

Lines in the Sand: Originality and Cheating in the Age of ChatGPT

Lauren Bialystok

University of Toronto

I used to assign an article to my educational ethics students called “The Shadow Scholar,” written by someone with the alias “Ed Dante.”¹ Dante was an academic writer for hire. He would take instructions for university assignments in the humanities and social studies, often anonymously via phone or email, and, for a fee, write a text from scratch that fulfilled the parameters of the course and degree for which it had been commissioned.

What made Dante’s services worth the money was that he provided clients with a text that was *original*. Farming out the entire assignment to a competent jack-of-all-disciplines, instead of copying and pasting from SparkNotes or Wikipedia, gave students a better chance of submitting a coherent text. Importantly, this quality also meant that their subterfuge was likelier to go undetected. In lieu of the choppy texts stitched together with excerpts from random places on the internet, impervious to voice or fluidity, Dante’s clients could submit an actual novel text written by one person. Academic integrity is measured in plausible originality, and our detection mechanisms, such as Turnitin, are honed to sniff out unattributed borrowing, and not much else.

Whoever Ed Dante is (or was), he’s obviously out of business now. Generative AI such as ChatGPT and Bing can spit out everything the Shadow Scholar did in a matter of seconds, free of charge, and probably with more adjustable parameters. That such uses of AI by students are unethical is irrefutable. However, the meaning of “cheating” has entered uncharted grey territory since the arrival of robotic Ed Dantes. We need to rethink academic integrity for a world in which textual analysis and writing are no longer the unique purview of human minds.

I want to make a modest contribution to this project by probing the notion of “originality” and the ethical work it performs in our educational practices and judgments. I will argue that “cheating” in the humanities depends on the existence of some benchmark for originality, and that this benchmark

proves to be perilously vague upon closer examination. Moreover, in a rapidly evolving digital world, we need to ask whether the educational aims that originality is supposed to guarantee are still the correct ones.

WHAT'S WRONG WITH CHEATING?

There are two main things wrong with cheating. First, it frustrates the aims of education. When students cheat, they don't learn what they are supposed to. There was a '90s sitcom in which a comically dimwitted character was repeatedly caught hiding cheat sheets under his desk during tests. At the end of one episode, he gleefully announced to his sister that he found the perfect way to escape detection: he hid the answers "in his head." The audience laughs because rather than cheating, he had learned what he was supposed to.

Second, cheating is unfair to others. Since students are evaluated in relative terms using common assignments and standards, (undetected) cheating disrupts the economy of assessment. Even when a compliant student's grades are not affected by the grades of a non-compliant student, we might think that moral indignation is appropriate, since the cheating makes a mockery of the good faith effort put in by others. In addition, we may think that cheating is unfair because it disrespects the teacher's time and effort.

I take these two concerns about cheating to be reasonably comprehensive and uncontroversial. In subject areas where essays have been the default form of learning assessment for decades, the purchased essay is the paradigmatic case of cheating. It frustrates the aims of education because the student must learn nothing relevant to the subject area to hire a Dante or other ghost writer: a simple business transaction substitutes for any engagement with the material. In this sense, it is a more consummate avoidance of learning than even a cheat sheet under the desk, which presumably requires some understanding to prepare. And it is maximally unfair to other students because Dante can reliably earn a B+ or better, giving the cheater a relative reward without going through the requisite motions.

Of course, ethically significant distinctions lurk in the details. When we say that cheating frustrates the aims of education, we must be presupposing certain aims of education. In the case of a classroom test, the aim is to learn the information that can be demonstrated under test conditions; in the case

of an essay, it is to express a novel and academically well-founded argument on a given topic. What if AI has changed, or should change, those same aims? Furthermore, when we say that cheating interferes with the learning that students are supposed to experience, we must refer to some fixed conception of “learning.” What if AI actually supports learning in the traditional sense, or provides a form of learning that we hadn’t anticipated? Large Language Models (LLMs) are so confounding for education in part because the very same use of technology can be hostile to familiar forms of education, yet educational in its own right. For example, a dialogue that a student undertakes with a chatbot about course material may supersede the learning that the teacher has provided, in which case they’re not learning “what they’re supposed to,” but the learning is not necessarily irrelevant.

Similarly, when we say that cheating is unfair to other students, there are variables to iron out, which may become even more wrinkly in the age of ChatGPT. It is unfair as long as at least some students abide by academic norms, though this need not be the case. A class full of dishonest students are not *pro tanto* harming each other with their respective cheat sheets. The unfairness of cheating is more obvious when students have differential opportunity to benefit from it; not everyone is able to afford to purchase a freshly written essay from Dante, for example, even if they are all unscrupulous enough to do so. The democratization of Dante’s services changes this. If all students have recourse to AI for the completion of academic work, we may still call it cheating, but its wrongness would not necessarily have to do with unfairness in the sense of inequality of opportunity.

In the paradigmatic case of a purchased essay, we can see how originality is a rough placeholder for academic integrity. When students write the essay themselves, as intended, they are avoiding both morally objectionable aspects of cheating, no matter how bad the result is; they are fulfilling the cardinal requirement of trying to do the work. The purchased essay, by contrast, is the antithesis of trying to do the work, and we take its unoriginality to be a stand-in for both not learning the material and not being fair to other students.

WHAT WE MEAN BY “ORIGINALITY”

Originality is a useful heuristic, then, for translating our moral objec-

tions to cheating into procedural requirements for academic work. Yet, just as the two dimensions of the wrongness of cheating are schematic and in need of interpretation, so too is originality a sketchy standard that masks important distinctions. There are two senses of “originality,” both of which have been implicitly used as a test for judgments of academic integrity. In the first sense, “original” refers to something that is numerically distinct from what already exists. That is to say, it is novel. Turnitin and similar cheating checks provide a brute assessment of whether, and how much of, a text matches things that already exist on the internet. If it can’t be found, it is presumed to be novel. This is the sense of “original” being used when a teacher website (sounding very much like an ad for AI cheating) explains: “What ChatGPT creates is an original work each time you ask it. It creates something new every time you ask, and its responses are not in the databases of plagiarism checkers.”²² Notice that the novelty criterion in no way correlates to the value of what is produced. A brand-new essay may be rubbish.

More robust than the criterion of mere novelty, the second sense of “originality” refers to something that has issued from a working mind—or, to translate into human terms what is now familiar from AI, the students need to have “generated” the work “themselves.” This comes closer to capturing the aim of “learning” toward which academic assessments are ostensibly oriented, and it may bleed into the related categories of “effort” and “creativity.”²³ We care that a student can come up with something that hasn’t been said before, not because that student’s opinion or contribution is so valuable in itself, but because it is evidence for the learning that we presume must have taken place in the course of its production. Roose calls this “the basic principle that the work students turn in should reflect cogitation happening inside their brains.”²⁴

In this way, the first sense of originality—and the recourse to now-obsolete software like Turnitin—is just a stand-in for the kind of originality we actually care about. If it is numerically original, we assume, it is likely to be the result of the student’s epistemic labour (these things come apart in the case of a purchased essay). But notice again that an idea can be the result of an individual’s own thinking without necessarily being numerically original. Historically, in fact, humans in different parts of the world have made the same mathematical

or scientific breakthroughs simultaneously, without being aware of each other. When it comes to student learning, it's not the uniqueness of the specific ideas that matters, but that the resulting work issues from processes that are proper to the thinker who is supposed to be doing the thinking.

Educators at all levels, especially in the humanities, take these assumptions for granted. To summarize what we typically pack into "originality," in order to be academically legitimate, written work has to be new, meaningful, related to the intended learning, and completed by the person who is taking credit for it. These ideas are rooted in well-established values from the Western Enlightenment, including individuality, rationality, and autonomy, which tacitly guide much of our education system. John Stuart Mill declared in his treatise *On Liberty* that "all good things which exist are the fruits of originality."⁵ Kant, in his answer to the question "What Is Enlightenment?," stressed "each person's calling to think for himself."⁶ And Rousseau announced in the *Lettres Morales* that "man's first idea is to separate himself from all that is not himself."⁷ There is no denying the philosophical precursors to our dependence on originality in Western education.

HOW FAR DOES ORIGINALITY GET US?

My purpose here is not to refute the value of originality overall, but rather to illustrate its shakiness as a guide to academic ethics, especially in the current context. I will briefly identify five ways in which reliance on originality is internally at odds with other presumptions and values at work in education, and liable to be further complicated through the use of AI.

THE HIGHEST FORM OF FLATTERY

It's an open secret that the norms of academic integrity are not universal. International students, in particular from East Asia, may be bewildered by the conception of originality applied in our institutions and inadvertently commit what we call plagiarism because it resembles what they have been trained to do.⁸ Whereas Western institutions view unattributed borrowing of others' words and ideas as theft, East Asian cultures are more likely to see them as a sign of humility.

The cultural contingency of originality as academic standard would be less problematic in a Western context were it not that we sometimes value its

precise opposite. Emulation is recognized in many domains as the precursor to novel creation, and hence also has a special role in pedagogy. This can be seen in the training of artists during European modernity, the same period in which Enlightenment values such as originality were emerging. Aspiring painters were expected to *copy* the great masters' works, down to the most minute stroke. This is why we have fake Vermeers. The fact that some art collectors obsessively try to distinguish authentic paintings by Vermeer from the faithful copies of his students points to our fascination with originality as well as the longstanding norm that students emulate, not originate.

Today, education in East Asian societies continue to put great emphasis on the pedagogical value of emulation, whereas Western society has increasingly developed educational norms that encourage students' self-expression and independent scholarship. Even so, copying teachers or experts still constitutes a significant part of our educational methods. The argumentative essay is in fact one of the few areas in which we expect students to display originality from a very early stage. Is linguistic expression categorically different from other types of academic outputs? Western academic norms imply that the ability to express oneself in original language signals the ability to form clear thoughts, and therefore evinces understanding. In the following sections, I suggest that the picture is more complicated than this.

BOUNDED BRAINS

The value of originality is predicated on views about the separateness of persons and the activity of thinking. Enlightenment metaphysics presume that thinking is the product of a bounded brain in a bounded body. We can identify the source of an original idea and give credit for it only insofar as we accept that my brain is separate from your brain, and each brain thinks for itself.

Technology is often posited as the enemy of originality in this sense. The metaphysical separation between selves as thinkers extends to a stipulated separation between humans and computers. Conjuring images of Descartes sitting at his fireplace in the 17th Century, the aforementioned teacher website croons, "If we close up the computers and ask students to write something on paper, we know it's being generated by their own brains." Yet to say that we know that some product is "being generated by their own brains" still makes

numerous assumptions about what is in students' brains and how it got there. Even the students sitting at a desk with a pencil and notebook are not inventing something new *ex nihilo*, with no traceable origins in anything else they've encountered—on the contrary. Indeed, our assessments require that students' outputs are mostly unoriginal: if they wrote something utterly novel and undervivable from what we had taught them, we would not consider the result worthy of academic reward. In this sense, low-tech writing is not ontologically distinct from writing that is produced on a computer, or, necessarily, with the assistance of AI. The method is itself no guarantor of originality, and “pure” originality is not what we are actually looking for.

Despite the endurance of the Cartesian thinker as the paradigm of originality, then, many aspects of our pedagogy suggest, and mounting scientific understanding confirms, that this is not actually how new ideas come about. Human brains, like neural networks, recombine huge amounts of data, finding patterns and gaps. Mark Twain is reported to have said in 1906:

There is no such thing as a new idea. It is impossible. We simply take a lot of old ideas and put them into a sort of mental kaleidoscope. We give them a turn and they make new and curious combinations. We keep on turning and making new combinations indefinitely; but they are the same old pieces of colored glass that have been in use through all the ages.¹⁰

This aligns with the common defense that AI isn't actually coming up with something new, but merely re-packaging the collective wisdom of humanity (or, depending on your perspective, all the garbage on the internet). There remains an important ethical distinction between human-generated text and text generated by ChatGPT; but the distinction is not obviously tracked by a simple judgment of originality, since human-generated text is also a form of repackaging. Any time we call a thought or a text “original,” we are drawing a line in the sand, which could in principle be drawn elsewhere.

OOPS, I DID IT AGAIN

Another concern about originality as the standard-bearer for academic integrity is that it presumes we are aware of what we are doing. It is unfair

to penalize someone who could not have known that they were violating an academic norm (hence the case of unknowing plagiarism by a student trained in East Asia is *prima facie* ethically different from the case of plagiarism by a student who has spent her whole life in the American education system). Yet we are not always aware of where our ideas, or even our words, come from, and we may overestimate the degree to which we deserve credit for things that we express. Just as we unconsciously start talking like our partners, or retelling anecdotes that we heard from our friends, academic expression echoes the influences we absorb from everything we learn. If we happen to be reading a lot of Kant, it shouldn't be surprising that the next time we go to write a paper, we write in convoluted multi-clause sentences with made-up compound nouns.

Judgments of academic integrity assume that people are either trying to be honest or not trying to be honest. Now, while it's implausible that someone could have "accidentally" copied twenty pages verbatim from another text and passed it off as their own, we should not dismiss out of hand the innocence of recycling a phrase or a particular combination of ideas that the mind had simply stored away without attribution. Psychologists have a word for this: "cryptomnesia" is "when a forgotten memory returns without its being recognized as such by the subject, who believes it is something new and original."¹¹ Insisting on originality may underestimate our psychological limitations and mistake sincerity for cognitive transparency. This discrepancy is likely to grow; already, algorithms surreptitiously prompt us to "think" of new things. We will not be able to keep track of where all our ideas came from as AI becomes increasingly sophisticated.

COLLECTIVE ORIGINALITY

The norm of originality and its individualist metaphysics sit uncomfortably alongside other Western educational values. We also celebrate co-produced knowledge, or what we might term "collective originality," in the same contexts where Enlightenment assumptions about individual rationality usually prevail. Students are required to do groupwork where they will presumably accomplish something that exceeds their individual potential. In this case, the aim of originality seems to be punted to the collective output of the group. They are still supposed to design a novel presentation, not just copy off YouTube; but they

are supposed to share in creating the novelty. On the one hand, groupwork echoes well-founded constructivist attitudes toward learning and may promote social-emotional development. On the other hand, even in this context, each student is still supposed to think for themselves. The familiar conundrum of how to grade groupwork crystallizes the difficulty of having it both ways.

As soon as we make any judgments about academic credit, we have applied some conception of the intellectual unit who is behind a given output. Originality is usually thought to track the generative activity of a single mind, providing the benchmark against which uncited work emerges as illegitimate. The very possibility of collaborative work acknowledges the inadequacy of the individualist metaphysics that gives rise to our norms of academic integrity. On the face of it, ChatGPT is just another manifestation of how the thinking process may be collective.

TRANSLATION, PLEASE?

Typically, judgments about what counts as original text presume that the work is happening within a single natural language. Ed Dante could only write in English. But before chatbots made their triumphant appearance, AI translators such as Google Translate had already upended our means for detecting unoriginal work. The phrase “bilingual plagiarism” refers to an academic paper that is presented (and sometimes published!) as something new when it is merely a translation of something that was written in another language.¹²

This creative form of plagiarism is the sharp end of a continuum that will become increasingly blurry with AI. We do not forbid the use of a dictionary to look up or translate words, either for English language learners or for native speakers; the goal, after all, is to increase our vocabulary and the precision of our linguistic expression. Such mini-translations do not even require citation for the most part. At what point does translation turn into plagiarism?

The answer hinges on the relationship between original thought and expression in a natural language. If we considered language an inert medium for conveying the contents of our ideas, extensive writing assistance, up to and including AI-generated text, would be permissible as long as it somehow reflected independently derived ideas. In the humanities, however, we tend to consider language to be an aspect of thought itself: students demonstrate their

understanding of philosophy, for example, through coherent writing about philosophical ideas. We help them to a certain point, but expect that work submitted for evaluation will be original in both thought and expression. Our attitudes toward editing are instructive here. We forbid students from letting other people write their work for them, but we send them to campus writing centres for help with the logic, organization, grammaticality, and fluidity of their writing—none of which ends up being cited. The arbitrariness of these determinations can also be seen in the differing norms between disciplines regarding co-authorship.

Introducing multiple languages further clouds this already indeterminate territory. I've had several international students submit work that they clearly didn't write, and insist, upon questioning, that they employed no illegitimate aid in its production. Rather, they fed my assigned readings into an AI translator, read the resulting text in their native language, wrote a response in their native language, and fed the response back through a translator to produce what they considered an "original" English text. Whatever is academically unethical about this approach, originality is only a brute guide to the ethical alternative. We want students to wrestle with texts (both understanding and producing them) within a natural language, even as we acknowledge that all students need the input of others to become better writers.

DRAWING NEW LINES

The preceding considerations point to our reliance on a largely stipulative definition of "originality" for safeguarding academic integrity. Such reliance is not surprising: in order to assign academic credit to an individual student or scholar, we need recourse to some standard of how the work was produced. However, as long as the metaphysics of originality are not fully stable, our distinctions between legitimate and illegitimate work will be at least somewhat contingent. Originality is a matter of degree, not a bright line, and it functions as a stand-in for the activity of mind that we think it represents.

While the concept of originality was already fraying at the edges, AI exacerbates the worries about its utility. As the adjectives "interactive" and "generative" suggest, AI prompts and responds to human thought in novel ways, inevitably chipping away at any romantic benchmark of a bounded brain

thinking for itself. It is more important than ever to define the scope of human originality if we intend to keep using it as a gatekeeper of academic legitimacy.

While completing this task is well beyond the scope of this paper, I will conclude by revisiting the two ethical problems with cheating and gesturing at how we might rethink academic integrity in the era of AI. Cheating is still wrong because it generally frustrates the aims of education and is unfair to other students, but our methods of defining and policing it need to evolve.

First, it is still a problem if students don't learn what they're supposed to. There are uses of AI that absolutely allow students to circumvent the intended learning, even while potentially earning high grades. However, originality is not a good proxy for this; numerical originality is available at the click of a button, and students who feed an assignment directly into ChatGPT may be learning nothing. Conversely, they may be learning a great deal through sequenced prompts and iterative writing processes that AI can facilitate, even if we consider the output less "original" than an in-class essay written with a pen and paper.

Many educators have therefore begun requiring that students attach records of the interactions with AI that were used to complete an assignment. Although this is no foolproof way of ensuring that students didn't still cheat, it takes seriously the possibility that an AI-enhanced assignment could be original enough to fulfill the aims of education. The assessment would need to depend not on some test of whether all the words and ideas came from the student's brain, but on whether there was evidence of adequate learning, likely measured in terms of the student's own improved understanding through the use of AI. Developing appropriate metrics for this kind of learning will take time.

Inseparable from this maneuver is the need to re-evaluate what students are "supposed" to learn in 2024. The tedious calculator analogy is simplistic but provides a clue. Producing minimally grammatical sentences and passable textual analyses is about as hard for AI as adding two digits is for a calculator. Yet, just as we still teach multiplication and division so that students can understand when and how to use calculators, students need to learn language and logic and higher-order thinking to use AI and interact with other humans. Our assessment methods need to take account of how easily AI-assisted cheating can go undetected, but it is far too soon to say that AI has obviated the need

for humans to learn these quintessential human abilities.

To return to the second way in which cheating is wrong, it is obviously a problem if students are doing their work in ways that are unfair to other students. AI can be marshalled toward such uses, too, but times have changed: ChatGPT is a free Ed Dante for everyone. We have no control over the availability of AI and how students choose to use it. It's important to remember that fairness depends on how students are being assessed and what rewards or punishments are contingent on their observation of academic protocols. We have more control over that. Given the magnitude of lifetime opportunities attached to educational outcomes, all of which track background social inequities, we must be increasingly vigilant about doling out or withholding academic rewards on the basis of fallible judgments of originality. It is unfair if some students beat the system and others languish due to our archaic standards of academic integrity. Future work in educational ethics should examine the consequences of applying particular standards of originality throughout the whole educational pipeline.

LLMs have thrown into disarray many of our longstanding, mostly Enlightenment-based presumptions in education, especially the humanities. In order to retain a meaningful category of cheating in the new technological landscape, we need to be more precise about the exact skills we want students to master, and what constitutes learning such skills as opposed to learning to outsource them. Though the unwelcome bombshell of ChatGPT's release has damaged academic integrity for now, perhaps it is not a bad thing that we are being called upon to clarify these ideas.

REFERENCES

1 Ed Dante, "The Shadow Scholar," *The Chronicle of Higher Education*, November 12, 2010, https://www.chronicle.com/article/the-shadow-scholar/?s-ra=true&cid=gen_sign_in.

2 Matt Miller, "ChatGPT, Chatbots and Artificial Intelligence in Education," *Ditch That Textbook*, December 17, 2022, <https://ditchthattextbook.com/ai/>.

3 "Origination" is sometimes used as evidence of "creativity" here, and it has been analyzed in terms of AI's capacity to create new artwork as well as gen-

erate language. Avery Slater, "Automating Origination: Perspectives from the Humanities," in *The Oxford Handbook of Ethics of AI*, eds. Markus Dubber, Frank Pasquale, and Sunit Das (Oxford: Oxford University Press, 2021).

4 Kevin Roose, "Don't Ban ChatGPT in Schools. Teach With It," *The New York Times*, January 12, 2023, <https://www.nytimes.com/2023/01/12/technology/chatgpt-schools-teachers.html>.

5 John Stuart Mill, *On Liberty* (Indianapolis, IN: Hackett, 1978).

6 Immanuel Kant, "An Answer to the Question: What is Enlightenment?" trans. Ted Humphrey (Indianapolis, IN: Hackett, 1992).

7 Jean-Jacques Rousseau, *Autobiographical, Scientific, Religious, Moral, and Literary Writings*, trans. Christopher Kelly (Hanover: University Press of New England, 2006).

8 Leelian Kong, "Plagia-what? How the Concept of Academic Integrity Splits the World in Half," *Study International*, July 28, 2017, <https://www.studyinternational.com/news/plagia-concept-academic-integrity-splits-world-half/>.

9 Miller, "ChatGPT, Chatbots, and AI."

10 Alexander Atkins, "There Is No Such Thing as a New Idea," *Atkins Bookshelf*, October 8, 2018, <https://atkinsbookshelf.wordpress.com/2018/08/10/there-is-no-such-thing-as-a-new-idea/>.

11 Wikipedia contributors, "Cryptomnesia," *Wikipedia, The Free Encyclopedia*, accessed November 1, 2023, <https://en.wikipedia.org/w/index.php?title=Cryptomnesia&oldid=1156667917>.

12 Carmel McNaught and David Kennedy, "Bilingual Plagiarism in the Academic World," in *Ethical Practices and Implications in Distance Learning*, eds. Ugur Demiray and Ramesh Sharma (Hershey, PA: IGI Global, 2009), <https://doi.org/10.4018/978-1-59904-867-3.ch020>