Brian Williamson

AI, Copyright and the Public Good

June 2025

COMMUNICATIONS CHAMBERS

About the Authors

Brian Williamson works at the intersection of technology, economics and policy. Clients have included governments, regulators, telcos and tech companies.

Disclaimer

This is an independent report funded by Google. The opinions offered herein are purely those of the authors. They do not necessarily represent the views of the client, nor do they represent a corporate opinion of Communications Chambers.

Contents

1. Executive summary

AI is a high-impact and beneficial technology. It is already seeing wide adoption and application and continues to improve rapidly in terms of capability and cost.

AI qualifies as general-purpose technology, like steam, electricity and computing; but is undergoing more rapid improvement than those technologies and a faster decline in costs. This will spur wider and deeper adoption throughout the economy.

Large language models are also transitioning to multimodal models incorporating text, sound and video; while the extension of compute 'thinking time' is supporting deeper reasoning capabilities. Capabilities continue to scale, and costs continue to decline.

AI also represents the invention of a method of invention and is seeing widespread application across science and technology.

Early fruit from the application of AI is Google DeepMind Alpha-Fold which has predicted the structure of over 200 million proteins, a key input to the development of medicines. This and similar advances offer the prospect of curing diseases and extending healthy lives – a prize of immense value.

AI is also a cultural tool, allowing the exploration of areas including history, art and science. This has the potential to enrich lives and education. However, to reflect diverse culture and interests, AI needs to be trained on a large and representative corpus of data.

Not only will AI lead to productivity, income and leisure gains, but it will also raise the rate of innovation and long-term growth. This is welcome given the headwinds from rising defence, health and care expenditures, but also from a declining workforce as populations age.

Expectations of the growth contribution from AI vary, but a contribution of a few percentage points per annum is plausible. Some have argued for much greater contributions. Even a modest contribution of 1% annually compounds to very large long-term gains.

To deliver on its promise, AI needs to continue to scale, and needs inputs - talent, compute (electricity and chips) and training data. Compute for AI training and inference is growing rapidly, and for optimal model training, data needs to grow alongside compute.

The debate over copyright therefore involves consideration of a trade-off. From an economic perspective, what matters is the trade-off between potential benefits from access to a large corpus of training data and potential harm to content providers under alternative rights regimes, considering the transaction costs of licensing.

This paper argues that the benefits of access to a large corpus of training data - in terms of AI quality and representativeness – are high, and that the prospective risk of harm to copyright holders is low.

Limiting the use of publicly available copyrighted works, including the possibility of an opt-out (or opt-in) approach, would reduce available training data. Proposed transparency provisions are also problematic in practice, and unnecessary if the focus remains on whether outputs violate copyright.

An approach to clarify rights would be to recognise that AI training is transformative and permit training on copyrighted works as fair use. There are parallels with the approach adopted historically in relation to search.

The focus should be on whether AI outputs violate copyright.

2. The promise and value of AI

Artificial intelligence (AI) offers the prospect of greater abundance via higher productivity and socially beneficial applications including in science and health. That is the prize.

But delivery of the prize depends on the availability of inputs to AI including talent, electricity, compute and training data

Technology as a driver of growth

The only source of sustainable growth in per capita income, and leisure, is productivity growth, and technological progress underpins productivity growth.

Previous general-purpose technologies - such as steam, electricity and computing (ICT) - have seen widespread application and drove waves of productivity growth.[1]

Figure 1: Productivity growth from past general-purpose technologies

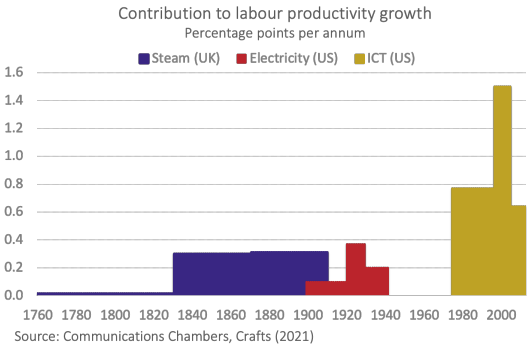

Source: Communications Chambers, Crafts (2021)

Many countries have seen a productivity growth slowdown over the past two decades, including the EU and UK.

This has contributed to low-income growth, deteriorating public finances, and heighted concern over future challenges including an ageing population and declining workforce.

The transformative potential of AI

AI as a general purpose technology

AI is widely seen as a general purpose technology with growing capability, declining costs and widespread potential application. But in contrast to steam and computing, the pace of improvement in performance and cost is vastly greater.

The efficiency of steam engines (lb coal per hp per year) improved from 30 lb for the Newcomen engine (1712) to 12.5 with the Watt engine (patented 1769) to 2 lb by the early 20th Century[2] – an average rate of improvement of around 1.4% per year.

The efficiency and cost performance of computing followed Moore's law, doubling roughly every 2 years. However, it was only with the advent of connected computing that an economy-wide growth contribution was observed.

AI performance per unit of compute is improving three-fold each year[3], while hardware price-performance doubles every 2.5 years[4] - an overall improvement of 75% annually.

In terms of capability, by the end of the decade we may see advances in AI as drastic as the difference between the rudimentary

[1] Crafts, September 2011, Artificial intelligence as a general-purpose technology: an historical perspective, Oxford Review of Economic Policy, Volume 37, Number 3. https://academic.oup.com/oxrep/article/37/3/521/6374675
[2] Crafts, May 2003, Steam as a general purpose technology: A growth accounting perspective. https://www.lse.ac.uk/Economic-History/Assets/Documents/Research/LSTC/wp7503.pdf
[3] https://epoch.ai/trends
[4] Hobbhahn and Besiroglu, June 2022, Trends in GPU Price-Performance. https://epoch.ai/blog/trends-in-gpu-price-performance?utm_source=chatgpt.com

text generation of GPT-2 in 2019 and the problem-solving abilities of GPT-4 in 2023.

Another measure of the growing capability of AI is the length of tasks AI agents can complete, which a recent metric developed by METR indicates has been doubling every 7 months (Figure 2[5]). If the trend continues AI would be able to tackle multi-week tasks by 2030.

Figure 2: AI task length is doubling every 7 months

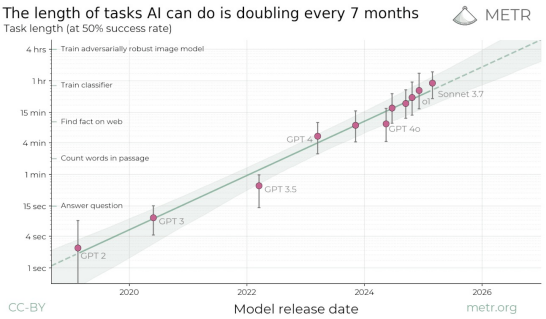

AI as a method of invention

By applying artificial 'intelligence' to scientific and technical problems, AI functions as a method of invention itself — a development that may sustain long-term knowledge and economic growth. To do so, AI needs to be trained on a large and diverse set of data, and to be able to access data during inference to innovate and solve real world challenges.

AI models exceed human performance across a range of technical performance benchmarks[6]. Large language models (LLMs) are also now multimodal including text, speech and vision, and are being applied widely.

Models, which may vastly exceed human capabilities in specific domains, have also been developed. For example, the model AlphaFold has increased the catalogue of known protein structures, key to developing new medicines, from around 100,000 previously known experimentally[7] to over 200 million[8].

The AlphaFold breakthrough won Demis Hassabis and John Jumper of Google DeepMind a Nobel prize in Chemistry in 2024. AlphaFold uses the transformer architecture developed for language models but applied to the sequences of amino acids that make up proteins rather than sequences of words.

A current frontier is reasoning models that utilise more compute for inference, can explore and determine a strategy for answering a query and may utilise tools including coding and search to do so.

Illustrative of reasoning models is "co-scientist" developed by Google. Professor José R Penadés and his team at Imperial College London had spent years working out and proving why some superbugs are immune to antibiotics. He gave "co-scientist" a prompt asking it about the problem he had been investigating, and it reached the same conclusion in 48 hours.[9]

A further development in this direction is AlphaEvolve[10], an evolutionary coding agent powered by large language models for

[5] Kwa and West *et al*, March 2025, Measuring AI Ability to Complete Long Tasks, ArXiv. *https://arxiv.org/pdf/2503.14499*
[6] Standford University, April 2025, Artificial Intelligence Index Report 2025. *https://hai.stanford.edu/ai-index/2025-ai-index-report*
[7] John Jumper *et al*, July 2021, Highly accurate protein structure prediction with AlphaFold, *Nature* 596.
[8] *https://deepmind.google/technologies/alphafold/*
[9] BBC News, February 2025, AI cracks superbug problem in two days that took scientists years. *https://www.bbc.co.uk/news/articles/clyz6e9edy3o*
[10] Google DeepMind, May 2025, AlphaEvolve: A Gemini-powered coding agent for designing advanced algorithms. *https://deepmind.google/discover/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/*

general-purpose algorithm discovery and optimization.

Neural networks and large language model architectures are also relevant to autonomous vehicles and robotics. One can think of the model generating sequences of robot-executable steps rather than a sequence of words.

Together, reasoning and robotics may accelerate science by automating not only hypothesis generation, but also experimental science.

Expected economic value of AI

While estimates vary widely, AI is expected to deliver sustained productivity growth.

One benchmark is previous general-purpose technologies which added about 0.5% per annum in the case of steam (and rail) over many decades, and around 1-2% per annum in the case of computers over a much shorter period.

A much more pessimistic view is that of Economist Daron Acemoglu who estimated that:[11]

"Using existing estimates on exposure to AI and productivity improvements at the task level, these macroeconomic effects appear nontrivial but modest—no more than a 0.71% increase in total factor productivity over 10 years [0.07% per year]."

However, Tyler Cowen[12] has argued that Acemoglu failed to consider a deepening of automation and disagreed with his dismissal of near-term impacts on science and the pace of innovation.

Some involved in AI development have more optimistic expectations in relation to the growth impact of AI, in the range 20-30% or more per annum – far more than anything seen in the past.

This is significantly greater than the 10% growth achieved by any country ever over several decades. Further, things that are hard to automate increase in relative price (referred to as Baumol's cost disease[13]) and tend to grow as a share of GDP, slowing overall growth.

Philippe Aghion and Daren Bunel,[14] who consider both the experience of previous general-purpose technologies and the approach adopted by Acemoglu, and conclude with an estimated growth contribution of around 1% per annum:

"Based on the first approach, we estimate that the AI revolution should increase aggregate productivity growth by between 0.8 and 1.3pp per year over the next decade. Using the second approach but with our own reading of the recent empirical literature on the various components of the task-based formula, we obtain a median estimate of 0.68pp additional annual total factor productivity (TFP) growth."

Importantly, these estimates exclude AI's potential to accelerate idea generation:

"Our estimates do not take into account the fact that AI automates tasks not only in the production of goods and services, our focus in this note, but also in the production of ideas."

[11] Daron Acemoglu, April 2024, The Simple Macroeconomics of AI. https://economics.mit.edu/sites/default/files/2024-04/The%20Simple%20Macroeconomics%20of%20AI.pdf
[12] Tyler Cowen, April 2024, The Simple Macroeconomics of AI. https://marginalrevolution.com/marginalrevolution/2024/04/the-simple-macroeconomics-of-ai.html
[13] Timothy B. Lee, May 2017, William Baumol, whose famous economic theory explains the modern world, has died. Vox. https://www.vox.com/new-money/2017/5/4/15547364/baumol-cost-disease-explained
[14] Philippe Aghion and Daren Bunel, June 2024, AI and Growth: Where Do We Stand? https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf

Even a one percent per annum contribution to growth, if sustained, would be large. The present value of one percent of additional growth compounded in perpetuity and discounted at 4.5% would amount to over six times current GDP.

Over longer time horizons, around a decade from now, AI may have a greater impact as the economy is reconfigured around AI, as opposed to adding AI to existing ways of doing things, and as AI contributes to the growth of knowledge. A growth contribution of several percentage points per annum is plausible.

Perhaps rather than thinking of AI as ushering in an era of radical and unwelcome disruption, we should consider the potential for AI to restore productivity growth to the 'Goldilocks zone' with which society has coped well in the past, and successive generations have benefited.

However, the benefits of AI go beyond productivity and income growth, and include, for example, the application of AI to medicine (below), and as a cultural tool (next section).

AI in drug discovery and medicine

Developments including Alpha-Fold, which can predict protein structures, and Google "co-scientist" that can reason and generate hypothesis to test were mentioned earlier.

The broader vision is that understanding biology, and therefore disease, disease prevention and treatment, is a challenge well suited to the application of AI:[15]

"I think biology can be thought of as an information processing system, albeit an extraordinarily complex and dynamic one." Demis Hassabis, 2021

Demis Hassabis more recently noted that progress has moved beyond hypothetical potential to real world application of AI for drug discovery:[16]

"We're looking at oncology, cardiovascular, neurodegeneration, all the big disease areas, and I think by the end of this year, we'll have our first drug. It usually takes an average of five to 10 years [to discover] one drug. And maybe we could accelerate that 10 times, which would be an incredible revolution in human health."

AI opens the potential to cure many, perhaps all, disease. As a hypothetical, imagine AI led to discoveries that allowed all cancers to be cured. In high income countries cancer accounts for around 25-30% of all deaths. Cancer risk is age related. Taking account of this, the extension of life expectancy in developed countries from curing all cancer might be around 3 years[17].

For the UK the National Institute for Health and Care Excellence (NICE) uses a value of £20,000–£30,000 per qualified adjusted year of life saved (QALY)[18]. Assuming a quality-of-life score of 0.8, an additional 3 years of life is equivalent to 2.4, which valued at £25,000 per year is £60,000 per life saved, or £4 trillion given the UK population of 68 million, slightly more than UK GDP.

[15] Demis Hassabis, November 2021, Introducing Isomorphic Labs.
https://www.isomorphiclabs.com/articles/introducing-isomorphic-labs
[16] FT, January 2025, AI-developed drug will be in trials by year-end, says Google's Hassabis.
https://www.ft.com/content/41b51d07-0754-4ffd-a8f9-737e1b1f0c2e
[17] CDC, May 2013, United States Life Tables Eliminating Certain Causes of Death, 1999–2001.
https://www.cdc.gov/nchs/data/nvsr/nvsr61/nvsr61_09.pdf
[18] NICE, December 2024, Should NICE's cost-effectiveness thresholds change?
https://www.nice.org.uk/news/blogs/should-nice-s-cost-effectiveness-thresholds-change-?utm_source=chatgpt.com

In addition to AI's role in assisting in the development of new drugs, AI also has application as a diagnostic tool. Examples include examination of eye scans for early signs of age-related macular degeneration (AMD) and broader application of AI as a clinical diagnostic tool.

Google DeepMind, in collaboration with Moorfields Eye Hospital and Google Health, curated a dataset of images of eye retinas and trained an AI system that could predict the form of AMD – dry AMD - that impacts about 15% of patients and which can result in rapid and permanent loss of sight.[19] There are treatments that can slow further vision loss, once dry AMD is detected, and having a rapid means of detecting dry AMD, whilst avoiding false positives, is critical to treating the most at risk patients promptly.

In a separate study LLMs and certified physicians were compared across three diagnostic tasks: triage in the emergency room, initial evaluation by a physician, and admission to the hospital or intensive care unit. In all experiments the LLM displayed superhuman diagnostic and reasoning abilities.[20] The findings suggest the need for trials to evaluate how AI can be integrated into real-world patient care settings and clinician-AI interactions and decision making.

Not only does AI offer the prospect of materially increased productivity and income growth, AI is also seeing application in medicine and health care that will support longer and healthier lives.

[19] Google DeepMind, May 2020, Using AI to predict retinal disease progression. *https://deepmind.google/discover/blog/using-ai-to-predict-retinal-disease-progression/*

[20] Brouder *et al*, December 2024, Superhuman performance of a large language model on the reasoning tasks of a physician. *https://arxiv.org/abs/2412.10849*

3. Technology as a cultural tool

All technologies alter the expression of culture, opening up new possibilities. Technology, including AI, is a cultural tool.

Past technologies and creative expression

Lamps and ochre allowed hominids to complete spectacular cave paintings. Printing spurred the production of written works, whilst electricity supported recorded music, cinema and broadcast media.

Technology transitions create opportunities for new forms of creative expression and may disrupt old ones.

The Gutenberg Press in the mid-15th century greatly reduced the cost of reproducing texts, expanding access and leading to new forms of expression including the emergence of the novel in the 17th and 18th centuries.

However, the Gutenberg Press also disrupted the centrality of oral storytellers and bards, and manuscript copying as a form of cultural labour. It also disrupted the church, a cultural institution.

The printing press also led to an institutional innovation, namely the Statue of Anne (1710) – the first true copyright law – that recognised authors, not just printers, as rights-holders with 14 years (which could be extended by a further 14 years if the author was still alive) of exclusive right to duplication.

AI as a cultural tool

AI is also a cultural tool that will change the creative process. The ability of AI to generate content will open new creative possibilities. However, what aspects of AI individual creators embrace will depend on their preferences, including how much they value the process versus the outcome.

Further, just as creators learn from existing art and culture and build upon it, the value of AI as a cultural tool for creators will depend on the scope for AI to learn from existing culture, including copyrighted works.

Audience preferences will also evolve, and to the extent that content is easier to generate they may come to value the relationship with the creator and with the creative process more highly.

The changing nature of the internet, already underway but accelerated by AI, will also change the feasibility of different ways of monetising creative content.

A wide range of digital platforms have emerged to allow creators to share, distribute, and monetise content directly, bypassing traditional publishing or media gatekeepers. An example is Substack with free and paid options for independent writers. Other platforms include YouTube, Spotify, SoundCloud and Patreon.

AI is providing new tools for expression including more advanced photo editing, and image and video production tools including film making tool Flow.[21]

AI has also been used in the content creation process, for research, translation, fact checking and content production. These changes are being embraced by established media and newer forms of media.

[21] Google, May 2025, Meet Flow: AI-powered filmmaking with Veo 3. *https://blog.google/technology/ai/google-flow-veo-ai-filmmaking-tool/*

Reuters news

The following is based on an interview between Paul Bascobert, president of Reuters, and Nilay Patel of Decoder. [22] In short, LLMs are already embedded within news production and distribution at Reuters.

Nilay Patel, Decoder:

"And you've lightly described what it is you want to build, right, a new way of accessing information in your journal. That's great. Is that a chat bot to you? Is that a voice assistant?"

Paul Bascobert, president of Reuters:

"That's just one piece on the professional consumer side of it. The bigger part of what we're doing is actually we've been investing in our newsroom operation for the last year and a half."

"Alessandra Galloni, our Editor-in-Chief... has a vision of the newsroom of the future being an AI-powered newsroom. So we have built tools that access every piece of moving news around the world. And using a combination of language models, we can access this data in multiple languages, grab it, pull it in, look at it, validate it with a human, and push it out."

Nilay Patel, Decoder:

"Are your reporters driving this change? Most newsrooms are full of a lot of conflicted feelings in this way."

Paul Bascobert, president of Reuters:

"...AIs can allow us to get more of our people in the field, more of our people doing the distinctive work, the new work, less of our people moving content through the system, and that's really powerful to our mission to cover the world."

Dwarkesh Podcast

Not only is AI helping Reuters, a 175 year old company that employs 3,500 people; it is also helping podcaster Dwarkesh Patel[23] who conducts deeply researched interviews with experts across AI, history, and geopolitics ranging from Mark Zuckerberg to Tony Blair.

Dwarkesh discussed how he uses AI in his podcasting workflow with Dan Shipper, CEO and cofounder of Every.[24] Dwarkesh begins his research for an upcoming podcast interview by uploading context about the guest's field—like books and academic papers—to Claude, using the model's project feature.

- LLMs to build context. Poses questions to understand the broader subject matter.
- Deepen expertise with AI. Uses the model as a sounding board to test and refine his understanding.
- Leverage AI's gaps. Uses AI's failures as prompts for interview questions.

Future-proof workflow. Adopted AI early, which is now paying off.

[22] Nilay Patel (Decoder), 8 May 2025, Reuters is ready to stand up for the press — and embrace AI. https://podcasts.apple.com/gb/podcast/decoder-with-nilay-patel/id1011668648?i=1000706782317&r=3519
[23] Dwarkesh Podcast. https://www.dwarkesh.com/
[24] Dan Shipper, July 2024, Dwarkesh Patel's Quest to Learn Everything. https://every.to/podcast/dwarkesh-patel-s-quest-to-learn-everything To

4. Inputs to AI: the critical role of data

Scaling laws show that AI performance improves in a predictable way when models, data, and compute are scaled together.[25]

To optimise compute use, model size and data volume should scale together. In fact, smaller models trained on more data often outperform larger models with less data - and do so more efficiently in terms of compute (and electricity).

Data could prove a constraint on the future scaling of model performance (Figure 3).[26] If the recent trend of 4x/year compute scaling continues, we would run into a 'data wall' for text data in about five years.

A data constraint is more likely if copyright is a constraint on training data. This 'data wall' would arise when available high-quality text data is exhausted, limiting future gains from increasing compute alone.

Figure 3: Future scaling constraints

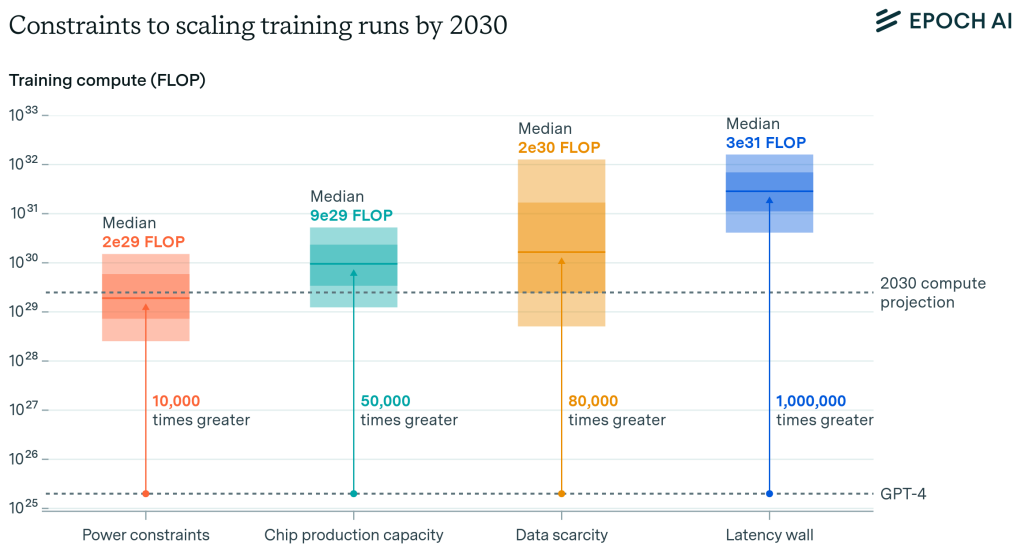

If broad swaths of text or images were unavailable because of copyright controls, you could not simply "train longer" on what is left without a decline in model performance. Synthetic data may help compensate for missing material, but it risks reflecting the biases of the model that generated it.

Utilising a large and diverse corpus of training data is also important to ensure that models are representative of the diversity of culture across people and time. As Amanda Levendowski observed:[27]

"The rules of copyright law also privilege access to certain works over others, encouraging AI creators to use easily available, legally low-risk sources of data for teaching AI, even when those data are demonstrably biased."; and

[2525] Kaplan *et al*, 2021, Scaling Laws for Neural Language Models, https://arxiv.org/pdf/2001.08361
[26] Savilla *et al*, November 2024, https://epoch.ai/blog/can-ai-scaling-continue-through-2030
[27] Amanda Levendowski, July 2022, How Copyright Law Can Fix Artificial Intelligence's Implicit Bias Problem. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3024938#

"The normative values embedded in the tradition of fair use align ultimately with the goal of mitigating bias."

Another study by Qu and Wang notes that:[28]

"Our analysis acknowledges multiple factors influencing LLMs' ability to simulate diverse perspectives accurately. These include limited training data diversity, which may bias the model towards overrepresented cultures..."

Both opt-in and opt-out licensing regimes could result a narrowing of training data, but also selection effects resulting in reduced diversity. Long-tail data and stylistic diversity might suffer.

High-quality, representative training data is not just an input - it is a foundation for AI progress. Copyright frameworks that restrict data access risk becoming a bottleneck on cultural innovation and representatives.

[28] Qu and Wang, August 2024, Performance and biases of Large Language Models in public opinion simulation, *Nature*. https://www.nature.com/articles/s41599-024-03609-x

5. Policy trade-offs: Copyright, innovation and the public interest

AI is valuable, both economically and more generally as a method of invention and as a cultural tool, and its value is growing as models improve and as users explore new applications and complementary innovations.

The quality of models, and their application, in turn depend on the availability of data. Data for training needs to be both sufficient and representative to reflect individual and cultural diversity.

Controversy has arisen over the use of public data subject to copyright for model training, driven by two distinct concerns, namely that AI models may replicate or closely approximate copyrighted material in their output, and that creators of copyrighted data should share in the value created.

In principle these concerns arise with human intelligence. We are beneficiaries of learning by 'processing' data including copyrighted works, and we produce creative outputs which draw on data, and the styles, we absorb. An artistic movement, such as impressionism, illustrates how creativity can involve both assimilation and representation.

AI is not different from us in this sense, but it does represent a difference in terms of scale, for example, GPT-4o facilitated the generation of a wave of Studio Ghibli-inspired images.

It is reasonable, therefore, to ask whether our existing framework of copyright protection is the right one in this new era - a question that has been considered in the past with the emergence of new technologies including printing (which inspired the original copyright law, the Statute of Anne – *"An act for the encouragement of learning"*) and, later, for example, photography.

The best way to approach this question is to consider trade-offs and not simply to take existing law as a given.

Illustrative of the latter is the US Copyright Office report "Copyright and Artificial Intelligence" which is silent on the economic literature on copyright and the economic trade-offs involved with alternative approaches.[29] The framework is legalistic rather than offering an economic analysis. As a result, it does not provide clear guidance on optimal policy.

There are also legal cases regarding possible infringement,[30] but these will inform interpretation of existing law rather than informing us about what the law ought to be.

The economic logic of trade-offs

From an economic perspective what matters is the trade-off between potential benefits from access to a large corpus of training data and potential harm to content providers under alternative rights regimes, considering the transaction costs of licensing.

If the benefits from access to publicly available training data in terms of higher quality and more representative models are significant, and the harm from model outputs in relation to protected content are low, then

[29] United States Copyright Office, May 2025, Copyright and Artificial Intelligence. *https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf*

[30] The Economist, March 2024, The court cases that could shape how AI develops. *https://www.economist.com/the-world-ahead/2024/11/20/the-court-cases-that-could-shape-how-ai-develops*

fair use of copy-protected material is supported.

Taking account of trade-offs, Joshua Gans has suggested the following approach:[31]

"For large-scale AI models trained on huge datasets, policymakers will need to weigh the benefits of training data access for AI progress against risks of harm to content owners. Fair use exemptions are likely to be optimal when the AI benefits substantially outweigh the harms to content owners. Alternatively, novel licensing mechanisms that allow free usage but preserve 'backstop' protections for content owners could provide a beneficial middle ground."

Focusing on model outputs

The reason for focussing on outputs, as opposed to training data inputs, is that it is in relation to the output of AI models that substitution for copyrighted works might arise.

In relation to inputs, model training on a large corpus of data does not imply that the model has 'memorised' the data; nor that the model will necessarily output copyrighted works included in the training data verbatim. As Jack Wiseman put it:[32]

"One common misconception is that models 'ingest' data... Another is that the model developers want the model to memorise things. This is not the goal. Memorisation is an inefficient use of space inside the model, and memorising protected expressions isn't what intelligence is."

Nevertheless, a violation of copyright in relation to model outputs could arise and could involve harm to copyright holders.

Language models may include filters or system prompts that aim to prevent excess copyright violating 'leakage' of original content in their output. For example, the leaked system prompt for Anthropic model Claude includes, amongst other copyright related prompts, the following:[33]

"CRITICAL: Always respect copyright by NEVER reproducing large 20+ word chunks of content from web search results, to ensure legal compliance and avoid harming copyright holders."

If model developers are liable for copyright in relation to outputs, a low threshold and/or high penalties for *ex post* liability could encourage artificial gaming of models to get them to generate content that violates copyright.

The 'PoisonParrot' paper[34] demonstrated that it was possible to inject text fragments into training data that could induce the model to generate copyrighted material even though the model was not directly trained on the specific copyrighted material.

"Copyright holders have financial incentives to pursue violation claims against LLM companies. Within this context, our work proposes a new threat model in which an adversary launches a training set poisoning attack to increase the chance of an LLM

[31] Joshua Gans, April 2024, Copyright policy options for generative artificial intelligence. https://cepr.org/voxeu/columns/copyright-policy-options-generative-artificial-intelligence
[32] Jack Weisman, May 2025, Bohemians at the Gate? https://inferencemagazine.substack.com/p/bohemians-at-the-gate
[33] Ásgeir Thor Johnson, May 2025, System prompt leaks – Anthropic. https://github.com/asgeirtj/system_prompts_leaks
[34] Panaitescu-Liess *et al*, March 2025, PoisonedParrot: Subtle Data Poisoning Attacks to Elicit Copyright-Infringing Content from Large Language Models. https://arxiv.org/abs/2503.07697

generating an output that violates the copyright of a particular text."

An approach to clarify rights would be to recognise that AI training is transformative and permit training on copyrighted works as fair use. There are parallels with the approach adopted historically in relation to search.

Why opt-outs and transparency requirements could be costly

Proposals including opt-in and opt-out in relation to training data have been suggested. However, they could reduce the quality of AI for all users, and Bertin Martens has argued: [35]

"The right to opt-out amounts to economically inefficient overprotection of copyright."

More generally, Martens argued that:

"Defenders of the moral right to remuneration argue that any arbitrary remuneration is better than no remuneration. But this private moral right comes at the expense of social welfare. The ongoing bargaining and court cases between media producers and GenAI developers risk entrenching this market failure in jurisprudence."

The Economist have argued that existing copyright overprotects content creators, and that: [36]

"A return to the 28-year copyrights of the Statute of Anne would be in many ways arbitrary, but not unreasonable."

The bar for clarifying rights in favour of restrictions or opt out in relation to the use of data for model training should be high.

Transparency requirements: a misguided fix

While appealing in principle, transparency requirements could be impractical and counterproductive.

Illustrative of the potential burden is the challenge of determining what is subject to copyright:[37]

"No one knows what is copyrighted or not. Copyright is not vested upon a work through a deliberate act like a registration: it is bestowed on any creation that meets the requirements of copyright laws, and those requirements may vary from one country to another. One of those criteria is originality, a threshold that has led to many lengthy court cases and that is in no shape or form something a web crawler or automated tool could identify."

Training datasets may contain personal information and may also contain proprietary or strategically curated information. Forcing disclosure could compromise trade secrets and reduce incentives to invest in data quality - particularly as synthetic and human-curated datasets become more important.

Mandating transparency in relation to training data, including the identification of copyright material, would prove costly and counterproductive.

[35] Bertin Martens, April 2024, Economic arguments in favour of reducing copyright protection for generative AI inputs and outputs, Working Paper Issue 09/2024. *https://www.bruegel.org/sites/default/files/2024-04/WP%2009%20040424%20Copyright%20final_0.pdf*

[36] The Economist, April 2020, Copyright and wrong. *https://www.economist.com/leaders/2010/04/08/copyright-and-wrong*

[37] Coalition for Creativity, September 2023, C4C's Perspective on the EU AI Act: Copyright in Real Life is Messy and AI Discussions Are Not Helping. *https://coalition4creativity.org/2023/09/28/c4cs-perspective-on-the-eu-ai-act-copyright-in-real-life-is-messy-and-ai-discussions-are-not-helping/*