



UvA-DARE (Digital Academic Repository)

Generative AI in Media & Journalism: Think Big, But Read the Small Print First

Helberger, N.

Publication date

2023

Document Version

Final published version

[Link to publication](#)

Citation for published version (APA):

Helberger, N. (2023). Generative AI in Media & Journalism: Think Big, But Read the Small Print First. Web publication or website, Medium. <https://medium.com/generative-ai-in-the-newsroom/generative-ai-in-media-journalism-think-big-but-read-the-small-print-first-375f2ecb1256>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

★ Get unlimited access to the best of Medium for less than \$1/week. [Become a member](#)



Generative AI in Media & Journalism: Think Big, But Read the Small Print First



Natali Helberger · [Follow](#)

Published in [Generative AI in the Newsroom](#)

18 min read · Jul 18, 2023





Image created with Adobe Firefly.

Terms of Use — typically hidden away at the bottom of a page, barely readable with the naked eye and notorious for their use of legalese— are where the use of generative AI systems is currently being governed. The Terms of Use of generative AI providers such as Open AI, Midjourney, Stable Diffusion, and others lay the ground rules for the responsible use of generative AI, instruct downstream users

how to mitigate potential safety risks and risks to the realisation of fundamental rights, and create binding legal obligations. Recent legal initiatives such as the draft AI Act are likely to increase the importance but also the weight of the Terms of Use and Usage policies because it is here that providers of generative AI instruct users how (or how *not*) to use the technology and distribute responsibility. At the same time, the AI Act is also likely to trigger closer scrutiny and trigger a debate on the fairness of some of those terms.

In the following post, we will critically and thematically review the Terms of Use and usage policies of five generative AI providers and recommend a number of terms to the attention of media organisations that are planning to use the models.^[1] We included proprietary and non-proprietary, open systems and platforms in the selection, including Open AI, Midjourney, Anthropic, Hugging Face, and StabilityAI.

Before we get into the specifics, here are some of our key take-aways:

- Before starting to use or integrate a particular generative AI model media organisations would do well to carefully read the Terms of Use. Different providers make their models available under different conditions.
- Of particular importance are the conditions on who is liable for the output of the model, who owns the rights to any prompts or contents shared with the model, the output, and to what extent the provider reserves the right to use any pictures, articles, and other content from a media organisation to further train the model. Some providers of generative AI also prohibit the use of their models for particular use cases, or stipulate additional conditions for the generation of news content (such as OpenAI).
- Some conditions are potentially problematic from the perspective of media organisations, and merit a broader discussion within the sector of what acceptable or unacceptable commercial practices are, also with view to the pending AI Act.
- Open-source models typically offer media organisations the possibility to exercise more control, have more transparency and adjust a model to their own needs and preferences, provided the necessary in-house skills are present, or there is a commitment to invest in such skills.
- Using generative AI responsibly includes choosing a responsible provider of generative AI. Additional information to look for includes any actions that a

particular provider has undertaken to address issues of bias, support for research, the lawfulness of training data, to reduce the ecological footprint of the model, and guarantee fair labour conditions of information workers.

Fair division of responsibility

The division of responsibility between providers of generative AI models and downstream users for compliance with the legal obligations under data protection law, copyright law or the pending AI Act is not always clear-cut and will require in many instances cooperation and a workable division of tasks. For example, for media organisations to be able to do a risk assessment or ascertain that content generated by an LLM does not infringe upon copyrights or data protection rights of others, they require information from the provider on how the model has been trained and what initiatives were taken to account for the rights of third parties. Vice versa, in order to be able to demonstrate that the copyrights of a media organisation or journalists have been infringed upon in the process of training the model, the media organisation depends on the cooperation of the developer of the model to prove that the latter had access to that particular piece of content. Ideally, the contractual agreements between providers and downstream users therefore include provisions on the mutual rights and obligations of parties, including issues around transparency, access to data, support, but also mutual assistance and cooperation when it comes to conformity assessment, risk management and legal responsibility. The need for a fair division of tasks and some form of cooperative responsibility is reportedly also an important point of departure in the European Parliament's compromise text, demanding a division of responsibility that reflects the different parties' levels of control.

An initial review of the Terms of Use of the five providers of generative AI reveals that the contractual reality is far away from this ideal of mutual assistance. Typically, the Terms of Use determine that downstream users bear full responsibility for making sure their use of the technology complies with legal requirements from e.g. data protection and copyright law or the AI Act and is otherwise responsible. All of the Terms of Use examined contain the extensive disclaimers of responsibility typically known from other software contracts. In addition, all include indemnification clauses obliging the downstream user to defend the provider against any claims that may arise from their use of the model, typically at their own costs. None of the terms analysed addressed the situation that safe use of the model or compliance with the applicable laws by downstream users

is not possible without the cooperation of the developers of the models themselves. None of the Terms of Use examined included any language of assistance, described a responsibility on the side of the providers of the models to cooperate or any rights to transparency on the side of downstream users. A notable exception was the Terms of Use from Hugging Face, committing to “all reasonable assistance and cooperation in connection with the defence and settlement of the claim, at its own expense”, albeit limited to a situation of litigation.

And in case of questions? Stable Diffusions FAQs are hopefully not symptomatic: “If you can’t find what you’re looking for, email our support team and if you’re lucky someone will get back to you.”

Dealing with dependency

A potentially problematic term from the perspective of downstream users is the *lack of guarantees regarding the continuation and support of the model*, particularly to the extent that a media organisation is deciding to rely on external technology instead of building their own model. Midjourney’s Terms of Use are most outspoken and warn all media organisations that want to rely on that model: “Please do not create any dependencies on any attributes of the Services or the Assets. We will not be liable to you or your downstream customers for any harm caused by your dependency on the service.” In other words, Midjourney reserves the right to change or even cease the service, without an obligation on the part of the developer to even issue a prior warning. While understandable from a business point of view, the fact is that seeing the resources needed to build a generative AI model, at least today only a few companies will have the capacity, finances, skills and data to do so. The rest will depend on the use of generative AI models by the few companies that are in a position to deploy their models. This does not only mean that large parts of the downstream market may end up building their futures on extremely unstable ground, it also means that without some safeguards to transparency, due notice and continuation, downstream providers will have little to no negotiation power, and find themselves in a constant state of economic uncertainty. The situation echoes earlier concerns about problems with the lack of transparency and predictability of some of the decisions of social media platforms, including the Facebook Newsfeed, and the consequences for those that depend on the platforms — a situation that has ultimately resulted in regulatory action in the EU in the form of the Digital Markets Act.

The lack of predictability and guarantees is potentially further exacerbated in a situation in which Terms of Use are used as an instrument to chill competition and alternatives. An example is the inclusion of non-competition clauses (which can be found e.g. in the Terms of Use of all the proprietary models examined). For example, Open AI does not allow to use the output of the model to develop models that compete with Open AI, and according to the Terms of Use of Anthropic, using prompts and results from the model to train other models will be considered fraudulent use in breach of the Terms of Use.

To the contrary, open-source models leave their users more possibilities to scrutinise the model and data sets, adjust existing open source models, and exercise control of their own. Doing so requires a certain level of in-house technical expertise, or the willingness to invest and build up such expertise. From the competition standpoint, encouraging investment in skills and alternatives, and reducing dependency by leveraging platforms for open-source models are worth serious consideration in case the lack of predictability and dependency on external providers is a concern.

Research, auditing and critical reporting

Importantly, and rather ironically, all proprietary models examined have extensive language that bans the reverse engineering or scraping of the model and its output — ironically because scraping the public internet is what has allowed training some of those models in the first place. Open AI does not allow reverse assembling, compiling, decompiling, translating or otherwise attempting to discover the source code or underlying components of models, algorithms, and systems; or the automated extraction of data or output, including scraping, web harvesting, and web data extraction. Similarly, Midjourney and Anthropic ban any attempts at reverse engineering and scraping.

Practically, this makes auditing and research into the system to understand its affordances and limitations difficult and potentially a breach of the Terms of Use, depending on how the provider will interpret these restrictions. Also, testing the systems for critical reporting about their functioning and investigative data journalism can potentially fall under the contractual restrictions. Clauses like these echo very similar concerns regarding the clauses of social media platforms that restricted access to data and that were the basis for attempts of platforms to shut down research projects ([here](#) and [here](#)). In Europe, these debates resulted ultimately in the adoption of Article 40 DSA, which however, only binds Very Large

Online Platforms and Very Large Search Engines in the sense of the DSA, and only benefits researchers, not journalists. The restrictive provisions on reverse engineering could be another important reason to choose an open source model, at least for downstream users that are not willing to use a model that they have no means to understand or audit, or that wish to invest in building in-house skills and capacity.

This is not to say that providers of generative AI do not engage in, and stimulate research. Anthropic, for example, makes research into responsible development of generative AI an explicit part of its mission. And OpenAI writes in its Charter: “We seek to create a global community working together to address AGI’s global challenges”. But it is the providers of the models that dictate the conditions of research. Besides some sporadic blog posts and white papers, providers like Anthropic and OpenAI seem to keep most research in-house. An important element of OpenAI’s strategy to addressing generative AI’s global challenges is conducting and publishing details on its alignment research and safety research, as well as running a researcher access program in which third parties can apply for credits to use the API. Grounding those models in (their own) critical research is certainly to be applauded, as is supporting third party research but it is also important to flag that the OpenAI research program or its Terms of Use do not give an independent right of access for research purposes. It is thus entirely up to the discretion of OpenAI to decide what research to support or not.

Managing Data Rights on Inputs and Outputs

Over the past months, we have seen controversies around the IP rights of the original creators of material that the models have been trained on, as well as concerns around data protection and the extent to which the models were trained on personal data without consent. The use of generative AI is as privacy-friendly as the data it has been trained on, and problems with the lawfulness of the content used to train the model can trickle down into downstream uses (e.g. in the case of training the model on contents that are protected by copyright law). Similarly, if equality and diversity are important editorial values, the ability to realise those values also critically depends on the external technology provider. Therefore, responsible use of generative AI starts with responsible procurement and selecting providers that can substantiate how in the development and subsequent deployment process, the rights of data subjects and right holders are taken into account.

Again, most providers are relatively silent on this point. Some Terms of Use mention the possibility for original right holders to file a complaint. Hugging Face offers a dedicated DMCA email address for rightsholders, and also Midjourney foresees a notice-and-take-down procedure. In both cases, the burden of proof that the model has been trained on copyrighted content without consent is with the respective rightsholder. It will be interesting to see how the European Parliament's recent suggestion to mandate a publicly available summary on the use of copyright-protected works for training the models may help rightsholders to arrest their claim.

While such a distribution of responsibility may be common practice elsewhere, in the case of generative AI and without transparency about the training data, making and documenting such a claim may impose insurmountable burdens on the right holder. Neither Midjourney, Open AI nor Anthropic volunteer any transparency about the training data, at least not on the website or Terms of Use. In contrast, Stable diffusion explains that the model is trained on the LAION dataset. Maybe most detailed and possibly the best practice so far (at least among the providers examined) is HuggingFace's approach to its open-source Bloom model. Hugging Faces provides detailed information about training data, but also foreseeable users, model objective and architecture, risks, and limitations.

Then there is the question of how the Terms of Use handle ownership in the data acquired *through* the use of the system. For example, who has the right to content that a media organisation uploads into the system, for example to make summaries or rewrite in a different format? Again, for media companies planning to use generative AI to (co)produce output this can be a relevant question, both in terms of their ability to control the publication of the output they create, as well as in terms of confidentiality of the prompts or content uploaded. Here, the picture is diverse. At one end of the spectrum is Stable Diffusion: "Images created through Stable Diffusion Online are fully open source, explicitly falling under the CC0 1.0 Universal Public Domain Dedication." In other words, users of Stable Diffusion's image generator must be aware that the images they create end up in the public domain by default. Similarly, Midjourney stipulates that non-paying users cannot claim any ownership, either in the prompts or in the output generated. Instead, Midjourney issues a CC Noncommercial 4.0 Attribution Licence, meaning that third parties are free to use and remix images and prompts, and by default, all content generated is publicly viewable. If Midjourney users wish to exercise more exclusive control, they need to sign up for a paid Pro Membership.

At the other end of the spectrum are providers like Open AI that assign users full ownership over prompts and output, including for commercial uses. Interesting is OpenAI's definition of when exactly a user owns the output. Whereas under copyright law, it is the creator of an original work that owns the IP rights, in the realm of OpenAI the owner of a prompt is the one who requested and prompted the system, unless someone else used the same prompt. Since it is impossible to know if other users used the same prompt this condition does not contribute to much legal certainty. In addition, the debate is still ongoing on how Open AI's "right to the prompt and its output" relates to competing claims of, for example, the original creators whose content the generated output may resemble.

At the same time, some of the Terms of Use try to carve out their own exception for training data. According to OpenAI's Terms of Use, for example, content collected from the API is not used for training the models in principle, but users can opt-in to share data. In contrast, in the case of web-based services such as ChatGPT and DALL-E, the default is that data input can be used to improve the model unless users opt out. In other words, if a media organisation would rather not like its content or prompt to be used to train the system (by default), it should access the model via the API. A variation on this is the Terms of Use from Anthropic or Hugging Face that in principle recognise that users retain all rights to prompts and results but that nevertheless reserve the right to use all **publicly available** prompts and results to train their systems (Anthropic) or those of others (Hugging Face). This could mean, for example, that content created with a generative AI model and published online could be used for training purposes. On the far end of the spectrum is Stable Diffusion, reserving the right to train their model on all kinds of data and offering no opt-in or opt-out for the LAION 5b model data because the model "is intended to be a general representation of the language-image connection of the Internet."

Defining responsible use of AI

Next to delineating rights and responsibilities in the downstream relationship, the Terms of Use of generative AI providers will also have a critical role in defining what responsible or irresponsible use of generative AI is, not only in relation to professional users but also end-users. Being at the centre of the debate around possible risks to society, arguably, generative AI providers are not only well positioned to identify and mitigate societally harmful uses of AI by both professional and individual users, they arguably also have a societal obligation to do so. Apart from that, terms of use and usage policies are also important means of risk management and governing the responsible use of AI.

Accordingly, the creators of the RAIL (Responsible AI Licence) have argued: “licenses could serve as a useful strategy for creating, expanding, and enforcing responsible behavioural norms given the limitations of self-regulation and governmental legislation. [...] we advocate for the use of licensing as a mechanism for enabling legally enforceable responsible use.” A case in point is Stable Diffusion that explicitly characterises usage restrictions as a tool to govern downstream responsibility: “Referring to the downstream responsible use, we added use-based restrictions not permitting the use of the Model in very specific scenarios, in order for the license to be able to enforce the license in case potential misuses of the Model may occur.” It is worth mentioning that Stable Diffusion also determines that the original use-based restrictions must also be included — ‘at minimum’ — in downstream derivative versions. And in a (leaked) policy response of Open AI to the AI Act, the company made repeated references to its terms of use as one of the ways in which Open AI addresses concerns around the responsible use of its model.

Having said so, the potential of Terms of Use to set legally enforceable conditions of responsible downstream use critically depends on the way those conditions are defined. The Code of Conduct of Midjourney is relatively short, for example, and can be summarised by its first rule “Don’t be a jerk.” In other words, Midjourney primarily appeals to the responsibility of end-users to respect the rights and interests of others and also leaves it to the imagination and responsibility of users to define what that means. In practice, this also means that legally enforcing these conditions will be difficult, if not impossible due to their vagueness. In contrast, all of the other providers examined provide more or less extensive attempts to identify prohibited uses and use cases. OpenAI, for example, includes a very detailed list of non-permitted uses and use cases. Particularly relevant for the media will be its guidelines for co-authoring content with Open AI: according to these guidelines, content co-authored with Open AI must respect Open AI’s terms of use, must not offend others, clearly disclose the role of OpenAI’s model, provide a disclaimer to users informing them that AI is being used and of its potential limitations. In practice, doing so can be difficult to the extent that Open AI does not share what the limitations of the model are, and prohibiting the use of generative AI to create content that offends others can potentially limit the fundamental right of the media to report content even if it shocks or offends.

Comparing the different instances of prohibited uses in the terms of use and usage guidelines, some identified use cases are more uncontroversially irresponsible than

others. There seems to be broad agreement, for example, that generative AI models may not be used to create unlawful content, and that deals with sexual abuse of minors, as well as hate speech, harassment, defamation, bullying, and discrimination. Interestingly, the list of non-permitted content reflects more or less broadly the types of content that are not permitted to publish or distribute on social media sites. A more nuanced picture arises around adult content. While OpenAI, for example, does not allow its models to be used to create adult content, Hugging Face does not banish synthetic adult content generally but requires securing access to such content via opt-in mechanisms.

Next to certain forms of content that may not be produced with generative AI, certain uses or use cases are prohibited. Here, the picture varies more broadly. And while some of the categories of prohibited uses are reminiscent of some of the “high-risk areas” that the European draft AI Act defines (such as the use in high-risk government decision-making, including law enforcement and criminal justice, migration and asylum), others are new. For example, Open AI does not allow the use of its models in the context of political campaigning, or for giving legal, financial or health advice. Anthropic, too, banishes the use in political campaigning and lobbying but also the use of generative AI to track or target individuals. Stable Diffusion adds the production of disinformation to the list of prohibited uses, but also using generative AI to make fully automated decisions and exploit the vulnerabilities of individuals. Anthropic created an additional category of “restricted business use cases” for which a special “Restricted User Agreement” is needed.

Finally, if the providers of OpenAI define the ground rules for the use of generative AI in society, they are not particularly democratic about this. None of the usage policies of the proprietary models explain how the providers came to the conclusion to ban certain uses, while permitting others, though some point more clearly to, and facilitate a role for the public to participate in that process. Or as Midjourney explains: “We are not a democracy. Behave respectfully or lose your rights to use the Service.” And Open AI simply stipulates: “Our policies may change as we learn more about use and abuse of our models.” In contrast to this top-down approach, open source models like Hugging Face’s invoke the community’s responsibility for developing standards of fair behaviour: “As we as a community develop better norms around training data and AI-related rights, we also ask that you take

reasonable steps to honour other people's wishes regarding systems trained primarily on their data.”

In Sum

In the discussions around the ethical and legal uses of generative AI and the need for defining some guardrails on what (ir)responsible use of generative AI means, there is currently a lot of attention on the role of formal regulators but also the emerging ethical guidelines that many media organisations are working on. So far, there has been relatively little attention on the Terms of Use of the providers of generative AI. And yet, it is the Terms and Use that are currently one of the most important sources of legal obligations between media organisations that wish to use generative AI, and the providers of these models. Also, the Terms of Use play an important role in defining what responsible use of generative AI is.

A review of the terms of use of five providers of generative AI models or platforms has shown that:

- There are certain conditions in the Terms of Use that are particularly relevant for media organisations that want to use generative AI. These include the provisions around the division of responsibility for the output generated, as well as the rights to input and outputs of the model. The conditions can differ between providers. Before using a particular model, media organisations probably want to carefully read what the Terms of Use stipulate about who is responsible in case content generated is inaccurate or infringes the copyrights and data protection rights of third parties (typically the media organisation), who owns the rights to any prompts or contents shared with the model, the output, and to what extent the model reserves the right to use any pictures, articles and other content from a media organisation to train the model, if there are particular labelling or other obligations when generating news content (like in the case of Open AI), if there is a possibility to contact the provider, and what happens if the model is changed or terminated.
- There are certain conditions in some of the Terms of Use that are potentially problematic from the perspective of a media organisation, including provisions that make auditing and critical reporting unlawful because of a ban on scraping or reverse engineering, the lack of any prior warning in case a model is changed or terminated, or provisions that place full legal responsibility for the output generated and the use of the model with the media organisation even in

situations in which a media organisation lacks the necessary information or cooperation from the provider of the model.

- Finally, there are issues that are typically not addressed in the Terms of Use but that media organisations may want to know about to be able make a responsible choice. This is because the responsible use of generative AI also means choosing a responsible provider. Examples include any initiatives on the side of the provider of the model to address concerns about biases, transparency about the underlying infrastructure (e.g. with which cloud service provider a particular model is affiliated, a question that could influence for example where data will be stored), the possibility to adjust a model to adhere to a media organisation's own editorial values, but also what a provider has done to reduce the ecological footprint of a model, address concerns around the working conditions of information workers or issues around the security of a model.
- Regarding the latter two categories, more discussion and potentially also cooperation among media organisations is needed to bundle and exercise sufficient negotiation power to address potentially problematic Terms of Use.
- Finally, with the pending AI Act, there is momentum to have a broader discussion on what “good commercial practices” are with respect to the provision and use of generative AI in the media. According to one of the newest proposals in the AI Act, for example, a contractual condition that is contrary to good faith and fair dealing or creates a significant imbalance between the rights and the obligations of the parties in the contract shall be considered unfair.

[1] The analysis of the Terms of Use took place in the months May and June 2023.

Generative Ai

Journalism

Terms Of Use

Ai Act

Values




Follow

Written by Natali Helberger

23 Followers · Writer for Generative AI in the Newsroom

More from Natali Helberger and Generative AI in the Newsroom



 Natali Helberger in Generative AI in the Newsroom

A Possible Future of News in 2043: FutureNewsCorp

Natali Helberger elaborates a scenario around a fictional AI-driven news organization—FutureNewsCorp.

Jul 2  21

