

RESEARCH PAPER

Automated plagiarism

Brian Martin

University of Wollongong, Australia

Submission date: 30 August 2025; Acceptance date: 17 November 2025; Publication date: XX XX XXXX

Copyright: © 2026, Brian Martin. This is an open-access article distributed under the terms of the Creative Commons Attribution Licence (CC BY) 4.0 <https://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

ABSTRACT

The advent of large language models (LLMs) has triggered concerns about various issues, one of which is plagiarism. In contrast with two long-standing types of human plagiarism, competitive and institutionalised, LLMs introduce a new type, which can be called automated plagiarism. There are many possible responses to it, including ignoring it, banning it, polluting it, exposing it, labelling it and suing for copyright infringement. The rise of automated plagiarism may stimulate thinking about alternative ways of allocating credit for human creative work.

Keywords

plagiarism, copyright, large language models, LLM, human plagiarism, automated plagiarism

DOI: 10.13169/Prometheus.41.1.0005

Introduction

Artificial intelligence, AI, enables machine production of text that mimics human text creation. The producers are called large language models (LLMs). Users of LLMs, such as ChatGPT, Gemini and Claude, quickly become familiar with their uncanny ability to answer questions, provide summaries of topics and produce what seems like human work, outputting flawlessly expressed text in seconds.

The advent of LLMs has generated concerns in several areas. Teachers can be dismayed when their students use these programs to generate assigned essays. Many authors are upset that their works are being used to train LLMs without acknowledgement or recompense. The now ubiquitous AI assistants used by Google and others, when they provide concise summaries of texts, either downgrade or omit links to sources, causing dramatic falls in click-through rates, undermining the economic model of some creators who depend on traffic from search engines (Garanko, 2025).

The focus here is on a different but related LLM issue: plagiarism. It is typically defined as the use of others' creations as if they were one's own, without appropriate attribution. In terms of written text, this means presenting others' ideas or words as your own without giving them suitable acknowledgement. The outputs of LLMs are generated from training data, usually without acknowledgement, which suggests that plagiarism may be involved.

Plagiarism is often conflated with copyright infringement, but they are not always the same. Plagiarism is a matter of appropriate acknowledgement of sources according to relevant conven-

CONTACT: bmartin@uow.edu.au

ACCEPTING EDITOR: Peter Drahos

tions, whereas copyright infringement is a legal matter. It is possible to plagiarise text that never was or is no longer under copyright. Also, some copyrights are held by plagiarisers, most egregiously when a supervisor takes credit for the work of a subordinate. Furthermore, copyrights can be bought and sold. When a plagiarising author assigns copyright to a publisher, the connection between the plagiarism and copyright infringement becomes tenuous. The plagiarism–copyright connection is discussed later.

To understand the relationship between LLMs and plagiarism, it is useful first to examine different types of plagiarism. This is covered in the next section. Then, to continue the examination, the following section deals with a range of possible responses to AI text production, introducing some additional considerations along the way, especially copyright infringement. The overall theme is that LLMs introduce a new version of plagiarism, called here ‘automated plagiarism’, which despite its novelty has connections with human creativity.

Plagiarism

Plagiarism is commonly thought of as presenting the words of others as if they were your own. An example is copying some dialogue from Shakespeare in one’s own play, but not mentioning Shakespeare. Already there are complications. If the dialogue is well known to the audience – ‘To be, or not to be, that is the question’ – then attribution is not necessary. The author knows the audience will recognise the text as Shakespeare’s and doesn’t need to acknowledge it explicitly. This is an example of how acknowledgement is context-specific. What is appropriate depends sensitively on the genre, style and audience expectations. At the opposite end of the spectrum is the person who is listed as the author of a book, each chapter of which is an exact reproduction of an already published paper by other authors, with no acknowledgement of their authorship.

Before continuing further, it is useful to distinguish several types of plagiarism, starting with those applicable to outputs. These are drawn from a range of writings about plagiarism (Mallon, 1989; LaFollette, 1992; Anderson, 1998; Harris, 2001; Sutherland-Smith, 2008; Weber-Wulff, 2014). In *word-for-word plagiarism*, passages, paragraphs or entire works from one author are presented by another as their own. When a student copies a passage from a book or paper, by retyping it or using cut and paste, but does not give the source, this is word-for-word plagiarism. When the student makes an attempt to hide the source by changing some words so the sequence of words is similar but different in detail, this can be seen as *paraphrasing plagiarism*. A careful scholar avoids word-for-word plagiarism by putting the passage in quotation marks and giving the source and avoids paraphrasing plagiarism by making sure the passage is not too close to the original, which is given as the source, desirably with a page number. Howard (1999) introduced the idea of patchwriting (close to what is seen here as paraphrasing plagiarism), arguing that for pedagogical purposes it should be reconceived as part of learning rather than cheating.

Another type is *secondary-source plagiarism* (Bensman, 1988: pp.456–7). Imagine finding a paper that gives several original sources. You write your own text but, to save time and effort, you do not bother reading or even looking at the original sources, instead copying the references from the paper that cited them. You have not plagiarised the paper’s text, but you have taken its sources without acknowledgement. This sort of plagiarism is difficult to detect. One telltale sign is a mistake in a copied citation. Although it is plausible that secondary-source plagiarism is widespread (how many biologists cite Darwin’s *On the Origin of Species* without looking at it?), it is little studied. Noted biochemist Erwin Chargaff (1976: p.324) referred to slabs of bibliographies ‘wafted in their entirety from one paper to the next’.

A more important and quite different variety is *plagiarism of ideas*. You write your own original text but use, without attribution, the ideas of one or more others. Again, context is important. To continue the biological theme, when writing about evolutionary theory, it is not necessary to cite Darwin or Wallace because their ideas are so well known. But if the idea is specific, it is a different matter. There are allegations of cases in which a scientist submits a paper with an original

finding and a reviewer uses the ideas in the paper to write one of their own, getting it published quickly while stalling on the review.

Three of these four types of plagiarism – word-for-word, paraphrasing and ideas – are commonly discussed in studies of plagiarism (e.g., Lee *et al.*, 2023). The fourth type, secondary-source plagiarism, has been added because it is especially relevant to LLMs. These four types refer to the output, to the text or other sorts of intellectual product involved.

There is also a different way of classifying types of plagiarism that refers to the author or agent involved. The canonical idea of plagiarism involves one person, the plagiariser, presenting the work of another person, the plagiarised, as their own. A common way of thinking is that this is a form of theft. It steals credit for authorship from one person, the plagiarised, while the credit goes, unfairly, to the thief, the plagiariser. This also applies when groups are authors. The metaphor of theft reflects and abets the sense of outrage attached to plagiarism: it is widely seen as cheating. In academia, it is a cardinal sin. This sort of plagiarism can be seen as competitive because the plagiariser gains credit at the expense of the plagiarised.

Teaching students not to plagiarise is considered important, and when they do, they may be penalised severely. In the wider academic environment, when a scholar is exposed as a plagiariser, it can be a source of great shame. Stealing from a fellow scholar, taking credit for their efforts, is a low act, one widely condemned. This also applies in some other domains; for example, when a politician in giving a speech draws on another politician's speech without acknowledgement.

However, politicians plagiarise in another way, when they have speechwriters. The speechwriter is the creator of the speech, or at least a major contributor, but the politician receives the credit. Seldom does a politician say, 'Now I'm going to read a speech, most of which was written by my staffer, Jane Smith'. The same misallocation of credit occurs when political staffers send letters to constituents, respond to emails, write papers in newspapers and make submissions to formal inquiries, all in the name of the politician. In some cases, the politician never even sees the work put out in their name. Despite these activities fitting the usual definitions, they are seldom called plagiarism. In contrast to competitive plagiarism, which is stigmatised and condemned, the routine misallocation of authorship by politicians can be called *institutionalised plagiarism*. It is an aspect of a well-established system, considered routine and unobjectionable, and seldom penalised or even mentioned (Martin, 1994; 2016).

Institutionalised plagiarism is common in many fields. In corporations, it is standard practice for documents, media statements and speeches to be composed by subordinates before being signed off or delivered by top managers. Nominal authorship is sometimes justified as taking responsibility for outcomes, though why this has to involve invisibilising the true authors is never explained. The same applies in most government departments. When an inquiry or complaint arrives from a citizen, it is typically given to a low-level employee to draft an answer, which is then passed up the hierarchy for approval until a response is issued in the name of the secretary of the department or some other high-level figure. Because this process of misrepresenting authorship is so common in corporations and governments, it has been called *bureaucratic plagiarism* (Moodie, 2006). However, this sort of plagiarism is also common in other settings. In small community groups, public statements, letters and submissions may be issued under the name and signature of the president or other office bearer rather than the author of the text. There is a good reason for this: statements have more credibility when their authors have higher status.

Another venue is research publications, especially in scientific disciplines. In many fields, especially in large labs, papers have numerous authors, many of whom have contributed little or nothing to the research. The leaders of some labs have their names on every paper produced by a team member, ending up with dozens of publications annually, which pumps up the reputation of the leader, thus increasing the odds of receiving research grants that fund the salaries of junior team members, including research students. These students and some junior scientists receive less credit than they deserve, compensated by the implicit expectation that they can, if successful, become exploiters in their turn, thereby reproducing the system. In some cases, supervisors blatantly take

partial or full credit for the work of research students and assistants (Tarnow, 1999; Martin, 2013; Macdonald, 2025).

Institutionalised plagiarism reaches its apogee in ghostwriting, in which an author, a ‘ghost’, does most or all the writing and is usually paid, while someone else receives most or all of the credit (Shaw, 1991; Schlesinger, 2008; Sismondo, 2018). Ghostwriters sometimes receive some limited credit, for example, thanks in a book’s acknowledgements. Ghostwriting is most obvious in the case of celebrities, whose autobiographies are well written in a way strikingly incongruent with their own compositional capacities. Donald Trump’s book, *The Art of the Deal*, was written by Tony Schwartz. In some cases, celebrities do not even read their autobiographies. It is sometimes claimed that ghostwriting is ethical or acceptable because the ghost is paid for their work. However, being paid does not negate being plagiarised. When students purchase essays from essay mills, it is never said to be acceptable because the mill employees are paid (*cf.* Sivasubramaniam *et al.*, 2016).

In summary, competitive plagiarism typically involves one person using the work of another without adequate acknowledgement, in a context in which doing this is seen as a serious transgression of relevant norms. Institutionalised plagiarism involves a person using the work of another without adequate acknowledgement, in a context in which doing this is normal practice and usually involves a person with more power taking credit for the work of someone with less power.

Do LLMs plagiarise?

LLMs are computer programs designed to replicate human language use, using sophisticated prediction processes (Zhao *et al.*, 2023; Naveed *et al.*, 2025). An extremely simple language model draws on the frequency of words and letters in its corpus to predict the next letter and word in a sequence. LLMs use more complicated processes: deep learning and neural networks. With transformer architecture, text is converted to numbers along with information about sequences and relations between words. In doing this, they draw on the texts on which they have been trained – the more extensive the training corpus, usually the more convincing the outputs. After parallel processing, numbers are decoded back to text. When the models are sufficiently complex, the processes by which any given outputs are generated become opaque. LLMs use their training data to predict what to say, but the internal steps by which they do this cannot be easily explicated or even understood.

Because LLMs are complicated prediction machines, they usually do not reproduce text from their training data verbatim. This is akin to a famous analogy: with enough monkeys typing for long enough, eventually they will reproduce one of Shakespeare’s plays – but this is infinitesimally unlikely. The same might apply to LLMs, except for two factors. One is that the chance of reproducing a fairly short piece of text is much larger than for an entire play. The other is that complete texts are held in LLM databases, and suitable prompts can induce LLMs to regurgitate them exactly. The implication is that LLMs can sometimes plagiarise word-for-word from their sources, though this is unlikely with more advanced models.

Most LLMs are programmed to rewrite text enough to avoid word-for-word reproduction. When an LLM generates the same sequence of ideas as found in a source, but with changes in expression, the result might be called paraphrasing plagiarism, except that the paraphrase may not be of a single source, but rather of multiple sources. When LLMs include references, this involves secondary-source plagiarism when the LLM obtains the reference from a source other than the original that it cites. Finally, LLMs draw on ideas expressed in the sources on which they are trained. When this is done without suitable acknowledgement, it amounts to plagiarism of ideas. Lee *et al.* (2023: p. 2) find that ‘machine-generated texts do plagiarize from training samples, across all three types of plagiarism’, referring to verbatim (word-for-word), paraphrasing and idea plagiarism.

When a student submits an essay written in whole or part by an LLM, without acknowledging the use of the LLM, this is plagiarism, closest to competitive plagiarism, since the aim is to obtain an advantage at the expense of another creator. However, an LLM is not a typical human

creator, but itself draws on the work of human creators. Is an LLM itself a plagiarist? Some writers who comment on LLMs and plagiarism seem to assume that only humans can plagiarise (e.g., Lemley and Ouellette, 2025), so when a student uses prompts to generate the draft of an essay that uses ideas without appropriate acknowledgement, it is the student who is the plagiariser. This is a reasonable perspective because there are fundamental differences between human creativity and LLM processes. Humans use reasoning processes, both conscious and unconscious, exercising judgement to achieve a purpose. In contrast, LLMs have neither consciousness nor purpose, but rather follow instructions to produce an output based on probability considerations. In this perspective, LLMs only mimic being plagiarisers because there is no intent, just automatic processing. They do not understand what they are doing in the way humans do.

Despite this difference, there is another perspective: to conceive of an LLM as a plagiariser because its outputs replicate the activities of human plagiarisers. This is in the tradition of actor-network theory (ANT), a framework for thinking about the world in which technologies are treated as ‘actors’ or, more generally, as ‘actants’, operating in conjunction with humans in a system of networks (Law, 1986; Latour, 1987; Callon *et al.*, 1988). Actants may include scallops and door-closers, among many other possibilities. We need not adopt the entire ANT apparatus in order to think of LLMs as potential plagiarisers, ones that are automated, in which case a student who submits an LLM-generated text without appropriate acknowledgement plagiarises from a plagiariser.

As LLMs become more sophisticated and powerful and are treated as having human-like intelligence, it is methodologically useful to attribute to them the capacity to plagiarise. When humans copy from LLMs, it is analogous to a politician giving a speech written by a speechwriter when the speechwriter plagiarises from someone else’s speech. In this scenario, an LLM is analogous to a speechwriter.

If LLMs can be thought of as plagiarising, is it competitive or institutionalised? LLMs are seldom competitive plagiarists: they do not take credit for others’ work for their own advantage because LLMs do not have human agency but rather are tools for others. For the same reason, neither are LLMs institutionalised plagiarists. Based on these considerations, it is reasonable to argue that LLMs involve a novel form of plagiarism, called here ‘automated plagiarism’. LLMs use the work of others, presenting it as their own without adequate acknowledgement. In a sense, it is similar to institutionalised plagiarism because using LLMs is normalised in many circles, but in such circumstances, it is the users who plagiarise from LLMs. The LLMs are also plagiarists, but in a different way.

Referring to LLMs as plagiarists is by analogy with human plagiarism, which also provides a reference point for what does not constitute plagiarism. Software for spellcheck and auto-complete can be run by AI, but just because automated processes are involved, that does not make it plagiarism. Checking spelling and grammar has traditionally been done by humans, including copyeditors, supervisors, colleagues and friends. Relying on them has never been considered plagiarism, so the same should apply to AI-driven systems. It is when LLMs do the same things as those humans who *are* called plagiarists, whether competitive or institutionalised, that it makes sense to call LLMs themselves plagiarists.

Responses

Given the concerns about LLMs – by teachers about student cheating, by authors about use of their works, by workers about their jobs – the question arises, should anything be done and, if so, what? This question can be approached from various angles; a common one is copyright infringement. Here the focus is on automated plagiarism, which provides its own special angle on LLM social issues. To provide a broad perspective on plagiarism in the age of LLMs wider than copyright infringement, the discussion here is structured around a range of responses to automated plagiarism that might be made by those concerned about it.

Ignore

One possible response is simply to ignore automated plagiarism, treating it as a non-issue in the same way that institutionalised plagiarism is seldom brought to anyone's attention. When students submit essays created by LLMs, teachers may accept or deplore this, or find alternative methods of assessment, but these responses focus on plagiarism by the students, which is stigmatised and is in the category of competitive plagiarism. This avoids addressing plagiarism by the LLMs, namely automated plagiarism.

Ban

A diametrically opposed response is to attempt to prevent the use of LLMs. Given their widespread use and the powerful forces promoting them, this might seem like a hopeless cause, but it is worth noting nevertheless. Banning is especially important in domains where originality is prized.

In an educational context, originality is valued because demanding it encourages learning and because when students fake originality, it is seen as cheating. Students found or assumed to use LLMs without permission may be penalised. When banning is seen as futile, teachers may switch to forms of assessment for which LLMs are not helpful. Research is a domain where originality is demanded and expected; it is a prime area where competitive plagiarism is unequivocally condemned because careers depend on the attribution of original work. Using an LLM to create a research paper is seen as wrong, as fraud, so attempts are made to ban this. However, other uses of LLMs seem more acceptable; for example, for improving expression (especially for authors writing in a second language), creating tables and reformatting a reference list. As LLMs become more widely used, banning them in research publications will become more fraught, especially when detecting LLM use is difficult. (On detection, see Pudasaini *et al.*, 2024.)

Block

Authors can seek to prevent LLMs from using their work for training purposes (Chapman, 2025). When text is online, this can be achieved by embedding code in the text or its host. However, this has two limitations: companies developing LLMs can find ways around blocks, and it is very difficult to organise widespread participation in blocking schemes. For example, everyone who posts on Facebook provides text useful to LLMs. To block this use would require action by Facebook's owner, Meta, which has its own LLM and thus no incentive to block use.

Poison

LLMs are trained on vast quantities of text, some of which is high quality, but much of which is not. When training data is contaminated, outputs may be compromised in various ways, adding to the well-documented problems of racism, abuse and fabricated outputs (Bender and Hanna, 2025). For those opposed to or disturbed by LLMs or who just want to cause mischief, a possible strategy is to poison the data pool by inputting corrupted material; the number of tainted samples needed to accomplish this may not be large (Souly *et al.*, 2025). Although there is no evidence that this has been attempted at scale, it has been suggested that LLMs will eventually train on data output by LLMs themselves, causing a feedback loop that will undermine quality (Casco-Rodriguez *et al.*, 2023). Any efforts to pollute training data could add to problems with outputs.

Expose, label and mobilise

Many if not most creators of text have little awareness about its use by LLMs. Creators in this context means everyone who composes text, even memos, shopping lists and personal messages, as

long as they end up online. People who use Gmail make their messages available for training. When people are aware that their words are being used by LLMs, and thus by large companies, they may become concerned. This can be because of the widely touted impacts of AI, including mass job loss. Opposition to its destructive impacts might lead to willingness to take action against AI, including LLMs. A crucial step in raising awareness is to expose how LLMs are exploiting the works of creators – and that nearly everyone is a creator of text.

To raise concern about LLM use of text, it is useful for it to have a striking and memorable label. The importance of labels is apparent in other novel expressions, such as climate change (or global warming), sexual harassment and genocide, in which a name for a phenomenon helps in coalescing concern. So far, there is no widely used name for LLM exploitation of texts, certainly nothing familiar in everyday conversation. To have a good chance of controlling LLM text use, it is also necessary to mobilise support. On other issues, this means organising campaigns that trigger greater awareness and support. If campaigns take off, the result can be a social movement in the tradition of the labour, feminist, anti-racist and peace movements. If an anti-AI movement emerges, LLMs might be restrained. So far, though, this is no more than nascent.

Sue

Many concerns about LLMs and authorship focus on copyright rather than plagiarism. In 2023, the *New York Times* sued OpenAI and Microsoft for copyright infringement based on their use of *Times* papers to train such LLM models as ChatGPT. Setting aside the many legal complications (e.g., Pope, 2024), the *Times* lawsuit gives voice to a widespread disquiet among authors about LLMs.

When authors feel their work has been ‘stolen’ – namely, others have taken credit for the products of their labour – one impulse is to sue. What law has been violated? The most obvious is copyright. In most countries, copyright in creative works lasts for a long time, typically 70 years after the author’s death. And copyright is wide-ranging. It covers business memos, personal emails, course syllabi, street directories and doodles on scraps of paper, as well as more obvious products, such as papers and books. Copyright also covers films, photos and computer code. Unlike patents and trademarks, no registration is required to obtain copyright: it is conferred automatically the moment creation occurs. This means that every draft of a novel, and every scribble, is copyrighted as soon as it is produced (Halbert, 1999, 2006; May, 2000; Vaver, 2000; Bellos and Montagu, 2024).

When someone uses copyrighted material improperly, this is called, in law, ‘copyright infringement’. However, some copyright owners think this sounds tame and prefer to call it ‘theft’. Owners of film rights like to call it ‘piracy’, even though the connection between piracy and copyright infringement is tenuous and is certainly not found in law (Loughlin, 2007). Another consideration is that while there is an overlap between plagiarism and copyright infringement, they are different in several ways (Stearns, 1992; Vaver, 2026). The trouble with the language of theft and piracy is that when copyright is infringed, the owner still has access to the creative work. If someone steals a computer, car or pair of shoes, the thief has the object, and the legal owner does not. Copyright is different, as are all forms of intellectual property. Critics argue that intellectual property is better thought of not as a form of property but rather as a ‘monopoly privilege’: the government guarantees the IP owner exclusive rights, in essence a monopoly, for the duration of protection (Drahos, 1996). As noted, the duration in the case of copyright is almost indefinite.

Authors may believe that copyright is to their advantage, a protection for an income stream, and for a small minority this is true. But the bigger picture is different. Most copyrights of value are owned by corporations, including Hollywood producers, software companies and big publishers, better termed ‘data companies’ (Lamdan, 2023). They receive the bulk of returns from copyright. It was always thus. Delving into the history of copyright reveals that it was never about benefits to authors but rather about protecting the profits of publishers (Rose, 1993; Bellos and Montagu, 2024).

Authors of books and papers often have to assign their copyright to the publisher, which thereafter has a stake in extracting maximum returns. For journal papers, this means being hidden behind paywalls, so works are available only to those who can pay or, more commonly, to students and staff at universities that can pay. One of the central problems with copyright is that it can be bought and sold, which means that those with the greatest market power usually end up with control and most of the income. Authors, aside from a few who are listed as the authors of bestsellers or textbooks, receive a pittance.

There's yet another issue. The official rationale for copyright, and intellectual property more generally, is not to provide an income for creators but rather to foster greater intellectual productivity. In the words of the US Constitution (article I, section 8), Congress has the power 'To promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries'. However, for copyright, the fostering of productivity through limited protection has been gradually subverted by the repeated extensions of the duration. The bulk of sales for most books are in the first six to twelve months of publication: if copyright were for a year or two, this would cover most sales. Extending copyright to decades after the author's death does little to foster greater productivity. After all, what author is stimulated to more feverish efforts by the promise that returns will continue for 70 rather than 10 years after their death?

With this background, it is time to examine the strategy of suing LLM companies for copyright infringement. This sounds appropriate on the surface because the companies are training their models on text created by a great many authors whose works are under copyright. However, there is another important distinction. Copyright normally covers only the expression of ideas, not the underlying ideas themselves. If an LLM draws too heavily on a single source, for example a substantial paragraph, then copyright may be infringed. But if the wording is changed sufficiently, then copyright does not apply. If the US Constitution were rewritten in a different style, with different words and sentence structures, then copyright would not be infringed, even if the Constitution were still under copyright, which it is not.

Imagine that a court rules that LLMs infringe copyright when they output text too close to the original. Reprogramming LLMs could ensure they do not do this. In the face of such a ruling, an LLM can be made copyright-compliant in the same way that a human could take care to avoid composing text not too similar to already published sources. Indeed, an LLM probably could do this more conscientiously, being trained on a larger volume of source text than any human could comprehend. Only in extremely narrow topic areas would there be any risk of infringement. This sort of technological workaround may not be sufficient because courts might invoke other principles: the *Times vs. OpenAI* case has many facets. The point is that copyright infringement is not a solid basis for challenging LLM use of sources.

In terms of plagiarism, LLMs can mostly avoid word-for-word plagiarism. What they do more effectively is plagiarise ideas. They take ideas from various creators, mix them together and produce a text that draws on the ideas without acknowledging the sources or, when LLMs provide references on request, without acknowledging secondary sources.

Support creators in other ways

Creators argue that they deserve both credit and compensation for their efforts. Credit on its own does not provide an income, but it can benefit careers. Scholars seldom earn much from royalties, but being the author of scholarly publications is the basis for obtaining jobs and promotions. For others, credit for authorship is valuable for obtaining speaking engagements, sponsorships, commissions and prizes. For some musicians, trying to prevent pirate recordings of concert performances is less important than the greater visibility they gain from these recordings, leading to larger live audiences.

What about journalists whose efforts are used by LLMs? A simple solution is to increase payments for stories and not bother with copyright. In any case, copyright fees seldom go directly to journalists; they go to publishers. A more radical option is the introduction of a guaranteed annual income, also known as a universal basic income (UBI) (Bidadanure, 2019). Many independent authors receive very little for their efforts. By having a liveable income, they could, if they wished, write to their heart's content and there would be no need to lock their words behind paywalls for the indefinite future. The alternative of a UBI raises all sorts of issues. Its relevance here is as a different model for rewarding creators, one in which outputs could immediately go into the public domain, the commons, thereby enabling greater productivity from others.

Create domains

Imagine the division of 'intellectual space' into two domains: one where credit is valued and the other where it is not assigned. In the credit domain, assigning correct authorship is important. An example is the publication of scientific papers, where credit is vital to careers. Some scholarly journals ask authors whether they use LLMs in their research and, if so, to specify exactly how. In the no-credit domain, authorship is treated as unimportant or irrelevant. In this domain, AI could operate unimpeded. An example is the AI summary provided by Google in response to a query. Readers are not expected to be concerned about the sources used by the LLMs in generating the summary. Another example is using LLMs to brainstorm options, which are then examined by humans, who are more interested in the options than in how they were generated.

Human and automated plagiarism of ideas

Some unlikely similarities between human text creation and LLMs can be found by looking at plagiarism, especially in giving credit for ideas. Humans are not very good at attributing the sources of ideas. Young children learn to speak and count, but no one expects them to give credit to the creators of language and mathematics. These domains are considered 'common knowledge,' part of the public domain, the heritage of a cumulative process of human creativity. Only when knowledge is more specialised would anyone expect credit to be given to sources. In everyday conversation, you might say, 'I read online about a new therapy'. In most cases, the author of the online paper is not the developer of the therapy. Attribution is not treated as important, except to bolster credibility: 'There was some study about it published in a medical journal'.

It is common for people to hear or read something, forget the source and think it was their own idea. The technical term for this is cryptomnesia, which is important in scientific discovery (Merton, 1973: pp.402–12). The point is that people hear and read lots of things; these become mixed up in their minds, so attributing credit for the origin of ideas is seldom accurate or complete. Even in science, where giving credit is important, most authors are not very good at it. In one of the few careful studies of the accuracy and comprehensiveness of scientific referencing, MacRoberts and MacRoberts (1986; 1989) find that only a small fraction of the influence on a scientific paper is captured by citations. There are all sorts of reasons, including ignorance of the literature, uncited informal influences and not referencing basic assumptions and background knowledge.

LLMs, like humans, are not very good at attributing their sources. When they are asked to give references for their statements, these are copied from some other source. This suggests that humans and LLMs are alike in a certain way: in making statements, they often draw from a wide range of sources but either do not know what these are or make inaccurate attributions. This parallel between humans and LLMs has limits. Humans think using concepts and expressions of them, often intertwined. LLMs, in contrast, operate using texts; they only seem to think. With few exceptions, humans writing text do not draw directly on the text of others. The exceptions are when they sit with a text before them while they compose their own, and when they have photographic or other

word-for-word memories. LLMs, in contrast, routinely draw on others' texts. But rather than reproduce them word-for-word, they usually rewrite and mix them to produce their own outputs.

The result is the same: systematic plagiarism of ideas, so pervasive that it seems like creative output. And sufficient modification of others' ideas is itself a form of creation. This is only to say that human creativity is necessarily built on the work of others. As Isaac Newton famously admitted, 'If I have seen further than others, it is by standing upon the shoulders of giants'. Humble human creators acknowledge that their works depend on the prior accomplishments of generations of others. In this context, to claim their outputs are solely their own and deserve to be legally protected from use by others is a sort of hubris. Yes, individual human creation does exist, but only as part of a broader social process in which every creator builds on a vast corpus of ideas and products that went before and provides the foundation and testing ground for new contributions.

Conclusion

Imagine a conscientious scientist working on a literature review, surrounded by piles of reprints or a folder filled with numerous PDFs. After making sense of the field, the scientist explains who has done what, putting each relevant reference in an appropriate context, carefully giving credit where credit is due. No relevant source is omitted, and no irrelevant or unimportant study is mentioned. This is an ideal picture, only seldom achieved, but nonetheless useful for a comparison with how an LLM would proceed. Digesting the text of all the PDFs, the LLM would create a summary of the field, drawing on the words used. If the LLM were tasked with providing a fully referenced summary, it would draw on the literature reviews of the papers in its system. If our ideal scientist's literature review is available, the LLM will use it for choosing references to cite.

The conscientious scientist grasps the meanings of the prior work in the field and seeks to acknowledge each relevant and important study. The LLM digests the texts of work in the field, producing a new text summarising these texts. If the LLM gives no references, it is a clear case of plagiarism of ideas, in other words, of making a new expression of prior expressions. If the LLM gives references, it is most likely a case of plagiarism of secondary sources, namely of listing references that were used by others. This is the basis for saying that LLMs commonly involve plagiarism, called here automated plagiarism to distinguish it from two human types, competitive and institutionalised. The conscientious scientist is careful not to copy text from prior papers, thus avoiding competitive plagiarism, and is careful to avoid having co-authors who did not contribute substantially to the research, thus avoiding institutionalised plagiarism. The scientist, though, might be tempted to use an LLM to prepare a draft of the literature review. This would be yet a different form of plagiarism, unless the LLM were acknowledged, as some journals now expect.

In everyday conversation and other informal situations, credit for ideas is seldom important. But where giving appropriate credit is deemed crucial, where people compete for credit – for example, in student work and academic papers – plagiarism is seen as a serious transgression, deserving penalties. Yet, in such domains, there is a major contradiction, perhaps hypocrisy: institutionalised plagiarism, when powerful or high-status figures take credit for the work of subordinates and this is treated as the norm. Enter LLMs, with their own type of plagiarism, neither competitive nor institutionalised in the usual senses. LLMs introduce automated plagiarism in many of their applications and have caused angst among creators whose works are being exploited without suitable credit or compensation. However, rather than being called plagiarism, the outputs of LLMs are variously used, admired and deplored. The main form of organised resistance relies on alleging copyright infringement, but this is a precarious basis for opposition because copyright is itself so poorly correlated with creative contributions.

Another option is to use the rise of LLMs as an opportunity to rethink human creativity, to get away from the idea that new text and new thoughts are the product of single creators who need credit for them for their livelihood, careers or fame. However, new models may involve completely different ways of rewarding creative work, for example, giving everyone a guaranteed income and

assigning all creations to the commons. Moving towards such alternatives might challenge economic inequality and the role of intellectual property in maintaining it. Automated plagiarism thus sits near the core of some of the crucial contradictions of the current economic order.

Acknowledgements

Thanks to David Bellos, Alex Montagu, David Vaver and two anonymous reviewers for valuable comments.

References

- Anderson, J. (1998) *Plagiarism, Copyright Violation and other Thefts of Intellectual Property: An Annotated Bibliography with a Lengthy Introduction*, McFarland, Jefferson NC.
- Bellos, D. and Montagu, A. (2024) *Who Owns this Sentence? A History of Copyrights and Wrongs*, Mountain Leopard Books, London.
- Bender, E. and Hanna, A. (2025) *The AI Con: How to Fight Big Tech's Hype and Create the Future We Want*, Bodley Head, London.
- Bensman, J. (1988) 'The aesthetics and politics of footnoting', *Politics, Culture, and Society*, 1, pp.443–70.
- Bidadanure, J. (2019) 'The political theory of Universal Basic Income', *Annual Review of Political Science*, 22, pp.481–501.
- Callon, M., Law, J. and Rip, A. (eds) (1988) *Mapping the Dynamics of Science and Technology: Sociology of Science in the Real World*, Macmillan, London.
- Casco-Rodriguez, J., Alemohammad, S., Luzi, L., Humayun, A., Babaei, H., LeJeune, D., Siahkoochi, A. and Baraniuk, R. G. (2023) 'Self-consuming generative models go MAD', *37th Conference on Neural Information Processing Systems*, available at https://research.latinxinai.org/papers/neurips/2023/pdf/Josue_CascoRodriguez.pdf (accessed December 2025).
- Chapman, S. (2025) 'Protect your site from ChatGPT in 2025: how to block LLM crawlers', *PrivacyJournal.net*, available at <https://www.privacyjournal.net/block-llm-crawlers/> (accessed December 2025).
- Chargaff, E. (1976) 'Triviality in science: a brief meditation on fashions', *Perspectives on Biology and Medicine*, 19, pp.324–33.
- Drahos, P. (1996) *A Philosophy of Intellectual Property*, Dartmouth, Aldershot.
- Garanko, J. (2025) 'Semrush AI overviews study: what 2025 SEO data tells us about Google's search shift', *Semrush Blog*, 23 July, available at <https://www.semrush.com/blog/semrush-ai-overviews-study/> (accessed December 2025).
- Halbert, D. (1999) *Intellectual Property in the Information Age: The Politics of Expanding Ownership Rights*, Quorum Books, Westport CT.
- Halbert, D. (2006) *Resisting Intellectual Property*, Routledge, London.
- Harris, R. (2001) *The Plagiarism Handbook: Strategies for Preventing, Detecting, and Dealing with Plagiarism*, Pyczak Publishing, Los Angeles.
- Howard, R. (1999) *Standing in the Shadow of Giants: Plagiarists, Authors, Collaborators*, Ablex, Stamford CT.
- LaFollette, M. (1992) *Stealing into Print: Fraud, Plagiarism, and Misconduct in Scientific Publishing*, University of California Press, Berkeley CA.

- Lamdan, S. (2023) *Data Cartels: the Companies that Control and Monopolize our Information*, Stanford University Press, Stanford CA.
- Latour, B. (1987) *Science in Action: How to Follow Scientists and Engineers through Society*, Open University Press, Milton Keynes.
- Law, J. (ed.) (1986) *Power, Action and Belief: A New Sociology of Knowledge?* Routledge and Kegan Paul, London.
- Lee, J., Le, T., Chen, J. and Lee, D. (2023) ‘Do language models plagiarize?’ in *Proceedings of the ACM Web Conference 2023 (WWW’23), May 1–5, Austin TX*, Association for Computing Machinery, New York.
- Lemley, M. and Ouellette, L. (2025) ‘Plagiarism, copyright, and AI’, *University of Chicago Law Review*, in press.
- Loughlin, P. (2007) ‘“You wouldn’t steal a car ...”: intellectual property and the language of theft’, *European IP Review*, 29, 10, pp.401–5.
- Macdonald, S. (2025) ‘In search of an author’, *Prometheus*, 40, 3, pp.123–5.
- MacRoberts, M. and MacRoberts, B. (1986) ‘Quantitative measures of communication in science: a study of the formal level’, *Social Studies of Science*, 16, 1, pp.151–72.
- MacRoberts, M. and MacRoberts, B. (1989) ‘Problems of citation analysis: a critical review’, *Journal of the American Society for Information Science*, 40, pp.342–9.
- Mallon, T. (1989) *Stolen Words: Forays into the Origins and Ravages of Plagiarism*, Ticknor and Fields, New York.
- Martin, B. (1994) ‘Plagiarism: a misplaced emphasis’, *Journal of Information Ethics*, 3, 2, pp.36–47.
- Martin, B. (2013) ‘Countering supervisor exploitation’, *Journal of Scholarly Publishing*, 45, 1, pp.74–86.
- Martin, B. (2016) ‘Plagiarism, misrepresentation, and exploitation by established professionals: power and tactics’ in Bretag, T. (ed.) *Handbook of Academic Integrity*, Springer, Singapore, pp.913–27.
- May, C. (2000) *A Global Political Economy of Intellectual Property Rights: The New Enclosures?* Routledge, London.
- Merton, R. (1973) *The Sociology of Science: Theoretical and Empirical Investigations*, University of Chicago Press, Chicago.
- Moodie, G. (2006) ‘Bureaucratic plagiarism’, *Plagiary: Cross-Disciplinary Studies in Plagiarism, Fabrication, and Falsification*, 1, 6, pp.1–5.
- Naveed, H., Khan, A., Qiu, S., Saqib, M., Anwar, S. et al. (2025) ‘A comprehensive overview of large language models’, *ACM Transactions on Intelligent Systems and Technology*, 16, 5, Paper 206, pp.1–72.
- Pope, A. (2024) ‘NYT v. OpenAI: The Times’s about-face’, *Harvard Law Review*, 10 April, available at <https://harvardlawreview.org/blog/2024/04/nyt-v-openai-the-timess-about-face/> (accessed December 2025).
- Pudasaini, S., Miralles-Pechuán, L., Lillis, D. and Llorens Salvador, M. (2024) ‘Survey on plagiarism detection in large language models: the impact of ChatGPT and Gemini on academic integrity’, arXiv:2407.13105v1.

- Rose, M. (1993) *Authors and Owners: The Invention of Copyright*, Harvard University Press, Cambridge MA.
- Schlesinger, R. (2008) *White House Ghosts: Presidents and their Speechwriters*, Simon & Schuster, New York
- Sismondo, S. (2018) *Ghost-Managed Medicine: Big Pharma's Invisible Hands*, Mattering Press, Manchester.
- Shaw, E. (1991) *Ghostwriting: How to Get into the Business*, Paragon House, New York.
- Sivasubramaniam, S., Kostelidou, K. and Ramachandran, S. (2016) 'A close encounter with ghost-writers: an initial exploration study on background, strategies and attitudes of independent essay providers', *International Journal for Educational Integrity*, 12, paper 1.
- Souly, A., Rando, J., Chapman, E., Davies, X., Hasircloglu, B. *et al.* (2025) 'Poisoning attacks on LLMs require a near-constant number of poison samples', arXiv:2510.07192v1.
- Stearns, L. (1992) 'Copy wrong: plagiarism, process, property, and the law', *California Law Review*, 80, 2, pp.513–53.
- Sutherland-Smith, W. (2008) *Plagiarism, the Internet and Student Learning: Improving Academic Integrity*, Routledge, New York.
- Tarnow, E. (1999) 'The authorship list in science: junior physicists' perceptions of who appears and why', *Science and Engineering Ethics*, 5, pp.73–88.
- US Constitution (1787) In *The Constitution of the United States of America: Analysis and Interpretation*. Library of Congress, Washington DC.
- Vaver, D. (2000) 'Intellectual property: the state of the art', *Law Quarterly Review*, 116, pp.621–37.
- Vaver, D. (2026) *Intellectual Property Law*, Third Edition, University of Toronto Press, Toronto.
- Weber-Wulff, D. (2014) *False Feathers: A Perspective on Academic Plagiarism*, Springer, Heidelberg.
- Zhao, W., Zhou, K., Li, J., Tang, T., Wang, X. *et al.* (2023) 'A survey of large language models', arXiv:2303.18223v13.