisely the kind of equilibrium that does not survive a large enough shock." (Fernández-Villaverde, 2026, Mar 27).

But he introduces "two principles that any economist would emphasize but are largely missing from the public debate about AI and higher education." (Fernández-Villaverde, 2026, Mar 27)

The first is marginal thinking. This is the kind of thinking that refers to the process of making decisions by considering the incremental or additional changes that result from a small, incremental change in a variable.

So, instead of binary thinking (universities will disappear, or not), we should be thinking about what is happening as a series of shifts in the decisions of the marginal agents: "(H)igher education may look very different in a few decades. Not because universities are going away, but because the marginal student, the marginal program, and the marginal institution will all face a different set of relative prices. And when prices change, behavior changes. The adjustment will begin at the margins and move inward from there. Markets eventually do their work, even in the department of anthropology." (Fernández-Villaverde, 2026, Mar 27)

The second economic principle he introduces, is heterogeneity, and it is especially relevant in light of the differences between the US and the European situation.

The appropriate unit of analysis when thinking about the future of higher education is not "higher education" as a single entity, but individual college-major combinations. With an American example: "A finance degree from Wharton is not equivalent to a philosophy degree from a small, non-selective liberal arts college in Vermont." (Fernández-Villaverde, 2026, Mar 27)

Readers will probably have noticed that in the graph pictured above, where Fernández-Villaverde assesses the impact of AI on different arguments for higher education, he coloured the "Signalling" argument blue, for "largely robust" to AI. But that is only true for highly selective elite institutions. Indeed, as he explains: "The signaling value of a STEM degree from an Ivy is untouched by AI, because the signal comes from admission and from the ability to complete the degree. By contrast, a degree in humanities from a non-selective institution is a weak signal, because both admission

and completion are easy." (Fernández-Villaverde, 2026, Mar 30)

Now, as to "highly selective elite institutions" (supposedly robust to AI impact) and only looking at the Humanities, in the top 15 of the QS World University Rankings 2026 for Arts and Humanities, four are in the UK, one in Canada, the rest in the US. Same picture in The Times Higher Education World University Rankings 2026 for Humanities, except that LMU Munich is added as fifteenth. (By the way, of those fifteen, only five, all American, offer Journalism degrees – noting that, of course, a PPE (Philosophy, Politics and Economics) degree from Oxford would be at least as good an entrance ticket to a journalism career as a Journalism degree. But that is another discussion)

So, rare exceptions notwithstanding, in Europe (and at American non-elite institutions) we have this combination in higher education: Easy admission + Very little sanctioned quality control, not at the level of the school, nor at the level of individual courses + Focus on outcomes instead of process.

And now AI is added to that equation. This is bad news for higher education institutions in general. We will now narrow down our reflections on journalism education. Warning: Even more bad news is coming.

Journalism education: More bad news

The problem of the institutional design of our Journalism schools is probably most sharply illustrated by comparing the impact of AI on two very different departments. This came up in a discussion with a colleague who teaches basic programming at an IT department. His first and most obvious problem is that students vibe-code their programming assignments and, hence, don't really learn to code anymore. Is that a problem, one could ask, in a world where AI can code as well as and surely faster than most programmers? Yes, he thinks, because we will, in the foreseeable future, still need humans who can assess the results of AI-coding and who are able to see the bigger picture. The ideal solution, if there would be enough time, would be to teach the IT-students more abstract reasoning.

There is a parallel with the challenges in journalism education. In a world where AI can read, analyse and write as well and surely faster than most journalists, and that anyhow is becoming more "digital" by the month, journalism students should at least acquire a decent level of IT-skills and probably even learn the basics of programming.

This discussion about teaching digital skills in journalism education has been going on since long before the current AI-hype. In a study and a survey of young German journalists, exploring the intersection of technological advancements in the media industry and the realities of journalism education, researcher Vera Katzenberger from Leipzig University found that "many journalists report insufficient exposure to critical digital skills such as search engine optimization, user metrics, data journalism, and coding (during their education)". (Katzenberger, 2024, p. 1) One of the main causes of this shortage, Katzenberger found, was that "(i)t appears to be quite demanding for journalism educators in the higher education sector to keep pace with the rapidly changing needs of the industry. There are reportedly not enough academically trained lecturers who can lead and teach interdisciplinary data literacy and coding courses in journalism programs at universities." (Katzenberger, 2024, p. 11)

Tellingly, as Bisiani et al (2025) found in their analysis of The State of the Data Journalism Survey 2021, seven out of ten surveyed practising data journalists stated they are self-taught. "Due to the limited availability of formal educational paths specific to data journalism ... online courses and certifications (are a) more popular development path than both higher education or workplace training." (Bisiani et al, 2025, p. 507-8) Findings like these confirm the fears about the declining value proposition of higher education that we noted in the analysis above.

So, simplifying (but only barely):

IT-students should learn to read and write.

Journalism students should learn to code.

Maybe. But we have now arrived at the real and creepy problem. In the current institutional design of our schools, with AI added and, again, only barely simplifying:

IT-students don't learn to program anymore.

Journalism students don't learn to read, analyse, and write anymore.
[Note that "read" and "write" are used in the broadest possible sense, so they also include "watch", "listen to", "film", ...]

We will leave it to our IT-colleagues to solve their problem (returning to them later), and investigate instead what the real value of a Journalism diploma is or should be, in order to argue that journalism education is hit especially hard by AI.

The signalling value of a diploma in Journalism implies at least two, but ideally three layers. Specifically:

The ***skill layer*** certifies: This student can write, research, verify, and use technical and other journalistic and storytelling skills.

The ***competence layer*** certifies: This student can do this consistently under academic conditions, and is ready to transfer her competencies to professional conditions.

Problem 1: Even before AI, this was already fuzzy. But at least, we could confidently certify that work had been mostly authored by the student, because the process of producing it was partially observable. Under AI-conditions, as noted, we can only certify that "this student can produce acceptable outputs under unknown levels of assistance."

So far so bad.

To be sure, the skills and competencies layers are and will remain ***foundational*** because (1) quite simply, they answer the question: Can this student produce journalism? And (2) our curricula are designed to clearly define, teach and assess these foundations.

Now let's for a moment assume a best case scenario where we could somehow credibly certify that work has been mostly and consistently authored by the student throughout her study programme, and has been duly assessed, so that her degree signals that she indeed masters these foundations. Do we then deliver the young journalists that media companies ***and their audiences*** want?

Media companies and their audiences will increasingly realize that what graduate students can produce with their skills and competencies, as foundational as those may be, can be produced just as well by AI. It is time we acknowledge that.

Unfortunately, in an AI-world, the answer will increasingly be a resounding NO. Media companies and their audiences will increasingly realize that what graduate students can produce with their skills and competencies, as foundational as those may be, can be produced just as well by AI. It is time we acknowledge that.

The only way out of this deadlock is to have a hard look at a third layer in the signalling value of our diplomas that has always been present somehow but is now set to become the most meaningful signal of expertise.

Where skills and competencies answer the question: Can you produce journalism?, this third layer answers the question: Is it worth reading, trusting, and remembering? It is the layer - let's call it craft - that makes a Journalism department more like an Arts department than, say, the IT department. Where competencies and skills are and will remain the ***foundation***, craft is the differentiation. A ***differentiation***, nota bene, not only as an individual Author, but also from AI. If anything, the audience in an AI-world will want more of this. In our students, craft will manifest itself in emerging judgment, voice, and originality; in storytelling instinct; in sense of timing and narrative tension; in sensing the audience and the moment.

If we agree that these elements are becoming the most meaningful signals of expertise, then we have:

Problem 2: Skills and competencies were, in principle, clearly definable, teachable, and assessable (but see Problem 1), with assessment heavily based on product. Craft and Authorship are, almost by definition, harder to standardize and were, consequently, always harder to teach and to assess. In an AI-world craft becomes more central but even harder to assess because it cannot be reliably inferred from polished outputs anymore. It shows up in large part in ***process***, not (only) in ***product***.

It is in this process vs product tension that we discover the heart of the challenge and probably also the first steps towards its solution.

For Santiago Schnell, provost and professor of mathematics at Dartmouth College, "(t)he deepest challenge of AI in education is therefore not academic integrity, though that problem is real. It is whether we will allow our schools and universities to ***define learning as the production of acceptable outputs***." (Schnell, 2026; my emphasis)

In the humanities and, a fortiori, in journalism education, this "acceptable output" is a "final text" (again, in its broadest definition, including video and audio). It is what matters most. So much so that teaching text production is, for a large part, how we define what we do. We defined learning as the progressive ability to produce publishable texts. And for good reason!

The implicit assumptions we used in this view on teaching and learning were:

The text is the product. It is what students turn in and what we grade.

This worked, because text was not only a "final product" but also a good shortcut, a valid proof of underlying skills and competencies such as research methodology, general knowledge, and reasoning and writing skills.

And then AI learned to do the same trick. (Read the master or bachelor thesis requirements of your school. By now, we know they are a joke).

But "(t)he text is not the product", German philosopher and social scientist Lisa Maria Herzog reminds us. "Generative AI, for all its problems, has one virtue: it forces us to rethink that assumption. ... We treat texts as products; they get bought and sold. ... There may be situations in which texts are really products in and of themselves ... (user manuals?) But in reality, texts are almost always something else." (Herzog, 2026)

Brendan McCord, founder of Cosmos Institute ("The Academy for Philosopher-Builders, technologists who fuse deep inquiry with technical skill to ensure that AI serves the human good"), points to what went wrong and where we should look instead.

"Students learn to have a thesis before they write, to know their conclusion before exploring their question, to treat writing as transcription of completed thoughts rather than the medium through which thought happens. … If AI can produce polished text instantly, we can no longer pretend that polished text is what matters. We must return to what Montaigne knew: the essay is a ***process, not a product***."
(McCord, 2026; my emphasis)

Of course, most journalists don't see themselves as essay-writers, and their master or bachelor thesis will often be the last piece of academic work they write, the last time they will use APA-citations. But what both kinds of writing – academic or essay-writing, and journalistic writing – have in common is the importance of the process, as craft (albeit different kinds of craft).

Teaching this process is not easy. We will come to that in a moment. But what we know for sure is that, for journalism students, AI will be most unhelpful in the learning of this process.

"AI writing is the technology's bleakest use case."

Tech journalist Rebecca Heilweil nails it with her boutade in a recent Fast Company piece: "AI writing is the technology's bleakest use case." (Heilweil, 2026) Or maybe we will touch the nerves of our students more with this advice from Cal Newport, American nonfiction author and professor of computer science at Georgetown University: "Don't let AI write anything for you. Writing is to cognitive health what steps are to physical health." (Goldstein, 2026)

In an hour-long interview with writer David Perell, journalist Ezra Klein makes his "Case Against Writing With AI", focusing on the "craft" (used eight times in the interview) "of deep and meaningful journalism." (Perell, 2025; podcast with transcript provided)

Ezra Klein, in Perell's words "one of the Internet's most influential journalists", is a columnist for The New York Times since 2021 and host of the podcast "The Ezra Klein Show" . Before that, he founded and was editor in chief of Vox, and was a columnist and editor at The Washington Post.

The interview is worth reading or listening to because it is one of those rare instances where one is able to eavesdrop on two master-craftspersons discussing the tools of their trade.

For Klein, journalistic, non-fiction writing is all about the process and the irreplaceability of work:

"By the time I am writing, I've done a lot of work. ... Writing is the end point of the process. ... That structural work had to happen, that mental work had to happen, those false starts had to happen. This is a reason … I think it is very dangerous to use ChatGPT in any serious way for writing."

"Just do the work that others skip," is the down-to-earth advice he would give to young students. And for journalists in particular: "Just do the reading ... Good journalists are the ones who actually read things like the Congressional Budget Office report instead of skimming the summary." Because "reading is the wellspring of all the rest of the work." (Perell, 2025)

The shortcuts AI seems to offer are based on a misguided metaphor of how our minds need to work when we want to (read to) write. "We sort of have this download model; I'm downloading the information. ... I'm not saying AI can't be useful or that it won't be useful, but I'm against shortcuts. You have to limit the amount of work you're doing; you can't read literally everything. But in some ways, I think it's more dangerous to think you've read something that you haven't than to not read it at all." (Perell, 2025)

There are no shortcuts for "spending time grappling with the text, making connections. It will only happen through that process of grappling." This is the false promise AI makes to our students: the process, the work, and the time spent are incompressible; or hardly. Because without process, work and time, the connections will not happen. "What knowledge is supposed to do is change you, and it changes you because you make connections to it." An AI-summary will not change you. And "a magazine article will change you less than a book will. A tweet will change you less than a magazine article will." And indeed: "There's no prompt that would have told me, "Just go read more books, man. You're not there. You haven't gone far enough." (Perell, 2025)

Nor is the solution, now widely implemented, to just find modes of assessment where students can't use AI. This will actually reinforce the view that the product matters more than process.

The sheer pleasure with which Klein and Perell talk about their craftsmanship is infectious. But it will not be sufficient to dissuade our students to take the shortcuts that AI seems to offer. As noted above, moralizing and blaming students is not fair and will not help. "Students are responding rationally to a system that has communicated ... that completion matters more than comprehension." (Robbins, 2026, January 31) Nor is the solution, now widely implemented, to just find modes of assessment where students can't use AI. This will actually reinforce the view that the product matters more than process.

If the problem is institutional, the solution has to be an institutional design where we (re)create the environment where teaching and learning craft as process are prioritized over just final output.

The Cloister and the Starship. A radical solution?

❧

One such solution, recently proposed by the British-American historian, writer and opinion journalist Niall Ferguson, is attractive because it is simple, while at the same time more radical than what most other education-thinkers and designers have been hinting at, and because of the catchy metaphor.

With The Cloister and The Starship, Ferguson suggests "five essential steps, to be taken with immediate effect" (Ferguson, 2025, June 17; partly paywalled – short summary on X (Ferguson, 2025, July 6), with the remark that "I can assure you no established university will do anything remotely like this."):

1. Create quarantined space in which traditional methods of learning can be maintained and from which all devices are excluded. Call this "the cloister".
2. Assume that all study outside the cloister will be done with the use of LLMs. Call this "the starship".
3. Inside the cloister, allocate time to
   a) reading printed books
   b) discussion of texts and problems
   c) writing essays and problem-sets with pen and paper, and
   d) assessment via oral and written examinations.
4. Require time in the cloister to be around seven hours a day, leaving time on board the starship, as well as vacations, for the use of AI.
5. Revise admission procedures to ensure the university attracts students capable of coping with the discipline of the cloister as well as the opportunities of the starship.

Here is how I think some of the steps should be amended to a) adapt the proposal for journalism education and b) make more room for teaching and learning smart AI-use than Ferguson allows for.

First, inside the cloister (3), journalism students should also read, watch, listen to and than analyse and discuss extensive anthologies of the best short- and longform journalism. Writing should be focused on short- and longform journalism, more than on essays. Students should practise interview skills, first by interviewing each other; then, outside the cloister walls, in real environments. Students should, of course, also look over the walls of the cloister to follow the news closely.

Second: For the starship, Ferguson is frankly not ambitious enough. A decent starship not just (2) "***assumes*** that all study will be done with the use of LLMs", but ***organises*** for a thorough and broad programme of technological learning. Here, as promised, we return to our IT-colleagues to finally design the real interdisciplinary curriculum that we should already have organised years ago. In interdisciplinary courses, journalism students and IT-students together acquire and discuss at least three levels of AI- and digital literacy: (1) Knowing AI- and other digital tools and how and when to ***use*** them smartly and in an advanced way in a wide range of journalism applications, from research to text analysis, data journalism and coding; in the starship, students should learn to *think like a computer*; (2) ***Understand*** the inner workings of AI-tools; (3) ***Judge*** the ethical issues and learn how to mitigate ethical and other risks.

Third: I would add a third environment to the Cloister and the Starship, where students merge what they have learned in the two other ones. In this environment, let's call it the Bridge, AI use actually reveals student competence instead of masking it.

For instance, as proposed by Brendan McCord: "What if, instead of turning in essays, students turned in evidence of their thinking? Document your dialogue with AI. Show where you questioned it, where verification changed your argument, where your thinking shifted. … The intellectual moves that used to happen invisibly now leave a traceable record." (McCord, 2026)

This is maybe also the environment where students could try their hand at what writer and educator Arnold Kling calls vibe-reading and vibe-writing, where students augment their reading and writing in a dialogue with AI. "Vibe-reading as I conceive it only starts with asking an AI for a summary. It is a more intricate back-and-forth process." (Kling, 2026)

On the Bridge, teachers and students (IT and Journalism) would also experiment with new, AI-enhanced, article formats. Or, strangely missing in current media, design AI-applications for readers.

The Bridge will probably be the environment that most resembles the real environment where students will end up in their professional life. Its goal should be: Augmentation without the atrophy.

The Bridge, by the way, will probably be the environment that most resembles the real environment where students will end up in their professional life. Its goal should be: Augmentation without the atrophy.

Fourth: Of course, with the above amendments, the seven hours a day in the Cloister proposed by Ferguson will have to be spread out more evenly among the three environments. Maybe, in percentages, 50 (Cloister); 30 (Starship); 20 (Bridge)?

Even with the amendments, this proposal has the advantage of simplicity. It is also radical. To be effective, the metaphor should not only be implemented in the institutional and curriculum design but also in the physical architecture of our schools.

Inarticulable expertise and accountability

In the diagnosis part of these reflection, we proposed that journalism was hit especially hard by AI because it depends heavily on authorship, credibility, trust in process, not just in output. In journalism and journalism education, therefore, craft becomes the most meaningful signal of real expertise. It is also harder to teach and to assess.

But if we want our Journalism diploma's to signal any value beyond "this person can produce acceptable outputs under unknown levels of assistance", and educate our students to become journalists that are worth reading, watching or listening to, we will have to find ways to do just that.

As already briefly noted, teaching and assessing craft will ensure that our students can differentiate themselves, as journalist, but also from AI. This differentiation will germinate in The Cloister, through reading and writing.

reading is the apprenticeship to the inarticulable craft and expertise of writing

It starts with reading, of course. Not in the first place because of the knowledge that is acquired in the process, essential as that is, but because reading is the apprenticeship to the inarticulable craft and expertise of writing.

The power of the images of craftsmanship and apprenticeship, which fit well into The Cloister metaphor, was already noted by Michael Polanyi, physical chemist before turning to philosophy in his fifties, who, long before the current AI-hype, provided arguments for AI's limits and, thus, for a way for a writer to differentiate herself from AI.

This is the relevant passage in Polanyi's *Personal Knowledge*:

"An art which cannot be specified in detail cannot be transmitted by prescription, since no prescription for it exists. It can be passed on only by example from master to apprentice. … To learn by example is to submit to authority. You follow your master because you trust his manner of doing things even when you cannot analyse and account in detail for its effectiveness. By watching the master and emulating his efforts in the presence of his example, the apprentice unconsciously picks up the rules of the art, including those which are not explicitly known to the master himself. These hidden rules can be assimilated only by a person who surrenders himself to that extent uncritically to the imitation of another." (Polanyi, 1958/2005, p. 55)

AI-specialist Chris Walker thinks that "if we take Polanyi seriously, and I think we should, then we can make a precise claim about AI's limits: if we cannot write down everything we know, then AI cannot learn everything we know. Not because AI isn't sophisticated enough, but because the knowledge doesn't exist in a form that can be digitized. … The foundations of expertise, Polanyi argues, are essentially inarticulable." (Walker, 2026)

It feels plausible. But maybe it is also one version of the "yes-but ..." reactions that we have tried to avoid, in which we define optimism in the AI-era as finding things that AI cannot do.

Alexander Maeder, professor for digital journalism at Stuttgart Media University, thinks that the inarticulability argument doesn't work well:

"I don't see why we could not try to make Polanyi's implicit knowledge explicit and then teach it to an AI. But more importantly, it is possible that the AI finds a different way of doing what we do. It is assessed by its output, and if the output is good, who will care how it was produced? If I'm right about this, the arts & crafts aspect of journalism may not be good enough to provide a new guideline to what we teach in

journalism school. Perhaps, we have to redefine what journalists have to do (and what they are accountable for). If they can use AI to get better at researching and writing, the audience can expect more. In other words, the Bridge might take up more time at university and the Cloister a bit less." (Maeder, 2026, personal communication)

But even in a world where "AI will probably write most news articles in the future", as Maeder seems to believe, he still thinks that students will have to be trained "to become more sensitive to different styles of writing and what effects they have on the audience." (Maeder, 2026, personal communication). It is difficult to see how this training can be provided without the deep apprenticeship of reading (a lot).

There is another feature, besides inarticulable expertise, where we can differentiate from AI, and that is particularly relevant for the journalists we hope our students will become: it is the accountability that their future audiences will expect from them.

This accountability also originates in the writing process. Here is philosopher and social scientist Lisa Maria Herzog (from the "text is not the product" article), with a quote that could well serve as envoi for hopeful and well-meaning journalism teachers:

"A key point of a good text, written by a person, is that it expresses a sense of that person standing by the words they wrote, of taking a stance because it matters to them: because they want to correct what they see as a fallacy or wrong position, because it is connected to certain interests or values, because they care. ... Learning to write, as a human, also means learning to understand what one cares about, and what one is willing to take a stance on. ... Maybe it is this attitude, the virtue of truthfulness and the courage to find the right words for what one really thinks, that our ways of teaching students should focus on much more? Then we can be quite sure that no AI will ever replace us." (Herzog, 2026)

Bibliography

Note about citing AI: APA suggests following format for citations from a chat with an AI-chatbot: (example: OpenAI. (2026, May 6). Brainstorming ideas for historical research paper [Generative AI chat].
ChatGPT. https://chatgpt.com/share/69fb4a40-cd44-83ea-9d54-40288e4d4e6f

I don't agree with this format. It suggests that OpenAI is the author of the chat, but there are two authors of this text: OpenAI and the who initiated the chat and wrote the prompts. Instead, I will use following format:

OpenAI, Keustermans, E. (2026, May 6). Thought experiment flaws / FM: learning from language alone [Generative AI chat].
ChatGPT. https://chatgpt.com/share/69562582-4c3c-8008-abc4-79bc493d486a

++++++++++++++++++++++++++

Billiet, J., Opgenhaffen, M., Pattyn, B., & Van Aelst, P. (2018). The battle for the truth. Fake news and disinformation in the digital media world. KVAB Position Papers.

Bisiani, S., Abellan, A., Arias Robles, F., & García-Avilés, J. A. (2025). The Data Journalism Workforce: Demographics, Skills, Work Practices, and Challenges in the Aftermath of the COVID-19 Pandemic. Journalism Practice, 19(3), 502–522.

Caplan, B. (2019). The Case against Education – Why the Education System Is a Waste of Time and Money.

Casteels, W. (2025). Het masker van AI: Ai-geletterdheid als gids bij maatschappelijke kansen en uitdagingen. Lannoo Campus.

Cruijsberg, B. (2026). Keeping Students in the Loop: Designing an Ethical AI Assistant for Fact-Checking Education in Journalism. Master of Science in Artificial Intelligence thesis.

Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. Science Advances, 10(28).

Eeman, E. (2025, February). #115—Lost in Model Land | LinkedIn. https://www.linkedin.com/pulse/115-lost-model-land-ezra-eeman-gkoge/

EJTA AI and Data Journalism Focus Group (2025). Recommendations for Integrating AI in Journalism Teaching. EJTA - European Journalism Training Association. https://ejta.eu/index.php/2025/06/26/ejta-publishes-recommendations-for-integrating-ai-in-journalism-teaching/

Eufactcheck.eu. (2019). About us. https://eufactcheck.eu/about-us/

Eufactcheck.eu. (2019). Eufactcheck Flowchart. https://eufactcheck.eu/about-us/eufact-check-flowchart/

Fernández-Villaverde, J. [@JesusFerna7026]. (2026, Mar 30). Which of the rationales I outlined last Tuesday for traditional higher education still hold up against AI? [Tweet] https://x.com/JesusFerna7026/status/2038599152427311308

Fernández-Villaverde, J. [@JesusFerna7026]. (2026, Mar 27). On Tuesday, I shared the twelve arguments for traditional higher education. Today, I will outline the framework I plan to use to evaluate them. [Tweet] https://x.com/JesusFerna7026/status/2037527546628022333

Fernández-Villaverde, J. [@JesusFerna7026]. (2026, Mar 24). My previous post on LLMs for self-study has sparked considerable debate about the role of "traditional" higher education. [Tweet] https://x.com/JesusFerna7026/status/2036450050251833515

Ferguson, N. (2025, July 6). AI's great brain robbery—And how universities can fight back [Substack newsletter; paywalled, with link to The Times article, also paywalled]. Niall Ferguson's Time Machine. https://niallferguson.substack.com/p/ais-great-brain-robbery-and-how-universities

Ferguson, N. (2025, June 17). The Cloister and the Starship [Substack newsletter]. Niall Ferguson's Time Machine. https://niallferguson.substack.com/p/the-cloister-and-the-starship

Ferguson, N. [@nfergus]. (2025, July 6). Here are my proposals for @uaustinorg. I can assure you no established university will do anything remotely like this. 4/5 https://t.co/rpDvHOeEfH [Tweet]. Twitter. https://x.com/nfergus/status/1941803908294209756

Frateur, W., and Marain, F. (2019). A fact check flow chart for the eufactcheck project (powered by EJTA). Working paper, AP University College, Antwerp.

Frateur, W., and Marain, F. (2019). Fact Check Flowchart V1.0. Train the Trainer presentation. https://eufactcheck.eu/wp-content/uploads/2019/03/factcheckflow20190322.pdf

FT Strategies. (n.d.). A deep-dive into AI Agents in news: Cutting through the hype. https://www.ftstrategies.com/en-gb/insights/a-deep-dive-into-ai-agents-in-news

FT Strategies. (2026). Future Newsrooms Study 2026: A global benchmark of how newsrooms are changing, what they are prioritising and where they are going next.

Goldstein, E. (2026, March 27). Is AI Making Us Stupid? Cal Newport Is Worried. https://www.chronicle.com/article/is-ai-making-us-stupid-cal-newport-is-worried

Google Cloud Topics. What is Retrieval-Augmented Generation (RAG)? (n.d.). https://cloud.google.com/use-cases/retrieval-augmented-generation

Gutowska, A. (2024, July 3). What Are AI Agents? | IBM. https://www.ibm.com/think/topics/ai-agents

Heilweil, R. (2026, April 10). AI writing is the technology's bleakest use case. Fast Company. https://www.fastcompany.com/91524634/ai-writing-technology-bleakest-use-case

Herzog, L. (May 2026). The text is not the product. Crooked Timber. https://crookedtimber.org/2026/05/12/the-text-is-not-the-product/

Hume, D. (2011). An enquiry concerning human understanding. (Original work published 1748) Selby-Bigge Project Gutenberg EBook #9662. https://www.gutenberg.org/files/9662/9662-h/9662-h.htm

Hume, D. (2006). An Enquiry Concerning Human Understanding (Lewis A. Selby-Bigge, Ed.). (Original work published 1748) https://www.gutenberg.org/ebooks/9662

Hume, D. (1978). A Treatise of Human Nature. Oxford University Press. (Original work published 1739)

IBM Research Blog: What is retrieval-augmented generation (RAG)? (2023). https://research.ibm.com/blog/retrieval-augmented-generation-RAG

International Federation Of Journalists. Global Charter of Ethics for Journalists. (2019). IFJ. https://www.ifj.org/who/rules-and-policy/global-charter-of-ethics-for-journalists

Katzenberger, V. (2024). Mind the gap: Re-thinking digital skills in journalism education in Germany. Media Practice and Education, 0(0), 1–15.

Kling, A. (2026, January 8). Actually-existing UATX [Substack newsletter]. In My Tribe. https://arnoldkling.substack.com/p/actually-existing-uatx

Leo XIV (2026). Magnifica Humanitas. Encyclical Letter.

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. arXiv Preprint arXiv:2506.08872, 4.

Liang, Y., Palmer Jones, M., Felps, P., Nwatu, I., Spiegel, J. (2025). AI helps publishers monetise audience, improve workflows. International News Media Association (INMA). https://www.inma.org/blogs/Generative-AI-Initiative/post.cfm/ai-helps-publishers-monetise-audience-improve-workflows

McCord, B. (2026, January 24). The rise of the answer machines. Financial Times. https://www.ft.com/content/6f843938-007d-4925-8753-6c94613a068d
Mounk, Y. (2026, June 16). Martin Wolf on Why Trump's Economic Revolution Never Happened. https://www.persuasion.community/p/martin-wolf-2

Oakley, B., Johnston, M., Chen, K.-Z., Jung, E., & Sejnowski, T. J. (2025). The Memory Paradox: Why Our Brains Need Knowledge in an Age of AI (arXiv:2506.11015). arXiv.

OpenAI, Keustermans, E. (2026, May 6). Thought experiment flaws / FM: learning from language alone [Generative AI chat]. ChatGPT. https://chatgpt.com/share/69562582-4c3c-8008-abc4-79bc493d486a

OpenAI. Marain, F. (2025). Fact-checking Claim Process. [Generative AI chat]. ChatGPT. https://chatgpt.com/share/6a4109d8-51c0-83ed-b144-49a8f7372a10

Pattyn, B. (2020). Waarom waarheid ertoe doet. Tijdschrift Voor Filosofie, 82(1), 93–139.

Perell, D. (2025, May 28). Ezra Klein: The Case Against Writing With AI [Substack newsletter]. David Perell's How I Write. https://howiwrite.substack.com/p/ezra-klein-the-case-against-writing

Perell, D. (2025, Mar 5). Tyler Cowen: Will AI Kill Writing? https://howiwrite.substack.com/p/tyler-cowen-will-ai-kill-writing

Polanyi, M. (2005). Personal Knowledge. Towards a Post-Critical Philosophy. Routledge e-book. (first published 1958)

Pollock, S. (2026). AI agents in the media environment. Retresco Blog. https://www.retresco.de/en/blog/ai-agents-in-the-media-environment/

QS World University Rankings for Arts and Humanities 2026. (2026, April 1). Top Universities. https://www.topuniversities.com/university-subject-rankings/arts-humanities

Robbins, H. (2026, May 23). How to make a mint off the coming higher ed contraction [Substack newsletter]. Anecdotal Value. https://hollisrobbinsanecdotal.substack.com/p/how-to-make-a-mint-off-the-coming

Robbins, H. (2026, March 27). What should a university AI czar do? [Substack newsletter]. Anecdotal Value. https://hollisrobbinsanecdotal.substack.com/p/what-should-a-university-ai-czar

Robbins, H. (2026, January 31). Why Film Professors Can't Fail Students Who Won't Watch Films [Substack newsletter]. Anecdotal Value. https://hollisrobbinsanecdotal.substack.com/p/why-film-professors-cant-fail-students

Robbins, H. (2025, February 9). It's Later Than You Think [Substack newsletter]. Anecdotal Value. https://hollisrobbinsanecdotal.substack.com/p/its-later-than-you-think

Schnell, S. (2026, April 13). Repairing the Ruins: Why AI Can't Replace Education. NCR. https://www.ncregister.com/commentaries/schnell-repairing-the-ruins

Schliesser, E. (2023). On Large Language Models' Delirium (with Hume and Foucault)—Crooked Timber. https://crookedtimber.org/2023/03/26/on-large-language-modelss-delirium-with-hume-and-foucault/

Shroff, L. (2025, December 1). The People Outsourcing Their Thinking to AI. The Atlantic. https://www.theatlantic.com/technology/2025/12/people-outsourcing-their-thinking-ai/685093/

Times Higher Education (THE). (n.d.). Best Arts and Humanities Universities 2026 | World Rankings. https://www.timeshighereducation.com/world-university-rankings/2026/subject-ranking/arts-and-humanities

Van Damme, D. (2026). Technologies Smarter, Humans "Dumber"? How Technological Advances Improve and Diminish Human Cognition and What This Means for Education. Boston, MA: Center for Curriculum Redesign.

Vissers, N. (2020). EUfactcheck: A pan-European project. Methodology and lessons learned.

Williams, D. (2025, Dec 4). AI Sessions #5: How AI Broke Education. [Podcast with Henry Shevlin - Transcript provided] https://www.conspicuouscognition.com/p/ai-sessions-5-how-ai-broke-education

Walker, C. (2026, February 8). Tacit Knowledge and the SaaSpocalypse [Substack newsletter]. Out of Sample. https://cpwalker.substack.com/p/tacit-knowledge-and-the-saaspocalypse

Wolf, M. (2026, June 3). A guide for the perplexed on AI. Financial Times. https://www.ft.com/content/7cbc3986-be9f-4568-a3e5-fbd3b987f2e9?syn-25a6b1a6=1

To conclude: An inconvenient experiment

In the introduction of this essay, I pledged not to use AI as a writing-aid. I kept that promise, and readers will have to take my word for it. But, given the topic, it only seemed natural to try an AI-experiment with the completed text.

So, ***after*** I completed the text and the revisions, I prompted ChatGPT with these two questions:

PROMPT: *I wrote a booklet with 'Reflections on AI in Journalism Education' (full text attached). Can you review it? Let's do this step by step. First: short general appreciation. Then: style and grammar. Specifically, look for sentences, expressions, … that don't feel native English.*

Here is ChatGPT's answer (after 22 seconds):

https://chatgpt.com/s/t_6a89a3a52e1c8191ba12ec21f0955c4f

I particularly liked this part of the reply: "I would not "AI-polish" the whole thing; I would keep the voice, but clean these:" (A list of what I prefer to call typos follows)

This was the second prompt:

PROMPT: *Now somewhat more elaborate. You are a panel composed of a journalism teacher, a practising journalist, a higher education specialist, an AI-specialist, and a head of a Department of Journalism. Give general comments and objections. Be critical.*

And here is ChatGPT's answer:

https://chatgpt.com/s/t_6a89a4ee7b148191a58e1909bae83da4

Interesting throughout. None of the arguments or objections are invalid or jejune; most arguments add to the ideas developed in this essay. Let the discussion continue.