

Legal Aspects of Generative AI

SESSION 1 OF THE DISCUSSION SERIES 'THE IMPACT WE GENERATE'

The AI, Media and Democracy Lab is hosting a series of discussions on Generative AI in the spring of 2023. This is a report of the first discussion, held on March 21st, which focused on legal aspects.



REPORT BY MEMBERS OF THE AI, MEDIA AND DEMOCRACY LAB

PANELLISTS PROF. PHILIPP HACKER, VIKTORIJA MOROZOVAITE, PROF. MARTIN SENFTLEBEN, PROF. NATALI HELBERGER, DR. OT VAN DAALEN AND NAOMI APPELMAN

AI-driven applications such as **Stable Diffusion, Dall-E, GPT -3, ChatGPT** and **Bard** have been receiving increased media attention in recent months. These technologies, commonly referred to as *Generative AI*, learn from existing data to produce new content, including audio, (realistic) images and art, text, code, etc. Though these innovations have sparked renewed enthusiasm and interest in the field of AI, their potential to disrupt and transform has also been met with concern. In an effort to demystify the use of Generative AI, the *AI, Media and Democracy Lab* sought to explore the pitfalls and possibilities of these technologies' transformative impact on the media and journalism industries.

On March 17th, the AIMD Lab organised a discussion event focusing on the legal and governance aspects concerning the use and deployment of Generative AI by media and journalists. The lab invited six legal experts, that is Prof. Philipp Hacker, Viktorija Morozovaite, Prof. Martin Senftleben, Prof. Natali Helberger, Dr. Ot van Daalen and Naomi Appelman, to the Institute for Information Law at the University of Amsterdam to discuss what they consider to be the most pressing legal challenges. Their respective panel talks were further specified through audience participation.

In this report, we synthesize our key takeaways and attempt to unearth key points of attention for legal academics and regulators, as well as media professionals. The report concludes with three overarching questions that we challenged the panelists with: a) Are the challenges faced by the media industries all that revolutionary, or do they present familiar questions and dilemmas in renewed ways? b) Are there any urgent issues currently missing from the legal debate?; and c) How can legal disciplines join forces to collaboratively contribute towards the creation and interpretation of a more robust regulatory framework wherein the concerns of media and journalism industries are met, democratic values guaranteed and citizen's fundamental rights maintained?

Befitting for an event pertaining to the law, a disclaimer. While this session addressed a wide array of legal views, not every legal field of relevance could be addressed. Finally, these takeaways should be seen as a point of departure for future reflection. They are the result of a collaborative brainstorm, subject to further finetuning.

The Need for Holistic Content Moderation Prior to Deployment

The development and deployment of large-language models (LLMs), upon which systems like Chat-GPT are built, require an enormous infrastructure. In the current geopolitical climate, the US has positioned itself as an LLM frontrunner, followed by China, whereas the EU represents the third largest actor.

While this might appear primarily as an economic and geopolitical issue, it raises significant questions regarding the EU's dependency on internationally-located hardware and software. More specifically, one can view technological actors as having control over the means of production and entertainment, and hence, as players with the potential to affect the leading societal narrative. For example, what data have been collected for the training of LLMs, as well as how those data have been annotated and upon which instructions, are questions with a considerable influence on LLM's functioning. Newsrooms should be aware of the normative values underlying the systems they use. These, however, are not always made explicit by the designers, nor are they always known by those who create these systems in the first place. Indeed, bias need not be explicit, but can be implicit.

As generative AI becomes ingrained into media and journalistic culture, these dependencies can significantly undermine the role of media and journalism as independent, trustworthy sources of information. Moreover, they potentially affect other values journalists hold dearly, such as the need for more diverse and pluralistic news content. Certain social narratives, and often those of invisible or otherwise marginalized communities, remain underrepresented, as their lived experiences may not be properly understood, captured, or translated by these systems. While this dependency on others might bring to the forefront issues of social and economic inequality, the latter issues are more broad and unlikely to be resolved by leveling the technological playing field. Indeed, it is unlikely to have these issues solved through technological solutions only. In this regard, traditional forms of news reporting and journalistic ethics remain crucial bulwarks within a democratic society: an active need for media and journalists to reach out to those who remain underrepresented.

Next to social and economic inequalities reflected through and enforced via data, the data themselves are often personal in nature. Is there a legitimate basis or ground upon which these data can be collected and processed en masse? We also need to consider the possibility of inversion attacks and prompt engineering to extract information from the input data, and possibly revealing sensitive

information in the process. Overall, we can describe these concerns as pertaining to moderating how AI systems produce content before the deployment.

The Need for Holistic Content Moderation Upon Deployment

There is, of course, also a need to moderate the content Generative AI systems produce upon deployment. Here, content moderation can be a valuable instrument to combat misinformation, fake news, and racist, discriminatory, or otherwise harmful content, such as the misgendering of a person or information – information that violates the artistic integrity of creators. Yet, many legal instruments that target content moderation, including the Digital Services Act (DSA), were conceived with human users and producers of content in mind. It is uncertain whether, and if so, under what circumstances, the DSA would apply to AI-generated content. Hence, an adaptation of the DSA's moderation instruments, such as trusted flaggers, audit mechanisms, and notice and takedown actions, to better capture GAI, could help repudiate some of these ills, and help media and journalists in their attempt to tackle both the misuse of their original content, as well as fight against misinformation and harmful content online.

Dependency and Competition

Introducing advanced technologies within the newsroom often requires financial means and expertise. While this is a luxury big media conglomerates can afford, it is often not given to individual, small, and medium-sized media actors. One popular narrative favouring the introduction of Generative AI models has been the potential of these technologies to level the playing field. This is a benefit we should not overlook.

Yet, here again, we might need to consider questions of dependence. The idea that media actors would need to act fast and adopt early in light of technological change could make the industry dependent upon technologies currently most developed. Moreover, generative AI depends on other technological infrastructures, particularly the Cloud. Hence, who ultimately benefits from AI's uptake, and who is competing in the field? Currently, the ultimate beneficiaries are not necessarily the perceived end-user but the more prominent players who dominate the market. If dependencies can be reduced by having alternative systems available, the barrier to entering the market must be lowered. Questions regarding dependence are

central to competition law: how can current competition mechanisms be instrumentalised in the relevant market, which would need to be defined to ensure, as a downstream effect, a healthy, digitised media landscape? And could the Digital Markets Act support competition goals? Yet, if the answer is yes, they can, what next? If we want to give smaller-sized actors access to the infrastructural means of production, what infrastructures should they be given access to: The underlying code? The underlying data? This, however, brings a new set of questions: what about the copyright of the initial creators? What about the privacy and data protection of the data that has been created? And again, if biases have been ingrained within these large data sets, do we not hegemonise a particular worldview through their integration into novel technologies?

In that regard, we could consider whether, through these regulatory mechanisms, we can increase the agency and control over these technologies on an individual level. Given these systems have been trained on the entirety of the internet and its users, would it not be appropriate to make sure everyone has access to the results, or be able to use these tools on their computer?

The Protection of Intellectual Property and Creativity

Within the media creative industries, the emergence of Generative AI has raised several questions regarding these systems' rightful appropriation of content and the creative relationship between human and technology.

On the output level, it is useful to bear in mind that news of the day, as well as general ideas, concepts and styles, do not fall within the scope of copyright protection. Hence the extraction and use of news stories through AI pose no immediate issue as long as the specific way in which that news story was initially expressed was not copied. Likewise, style imitation as such is normally not enough from a copyright perspective to trigger protection. Copyright claims will only have bite when there is evidence a system has taken over individual creative choices of the journalist and reproduced them. Yet, proving this has been the case, will be quite tricky given the black-box nature of these systems.

From an input perspective, it can be difficult to attribute the creative and artistic expression requirements typically linked to copyright protection. This relates to the previous point: we do not always know what types of data and works are used, and how they are used in order to arrive at a specific result.

To address these issues, we would need to adapt the rules on evidence for infringement. Either the law would need to reverse the burden of proof, whereby the system developer must show a certain work has (not) been used for the planning process or we would need to start working with probabilities and assumptions rather than hard evidence. Perhaps then, if we can illustrate that there is a sufficient degree of probability that a pre-existing journalistic work was used as source material, it could trigger remuneration claims. While the media industries could invoke copyright legislation to prohibit the mining of their works, they would also not be able to get remuneration due to the current design of the text and data mining exemption in the Copyright in the Single Market Directive.

Do We Want Change, And When Do We Want It?

Though Generative AI systems provide both opportunities and challenges for journalism and media industries, a more fundamental question that should be asked is whether these systems should be adopted in the first place.

As part of the discussions regarding the proposed AI Act, voices have claimed that these systems should be considered high-risk. This however has been greeted with concerns of overregulation and the potential to stifle innovation. However, even if these systems are considered high-risk, it might be too little too late. The AI Act is focused on select areas or domains in which the use of AI is considered high-risk, such as education or law enforcement. Yet, Generative AI can be deployed in many areas. The risks associated with Generative AI, such as the aforementioned inequalities or copyright issue, would trickle down into a myriad of possible and difficult-to-predict future applications. Data quality standards might work well if one governs specific usecases, but might be more difficult to develop where general-purpose AI is involved. The true risk might not only be the application of GAI in certain areas but the pure massive scale of application. In this regard, the media sector, and any sector for that matter, including its individual members, have its own specific risks: there is a diverse set of actors with different skill sets and media literacy, with very different organisational embeddings and safeguards. In this regard, lessons could be learned from the DSA. For instance, we could introduce a monitoring mechanism to gauge and map systemic risks.

This also raises questions concerning the responsibility of media and end-users in their use of AI. The AI Act establishes the need for transparency in case people communicate or consume AI-generated content. Yet, in order to comply with this

obligation and other potential obligations under the AI Act, media actors may be dependent upon information they receive from the developers of these systems. Hence, there is a cooperative responsibility. We need to find better ways to divide respective responsibilities and increase cooperation to ensure that different players throughout the value chain ensure these systems are deployed in a fair and reasonable manner. Yet, when we allocate these responsibilities to both public and private actors, there is a need to reflect on how these relationships should be established – if we want to give these responsibilities to private actors in the first place. For instance, could we foresee a future where private actors can determine, through licensing schemes, for which purposes their systems may not be used? And if so, would this be desirable?

Responsibilities pertaining to data are not only about their reliability and the biases they exhibit, but should also pertain to their creation. Data must be labelled for example, and the conditions under which these labelling processes occur are often dire. These efforts tend to be outsourced to low-income countries. The precarious nature of these labelling tasks not only extends to the remuneration labourers receive. The act of labelling necessitates people to become exposed to harmful, illegal and offensive content that can cause psychological harm. There is a shared responsibility when we partake in the uptake of these systems: media actors should be aware that the values they represent in their practice might not extend to the technologies they use.

In recent months, scholars and industry actors have called upon the need to reflect on AI's impact and considered the need for moratoria regarding the development and deployment of AI systems. These risks are not new. The landmark paper ["On the Dangers of Stochastic Parrots"](#) by Bender and others is a must-read primer regarding the potential implications of Generative AI. The risks that have been raised also pertain more foundationally to the democratic values the media industries look to safeguard. The democratisation of Generative AI technologies increases the possibility for bad actors to manipulate readers or, to easily produce misinformation. We should therefore map foreseeable risks while creating a monitoring process to assess and address (hitherto unknown) future risks. As part of this exercise, we should consider what legal tools are currently available, the possibilities they offer, and the limitations they have, as well as consider future regulatory action where needed.

Questions From the Audience

Multiple audience members shared questions related to intellectual property and generative AI, underlining the urgent need for insight in this particular area. One question, in particular, addressed the complex relationship between journalism, content used by AI models and copyright: 'The fact-based nature of news and journalistic pursuit for factualness may contradict the originality and creativity requirement for copyright, especially when automated news is scaled up by algorithms. How can automated news be better protected, so that press publishers are motivated to innovate and produce quality journalism?'.

The answer to this question was not singular in nature. On the one hand, Article 4 of the Digital Single Markets (DSM) Act allows authors to reserve copyright, ideally in machine-readable form. This would mean that a crawler, looking for the latest news, would not use that journalistic text when coming across their reservation of rights. However, developers cannot ensure that they are actually identifying all opt-outs due to a lack of standardisation. In addition to that, it is debatable whether this is a desirable outcome in the first place. A preferable scenario could be that their text is mined and a source of revenue comes back in return. This is an ongoing discussion and a standard to ensure money flows back directly to the original journalists that have created the news story has not been established either. Alternative mechanisms could focus on providing evidence of copyright-relevant use, for example by a probability-based assessment of the use of source material. The problem from a copyright perspective is always that news of the day as such remains free. The mere use of the news story, the news facts as such, does not have copyright relevance. It is required to further open the 'black box' of the model.

Joint Discussion

The last 20 minutes of the event were dedicated to analysing 3 core questions in order to establish a common research agenda: are there novel problems in generative AI? Do we focus on the right or wrong issues at the moment? What kind of collaborations are needed between various legal fields? Please find the main remarks of the panellists below.

Firstly, all panellists agreed that tackling topics related to generative AI will require cooperation between different legal fields. As one person said, the 'main challenge is to appreciate the whole of what these systems are and finding a way to balance

that'. That will be difficult and will require much collaboration (from the perspective of intellectual property, competition law, non-discrimination, equality, data protection, etc.).

Secondly, fairness was underlined as a concept that should be explored in more depth. New dimensions of fairness could be helpful in this context ([data](#) extraction fairness in particular). How fair is it to collect all this data from the internet, train models and economically benefit from this? This is a narrative that would require research from legal and non-legal perspectives.

Thirdly, from a competition angle, the market should be constantly observed, looking at the main players and what kind of partnerships are being established. The DMA only applies to gatekeepers who have core platform services. Generative AI (or AI in general) is not mentioned. There is a possibility to extend the notion of core platform services to generative AI. However, this would require the Commission to conduct a market investigation running for up to 18 months (starting from the date of the official decision opening the market investigation), in line with Art. 19(3) and Art. 16(3), point (a) of the DMA. From a media perspective, how would a concentration of power in generative AI affect journalists' activities?

Fourthly, even if a small percentage of people using generative AI will use it for disinformation purposes, it could have a big impact. One way of tackling this problem, which has been proposed by one of the panellists (in addition to regulation, human oversight, more responsible user behaviour, etc.) involves AI that recognises disinformation. If only few had access to those technologies, and therefore the ability to defend themselves from disinformation, this would be highly discriminatory. A question that arises from this is whether media and (especially smaller or local) journalists will have access to such tools.

Fifthly, while perhaps not strictly media-related, the environmental dimension of generative AI was mentioned as highly problematic.

Sixthly, it is urgent to tackle the question of responsibility within the generative AI value chain. To do so, provisions such as Art. 26 GDPR on joint controllership or certain CJEU cases (for example, Fashion ID Case C-40/17) may be helpful. Approaching this correctly is crucial for journalists and media organisations in order to understand and rightfully assign their responsibilities for using, developing, and deploying generative AI-related systems.

Finally, we are entering a new area of machine-mediated thinking and perception of reality. These generative AI technologies will enter our routines of creating new work. If a journalist uses it as a template for a news story with information generated by a machine, this is a story built on machine data analysis (not entirely the journalist's work). If a press publisher increasingly relies on the process of automated creation (for example, to reduce expenses linked to sharing profits with journalists) then news sources are populated with machine-generated data. All of this will have an impact on how people perceive the world. It is important 'not to get lazy' due to those systems and not to lose the potential to quickly and effectively react to societal challenges.

Concluding Words

The dialogue initiated by this event made apparent that a more nuanced and cooperative approach toward the regulation of Generative AI systems is needed. In order to arrive at a sufficiently robust and future-proof regulatory framework, a multi- and interdisciplinary dialogue is called for. We can take up this task with optimism. Though complex due to coordinated efforts, we believe that each respective domain can be instrumental to support the media and journalism industries and bolster their function as the fifth estate, a public watchdog, a source of entertainment, etc. Even more so, by establishing a mutual understanding between disciplines, we will start to see not only how each domain can contribute to this overall framework from their respective "area of governance", but can help other legal areas flourish in their realisation of this common goal.

This, however, will also require legal scholars and regulators to cross their own disciplinary boundaries. Sufficient attention must be paid to the broader social and economic impact of Generative AI. The social and economic environment in which these systems are embedded largely determines their functioning. And while Generative AI systems take meaning from their environment (e.g., they draw from existing information and knowledge, including societal biases) during the design phase, they co-produce meaning upon deployment and interaction (e.g., the generation of news content, including, unfortunately, the potential to produce false or misinformation). Given the central role the media and journalism industries play as a source of information, entertainment and democracy, this mediation role of technology should be adequately understood. Hence, it is advised that the media and journalism industries pay attention to legal developments in the fields so they can understand what they can and cannot do with these technologies. At the same time, legal scholars and regulators should make a conscious effort to help translate

these legal limitations and opportunities and help the industry navigate the complexity of the law.

In case you are interested in what a Generative AI system might take from our event, we kindly refer you to this blog post of our colleague and visiting scholar Prof. Dr. Nick Diakopolous, who used our event as a little experiment to answer the question: "[Can GPT-4 help cover events?](#)".

Finally, we would like to thank our panellists, Prof. Philipp Hacker, Viktorija Morozovaite, Prof. Martin Senftleben, Prof. Natali Helberger, Dr. Ot van Daalen, and Naomi Appelman for making this discussion on the legal aspects of Generative AI possible. Moving on from this successful first session, the *AI, Media and Democracy Lab* will host the second session expanding on [News Media and Generative AI](#) on April 18th. Aiming to explore the more practical challenges of Generative AI, we will answer questions regarding how the implementation of these technologies impacts news reporting and how they will alter the process and product of news media and journalism. The final session on May 30th will focus on the [Societal Impact of Generative AI](#).