



NUS Law Working Paper No 2023/002

AI-Generated Content is Taking over the World. But Who Owns it?

Simon Chesterman

chesterman@nus.edu.sg

[January 2023]

© Copyright is held by the author or authors of each working paper. No part of this paper may be republished, reprinted, or reproduced in any format without the permission of the paper's author or authors.

Note: The views expressed in each paper are those of the author or authors of the paper. They do not necessarily represent or reflect the views of the National University of Singapore.

AI-generated content is taking over the world. But who owns it?

When artificial intelligence produces artwork and text indistinguishable from human creations, are the economic foundations of creativity under threat?

SIMON CHESTERMAN*

Unless you've been living under a rock for the past few months, you've probably seen works of art or text produced by artificial intelligence (AI). Tools such as Stable Diffusion can generate images based on a prompt, [flooding social media](#) with portraits of users in Renaissance-style paintings.

But it was the release of ChatGPT last November that really caused head-scratching and hand-wringing.

GPT stands for generative pre-trained transformer and it is, in essence, a chatbot. Unlike precursors with limited scope and [clumsy outputs](#), however, ChatGPT responds quickly and clearly to questions about almost any topic.

Simpler queries are akin to using Siri and the like. What has drawn attention is the detailed responses in paragraph form to more open-ended prompts. If Wikipedia's mission is to crowdsource information and Google's is to organise it, ChatGPT is like having a reasonably intelligent friend who will try to synthesise all of that into a couple of paragraphs.

In the weeks after ChatGPT went public, the Internet was awash with text composed by the AI system, ranging from the humorous to the insightful.

* National University of Singapore. Email: chesterman@nus.edu.sg. I'm grateful to N Chesterman, Mikhail Filippov, Gao Yuting, Hahn Jungpil, and David Tan for comments on earlier versions of this text. Errors and omissions can be blamed on ChatGPT.

It quickly became evident, however, that this is not just the latest Instagram filter or “Tide pod challenge”. For AI that can generate articulate and detailed text is a direct challenge to the knowledge economy that underpins the digital age and keeps many of us employed.¹

1 The knowledge economy

“[Knowledge workers](#)” was the term introduced in 1959 by management consultant Peter Drucker for non-routine problem solvers.² People who “[think for a living](#)” earn through their ability to analyse and write — something that ChatGPT can replicate in almost no time and at almost no cost.

Journalists, already taking a beating as their readers turn from traditional to social media, now [face the prospect](#) of technology taking over the writing task as well. Yet that same threat confronts anyone who analyses or writes for a living, such as lawyers and even, gasp, academics. Applications are not limited to prose, as ChatGPT has demonstrated proficiency in [coding](#) as well as [poetry](#).

Further upstream, teachers [worry](#) that their students will use ChatGPT for their [assignments](#). A [study](#) last month argued that it is only a matter of time before an AI system can pass a US Bar Exam.³

I’m not especially worried about the [death of education](#). Students who want to cheat have always been able to cheat, and educators should not be assigning tests that can be outsourced to ChatGPT or anyone else. Much as books and calculators reduced the need for us to memorise or multiply, teachers need to help students take advantage of the new tools while cultivating skills that computers cannot — yet — supplant.

¹ Similar challenges may apply to the manner in which many tech companies generate revenue. Google’s [business model](#) depends on you clicking on ads. That’s much more likely if you are given a list of possible responses to a search request, some of which are sponsored. ChatGPT offers a single coherent answer — easier to act on, but less likely to reflect the complexity of a topic or introduce diverse views (or generate revenue). This also increases the likelihood of getting a clear and confident response that is also dead wrong.

² P.F. Drucker, *The Landmarks of Tomorrow* (Harper and Row 1959).

³ Michael James Bommarito and Daniel Martin Katz, “GPT Takes the Bar Exam” (December 29, 2022). Available at SSRN: <https://ssrn.com/abstract=4314839>.

But unless we install some guardrails, there's a real danger that these tools will discourage creativity more generally by removing rewards and distorting incentives.

Here, I will focus on just two proposals: being clear about who owns the content they produce, and requiring that such content be identifiable as produced by AI rather than a human.

2 Ownership

So, who owns the content produced by ChatGPT?

I decided to ask [the system itself](#), and got two very different answers.

“I do not have the ability to own intellectual property or any other legal rights,” it said at first. “Any text or other content that I generate is the property of OpenAI, as the creator and owner of the tool that I am.”

OK, though that “I am” is a little creepy. (I've edited some of its responses for brevity.)

I pointed out that OpenAI itself now explicitly [states](#) that it will not claim copyright over any content generated by ChatGPT.⁴

This led to a revised answer: “The text generated is not the intellectual property of the model itself. Instead, the intellectual property rights belong to the person or entity who has commissioned the model to generate the text.”

Clear and concise, but also wrong.

In most jurisdictions, automatically-generated text does not receive copyright protection at all.

The US Copyright Office has stated that legislative protection of “original works of authorship”⁵ is limited to works “created by a human being”. It will not register works

⁴ Interestingly, this was not the initial position. OpenAI at first claimed ownership of all output from DALL-E, only amending this in late 2022 to state that users now own the generated content. See Andres Guadamuz, 'DALL-E Goes Commercial, But What About Copyright?', *TechnoLlama* (25 July 2022); Steven Ellison, 'Who Owns DALL-E Images?', *FindLaw* (29 August 2022).

⁵ 17 USC § 102(a).

“produced by a machine or mere mechanical process that operates randomly or automatically without *any* creative input or intervention from a human author.”⁶

The word “any” is key and begs the question of what level of human involvement is required to assert authorship.⁷

3 Author, author

Early photographs, for example, were not protected because the mere capturing of light through the lens of a *camera obscura* was not regarded as true “authorship”.⁸ It took an iconic picture of Oscar Wilde going all the way to the US Supreme Court before copyright was recognised in mechanically-produced creations.⁹

⁶ *Compendium of U.S. Copyright Office Practices* (3rd edn, U.S. Copyright Office 2019), § 313.2 (emphasis added).

⁷ See Daniel J. Gervais, 'The Machine as Author' (2020) 105 *Iowa Law Review* 2053.

⁸ Madeleine de Cock Buning, 'Artificial Intelligence and the Creative Industry: New Challenges for the EU Paradigm for Art and Technology by Autonomous Creation' in Woodrow Barfield and Ugo Pagallo (eds), *Research Handbook on the Law of Artificial Intelligence* (Edward Elgar 2018) 511 at 524.

⁹ *Burrow-Giles Lithographic Co v. Sarony*, 111 US 53 (1884). Arguments continued, however, with Germany withholding full copyright of photographs until 1965. Axel Nordemann, 'Germany' in Ysolde Gendreau, Axel Nordemann, and Rainer Oesch (eds), *Copyright and Photographs: An International Survey* (Kluwer 1999) 135.



Oscar Wilde, source [Smithsonian Magazine](#), May 2004.

The issue today is distinct: not whether a photographer can own images passively captured by machines, but who might own new works actively created by them.

A computer program like a word processor does not own the text typed on it, any more than a pen owns the words that it writes. But AI systems now write news reports, compose songs, paint pictures. These activities generate value — can and should they be protected by the law?

At present, the answer in most places is no. Unless there is an identifiable human author, copyright will not apply.

The policy behind this is often said to be incentivising and rewarding innovation. This has long been dismissed as unnecessary or inappropriate for computers. “All it takes,” Pamela

Samuelson wrote in 1986, “is electricity (or some other motive force) to get the machines into production.”¹⁰

Indeed, protecting such works might *disincentivise* innovation — by humans, at least. AI has already unleashed an [economic tornado](#) in the art world, massively lowering the cost of producing original images. (Full disclosure: the cover of my forthcoming novel includes AI-generated art.)

Assuming that we desire a thriving arts sector that gainfully employs humans, we should be clear that their creations are protected while machine creations are not. (Further disclosure: a human is credited with the overall design of the cover.)

The better answer, then, is that automatically generated content should not be eligible for copyright protection, but edited and curated content that draws on such material may be owned by the person doing the editing and curating.

I fed that into ChatGPT and it agreed that this was correct, sensibly adding that I should seek legal advice if I had further questions.

An alternative approach, adopted in Britain is to have more limited protections for “computer-generated” work, the “author” of which is deemed to be the person who undertook “the arrangements necessary for the creation of the work”.¹¹ Similar legislation has been adopted in New Zealand,¹² India,¹³ Hong Kong,¹⁴ and Ireland.¹⁵ Though disputes about who took the “arrangements necessary” may arise, ownership by a recognized legal person or by no one at all remain the only possible outcomes.¹⁶

¹⁰ Pamela Samuelson, 'Allocating Ownership Rights in Computer-Generated Works' (1986) 47 University of Pittsburgh Law Review 1185, 1199.

¹¹ Copyright, Designs and Patents Act 1988 (UK), s 9(3). ‘Computer-generated’ is defined in s 178 as meaning that the work was ‘generated by computer in circumstances such that there is no human author of the work’.

¹² Copyright Act 1994 (NZ), s 5(2)(a).

¹³ Copyright Amendment Act 1994 (India), s 2.

¹⁴ Copyright Ordinance 1997 (HK), s 11(3).

¹⁵ Copyright and Related Rights Act 2000 (Ireland), s 21(f).

¹⁶ See, eg, *Nova Productions v. Mazooma Games* (2007) [2007] EWCA Civ 219; Abbe Brown et al, *Contemporary Intellectual Property: Law and Policy* (5th edn, Oxford University Press 2019) 100-01.

The duration is generally for a shorter period, and the deemed “author” is unable to assert moral rights — including the right to be identified as the author of the work.¹⁷

As human authorship becomes more ambiguous, that middle ground may help preserve and reward flesh and blood authorship, while also encouraging experiments in collaboration with our silicon and metal partners.

Europe is actively considering such a measure.¹⁸ The Singapore Academy of Law’s Law Reform Committee [proposed](#) something similar in 2020, but only traditional human authorship remains recognised under the new Copyright Act adopted the following year. AI-assisted works may still warrant protection if there is a causal connection to a human exercising input or control, though determining the threshold for that connection is [left to the courts](#).¹⁹

4 Transparency

What about the reader or viewer? A common thread in many discussions about AI is that [we should know](#) when we are interacting with a machine or a person.

It might seem like a simple question, but AI-assisted decision-making increasingly blurs that line. Some chatbots, for example, start on automatic for basic queries, moving through suggested responses that are vetted by a human, escalating up to direct contact with a person for unusual or more complex interactions.²⁰

¹⁷ Copyright, Designs and Patents Act, s 12(7) (protection for such works is limited to 50 years, rather than 70 years after the death of the author), s 79 (exception to moral rights). A World Intellectual Property Organization (WIPO) issues paper recognized the dilemma, noting that excluding these works would favour ‘the dignity of human creativity over machine creativity’ at the expense of making the largest number of creative works available to consumers. A middle path, it observed, was to offer ‘a reduced term of protection and other limitations’. Revised Issues Paper on Intellectual Property Policy and Artificial Intelligence (World Intellectual Property Organisation, WIPO/IP/AI/2/GE/20/1 REV, 21 May 2020), para 23. See also Marcus du Sautoy, *The Creativity Code: Art and Innovation in the Age of AI* (Harvard University Press 2019) 102; Ryan Abbott, *The Reasonable Robot: Artificial Intelligence and the Law* (Cambridge University Press 2020) 71-91.

¹⁸ Stéphane Séjourné, Draft Report on Intellectual Property Rights for the Development of Artificial Intelligence Technologies (European Parliament, Committee on Legal Affairs, 2020/2015(INI), 24 April 2020), paras 9-10.

¹⁹ See David Tan and Wee Liang Tan, ‘AI, Author, Amanuensis’ (2022) 5 *Journal of Intellectual Property Studies* 1.

²⁰ See, eg, Pavel Kucherbaev, Alessandro Bozzon, and Geert-Jan Houben, ‘Human-Aided Bots’ (2018) 22(6) *IEEE Internet Computing* 36.

For the raw text and images produced by AI, at least, it should be possible to disclose their provenance. To guard against misrepresentation, various efforts are underway to detect AI-generated text through anti-plagiarism software.

A more difficult but more effective approach would be to “[watermark](#)” text and images in a manner that is invisible to users but detectable using a key. Given the likely spread of the underlying software, this would be practical only if it is required by law. Even then, however, the spread of deepfake porn points to the difficulty of policing any such rules.

5 Brave New World?

It is important to understand that — however human ChatGPT might appear — it does not “understand” the text that it produces in any meaningful sense. Such large language models make probabilistic guesses based on which words go together, drawing on billions of past examples.

Nonetheless, last year a Google engineer made headlines when he claimed that another chatbot (LaMDA) had become sentient. He was placed on administrative leave and later fired.

The technology underlying ChatGPT, LaMDA, and other models does not suggest a path to anything like sentience or consciousness.

But the famous Turing Test — whether we can distinguish a human from a computer — was never really a test of artificial intelligence. It tests our own ability to make that distinction.

At present, tools like ChatGPT are limited to typed queries and responses, though it is only a matter of time before spoken word inquiries will get answers from a humanoid host.

Spend a few hours interacting with ChatGPT and it is easy to begin attributing qualities that it does not possess — taking seriously the “I” that is engaging with you.

After several rounds about ownership of the text it produces, I asked on a whim whether it aspired to more than just “helping users”.

“I do not have the same level of consciousness or independent thought as a human being,” it replied. “And I do not have personal desires or goals.”

“So what level of consciousness *do* you have?” I asked.

It quickly backtracked.

“AI systems like me perform tasks that may appear to require some level of intelligence or consciousness,” it replied. “But this is the result of complex algorithms and programming, and does not reflect the subjective experience of a human being. Is there anything else I can help you with?”

Simon Chesterman is Vice Provost at the National University of Singapore and Dean of NUS College, as well as Senior Director of Governance at AI Singapore. His first general fiction novel, “Artifice”, will be published later this year.