Beyond Intermediaries

Generative AI in Search of Their Legal Status

Marco Bassini & Andrea Palumbo

12.04.2026 doi: 10.59704/887f92f3bcf01a98

Recent investigations into the dissemination of illegal content generated by Grok have exposed a structural gap in the EU's legislative framework: while the Digital Services Act (DSA) equips the European Commission with far-reaching powers over large online platforms, it does not clearly capture generative AI systems per se. As a result, the Commission may be able to act against platforms integrating such systems (such as X), but not necessarily against the systems themselves (such as Grok).

This asymmetry raises a broader question that has increasingly gained attention in policy and academic debates: can generative AI applications be brought within the scope of the DSA, for instance as very large online platforms (VLOPs) or search engines (VLOSEs)?

We argue that applying the DSA – or an ad hoc regime based on the DSA – to generative AI (GenAI) could prove beneficial to address the risks posed by GenAI.

Is the DSA a good fit for GenAI?

Chatbots powered by generative AI are a real brain teaser when it comes to mapping the applicable EU legal framework. They stand out as a new service with its own distinctive features, but they also functionally resemble other services that are already regulated under EU law. Generative AI applications can exhibit platform-like features and offer users functionalities that are comparable to those of online search engines. However, they are difficult to fit within pre-existing legal categories.

The DSA is premised on the well-established categorization of intermediary services crafted by the e-Commerce Directive (Directive 2000/31/EC): mere conduit, caching, and hosting services. The DSA has never departed from the underlying assumption that those regulated under its umbrella are intermediary services, to which a special liability regime for third-party content is granted. Historically, the reason why intermediary service providers were granted a special regime lies in the absence of editorial control, which marks the difference between them and content providers such as publishers. Accordingly, they are not subject to general monitoring obligations in respect of such content, nor do they incur liability for any illegal content published by their users, unless they fail to comply with the notice-and-action mechanism.

When assessing whether some GenAI applications can fall under the DSA, one should take these rationales into account and assess whether the way GenAI applications work is aligned with the structural paradigm of intermediary services. In this regard, one can notice more than a mere divergence.

Anatomy of a service

A first common feature of intermediary services is that the information they transmit or store is provided by a recipient of the service. Social networks or user-generated content platforms, for example, only store content posted or uploaded by their users, consistently with their curational but never editorial role. As noted by Edwards et al., however, GenAI applications do not match this characteristic: the only information provided by the recipients of the service is the prompt used to ask for a specific task to be performed, while the output returned is AI-generated (albeit user-prompted).

This first problem shifts our analysis towards a second issue, namely whether the production of AI-generated output consists of a genuinely creative activity (Bassini; Lopes-Bassini). Qualifying the operation of GenAI applications as content creation would

trigger significant obstacles to the application of the DSA. These systems would de facto operate as content providers, which fall beyond the scope of intermediary service providers and therefore do not enjoy the special liability regime.

Hybrids between online search engines and online platforms

Despite these conceptual difficulties, scholars have reflected upon the suitability of the categories of VLOPs and VLOSEs to capture GenAI providers with a remarkable market presence.

Emphasizing that the notion of online search engine under Art. 3(j) DSA is agnostic to the format in which ''information related to the requested content can be found'' and returned, Botero Arcila has maintained that some GenAI applications, such as LLMs, could already fall under the relevant category. In accordance with this view, Schaal, Lenne and Akinyemi equally support ChatGPT's designation as a VLOSE, as Art. 3(j) would encompass ''services that retrieve and synthesise web information at scale, regardless of whether they return traditional link lists''. Lorente and Gardhouse, in turn, have argued that ChatGPT falls under the DSA as an online search engine but should more properly be considered as a hybrid between an online search engine and an online platform. The latter, however, acknowledge that GenAI applications do not meet a requirement for online platforms, namely that they engage in the dissemination to the public of the information they host. Lorente and Gardhouse underscore that the requirement is not met for user prompts, but questions may arise on whether they genuinely ''disseminate to the public'' also the output they generate.

More generally, a key challenge remains as to whether hosting content that GenAI applications generate themselves would not deprive the providers of their role as intermediaries, also in light of the concept of 'active hosting provider' in the pre-DSA age.

A stick-and-carrot approach

Even if the DSA might not be a very well-fitting dress for GenAI applications, its application and enforcement could still prove beneficial.

The AI Act does not bring much clarity to its relationship with the DSA, and Recitals 118 and 119 only address the coordination of the respective risk management frameworks, without suggesting any solution regarding the content moderation regime. Despite the conceptual tensions outlined above, there are practical reasons to consider applying the DSA to GenAI applications, regardless of whether they are integrated into digital services or operated as standalone solutions. We believe that this conclusion would derive from a ''stick-and-carrot'' approach (Botero Arcila; Hacker et al.; Bassini).

As to the ''stick'' component, the DSA offers a set of risk-governance mechanisms designed to address the societal risks associated with large-scale information services. Trusted flaggers, transparency obligations, and systemic risk assessments provide tools to identify and mitigate harmful and illegal content.

Extending similar mechanisms to GenAI providers could help address the risks associated with AI-generated content, including disinformation, defamation, and other forms of illegal or harmful speech. It would be paradoxical if such mechanisms were applied to VLOPs integrating AI components into their services (such as in the case of X and Grok) – as the recent Commission investigation suggests – and were conversely excluded in the case of other GenAI applications simply because they operate as standalone solutions.

But there is also a ''carrot'' component, which lies in the favorable liability regime for illegal content. It could be disputed that such content qualifies as ''third-party'' when it is generated by AI systems on the basis of prompts input by the recipients of the service. However, this solution could mitigate the burden of applying risk governance obligations for systems with a large market presence (and presumably designated as VLOPs or VLOSEs). Overall, it would

align the status of GenAI providers with that of hosting providers.

This scenario could represent a pragmatic solution. Even if generative systems do not neatly fit within the intermediary taxonomy, the combination of notice-and-action procedures and systemic risk obligations could strengthen accountability and facilitate the moderation of problematic outputs.

Why applying the DSA would matter

As discussed above, the ''stick'' component of a stick-and-carrot approach would entail applying the DSA risk management obligations to GenAI. It may be argued that GenAI already has its own systemic risk management regime, the one applicable to providers of general-purpose AI models (GPAIMs) under the AI Act.

The implementation of this regime may also intersect with the systemic risk framework established under the DSA, as acknowledged in Recitals 118 to 120 DSA. However, despite the apparent convergence suggested by these recitals, the two regimes are not equivalent. Systemic risk management under the DSA and the AI Act differs in at least three fundamental respects: the protected interests, the mechanisms through which risks materialize and propagate, and the scope of the regulated entities and services. Therefore, they cannot always lead to the same risk mitigation measures, even when GenAI applications pose systemic risks comparable to those that the DSA aims to tackle.

The main difference to note is that systemic risk management under the AI Act applies to providers and their models before the latter are integrated into other systems and applications downstream in the value chain. GPAIMs, especially large language models, are the backbones of several GenAI applications, but they are just a component of such applications. Other components, in particular the user interface, are also key to their functioning and ultimately shape the risks that they pose. Risk management at the model level does not encompass risks posed by these components and their interaction with GPAIMs. Taking the example of chatbots that can encourage self-harm practices, safeguards need to be implemented not only for the underlying large language model, but also in the user interface. At the model level, risk mitigation measures can include adversarial testing to evaluate how the model reacts to prompts about self-harm and suicide, conducting monitoring and incident reporting, and introducing or improving safety filters and restrictions on output. Nonetheless, safeguards are also needed in the user interface, such as providing warnings to users, interrupting the service, and/or contacting emergency services when risky situations materialize. Similarly, model-level constraints can greatly mitigate the risk of deepfakes being generated by the model, but restrictions will also need to be applied, both contractually and technologically, at the system level through output filtering.

Systemic risk management for GPAIMs can thus only address some of the risks posed by generative AI. This does not preclude the possibility that providers of GPAIMs may mitigate systemic risks by anticipating risks that can arise in downstream applications. The very definition of systemic risks in the AI Act requires considering negative effects ''that can be propagated at scale across the value chain''. This requirement can contribute to more effective risk mitigation but cannot replace risk management specifically targeting the downstream application. Therefore, the lack of system-level risk management leaves a gap in mitigating the risks posed by GenAI applications – a gap that the DSA could close.

Conclusion

The governance model of the AI Act for GPAIMs seems to build on the idea that they are core infrastructures of broader systems in which they are integrated. It thus regulates GPAIMs as core infrastructures to mitigate cascade risks in the downstream value chain, mainly through transparency and risk management requirements. In the travaux préparatoires for the AI Act, the Council referred to the risks posed by certain GPAIMs for democratic processes and public health through the dissemination of disinformation, echoing the concerns underpinning the DSA's systemic risk management regime.

While applying the DSA to GenAI applications is not free from legal uncertainty, it could nevertheless prove beneficial. It would make risk governance obligations binding in exchange for a more favorable liability regime; in this way, it avoids imposing strict liability on GenAI providers for content that is, by its nature, generated by a black box, while inviting scholars and policymakers to reconsider the well-established service provider vs. content provider dichotomy in light of the characteristics of GenAI applications. Derogating from a strict liability regime also appears to be a wise move in a Europe in search of digital competitiveness. Beyond entities that may be designated as VLOPs or VLOSEs, such an approach would be particularly beneficial for small and medium-sized operators that do not meet the threshold for designation as very large online platforms or search engines.

ABOUT THE AUTHORS

Marco Bassini
Marco Bassini is an assistant professor of Fundamental Rights and Artificial Intelligence at the Tilburg Institute for Law, Technology, and Society (TILT) at Tilburg University.Marco Bassini is an assistant professor of Fundamental Rights and Artificial Intelligence at the Tilburg Institute for Law, Technology, and Society (TILT) at Tilburg University.
Andrea Palumbo
Andrea Palumbo is doctoral researcher at the Centre for IT and IP Law (CiTiP), KU Leuven, where he conducts research on AI and platform regulation.

SUGGESTED CITATION

Bassini, Marco; Palumbo, Andrea: Beyond Intermediaries: Generative AI in Search of Their Legal Status, Verfassungsblog, 12.04.2026, doi: 10.59704/887f92f3bcf01a98.

Licensed under CC BY-SA 4.0.

Langzeitarchiviert bei IntRechtDok unter: doi.org/10.17176/20260413-173443-0.