

SPECIAL FEATURE: AI AND LITERARY TRANSLATION

# Copyright in light of generative AI

*Gregor Stojin*

## Inputs and outputs

In September 2023 a group of prominent authors, including George RR Martin and John Grisham, filed a lawsuit at the federal court in Manhattan, New York, against OpenAI, creators of ChatGPT, claiming their copyright was infringed to train the system and accusing them of “systematic theft on a mass scale.” Similarly, during the course of this year, digital artists sued developers of image generators Midjourney and Stability AI for training their models on their graphic work, and there were also lawsuits by programmers who allege their code was used in a similar manner by Microsoft, OpenAI and others in training Copilot, an artificial intelligence (AI) tool used to assist in code writing.

The rapid development of AI over the past few years has spurred the need for social rules and society to adapt to the new reality, both due to impact and disruptions on the existing positions, and the second-order effects which require reevaluations of feasibility of legal and economic regimes which formed the existing rules in the first place. Some of these issues relate also

to copyright and other intellectual property protection mechanisms.

Initially, most discussions on the reevaluation of Intellectual Property (IP) legal frameworks have been devoted to protecting AI systems themselves, their products and their outputs. Questions regarding output include: are works made by AI copyrightable, if so, who owns the copyright, what is the level of human input required to satisfy the threshold for creative input, or even, can AI systems be considered authors and awarded copyright protection?

The recent proliferation of a new generation of technology, generative AI, has strengthened the salience of such questions but also sheds light on the opposite perspective. Many ambiguities were made visible on the other side of the equation i.e. in the input data used for training.

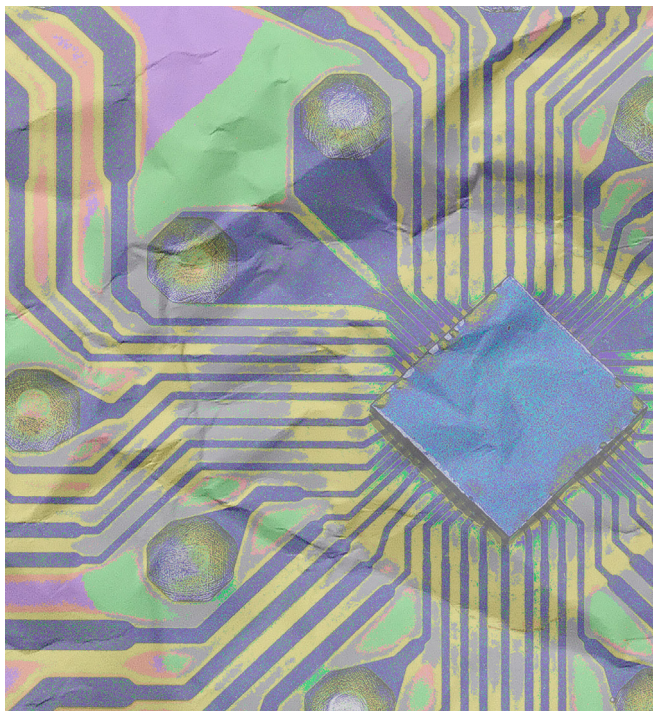
The former issues regarding output primarily address the potential rights stemming from the novel uses of the technology. The latter issues regarding input address the interests and rights



of existing human authors. The former group of issues shares many similarities with discussions relating to the benefits and potentials of the new technology and is often advocated by the same representatives of the wider tech community. Representatives of the latter group are typically fragmented and often unaware of their role in the dynamics of technology development, and at best, in a defensive role. These positions appear to be reflected in parallel legislative initiatives as well.

### What makes generative AI so special?

AI as a technology depends on three key components: data, algorithms, and computing power. Although AI has evolved through various forms since the mid-20th century, it has gained unprecedented momentum in the past decade mainly due to the exponential growth in the amount of available data.



By Kaisa Ranta, based on a  
photo by Mister\_fr CC-SA

The increasing digitisation of our lives contributes to more data being created every two years than humanity has produced in its entire history prior to that. Texts, images, sounds, videos, our use of mobile and other electronic devices, and the signals received by sensors on cars and various other devices can all be used as input for training AI models. Data represents an extremely broad category and also includes more complex information, such as unstructured documents and other works, including literary writings. It is the main ingredient in the development of AI models, and the fuel for their further growth.

Algorithms are methods of processing such data. They are abstract methods, which means that they cannot be directly protected by intellectual property mechanisms such as copyrights or patents, and the ability to use them primarily depends on available knowledge or human resources.

Many algorithms used in artificial intelligence systems have been known and used in statistics for a long time. Nevertheless, there have been significant advances in processing methods over the last decade. Initially requiring supervised training of pre-arranged data, machine learning abilities have progressed through large neural networks and generative adversarial networks toward advanced capabilities that allow self-supervised learning on previously unprocessed data. This, however, requires speed.

Fast processing of large amounts of training data allows neural networks to identify connections and relations

between all data elements in the network, recognise patterns, dependencies and statistical regularities often hidden to humans, and learn from them directly. This is used to develop parameters, which are essentially variables determined from the given training data. Eventually, they function as detailed instructions on how to use algorithms in connection with particular input data within set degrees of probability, decoding it and encoding new outputs. Through this, pre-trained models are created, which can be used for various types of tasks and fine-tuned further to create original and diverse content and synthesise data that resembles the examples from which they were trained.

Efficient training and running of such systems largely depends on their computing or processing capacities. In comparison with data, the capacities of processors are growing relatively more linear, as they are limited by many physical laws – from the size and speed of the chips to the geostrategic implications of their manufacturers' locations, which influence their availability and, consequently, competitiveness. In addition to this, they also consume substantial energy and human resources for setting up, running, cooling and maintenance. As they significantly depend on the available financial resources, they influence concentrations of large model development.

Increasing capacities of all three components allow for the creation of ever larger and more powerful pre-trained models, whose complexity is often expressed in the number of parameters and dataset size. For example, while the

first version of OpenAI's GPT was based on 117 million parameters, having been trained on a dataset of about 4.6 GB of raw text, and the second 1.5 billion with a dataset of 40 GB of filtered text, version GPT3 already used 175 billion parameters with a dataset of 570 GB filtered from 45 TB of plaintext. Interestingly, information on the GPT4's parameter size was never formally disclosed but is estimated at 1.7 trillion, and the dataset size and provenance remain unknown.

## “Advances in natural language processing allow for human-like quality levels of translation”

Growth of parameters above a certain critical point seems to trigger unexpected emergent abilities, which are not present at a smaller volume. This has, among many other functionalities, led to significant advances in natural language processing, allowing for human-like quality levels of translation, summarisation, stylisation, and overall content generation.

### Access to data without burden

While algorithms, computing power and their utilisation are generally freely available on the market and depend mostly on monetary resources, data is often exclusive and subject to various



restrictions, depending on its source and protection mechanisms that may apply. The provenance and legality of the materials used for high-volume and high-quality datasets for training and especially fine-tuning AI models are relevant questions. It seems, however, that the lack of developers' transparency makes it difficult, if not impossible, to answer them satisfactorily.

The importance of data, with priority on accessibility and availability, has been at the forefront of European legislators for many years. Such an approach has contributed, among other measures, to new rules for free access to data through different legal instruments, especially in the public sector with requirements for re-use of its data. Recently, this was also expanded to private individuals, as stipulated by the newly applicable [Data Governance Act](#), which encourages data altruism. The aim was to provide as much data as possible for development, stressing the importance of facilitating access to data and data sharing, open standards and open-source technology to encourage investment and boost innovation.

The revision of the EU copyright directive is an example of another such measure, as it introduced an important exception to the previously strong and exclusive rights of copyright holders and significantly expanded the ability of data users to process previously protected material through an opt-out principle. As an EU Commission Study on copyright and new technologies noted in 2022, "The use of protected content as AI-training data may involve certain protected acts, which require the rightsholders prior consent –

unless they are exempted under one of the copyright exceptions. The newly introduced exceptions for text and data mining (TDM) may relieve the developers and users of AI solutions in the cultural sector of this burden [*sic*]."

The opt-out principle, for example, allows copyright holders to reserve their rights in light of the newly introduced TDM exception, but the process and effectiveness of this remain unclear. The Commission Study did acknowledge, that (*inter alia*) it may prove difficult to verify compliance with the opt-out, as TDM processes are mostly invisible to the public, carried out without prior information, and there is no obvious legal basis to request access to the process or to force the AI solution provider to demonstrate that the protected content has not been used.

Another illustrative demonstration of how policymakers understand and assess such data can be seen in the progress of the currently negotiated EU Artificial Intelligence Act (AIA). The European Commission's original proposal from April 2021 touched on copyright only in passing, and even then in the context of the protection of the developers' IP rights. The Council of the EU also did not take a more specific view of it in November 2022. Significantly, however, it did introduce a new category of 'general AI', which would be subject to a significant reduction of the expected compliance obligations. The European Parliament's position in June 2023 may have had a certain advantage of hindsight since it was formulated after the public's exposure to generative tools in late 2022 and early 2023. Consequently, the Parliament

envisaged significantly higher compliance requirements for a new category of ‘foundation models’, which could, otherwise, easily fall within the Council’s ‘general AI’ category as a subset. Among many other new requirements proposed for foundation models, transparent information on the copyrighted data used in AI models’ training particularly stands out.

## *“Many issues relating to AI’s impact on creative work may not belong in the domain of copyright”*

Whether such obligations for foundation models will actually be used in the final text of the upcoming Regulation remains to be seen. The trialogues, the final phase of the negotiations, are currently underway and are expected to conclude by the end of this year (2023). The differences in the initial approaches of the three key institutions point both to the rapid changes in the technological field, and this makes it difficult for legislators to foresee all the implications in time and demand the agility and adaptability of the appropriate legal framework. They also indicate the complexity of the interests

and relationships of all stakeholders in the long chain of new technologies.

These changes also indicate the need to reevaluate some of the principles related to TDM. Increasing the levels of legal access for TDM indeed provides an important ingredient for emerging technology, but its second-order effects, primarily the eventual disruptions in the creative industries and culture, might indicate a costly imbalance. The problem echoes one of the findings of the Ad Hoc Committee on AI at the Council of Europe (CAHAI), which assessed the feasibility of a legal framework for the design, development and use of AI in line with standards on human rights, democracy and the rule of law (and whose successor, Committee on Artificial Intelligence (CAI) is currently negotiating an international treaty on AI). While we do not have a legal vacuum when it comes to AI, procedural and substantive gaps in current legal frameworks limit the effective protection and enforcement of existing rights due to the specifics of the new technologies. Similarly, the rights of the existing copyright holders may seem protected in substance but are, in fact, limited in actual scope or effective procedural enforcement mechanisms. The impacts of AI clearly show the need to reevaluate the efficiency and balance of the existing rules.

### **Issues beyond copyright**

Many questions raised by the lawsuits against generative AI developers go beyond the scope of copyright law. The style of a particular artist or facts presented in the works are traditionally not protected by copyright, yet they are important themes of contentions. Automated generation of similar works





**Gregor Strojin** is Vice Chair of the Committee on Artificial Intelligence at the Council of Europe (CAI). He is a bar-qualified lawyer (2007), a senior advisor to the President of the Supreme Court of Slovenia and also works with many other organisations at the crossroads of technology, information and law. Strojin has an LLB from the University of Ljubljana (2000) and an LLM from Chicago-Kent (2002) with a specialisation in IP and IT law.

**Gregor Strojin**  
**Photo: Private archive**

impacts authors' interests and society. Generative AI can, for example, change a text's style by using techniques such as style transfer and text synthesis, effectively bypassing many protected aspects of individual copyrighted work. As generative AI produces works that can often be indistinguishable from the work created by humans, it can be and already is used as an adequate substitute in some areas that primarily depend on creativity and originality of the output rather than its quality. After all, quality may just as well be a matter of subjective taste, or something that can be evaluated on the basis of measuring feedback from content providers' end users. This effectively disrupts not only the underlying business models of the creators of the works which were used for training the AI models but can eventually displace human-authored content on a much larger scale, affecting culture.

Centuries ago, modernising societies started developing intellectual property regimes as incentives for creating and disseminating knowledge, innovation and creativity attuned to particular economic policies and philosophies,

primarily humanism. As legal constructs, they can significantly differ between countries. Even more importantly, they are also subject to periodical changes. They adapt over time as technology progresses, and are bound to do so in the future. However, the direction of future legislative changes is always uncertain, and can just as well depend on principles or on different perceptions, priorities and lobbying abilities.

In a public response to the pending lawsuits, OpenAI said it respected the right of authors, and "believed they should benefit from AI technology". That does not seem to answer the plaintiffs' claims either directly or indirectly, but the benefit from AI is an important notion. Many issues regarding AI's impact on creative work, or work otherwise, may not belong in the domain of copyright or intellectual property. However, they do raise questions on what future incentives for creativity and innovation should be, how they might be balanced between stakeholders, who will benefit from the AI, in what share, and, not to be overlooked, at what cost.

