



Strasbourg, 10 June 2026

CDDEM(2025)18

STEERING COMMITTEE ON DEMOCRACY (CDDEM)

Generative AI and Democracy

**Study on advantages and risks of the use of generative artificial intelligence in
public debate and democratic processes**

Prepared by

Prof. Dr. Dr. h.c. Natali Helberger & Prof. Dr. Claes de Vreese, Gionata Bouche, Agustin Ferrari
Braun, Dr. Max van Drunen, Dr. Simon Kruschinski, Dr. Nicolas Mattis, Dr Sophie Morosoli, Dr
Laurens Naudts, Dr. Charis Papaevangelou, Dr. Theresa Seipp, Dr. Fabio Votta, Dr. Teresa Weikmann
(University of Amsterdam, AI, Media & Democracy Lab and GESIS - Leibniz Institute for the Social
Sciences, Cologne (Dr. Simon Kruschinski))

Generative AI and Democracy

**Study on advantages and risks of the use of
generative artificial intelligence in public debate
and democratic processes**



Table of Contents

Executive summary	5
1. Introduction	7
2. GenAI & democratic institutions	10
2.1 Journalism	10
2.1.1 The Transformative Potential and Applications of GenAI	10
2.1.2. How generative search challenges journalism	13
2.1.3 Citizens' Perspectives on the Risks of the Use of GenAI in journalism	14
2.1.4 The Rise of AI-driven Disinformation	17
2.1.5 Structural and Democratic issues: Digital Dependency and Power Asymmetry ...	18
2.1.6 Conclusion	20
2.2 Generative AI in Political Campaigning and Democratic Practices	21
2.2.1 Overview on Gen AI Deployment in Political Campaigns across the World and Technological Dependencies	22
2.2.2 AI-Assisted Campaign Activities	23
2.2.3 AI-Assisted Communicative Strategies	25
2.2.4 AI-Assisted Personalisation of Political Messaging	27
2.2.5 Gen AI for Political Information Seeking and Civil Society Engagement	28
2.2.6 Citizens' Attitudes towards GenAI use in Political Campaigns	30
2.2.7 Conclusion	32
2.3 Infrastructure	34
2.3.1 The politics of infrastructure	34
2.3.2 Control over the infrastructure needed for GenAI	35
2.3.3 Infrastructural dependency and democratic backsliding	37
2.3.4 Conclusion	38
3. Implications for human rights, political participation and democratic principles	39
3.1 Adding the human rights dimension: GenAI from the perspective of the rights to equality, non-discrimination and political participation	40
3.1.1 Inequalities and structural implications	40
3.1.2 The potential of non-discrimination law to address these concerns	43
3.1.3 Conclusion	44
3.2 GenAI & freedom of expression	44
3.2.1 Adding the human rights dimension: GenAI from the perspective of the right to freedom of expression	45

3.2.2. The potential of freedom of expression doctrine to contribute to the democratic governance of AI	49
4. Council of Europe legal framework.....	51
4.1 Situating GenAI in Council of Europe Recommendations	51
4.1.1 CoE recommendations and the role of GenAI in democratic society	52
4.1.2 Conclusion	55
4.2 Framework Convention on Artificial Intelligence.....	56
4.2.1 Human rights and democracy as point of departure	56
4.2.2 Risks to human rights and democracy.....	57
4.2.3 Positive obligations to protect the realisation of human rights and democracy	57
4.2.4 Public interest innovation	58
4.2.5 Democratising AI Governance.....	59
4.2.6 Conclusion	60
5. Overall conclusions and recommendations	61

Executive summary

New technologies offer new opportunities for the realisation of fundamental rights. This is also true for Generative AI (GenAI). GenAI tools can enable new forms of creating and disseminating news content and political communication in ways that are more accessible, engaging and inclusive, opening new avenues for realising human rights such as freedom of expression, non-discrimination and equal opportunities to participate. The technology can expand channels for voter information and engagement, and offer new ways to make information more relevant, responsive and available across languages and formats, including for organisations with limited resources. Moreover, AI-powered translation and accessibility features create meaningful opportunities to reach marginalised communities, adapt content to diverse cognitive and physical needs, and strengthen multilingual outreach.

Seizing those opportunities and creating the condition for societal innovation will not only require managing the manifold risks for the realisation of fundamental rights but also creating the broader institutional environment for public interest innovation. Drawing on state-of-the-art research in communication science, digital journalism, and law, the report also identifies key concerns at the level of individual citizens, e.g. concerns around the respect for privacy, amplification of biases and the spread of misinformation. It points to the struggles of democratic institutions, including threats to the societal and economic viability of journalism in the face of competition with the virtual assistants and AI overviews of large technology companies, or the use of synthetic content to intimidate and manipulate voters. And it calls for attention for the overall shape and structure of the underlying informational infrastructures and power dynamics, which are often less visible to policy makers but that ultimately provide the material, digital and organisational conditions under which information is produced, disseminated and exchanged, and are the source of considerable economic and political power.

The key finding of this report is that it is not enough to make sure that the technology is safe and no risk to human rights and public values. GenAI's impact and potential for the realisation of fundamental rights depends on the societal, organisation and regulatory context in which it operates, the possibility for institutions to make value-driven choices and regulators that actively promote public innovation where the interests of society as a whole outweigh purely economic interests of a concentrated handful of technology companies. This report has made the argument that GenAI and its implications for human rights and democracy must be seen in the context of the configurations of infrastructural and political power, and who is able and entitled to decide the course and character of innovation.

With its focus on human rights, democracy and the rule of law, the Council of Europe is particularly well positioned to realise such a more societal perspective on AI innovation and provide an important counter-narrative to more profit-oriented narratives of technology innovation.

To this end, the report makes a number of recommendations to the Council of Europe:

1. Provide substantive guidance and standardisation

As a human rights organisation, the Council of Europe is particularly well prepared to develop more substantive guidance and identify, concretise and monitor risks to democracy, human rights and political equality, be that in the way GenAI is used in democratic practices like elections, the more structural implications for democratic institutions (such as the free press) as well as the democratic governance of AI. Next to a focus on risks, the Council of Europe can also usefully develop guidance how AI can advance democracy and human rights and develop a value-driven interpretation of “public interest innovation” to inform States, companies and standardisation processes

2. Incentivise and enable public interest innovation

Next to the risk-based approach to AI that is also enshrined in the Council of Europe’s Framework Convention, the Council of Europe could develop guidance on policies and governance approaches that enable and incentivise public interest innovation. It could help identify the democratic and human rights implications of procurement choices, and develop narratives to prioritize value-oriented innovation, open-source datasets, decentralisation and interoperability above the promotion of scale and innovation driven by a concentrated number of technology companies.

3. Strengthen democratic resilience and ability to govern

The Council of Europe, and the CDDEM specifically, could have an important role in exploring mechanisms and regulatory solutions to map and remedy existing dependencies and identify mission-critical elements of a European digital ecosystem, for example through monitoring the use of GenAI in elections, stimulating the development of public interest information infrastructures and strengthening the ability of states to develop democratic and inclusive ways of governing GenAI. Next to digital autonomy, empowering public accountability actors like civil society, equality bodies, academics or investigative reporters to participate in the governance of the digital transformation are critical conditions for strengthening democratic resilience. Finally, democratic resilience also means enabling more meaningful citizen engagement, including of otherwise often marginalised groups, for example through transparency, literacy and meaningful participation in decision making.

Human rights law and also the realisation of the Council of Europe’s Framework Convention requires to not only manage the negative implications and risks from AI, but also emphasise the importance of actively promoting a positive contribution of innovation to the realisation of human rights, democracy and the rule of law.

1. Introduction¹

With the ability to analyse huge amounts of data, produce high-quality text, images and audiovisual content, and pioneer new forms of interaction with users through Virtual Assistants or Generative Search, the potential of GenAI to impact core democratic institutions and practices is enormous. This is because of the central role that communication plays in our democracy. Democracy depends on being able to share and access information, express one's viewpoint and agree with or contest those of others, activate, engage, congregate and call for political action. So far, this power of speech was "what distinguishes us from other animals and, this far, from computers" (Garton Ash, 2016, p. 73). With GenAI, this role of language as boundary condition between computer and human is eroding rapidly. GenAI, as a "language synthesiser" can directly feed into all these activities of communicative engagement previously reserved to humans, and in so doing make or break a democratic discourse, an autonomous political choice or a trustworthy public sphere.

GenAI (GenAI) is "artificial intelligence (AI) that can create original content such as text, images, video, audio or software code in response to a user's prompt or request" (Stryker, Cole & Scapicchio, 2025). Examples are systems like ChatGPT, Gemini, Claude, or Grok. These systems are powered by Large Language Models (LLMs, such Open AI's GPT or Meta's LLaMA), also referred to as foundational models in the sense of an AI system with broad capabilities that can be adapted to a range of different purposes. Their multi-purpose nature distinguishes GenAI systems from other AI applications that are often designed for a specific purpose. GenAI systems also become increasingly multi-modal, in the sense that they can operate with, and create text, images, videos or audio. Part of the success and rapid adoption of GenAI is its ease of use, often presenting an interface that is similar to the ones we already know e.g. from search engines. Under the hood, however, GenAI works very differently from a search engine.

GenAI's output is in the first place the result of a statistical calculation that enables the system to predict the next word in a sentence. The statistical character of GenAI communication is important to understand as it means that the system does not know or reflect on what it is saying. Accordingly, GenAI has been described as an entity for "haphazardly stitching together sequences of linguistic forms [...] according to probabilistic information about how they combine, but without any reference to meaning" (Bender et al., 2021, p. 617). This means that GenAI can generate high quality text, images, audio or video that mimic speech produced by humans in a highly convincing manner, with a complete lack of understanding or any sense of truthfulness. Because of the ease and speed with which GenAI systems can produce content, it significantly lowers the thresholds for certain communicative practices, such as personalising content to adjust to individual users' characteristics or preferences. The technology can also enable new forms of engagement (e.g. in the form of adaptive Virtual Assistants), a higher level of autonomous action, as well as scaling up communicative practices. These characteristics distinguish GenAI-enabled from other communicative technologies, like conventional recommender systems or search engines. GenAI can hypercharge digital

¹ We would like to thank the the members of the CDDEM, the Committee's secretary Urska Umek and the participants of the Generative AI & Democracy closed discussion event organised on October 28th 2025, by the AI, Media & Democracy Lab for their valuable feedback on this study.

communicative practices, such as the spread of synthetic disinformation across existing social media platforms.

GenAI and its implications for democracy cannot be seen in isolation from the wider context in which the technology operates. Though policy discourses often focus on the technology itself, the performance of GenAI is the result of a complex web of human and societal factors and resources. These include the production of content by creators and citizens, content that is then used to train the models. GenAI draws on the work of data labellers to prepare the content for training or inform content moderation practices, the red-teamers and fact checkers, the expertise and instructions of those training the models. Then there are the natural resources needed to build chips and cool data centres, and that power the necessary computational infrastructures. GenAI systems and models are also shaped by internal managerial practices and business models, as well as the engagement of countless individuals with the output as new input for training. Last but not least, there is the role of the ethical codes, community guidelines and regulatory frameworks that steer the course of innovation and the political processes behind those governance frameworks. As such, GenAI is more than just yet another technology. It is part of a societal process of extraction, transformation and re-distribution of human and natural resources, and ultimately, power.

Like with social media, the development of GenAI is being led primarily by a handful of very large technology companies, many of whom have ties or are identical with the operators of the leading social media companies. The development of the underlying computational infrastructures is led by a concentration of highly influential hyper-scalers, such as Amazon, Microsoft, Google and OpenAI, with strong incentives to use GenAI to further consolidate their control over the global provision of services. Further, the revolution in GenAI comes at a time when technology companies increasingly also play a political role and shape geo-political powerplays. The possibilities of GenAI for communication aside, the resulting new forms of concentration of technological, economic and communicative power have their very own implications for democratic processes and the realisation of human rights.

This broader, socio-technical perspective will also inform the analysis in this report. Taking a socio-technical approach means that we situate GenAI as technological elements (e.g. in the form of a virtual agent or recommender system) within a broader value chain of actors (including data providers, cloud infrastructure providers, application providers, downstream users, etc.) and (un)democratic practices. Doing so will allow for a more nuanced discussion of possible interventions and intervention points, but also allows us to explain how dynamics within this value chain ultimately affect the realisation of human rights and democratic processes. In particular, we will touch upon the problem of said power imbalances, infrastructural capture, digital sovereignty and how AI can be a technology of freedom, but also of autocratic control.

We will do so using three case studies. The first relates to the use of GenAI in journalism. According to the latest Reuters Trends and Prediction report (Newman & Cherubini, 2025), 90 % of the editors, CEOs, and digital executives interviewed globally say that newsrooms are being fully or somewhat transformed by GenAI. Newsrooms around the world experiment with integrating GenAI along the entire news value chain, from research and data analysis, translation, transcription and fact-checking, to generating personalised recommendations, new forms of news aggregation and communicating with the audience through chatbots and

virtual assistants. Insofar, the use of GenAI in journalism is instructive as it demonstrates how technological innovation can affect the functioning but also trust in and legitimacy of a core democratic institution, and be used to realise and advance the constitutional right to freedom of expression. At the same time, the example of journalism is also relevant to demonstrate how the above-mentioned power dynamics can affect the sustainability and future conceptions of the role of journalism in the information ecosystem.

The second case study is the use of GenAI in electoral processes and political campaigns. This case study enables us to highlight the way in which GenAI can shape communicative actions between citizens and political parties (e.g. in the form of personalised political advertising), empower citizens to play their democratic role as voters, or (malicious) third parties to weaponise the technology for un-democratic purposes. It is here that the ability of GenAI to produce language as the raw material of democratic practices, and the resulting implications for democracy are most pronounced, as the underlying dynamics of power, inclusion and exclusion play out.

The third case study is the role of GenAI as part of the public digital information infrastructure that underlies all processes of digital communication. While the use of GenAI in communicative practices is often approached from the perspective of a very concrete action and relationship between citizens and democratic institutions (informing oneself or others, persuasion and engagement), the infrastructural implications are far less visible but impact democracies on a far deeper, more pervasive level. The current discussions around digital autonomy are case in point. As digital communication depends on, but is also shaped by the underlying technological infrastructures, the question of who develops and controls those infrastructures, to which end, and with an eye to which values has the potential to fundamentally affect and alter the conditions of public communication and exercise of human rights, up to a point where the independence of democratic institutions such as the press is at stake, and the concentration of economic and political power raises to a level that is no longer compatible with the core principles of democracy.

The overall goal of this report is to signal advantages and risks that the use of GenAI (GenAI) can bring for democracy, with a particular focus on the media and electoral processes, and develop recommendations for priorities for the Council of Europe's Steering Committee on Democracy (CDDDEM).

The report will be grounded in the Council of Europe Framework Convention on Artificial Intelligence, Human Rights, Democracy and the Rule of Law (Framework Convention) and the recommendations and other instruments of the Council of Europe, but also look into interfaces with other regional/international attempts at GenAI governance, with a special focus on the European Union. The report will be research-driven, in the sense that it will use and translate the insights of state-of-the-art research in, *inter alia*, communication science, digital journalism and law, into suggestions for strategy and prioritisation. Seeing the limited scope of the report, the goal is not to give a comprehensive overview of the state of the art, but signal trends and key developments. The goal of the report is to provide a first overview of important topics to consider.

2. GenAI & democratic institutions

Following an outline of the general societal, technological and economic dynamics in the Introduction, this section will zoom in on the way GenAI is enabling, shaping, changing or threatening democratic institutions and processes. To do so, this section will provide a synthesis of existing research on the implications of GenAI for (1) journalism, (2) elections, and (3) digital infrastructures.

2.1 Journalism

2.1.1 The Transformative Potential and Applications of GenAI

Journalism is a cornerstone of liberal democracies, as it is crucial for facilitating an informed public, fostering public debate, and holding those in power to account. Because of its democratic importance, journalism also serves as an excellent case study for assessing the transformative potential of (generative) AI and its democratic and societal implications – especially so since audiences, across different countries, are rather critical about the impact of GenAI on journalism (see Figure 1).




Figure 1. Feelings towards Gen AI adoption in Journalism in the US, the Netherlands and Denmark. From Morosoli et al. (2025).

A look at current newsroom practices shows that GenAI is already transforming journalistic workflows. Although the use varies between news organisations, GenAI increasingly permeates the entire news value chain, from news gathering and content production to potentially personalised distribution (Beckett & Yaseen, 2023; Diakopoulos et al., 2024a; Mattis & de Vreese, 2025). Concrete use cases include but are not limited to automated trend analyses during news gathering, summarisation, transcription, headline optimisation and even the generation of entire articles during news production, or search engine optimisation and personalisation during news distribution (Cools & Diakopoulos, 2024). Indeed, in a 2024 survey among news workers by Diakopoulos et al. (2024), 70% of the respondents indicated having already used GenAI tools for textual content production, and close to 80% said they'd aspire to do so in the future (Diakopoulos et al., 2024b).

When asked about their motivations for using GenAI, news professionals often cite its ability to increase workflow efficiencies. This, they argue, might help free up journalists for more meaningful tasks, as GenAI automates the more menial aspects of journalistic work and augments creative tasks such as story discovery and idea generation (Cools & Diakopoulos, 2024; Diakopoulos et al., 2024b). In addition, GenAI also facilitates the creation of "fluid or modular journalism," where the content and presentation adapt to individual user preferences (Caswell, 2024; Cools & Diakopoulos, 2024; Jones & Jones, 2019). Among other things, this could help make news more relevant, diverse, and accessible (Lin & Lewis, 2022), for instance, by creating understandable summaries or by generating audio versions of written articles. At times, synthetic content can also be a means for citizens and journalists to report issues that would otherwise be censored. For example, in Venezuela, journalists have turned to GenAI to create virtual journalistic anchors as a means of protecting their identities from prosecution (Phillips & Torres, 2024). Another potentially worthwhile use case that news organisations are already exploring is conversational agents that can interact with readers and help them find relevant information.

In other words, there is reason to hope that GenAI will help improve various aspects and accessibility of journalism. However, it is important to note that this feeling is not necessarily shared by audiences. Indeed, survey results from a six country study (US, Japan, South Africa, Brazil, Denmark and the Netherlands) from April 2025 indicate considerable ambivalence and variation in the extent to which US, Danish and Dutch citizens see benefits in hyper-personalised news (see Figure 2, Panel B). The most prevalent risk across countries was the risk of hyper-personalisation negatively impacting data privacy and the risk of it facilitating the spread of misinformation. Participants saw the most benefits in hyper-personalisation, improving the user experience, also by providing more personally relevant news. Moreover, the same survey also points to a somewhat limited appetite for AI (co-)created news in both countries (see Figure 2, Panel A and section 2.1 b). Overall, participants were the least willing to share this type of content, with Denmark being the most sceptical. Across countries, individuals would also take AI generated news with a grain of salt, where Dutch citizens were the least willing to trust it. When it comes to consuming AI driven journalism, we find the most acceptance. People in Denmark and the U.S. are willing to read this type of content.




Figure 2. Engagement with AI-(co-)created news (A) and perceived benefits of hyperpersonalisation (B) in the US, the Netherlands and Denmark. From (Morosoli et al., 2025)

Somewhat paradoxically, however, we also see that increasing numbers of people use GenAI for information seeking and frequently encounter AI features such as Google’s AI Overview in their daily life (Simon et al., 2025). Here, GenAI companies are introducing a new layer of gatekeeping for news accessibility through chatbot interfaces, answer engines, and smart assistants integrated into mobile operating systems. This shift enables more fluid and participatory forms of news consumption (Caswell, 2024) but is already having economic and visibility impacts on publishers, particularly in the Anglo-Saxon media markets (Maher, 2025). This has important implications for the news media’s ability to exert the gatekeeping and agenda setting power that has traditionally been ascribed to it (see also section 2.3).

Moreover, greater public reliance on LLMs as access points to news also poses threats to the economic viability of journalism. That is, while these LLMs are often trained on journalistic content, they frequently fail to mention news sources adequately or factually misrepresent their content. For instance, comparing eight popular AI search engines (incl. ChatGPT,

DeepSeek, Gemini, Perplexity, and others), Jaźwińska and Chandrasekar (2025) find that all of them frequently fail to retrieve the correct article, publisher, link, or a combination thereof. Importantly, not even content licensing deals provided news publishers with the security that their content would be adequately cited (Jaźwińska & Chandrasekar, 2025). A recent study among 22 international public broadcasters confirms this, finding that on average, 45% of AI responses have at least one significant issue such as hallucinations or problems with sourcing (EBU, 2025). When it comes to accuracy, this research finds that roughly half of all tested queries had some and around 20% of queries had significant accuracy issues. Given that LLMs are increasingly used as direct access points for news (Simon et al., 2025), these inaccuracies threaten to further compromise the economic sustainability of journalism by misrepresenting the origin or nature of journalistic claims.

2.1.2. How generative search challenges journalism

GenAI companies are introducing a new layer of gatekeeping for news accessibility through chatbot interfaces, answer engines, and smart assistants integrated into mobile operating systems. The share of traffic flowing through these channels is steadily increasing, as users appear to prefer receiving synthetic, ready-made responses over navigating traditional hyperlinks (Reiss et al., 2025; Simon et al., 2025). This shift is already having economic and visibility impacts on publishers, particularly in the Anglo-Saxon media markets (Maher, 2025).

The long-standing feud between the publishing industry and major American technology companies has brought public attention to the existential threat facing the financial viability of digital journalism. This culminated in legislative interventions across several jurisdictions, mandating fair remuneration for the dissemination of news content on digital platforms. GenAI brings another structural reconfiguration of the online information environment extending beyond the influence once exerted by traditional search engines and social media platforms (see also section 2.3.2). The implications are both experiential, for how citizens encounter and consume news, and political, in terms of shifting power relations across the media ecosystem. The notion of “news as conversation” signals a move away from artisanal, item-based production toward more fluid and participatory forms of news consumption (Caswell, 2024). The one-stop, low-effort, and interactive mode of information delivery offered by generative systems is increasingly appealing in an age of digital information overload. Unlike earlier intermediaries, AI companies are now directly involved in the production of content, reaching deeper into the core functions and values of the journalistic institution (Guzman & Lewis, 2024; Lewis et al., 2025). The latest step in this process bypasses news organisations entirely: platforms such as OpenAI’s Sora and Meta’s Vibes offer feeds of AI generated content, potentially signalling a future in which the space for journalistically produced content on platforms shrinks further and further.

The news industry’s relationship with these AI companies is ambiguous. Traditional news publishers, and in particular large players, still wield considerable influence over the public agenda, a position which online intermediaries are contesting (Simon, 2024a). Yet, the public narrative surrounding these relationships is shifting from one of initial confrontation to strategic collaboration. Several publishers now partner with AI companies to train models or enhance contextual output through retrieval-augmented generation and online search (Tobitt, 2025). In return, AI companies offer remuneration and content visibility. These collaborations could be seen as institutionalising a new form of reciprocity between journalism and AI. While

generative systems shape journalistic practice, business models and consumption patterns, journalism informs AI development through the provision of data and ethical grounding for its communicative use (Guzman & Lewis, 2024, p. 349).

However, these trends carry latent risks. By helping improve technologies that can replicate and at some point replace an expanding number of journalistic functions, publishers may be undermining their own autonomy and future sustainability (Simon, 2024a). Systematic dependencies are likely to arise given GenAI companies' superior financial power and control over infrastructure, spanning from hardware to user interfaces (Sjøvaag et al., 2024). Moreover, as generative models represent "epistemic mutations in computational ways of knowing" they reshape not only the circulation of knowledge but also the political models underpinning it (Amoore et al., 2024). Attention must therefore be paid to the normative outcomes of these mutations, especially concerning visibility, representation, and user agency. While some readers can influence AI outputs through prompt engineering, others, particularly less technically skilled or politically aware users, may remain confined within these tools' default informational frames (Gillespie, 2024, p. 3). The implications for media pluralism are at this stage theoretical, yet plausible in the long-term: smaller outlets and independent journalists are unlikely to secure equitable treatment or remuneration, deepening inequalities in exposure and influence and further distorting media market dynamics (see also section 3.2.1.).

2.1.3 Citizens' Perspectives on the Risks of the Use of GenAI in journalism


Alongside the potential benefits of GenAI (Cools & de Vreese, 2025), its rapid integration in journalistic processes presents significant challenges that threaten journalistic integrity and citizens' trust in information (Mattis & de Vreese, 2025). GenAI is not merely a new tool that augments journalistic work. Depending on how it is used, GenAI "intervenes in the core creative processes of journalism, challenging traditional norms of authorship, originality, and professional identity" (Lewis et al., 2025, p. 1). This could have important implications for shared journalistic norms and values down the line and may, at least in some scenarios, undermine journalism's integrity, public perception, and its ability to perform its key democratic functions, such as combatting the spread of disinformation with accurate and high-quality reporting (Mattis & de Vreese, 2025).

And indeed, citizens are expressing diverse concerns and risk perceptions about GenAI and its impact on their lives. Data from a representative sample of Dutch citizens (N = 1545) shows that the three biggest risks connected to GenAI are (1) the fear that GenAI can be used to make people think that they interact with a human; (2) the perception that it is dangerous if everyone, including non-experts, has access to GenAI tools; and (3) that GenAI poses a risk to privacy because of the data it's trained on (see Figure 3). More specifically, when asked about the risks associated with the use of GenAI in journalism, evidence from focus groups with Dutch citizens indicates that citizens expect news organisations to be transparent about their use of AI at all times and desire a certain degree of accountability or responsibility from these organisations (Morosoli et al., 2025). These findings are not unique to the Netherlands, but echo findings in, for instance, various African (Gondwe, 2025) and European countries (Kieslich et al., 2024a; Mitova et al., 2024a). For example, Gondwe (2025) finds that citizens from Ghana and Nigeria see transparency about the use of AI in news creation as a crucial element for maintaining trust in the media. In Europe, Mitova and colleagues (2024b) conclude

that across four countries (the Netherlands, Poland, Switzerland, United Kingdom), the desire for transparency in the context of news recommender systems is relatively high, the more citizens were concerned about their privacy. With the help of scenario writing, Kieslich et al. (2024b) evaluated the effectiveness of transparency obligations in the news environment among different stakeholders in the Netherlands. Overall, the majority of respondents across stakeholder groups (content creators, technology developers, and news consumers) welcomed transparency obligations and expected some level of effectiveness, particularly linked to preventing the spread of misinformation and enhancing journalistic quality. However, some respondents expressed doubts about the practical enforcement and overall impact of transparency obligations, noting that they might not significantly alter the behaviour of news consumers or content creators.

Considering the above, we observe a mismatch between what citizens want and what they seem to be informed about. Citizens have strong desires when it comes to transparency and are opinionated in their risk perceptions directly involving AI. However, they do not seem to comprehend the deeper structural dependencies, which will lead to institutional shifts. These remain largely invisible or are treated as secondary. Without a more comprehensive awareness of these interdependencies, it is challenging for policymakers to create as meaningful perceived policies.


Figure 3. Risk Perceptions toward generative AI in %



Given the general concerns about the involvement of GenAI in journalism, we asked in the same survey with Dutch citizens (N = 1,545) as mentioned above, to what extent they trust the use of AI in the news production. We were particularly interested in how much individuals trust AI-generated news, news written by human journalists, and journalists who use AI in their work. We also asked them how credible they think the respective news outputs are. Results reveal a clear preference for human-authored content over AI-generated content across trust and credibility dimensions (see Figure 4). The means of trust in AI-generated news and trust

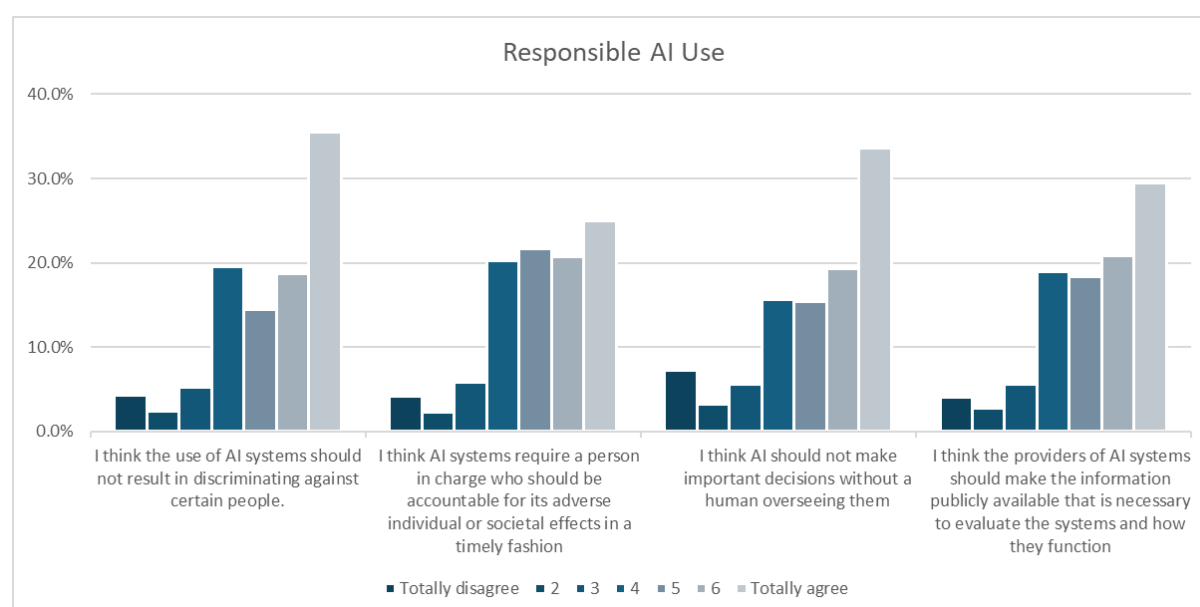
in journalists using AI are relatively low, while the reported trust in human journalists is higher. Credibility perceptions follow a similar pattern: AI-written articles are rated the least credible, followed by human–AI collaborative articles, and human-written articles are rated highest. The perceived impact of AI on journalism quality is slightly above neutral. These findings are not unique to the Dutch case. In a recent report on GenAI and the news, the Reuters Institute found across six countries (Argentina, Denmark, France, Japan, the UK, USA), citizens are the least comfortable when news is generated completely by AI and the most comfortable when news is produced by human journalists (Simon et al., 2025).

Figure 4 Reported levels of trust in %



To (re)build trust in the form of a meaningful relationship between citizens and news outlets, ethical and responsible use of GenAI is essential. But what do citizens consider ethical AI use? Our findings in a Dutch context highlight that a large majority agrees that AI should not discriminate against certain people, with 35.5% totally agreeing ($M = 5.35$), Figure 5. Similarly, there is substantial agreement that AI systems require a person in charge who should be accountable for adverse effects, and that AI should not make important decisions without human oversight is also strongly supported, with 33.6% totally agreeing, pointing toward a strong need for accountability. Finally, there is a clear consensus that AI providers should make necessary information, such as training data, publicly available to evaluate the systems, with 29.5% totally agreeing. These citizen-centred findings are consistent with the perceptions of media professionals. Diakopoulos et al. (2024a) surveyed 292 individuals working in the media industry and similarly found a lack of human oversight, false information, and bias as the most pressing concerns among media professionals. Overall, these results underscore the importance of fairness, transparency, and accountability in the deployment of AI systems across populations of citizens and professionals.

Figure 5. Ethical Concerns and Responsible Use of AI



Taking all the empirical evidence together, we argue for two things to mitigate the erosion of trust and combat the increasing fear of AI-generated mis- and disinformation: Meaningful transparency and AI literacy for citizens and journalists alike. Meaningful transparency in AI-driven journalism should go beyond a simple AI label (Piasecki et al., 2024). Transparency should explain why AI was used, what the system did, what data and editorial judgments shaped the output, and what uncertainties or limitations remain. Next to useful transparency, AI literacy efforts for both citizens and journalists should be developed and enforced (Deuze & Beckett, 2022; Stamboliev, 2023). These efforts should emphasise empowering citizens and journalists to question, resist, and influence how AI is designed and deployed. Thus, addressing access and power relations, not just teaching technical explainability (Stamboliev, 2023).

2.1.4 The Rise of AI-driven Disinformation

A central concern is that GenAI will create and expose citizens to false and misleading information, which they cannot recognise as such, leading to a growing distrust in information and journalistic processes (Morosoli et al., 2025). Specifically, the rise of GenAI has intensified concerns about its misuse for disinformation – false information created with the intention to deceive – as well as misinformation, meaning the accidental spread of falsehoods. For example, GenAI tools, such as Sora, can be used to produce images depicting events that never happened, whether in the form of deepfakes or synthetic photographs (Weikmann & Lecheler, 2023) (see also section 2.2 on the use of GenAI in political campaigning). Here, we find that the experience of being deceived by such forms of visual disinformation can lead to a general mistrust towards visual media, and therefore anything citizens may see in the media (Weikmann et al., 2025). For journalistic practices, two key concerns arise: (1) the increased pressure to identify AI-generated disinformation in order to avoid becoming accidental disseminators, and (2) the risk of becoming inadvertent producers of misinformation due to reliance on AI for research or content creation.

First, because of its close resemblance to original information, synthetic media contributes to an “authenticity crisis,” in which both audiences and journalists struggle to distinguish between human-created and AI-generated content (Dodds et al., 2025b). This raises the question of whether journalists trust their own ability to navigate this challenge. A survey we conducted amongst disinformation experts of the European Digital Media Observatory (EDMO) suggests that journalists generally feel competent in handling disinformation in the age of GenAI. However (see Figure 6), they report significantly lower confidence than fact-checkers, and roughly the same level of competence as academics studying the issue (Weikmann et al., 2026).




Figure 6. Bar plot representing self-perceived competence in dealing with AI-generated disinformation across expert groups, including group averages on a scale from 1 to 7 and significant differences (* = $p < .005$).

Second, as journalists become more reliant on GenAI in the news production process, AI systems may introduce errors, biases, and “hallucinations” into journalistic content. For instance, a recent report by the European Broadcasting Union (EBU, 2025) shows that AI assistants consistently misrepresent news content when used for information seeking. This can become an issue not only if citizens use AI assistants to inform themselves, but also if journalists use them for research purposes. On the one hand, potential shortcomings of LLMs may provide users directly with inaccurate or biased content. On the other hand, the quality of journalistic content may be affected if it is based on erroneous AI-reliant content. In line with this, our research shows that large-scale use of GenAI for content creation increases the risk of spreading misinformation and so-called “AI slop” – low-quality, mass-produced content – which can drown out credible journalism (Dodds et al., 2025a). This, if left unchecked, presents a direct threat to credibility.

2.1.5 Structural and Democratic issues: Digital Dependency and Power Asymmetry

The preceding sections have demonstrated that with the growing deployment of GenAI, there is yet another digital transformation of journalism unfolding. It also showed that as of now there is some discrepancy of what is possible and what is, from the perspective of the audience, desirable. While professional ethics, editorial oversight and transparent and honest engagement with the audience can go some way to address these concerns, there are also

clear limits to what news organisations can do to protect and promote a healthy information ecosystem. This is because their limited influence over the way the leading GenAI tools are developed, as well as deeper structural dependencies and power asymmetries.

Structural dependencies as such reinforce power imbalances where news organisations have limited capacity to influence or even understand the technologies they rely on (Dodds et al., 2025b). AI systems offer limited transparency into their inner workings, biases or decision-making processes (Piasecki et al., 2024; Simon, 2024b, p. 4). Therefore, in order to be able to address issues of trust and transparency, news organisations rely on cooperation from technology providers. At the same time, GenAI further intensifies and reinforces concentrations of technological, economic, and communicative power (see also section 2.1.2). News organisations are growing increasingly dependent on a handful of, sometimes unaccountable, Big Tech companies who not only influence nearly every stage of the news value chain, but also drive technological innovation, control key computational infrastructures, and shape the legal and ethical frameworks that govern the industry (OSCE, 2025). These companies have amassed extensive powers, with news media relying on them for nearly every stage of news production, from access and use of relevant tools to the underlying infrastructure. This reliance not only exacerbates existing dependencies, and creates new ones in production, it poses significant risks to human rights and values, particularly media freedom and pluralism (Simon, 2024b; van Drunen, 2025).

These dynamics may risk "infrastructural capture," where the news media becomes incapable of operating sustainably without the digital resources provided by external technology companies (Nechushtai, 2018). GenAI further contributes to that, as dependency is no longer confined to distribution channels like social media platforms, but a fundamental reliance that is structural, technological, and infrastructural in nature emerges (Braun, 2025). For instance, news organisations heavily rely on cloud services from Big Tech companies like Amazon Web Services (AWS), to host their data and run their operations, effectively outsourcing their core infrastructure (Braun, 2025). This risks lock-ins, where high switching costs and technical integration make it difficult for news organisations to change providers, even if terms become unfavourable or misaligned with journalistic values. Accordingly, such infrastructural entrenchment gives technology companies immense power, as they control the underlying architecture on which journalism is increasingly built. Although the report addresses infrastructural dependency in more depth in a later Section 2.3, it is important to note here that such dependency may undermine media independence and jeopardise media freedom and pluralism (van Drunen, 2025).

These threats and power concentrations reflect deep structural power imbalances between media and Big Tech. High barriers to in-house tech development give the latter a significant advantage, while news media remain limited in data, computing power, capital, and technical expertise (Gerbrandt & Phoa, 2022). In this context, the size and available resources of a news organisation play a crucial role in determining the degree of its dependency, as well as its ability to remain independent, sustainable, and resilient. While for smaller newsrooms, reliance on tools by Big Tech is often a necessity imposed by resource constraints, larger news organisations often adopt Big Tech tools as a strategic preference, bypassing alternative, open-source options. In fact, some large news organisations appear to increasingly adopt "platform-like" features and strategies, prioritising resource and media consolidation instead

of supporting a more democratised approach, where small media and tech start-ups stand a chance (Meese & Seipp, 2025).

Accordingly, (Gen)AI adoption is widening the divide between large, well-resourced news organisations and their smaller counterparts (Eder & Sjøvaag, 2025). Thus, there is a tangible risk that AI will empower already dominant media and tech companies to become stronger, further marginalising small outlets and harming media pluralism (Sjøvaag et al., 2024). These developments trigger a trend towards consolidation and the favouring of news mergers to get necessary scale, as observed in recent media mergers. Such mergers are often supported as the mere solution for smaller media to not disappear and accelerate the formation of "news deserts" (Blagojev et al., 2023; Borchardt, 2022). While this may hold some truth, the long-term consequences of increasing media consolidation and mergers may ultimately weaken journalism as a whole, particularly local and investigative reporting, which serve as essential democratic watchdogs.

2.1.6 Conclusion

It is clear that GenAI is impacting journalism, a key pillar of democracies, both in terms of production, organisation, content, and the relationship with audiences. Indeed, because of its democratic importance, journalism is such a key case study for assessing the potential of GenAI and its democratic and societal implications.

At a general level, it is clear that at the time of writing most citizens view AI in journalism with more scepticism than enthusiasm. It is also clear that citizens have ambivalent views with a mismatch between what citizens want and what they know. Citizens, for example, ask for transparency and they are opinionated in their risk perceptions. However, they do not seem to grasp the deeper structural dependencies. These remain largely invisible or are treated as secondary.

The potential opportunities of the use of GenAI in journalism are offset by the relative distrust and concerns of the audience. In part, there is a role for journalism to innovate *above with* the audience. Similarly, there is a role for regulators to make sure tech is safe but at the same time, a certain level of distrust and critical awareness is also critical, taking into account the broader trust crisis and difficulty of distinguishing trustworthy from non-trustworthy content.

At a macro-level, we might observe effects of GenAI that diminish trust in the information ecosystem in general. This is something to consider in combination with the impact of GenAI on other democratic areas, such as politics (see section 2). GenAI is also shaping new tech-driven business models that compete for the audience, and threaten economic viability of journalism. This dynamic has reminiscence of the dependency of journalism on social media and as such GenAI is therefore a technology that also affects the journalistic market and business model.

Ultimately, the changing production, content, and audience relations, as well as the new challenges in terms of business models, viability, and dependency on specific companies, leads to a situation that requires political choices and steering innovation in directions that promote freedom of expression and benefit the audience. However, as long as macro-level effects are invisible, this is unlikely to get the political attention that it merits. It is imperative to address these challenges in a timely and forceful manner to create results.

2.2 Generative AI in Political Campaigning and Democratic Practices

GenAI, through a growing set of applications, is transforming political campaigning worldwide, as shown in recent elections from India to Indonesia, Europe, Mexico and the US (see Figure 7). The capabilities and rapid diffusion of GenAI have triggered public, media, and academic discussions about its risks and opportunities for political actors, voters, and democracies. During the 2024 super election year, GenAI was perhaps not a game changing technology that single-handedly swayed elections and undermined democratic processes (Kapoor & Narayanan, 2025; Simon & Altay, 2025) as predicted by media outlets such as the Washington Post (Verma & Zakrzewski, 2024) (see Figure 7). Yet the broad availability of these tools meant they became embedded in everyday campaign practice, primarily amplifying existing dynamics and introducing a new layer of operational and regulatory challenges.




Figure 7. Examples of political AI use in Indonesia, Türkiye, India, and the United States. (for an overview of AI use during the 2024 election, see (Elliott, 2025).

Understanding the implications of GenAI in electoral campaigns requires examining both how political actors deploy these technologies and how voters experience and respond to AI-generated campaign content. GenAI's transformative potential extends across multiple dimensions: it substantially lowers barriers to content creation, enables more sophisticated strategic communication including visual manipulation, facilitates large-scale personalisation of political messaging, and expands channels for voter information and engagement. These capabilities carry both opportunities and risks for electoral integrity, political participation, and democratic processes.

From the campaign perspective, this section first examines the deployment of AI in political campaigns (section 2.2.1) and its transformative potential regarding operational activities (section 2.2.2), communicative strategies (2.2.3), personalisation (2.2.4), and the use of AI(-chatbots) for citizen engagement and voting advice (2.2.5). From the citizen perspective, we then analyse how citizens evaluate the use of AI in political campaigns (2.2.6) and discuss our results considering regulatory and societal challenges.

2.2.1 Overview on Gen AI Deployment in Political Campaigns across the World and Technological Dependencies

The deployment of GenAI in political campaigns varies significantly across national contexts, shaped by an interplay of factors operating at three distinct but interconnected levels (Dommett et al., 2024; Kruschinski & Votta, forthcoming): at the systemic level, a nation's electoral arrangements, party structures, and media landscape create the strategic environment within which campaigns operate and adopt AI technologies. Regulatory frameworks, including data protection statutes, campaign finance laws, and AI-specific legislation, establish varying legal parameters that constrain or enable technological deployment across different jurisdictions. At the organisational level, individual campaigns' decisions about AI implementation reflect their available financial and human resources, technological sophistication, ethical commitments, and ideological orientations. These multilevel factors interact dynamically to produce distinct patterns of AI adoption that differ between countries and even among campaigns within the same national context.

Global patterns of GenAI use in recent elections show how context shapes adoption (see (Elliott, 2025): in the United States, a permissive regulatory environment and large budgets have enabled early, wide-ranging deployment, from message, image, and video generation to chatbots, automated calling, and data-driven voter outreach. India has used GenAI to manage linguistic diversity by creating demographically tailored materials and live translations of speeches. In Pakistan, Imran Khan's campaign relied on AI voice-cloning to let him address supporters from detention. In the EU's more regulated, multi-party systems, uptake has been more cautious and experimental, though actors at the ideological fringes have used AI for satirical or negative attacks on opponents. At the same time, foreign influence operations from Russia and China have used GenAI to produce targeted social content, spoof news sites, and flood discussions with synthetic commentary (OpenAI, 2024).

The multilevel factors on the system, regulatory, and organisational level shaping GenAI adoption create dependencies on AI technology that vary inversely. Resource-constrained parties and campaigns, especially in systems with stringent spending limits, treat accessible "off-the-shelf" AI tools as a cost-efficient substitute for human labour. Well-funded campaigns with established technical infrastructure can instead adopt GenAI more selectively and integrate it into existing workflows. These resource effects are constrained by regulation, since stricter data-protection and AI-governance frameworks (for example in the EU) limit how parties can leverage AI tools regardless of budget, while more permissive environments expand the adoption space for both well-resourced and resource-poor actors. Political campaigns are at the same time moving closer to the kind of infrastructural dependency that has been observed in other domains (see Section 2.3. for broader infrastructure discussion, and 2.1.5 for the structural version of this problem in journalism). Parties once treated digital technologies as cyclical infrastructure that was activated mainly during election periods. The

rise of permanent campaigning and professionalisation of campaign staff increasingly blurs this distinction (Dommett et al., 2021). Once AI is adopted to compensate for limited resources, its productivity gains can make traditional methods appear economically unviable. This creates a dynamic of technological escalation. If some actors demonstrate that GenAI improves content production and voter engagement, other actors face competitive pressure to adopt similar tools or accept a strategic disadvantage. Under such pressure, efficiency and visibility needs, including the production of synthetic content that is optimised for platform algorithms, can start to override any ethical considerations. Dependency then forms less from party preference than from platform affordances, costs, and rules that shape what is technically viable and made visible.

The next subsection shows how these infrastructural and resource-dependent patterns translate into concrete GenAI-supported campaign tasks, communicative strategies, and voter engagement applications.

2.2.2 AI-Assisted Campaign Activities

GenAI technologies can assist campaigns with operational activities (Kruschinski et al., 2025). Language models like ChatGPT, Claude, and Gemini can handle repetitive work including research tasks, translation services, content summarisation, and text editing. They also facilitate message development for voter outreach across conventional media channels and digital platforms. Thus, GenAI can enhance campaign productivity, particularly advantaging organisations and candidates operating with limited budgets and can facilitate the entrance of new political movements into the public sphere (Dommett et al., 2024; Foos, 2024). Content development and multimedia production represent substantial financial and temporal investments in campaign operations, typically requiring contracted specialists offering expensive services. Given the resource constraints most campaigns face, AI tools enable campaigns to deploy their human resources, finances, and time more effectively. This productivity improvement may help level the playing field in electoral competition by allowing less-resourced actors to enhance their campaign sophistication and messaging quality (Kruschinski et al., 2025).

Empirical analysis in our CampAlgn Tracker project examined nearly 70,000 posts and advertisements from Facebook and Instagram published by over 3000 political actors including all major and minor parties and respective candidates in Germany during the four weeks preceding election day (Kruschinski & Votta, 2025) (campaigntracker.de). In the Netherlands, we tracked over 3000 accounts belonging to Dutch political parties and candidates as well as third-party accounts (e.g., commentators, influencers, fan pages, meme pages) who posted 44,000 pieces of content on Facebook, Instagram, TikTok, and X, as of time of writing (Votta & Kruschinski, 2025) (campaigntracker.nl). Using a semi-automated detection methodology combining commercial AI detection systems with trained human validation, we find that 1.4% includes visual AI generated content in the German election, with similar rates in the Netherlands (see Figure 8). However, this modest adoption rate masks significant variations across different types of political actors and the reach that these AI-generated visuals achieved.




Figure 8. Share of visual GenAI posts in the 2025 German election campaign. Labels inside the bars show the percentage of GenAI posts and, in parentheses, the number of GenAI posts over the total number of posts for that group. Minor parties are only displayed in the aggregate (Kruschinski & Votta, forthcoming).

Two distinct patterns of adoption emerged across both Germany and the Netherlands (see Figure 8 for Germany). Smaller parties with limited resources demonstrated higher rates of AI usage, presumably leveraging the technology to overcome resource constraints and produce content more efficiently (Kruschinski & Votta, forthcoming; Votta & Kruschinski, 2025). Additionally, parties with fewer normative constraints regarding AI deployment, frequently positioned on the far-right of the political spectrum (like the AfD in Germany, or the PVV in the Netherlands), showed elevated adoption rates. For these actors, AI-generated content sometimes served to provoke and shock audiences, testing boundaries of acceptable political discourse. In contrast, well-resourced, established parties operate under formalised decision-making procedures and risk-averse organisational cultures, typically integrating new technologies within existing routines and facing greater reputational risks from perceived misuse of controversial technologies. In contrast, less institutionalised parties possess organisational flexibility and face fewer traditional constraints.

Further, there is a substantial risk of information saturation and pollution when AI is used to generate political content at scale, as even individual bad faith actors can now produce volumes of content that would previously have been unrealistic with limited personnel (Simon & Altay, 2025). The critical issue herein lies not in the quality or factualness of AI-generated texts, images, or videos, but rather in their ease of production. This capacity to produce content at very low marginal cost has been termed "AI slop" and increasingly hampers citizens' ability to separate meaningful information from noise, a problem compounded by the fact that many platforms and campaigns fail to disclose AI-generated content as such (Kruschinski & Votta, forthcoming). Only 13 percent of AI-generated content carried any form of disclosure label during the 2025 German federal election. When labelling occurred, it was predominantly voluntary disclosure by parties themselves rather than platform-initiated detection and flagging. Similar but also divergent trends are observed in the Dutch election campaign: half or more of the AI content remains unlabelled but now the majority of Dutch AI content, if it is

labelled, is labelled by the platforms as opposed to party actors. This marks a significant change in platform behaviour between the German Federal Election in February 2025 compared to the Dutch Parliamentary Election in October 2025. Consequently, substantive political communication may be drowned out, the quality of public discourse may deteriorate, and voters may struggle to navigate their political decision-making processes (Jungherr, 2023). Indeed, in the recent Dutch election campaign two members of the Dutch Parliament systematically deployed large volumes of AI-generated material, including discriminatory imagery targeting minorities and women through an anonymous social media account, thereby achieving engagement levels that exceeded those of established national politicians working with full campaign teams (Hofman & Veerbeek, 2024; Votta & Kruschinski, 2025). This represents the flipside of the potential equalising benefits: tools that lower barriers for smaller campaigns also lower them for bad-faith actors seeking to harm the information environment.

2.2.3 AI-Assisted Communicative Strategies

Second, AI facilitates more possibilities in communicative strategies (Jungherr, 2023). Visual AI systems such as Midjourney, DALL-E, Nano Banana, and Sora enable campaigns to generate imagery and video content for various applications from campaign posters and television advertisements to websites and social media feeds. These tools offer flexibility in visual approach, ranging from documentary-style realism to artistic interpretation, and from neutral presentation to emotionally resonant messaging (Sanseviero et al., 2025). This versatility allows campaigns to frame issues and individuals strategically according to their communication objectives, such as persuasion, negative campaigning, get-out-the vote messages or donation solicitation (see Figure 9).



Figure 9. Examples of different types of political AI use in the 2025 German election campaign.

GenAI also enables the malicious use of photorealistic AI-generated images and video that can erode both factual accuracy and political trust (Simon & Altay, 2025; Weikmann et al., 2025; see also section 2.1.4). Because deepfake techniques now reproduce faces, voices, and gestures with high fidelity, they create an “authenticity crisis” similar to that previously described for journalism (section 2.1.4): it becomes harder in electoral settings to determine whether a seemingly real clip is in fact synthetic. This uncertainty can be exploited by bad-faith actors to manufacture controversies or to attribute inflammatory statements to politicians, with downstream risks including violations of personality and image rights, the intentional circulation of false political information, and reputational harm (Laude & Daum, 2024; Weikmann & Lecheler, 2023). These capabilities are now more accessible than ever, in ever-increasing sophistication. We have seen for example in the Irish presidential election how AI-generated videos that convincingly impersonate trusted broadcasters and political candidates can circulate long enough to mislead voters and force platforms and regulators into reactive mode (Ryan, 2025). Drawing on our CampAlign Tracker research, analysis of visual GenAI content used in Germany and the Netherlands showed that the vast majority of it relied on photorealistic aesthetics rather than clearly stylised or illustrative approaches (see Figure 10; Kruschinski & Votta, forthcoming; Votta & Kruschinski, 2025). This preference for realistic depiction suggests strategic efforts to harness the evidentiary weight traditionally associated with photography, though it simultaneously heightens risks of deception and erosion of trust in visual political communication.

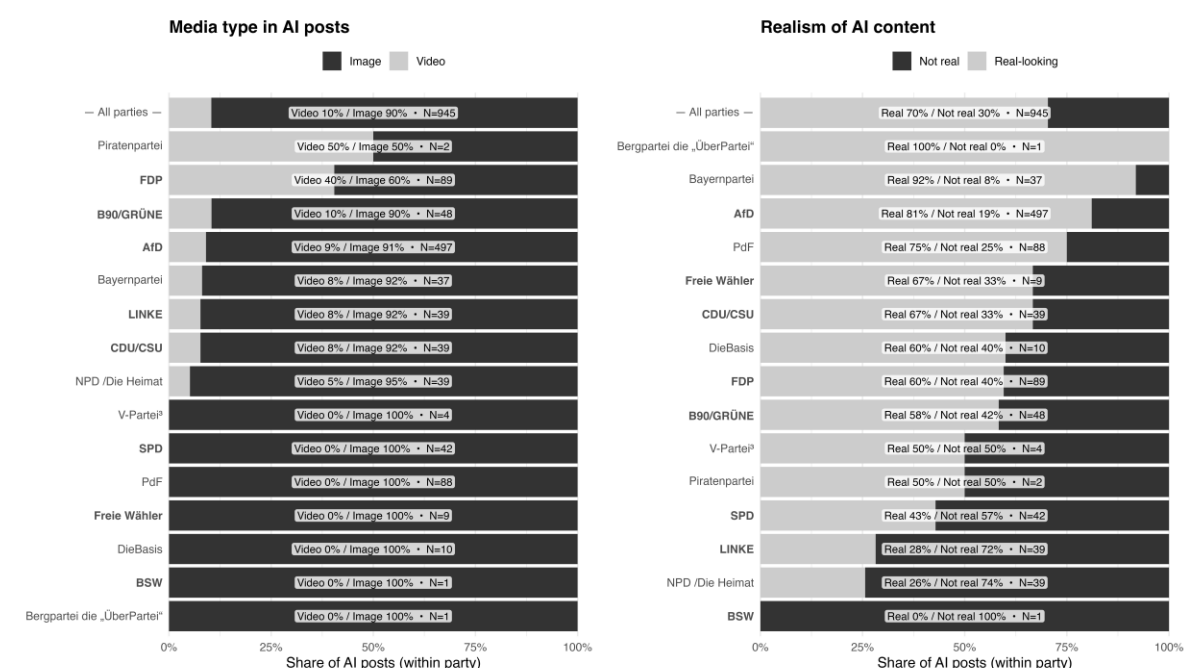


Figure 10. Media type and realism of visual GenAI posts in the 2025 German election campaign. Labels inside the bars report, for each party, the share of AI posts that were videos vs. images (left panel) or real-looking vs. not real-looking (right panel), followed by the total number of AI posts for that party (Kruschinski & Votta, forthcoming).

However, the photo-realistic content predominantly depicted ordinary citizens, supporters, and generic members of the public rather than political elites, both in Germany and the Netherlands. When deepfakes of known political figures appeared, they typically comprised lead candidates in Germany including Friedrich Merz (CDU/CSU), Alice Weidel (AfD), and in

the Netherlands, Geert Wilders (PVV) and Frans Timmermans (GroenLinks-PvdA). A prominent example of a deepfake of Frans Timmermans involved him getting arrested, which prompted a series of death threats against him and subsequent criminal complaints against the perpetrators (Feenstra & Sabel, 2025).

Our empirical findings provide further evidence of how these dynamics manifest in practice. Right-wing and far-right actors emerged as primary producers of negative AI content in both Germany and the Netherlands (Kruschinski & Votta, forthcoming; Votta & Kruschinski, 2025) who frequently deployed AI-generated imagery to criticise rival parties and societal groups including minorities, often characterised as criminals. This represents a concerning application of the technology for exclusionary messaging that reinforces harmful stereotypes. Major parties demonstrated greater restraint, and when employing AI for criticism, focused on policy differences and political accountability rather than personal attacks or stigmatisation. Furthermore, analysing the engagement patterns of AI content reveals a troubling pattern: while AI-generated visuals constituted a small fraction of total content, they generally achieved higher median engagement compared to non-AI posts in Germany (Kruschinski & Votta, forthcoming).

Further, AI generated visuals enable a more insidious communication strategy: the mere existence of deepfake technology provides a convenient escape route for political actors caught in genuine scandals, they can now claim that authentic recordings, images, or videos documenting their misconduct are AI-generated fabrications, exploiting public uncertainty about what is real to evade accountability. This phenomenon, known as the *liar's dividend*, means that those who engage in actual wrongdoing paradoxically benefit from the proliferation of synthetic media (Schiff et al., 2025). Moreover, when citizens lose the ability to reliably differentiate genuine content from manufactured material, pervasive scepticism about political information and public figures can take root (Jungherr, 2023; Weikmann et al., 2025).

2.2.4 AI-Assisted Personalisation of Political Messaging

Gen AI enables more large-scale personalisation of political messaging. By processing substantial volumes of information, including social media engagement, constituent correspondence, and demographic data, campaigns can identify public sentiment and voter priorities. This analytical capacity supports the development of customised communications that address specific constituent concerns, preferences, and ideological positions. Such tailored messaging can strengthen the resonance of political communication and stimulate civic engagement (Foos, 2024). Additionally, AI-powered translation and accessibility features may enable campaigns to connect more effectively with marginalised communities and citizens from diverse linguistic backgrounds. For example, both in media and in political campaigning, generative tools could be used to translate content to accommodate people's diversity in both cognitive and physical capacities and needs, both young and adult. Speech-to-text technologies for instance could help in transcribing audio-visual programming for people with hearing loss. Indeed, AI tools are already leveraged for multilingual voter outreach in India and the United States (Bhattacharya, 2024).

However, the flipside of adaptive voter outreach is that the very same techniques used to broaden inclusion can also serve to deliver ever more granular, opaque, and affectively charged messages to tightly defined audiences, setting the stage for large-scale

personalisation that blurs the line between persuasion and manipulation. AI-enabled large-scale personalisation of political messaging introduces risks of societal fragmentation, inauthentic coordinated behaviour during elections, and stronger incentives to push polarising content (Simon & Altay, 2025). When AI systems increasingly deliver information that reinforces voters' pre-existing viewpoints, the possibility for constructive dialogue across perspectives diminishes. This dynamic becomes particularly pronounced on social media platforms, where algorithmic systems can direct personalised AI-generated messages toward audiences predisposed to accept them. Such targeting can amplify extreme viewpoints and deepen political polarisation (Jungherr, 2023). The capacity of GenAI to produce content at scale also presents threats to political representatives. Field experiments demonstrated that AI-generated advocacy letters were nearly indistinguishable from human-written correspondence, suggesting malicious actors could distort legislators' understanding of constituent priorities (Kreps & Kriner, 2023). Beyond direct influence on officials, AI systems can shape democratic outcomes by swaying voters through personalised messaging and microtargeting (Zuiderveen Borgesius et al., 2018). The convergence of low-cost content generation with AI-driven personalisation creates what researchers characterise as a potential "manipulation machine" (Simchon et al., 2024), raising fundamental questions about informed democratic choice and electoral integrity when scaled AI production combines with targeted and algorithmic distribution systems prioritising personalised material for audiences more susceptible to them.

2.2.5 Gen AI for Political Information Seeking and Civil Society Engagement

Gen AI systems also offer new pathways for voter information and civil society engagement. While initially said to be unavailable in the large models in 2024 (but see (Nieuwsuur, 2024), in 2025 these guardrails, if they were ever there, are gone and they are used as alternatives to more established voting advice applications (VAAs). This development presents both opportunities and challenges. On one hand, chatbots and LLMs can provide valuable services such as summarising complex policy issues, synthesising consultation responses, and making political information accessible at different levels of literacy and depth. This capability could ideally improve responsiveness and help citizens engage with policy content more effectively. Several applications of chatbots have also been tried that can lower barriers to political knowledge by allowing citizens to directly 'chat' with a party manifesto, ask questions to digital clones of candidates, or debate a fictional supporter of an opposition party (Argyle et al., 2023; Schiele et al., 2024), even reducing conspiratorial beliefs through conversations (Costello et al., 2024).

On the other hand, these chatbots carry significant risks, including the misattribution of party viewpoints, the spread of misinformation, and the introduction of noise into the electoral information environment. Empirical testing by the Dutch Data Protection Authority (AP) revealed systematic bias in chatbot voting recommendations. Across four popular chatbots, over 56% of recommendations favoured just two parties (PVV or GroenLinks-PvdA), with one chatbot recommending these parties in over 80% of cases, regardless of user input. Meanwhile, other parties were rarely recommended even when user positions aligned perfectly with their platforms (Autoriteit Persoonsgegevens, 2025). The scale of potential exposure to these risks is concerning: survey research among Dutch citizens found that

approximately 1 in 10 voters reported they were likely to ask AI chatbots for voting advice, with highest numbers among 18-34 year olds (17%) to 6% among those 55 and older (de León et al., 2025). For future development, this may show how AI chat interfaces could become de-facto gatekeepers for electoral information in much the same way as AI-driven news search and overviews (see also section 2.1.2).

Beyond voting advice applications, GenAI offers potential for citizen participation and civil society engagement. GenAI tools can synthesise thousands of citizen submissions in public consultations, making participatory governance more scalable and accessible (OECD, 2025). For instance, deliberative democracy platforms can employ AI to process citizen input on policy proposals, identify common themes across diverse viewpoints, and facilitate structured dialogue between participants with different perspectives (McKinney, 2024). These applications have shown promise in reducing polarisation by helping citizens find common ground through AI-mediated discussions that highlight shared values rather than divisive positions (Argyle et al., 2023; Tsai et al., 2024). AI-assisted tools enable smaller organisations to participate more effectively in policy debates previously dominated by well-resourced actors, potentially democratising civic engagement (Bonin et al., 2025; Guler et al., 2024). However, as indicated in Section 3.1.1, in certain cases, a reliance on GenAI can flatten diversity and drown out minority voices to the favour of dominant majorities or popular, including polarising, narratives (See also: Lazar & Manuali, 2026). Moreover, public acceptance remains a significant challenge. Experimental evidence from Germany reveals a clear "AI penalty": citizens express significantly lower willingness to participate in AI-facilitated deliberation compared to human-facilitated formats, with participation dropping by 7.4 percentage points and perceived quality ratings falling nearly 10 percentage points (Jungherr & Rauchfleisch, 2025). This scepticism is particularly pronounced among those who perceive greater AI risks and lower societal benefits, potentially creating a new "deliberative divide" that cuts across traditional participation gaps. Finally, for citizens and civil society actors, GenAI could help lower the barrier to engage in political communication within (online) participatory structures, like public consultation mechanisms, aimed to decrease the constituency's distance to the legislator and render decision-making more democratic. However, even if equal access to the technology can be guaranteed, it remains uncertain whether over time the use of GenAI in these contexts will transform or instead deepen (existing) political inequalities (Danaher, 2024). Moreover, and as aforementioned, GenAI can be leveraged by bad-faith actors to "generate false 'constituent sentiment' at scale by effortlessly creating unique messages taking positions on any side of a myriad of issues (Kreps & Kriner, 2023, p. 123)." This makes it harder for public institutions to distinguish which public comments, whether co-produced by AI or not, "genuinely represent the preferences of interested stakeholders (Kreps & Kriner, 2023, p. 125)."

Summing up, as AI adoption increases and parties recognise potential engagement advantages, democratic discourse faces the risk that synthetic content increasingly displaces substantive political communication, with parties competing not through policy argumentation but through AI-generated emotional appeals and strategic negativity optimised for viral spread rather than informed deliberation. In view of these opportunities and challenges, political actors must critically ask themselves to what extent the strategic use of GenAI in their campaigns is justifiable.

The positive use-cases mentioned throughout this section can only manifest their full potential if the risks to political communication are sufficiently addressed. Indeed, it would be antithetical to the goal of furthering inclusion and citizen empowerment through GenAI if those applications would simply reproduce societal prejudices. GenAI's democratising potential extends beyond electoral campaigns to everyday democratic participation. When deployed responsibly, these tools can lower barriers to civic engagement by helping citizens navigate complex policy documents, participate in consultations regardless of language or education barriers, and organize collective action more effectively. Moreover, it should be understood however that the abovementioned scenarios do not address the underlying structural social and political inequalities at their root, and hence, should not be interpreted as positioning GenAI as the solution for political inequality. To truly increase the political visibility of marginalised or vulnerable communities, civil society's position should be strengthened, as civil communities offer space to articulate experiences and perceptions that are not part of dominant (political) discourse and can help understand and expose various forms of injustice in state and economic power (Young, 2002, p. 155). Moreover, as everyone can use these tools, including powerful actors, the benefits they offer smaller parties and grassroots movements, might soon subside if their overall use as a tool for communication becomes more commonplace.

2.2.6 Citizens' Attitudes towards GenAI use in Political Campaigns

Public perception and evaluation of GenAI in political campaigns influence both the strategic options available to political actors and the design of regulatory frameworks. However, substantial uncertainty remains about citizens' attitudes toward AI usage in elections.

Population surveys reveal that while the term "GenAI" enjoys broad recognition, most people lack deeper understanding of the technology (Simon et al., 2025), and this knowledge gap extends to its application in political campaigns (Kruschinski et al., 2025). In Germany, most respondents have not consciously encountered AI-generated content in election advertising and remain uncertain about whether and how parties deploy these tools, suggesting that public debate has outpaced actual awareness. Cross-national surveys demonstrate that risk perceptions dominate attitudes toward AI in political applications (see Figure 11). Similar to patterns observed in journalism (see report section 2.1), where many express scepticism about AI-generated news and fear compromised trustworthiness and transparency (Simon et al., 2025), German voters associate AI-driven campaigns with diminished credibility of political messages, voter manipulation, and threats to democracy (see Figure 12; (Kruschinski et al., 2025). First, acceptance varies by application: technical support functions such as spell-checking or translation receive broad approval, but as AI's role extends into content creation, particularly in voter communication, resistance intensifies. The use of AI to generate press releases, social media messages, and especially realistic images and videos encounters the strongest opposition.

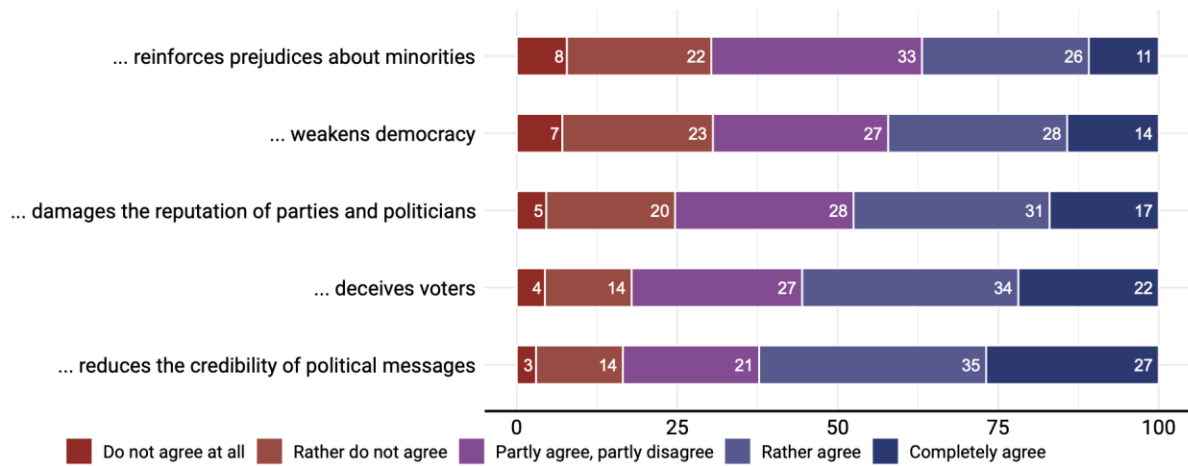


Figure 11: Risk perceptions of AI-driven Campaigns *Note.* N = 1991; *question.* "If parties and politicians increasingly use AI to create political advertising in the future, it..." *Source:* (Kruschinski et al., 2025).

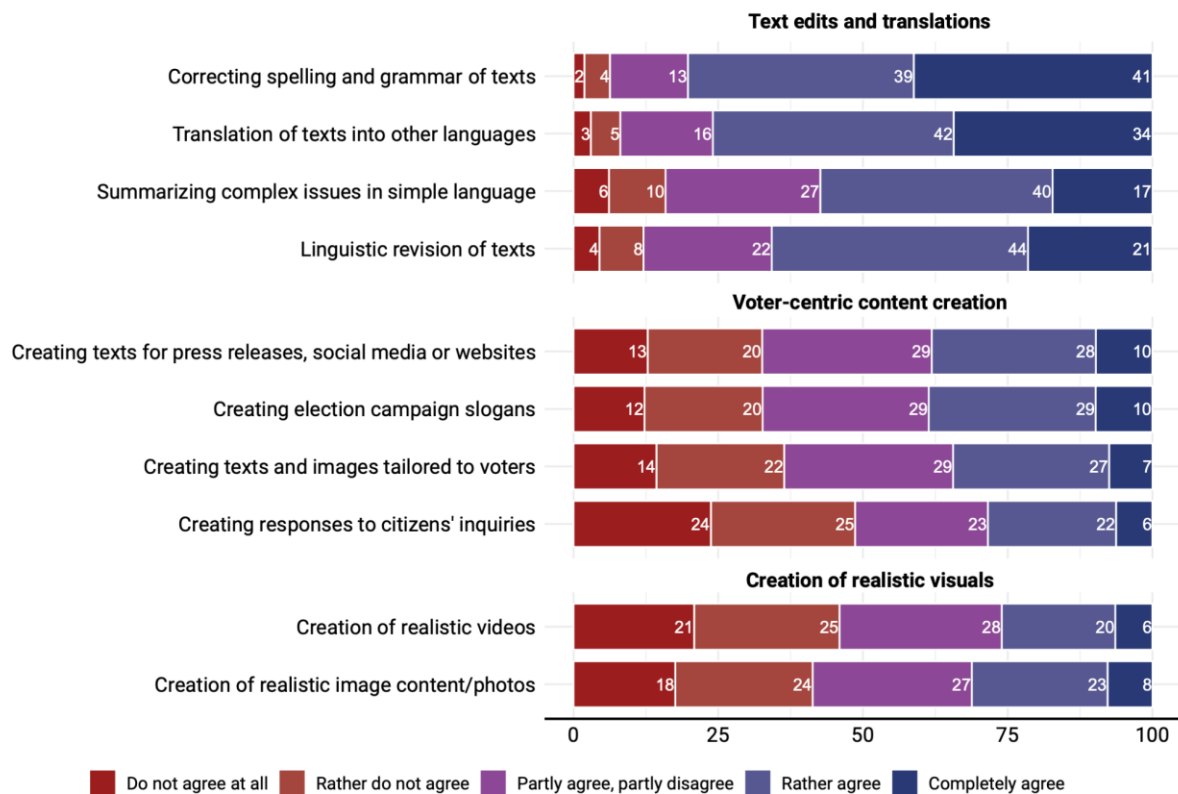


Figure 12: Acceptance of AI in political communication by task (in %) N = 1991; *question.* "Please indicate how acceptable you consider the use of AI for certain tasks."

Understanding what drives these attitudes reveals complex patterns. Individuals confident that their political actions matter (high external political efficacy) and those curious about technology perceive fewer risks and show greater acceptance across AI applications. Conversely, people more engaged with politics generally recognise greater risks, likely because they better grasp the political processes AI might disrupt. AI knowledge produces a paradox: informed citizens accept AI for simple tasks like translation but reject complex

applications such as generating realistic imagery, suggesting discriminating judgment rather than wholesale acceptance or rejection.




Figure 13: Public opinion towards regulating the use of AI in political communication. Note. N = 1991; question: "How useful do you find the following proposals?"

Regarding regulation, mandatory transparency emerges as the central public preference (see Figure 13; Kruschinski et al., 2025), which mirrors citizen expectations in journalism (see 2.1.6). Most Germans support labelling requirements for AI-generated campaign content, particularly realistic depictions of people or events, and find AI use more acceptable when transparently disclosed. Stricter measures such as voluntary abstention or outright bans receive less support. Risk perception drives regulatory preferences: those viewing AI as dangerous favour stronger interventions, especially AI prohibitions. However, citizens with greater AI knowledge and technological curiosity support transparency measures but oppose bans, preferring information to make their own judgments. Those confident in political system responsiveness (external efficacy) oppose both disclaimers and bans, perhaps viewing AI as less threatening when trusting democratic institutions to self-correct. These heterogeneous attitudes shaped by trust, knowledge, and technological openness reveal that building legitimate AI regulation requires accommodating diverse perspectives rather than seeking uniform consensus.

2.2.7 Conclusion

Overall, current empirical evidence regarding the campaign perspective suggest that GenAI is still being used rather experimentally by established political parties in the Netherlands and Germany. In contrast, extreme, smaller, and resource-poorer parties are strategically relying on GenAI because they set hardly any ethical boundaries for its use and want to offset their resource-related competitive disadvantage to some extent. As AI adoption increases and parties recognise engagement advantages, democratic discourse faces the risk that synthetic content increasingly displaces substantive political communication, with parties competing not

through policy argumentation but through AI-generated emotional appeals and strategic negativity optimised for viral spread rather than informed deliberation.

While democratising benefits exist for under-resourced candidates and parties, the current balance tilts toward risks for democratic discourse integrity. The pattern with established parties exercising caution while fringe actors experiment intensively suggests an intervention window before practices become entrenched. Comprehensive regulatory frameworks with robust enforcement mechanisms are essential for preserving informational foundations of democratic decision-making as AI capabilities advance toward increasingly indistinguishable synthetic content. The 2025 Dutch and German elections demonstrate that voluntary self-regulation proves insufficient, platform governance remains weak, and comprehensive approaches addressing transparency, accountability, and positive obligations to support democratic participation require priority attention. Yet, these findings also underscore the contingent nature of GenAI's political impact: its effects depend less on the technology itself than on the social, organisational, and regulatory contexts in which it operates.

Regarding empirical results for AI acceptance in the public, societies are confronted with a broader question about the boundaries between automation and authenticity in democratic communication as campaigns integrate GenAI into routine operations. Public attitudes reveal that citizens have a greater risk perception about the use of AI in political campaigns but distinguish between supportive and manipulative uses with acceptance varying systematically: technical support functions like spell-checking receive broad approval, while AI-generated imagery and videos encounter strong resistance. Further, mandatory transparency for AI emerges as the dominant regulatory preference, with most Germans supporting labelling requirements for AI-generated content instead of outright bans. However, risk perception, acceptance and regulatory preferences differ across citizen characteristics. AI knowledge produces discriminating judgment rather than uniform positions: informed citizens accept AI for simple tasks but reject complex applications like generating realistic depictions, suggesting sophisticated evaluation of specific risks. Political engagement correlates with heightened risk perception, as politically active citizens better understand the democratic processes AI might disrupt. Regulatory attitudes divide along knowledge and political efficacy lines: those with greater AI literacy favour transparency over prohibition, while those with high external political efficacy oppose both disclosure mandates and bans.

These heterogeneous and knowledge-dependent attitudes demonstrate that legitimate AI governance cannot rely on uniform public consensus but must accommodate diverse perspectives shaped by technological understanding, political trust, and engagement levels. This selective acceptance offers both a normative compass and a policy roadmap: transparency, accountability, and human oversight remain critical prerequisites for legitimate deployment.

Current empirical evidence thus points toward a double imperative. On the one hand, regulators and platforms must safeguard electoral integrity through enforceable transparency and labelling regimes. On the other hand, policymakers, educators, and civil society actors must cultivate citizens' media and AI literacy, enabling them to discern synthetic content and evaluate political communication critically. Future campaigns will likely blend human and machine creativity ever more seamlessly, demanding that democratic systems adapt not by

rejecting technological innovation but by embedding it within robust ethical and institutional guardrails.

In this light, GenAI does not mark the beginning of a new era of political communication but rather a new phase in the already existing struggle that seeks to align technological innovation with democratic responsibility.

2.3 Infrastructure

The deployment of GenAI in journalism and other sectors is underpinned by a vast and multilayered technological infrastructure. Crucially, many of these layers, including computing power, data storage, access to AI models with billions of parameters, and cloud services, are operated by the same firms that own leading GenAI tools, search engines, marketplaces and social media platforms referenced in our case studies. Thus, the concentration of ownership and control of our digital information infrastructure in the hands of untransparent, profit-oriented technology companies raise serious concerns about democratic accountability and the trustworthiness and quality of the digital information environment.

2.3.1 The politics of infrastructure

Infrastructures carry weighty (geo)political implications, in particular regarding critical topics like (digital) sovereignty, strategic autonomy, and democratic deliberation. For that reason, it is imperative to take stock of the importance of infrastructures and their politics, confronting them as a vector of power. Yet, infrastructure is a slippery term to take hold of. Conceptually, they can be understood as “built networks that facilitate the flow of goods, people, or ideas” (Larkin, 2013, p. 328), connecting and organising various digital, economic, social and political dimensions of our society.

Practically, on the one hand, infrastructure refers to the material and technical substrates required to build and run large-scale AI systems, such as data centres, semiconductor chips, underwater sea cables, and energy grids (Mayer, 2023). On the other hand, it often encapsulates more intangible dimensions, such as software architecture, hardware specificities and organisational structures that coordinate production, such as Content Management Systems (CMS), Application Programming Interfaces (APIs) and Software Development Kits (SDKs) (Pybus & Côté, 2024). By extension, talent, technical expertise, and knowledge are also considered critical resources for the development and deployment of infrastructure in this context (Rikap, 2021). This report adopts a synthetic and integrated approach considering infrastructure as the means and conditions of production in a given, often industrial, context. In short, ‘infrastructure’ encompasses a wide and varied set of technologies and conditions that facilitate communication. Concretely, as the previous sections have identified, where political actors often purchase platform-provided AI tools as they need them (2.2.5), journalism’s reliance on AI infrastructure extend to, for example, the cloud systems on which media organisations’ own AI systems are deployed (2.1.5).

The implications of the politics of infrastructures for democracy become clearer when examined within our broader information ecosystem, which may be conceived as a “complex and interconnected set of systems that [determine] who can receive or share what kinds of information and in what contexts” (Mansell et al., 2025, p. 75). Crucially, this set of complex systems comprising our information ecosystem is not neutral. Instead, it is conditioned by the interconnectedness of the political-economic structures, including platform ownership configurations (Leerssen, 2025) and the sociotechnical infrastructure upon which it is formed.

A recent example underlines the point: in January 2025, right after the re-election of Donald Trump, Mark Zuckerberg, Meta Platforms’ CEO, announced the simultaneous termination of its third-party fact-checking network in the US, the re-introduction of political content in users’ recommendations, and a relaxation of content policies risking the exacerbation of hateful online behaviour toward marginalised groups (de Brún & Sawyeddollah, 2025). These choices, made by a dominant infrastructure provider and operator, affect the conditions and quality of information in users’ feeds, as well as the ways in which news media organisations produce, distribute and monetise their content.

Therefore, in this context, infrastructures – from clouds and servers to software architecture and solutions – emerge as the constitutive backbone of our digital information ecosystem. They provide the material, digital and organisational conditions under which information is produced, disseminated and exchanged, and are the source of considerable economic and political power. Together, the implications for democracy are fundamental. As such, it is crucial to keep a materially grounded perspective when discussing infrastructures, taking stock of the politics that shape them, including power dynamics in their governance and ownership structures.

2.3.2 Control over the infrastructure needed for GenAI

Dominant technology companies have become best suited to “exploit the affordances of digital information and communications networks and supply infrastructures that facilitate particular types of interactions” (Cohen, 2019, p. 41). The development and deployment of GenAI has been marked by an infrastructural paradigm and a political economy where a handful of American and Chinese actors have disproportionate control over critical technology (Balayn & Gürses, 2024; Whittaker, 2021). This concentration of infrastructural power means that the values embedded in these systems (e.g., efficiency, scalability, profitability, surveillance, data extraction, etc.) are effectively engrained in the services supported by them. As already discussed in Section 2.1.5 on structural dependencies, Big Tech firms like Microsoft, Amazon and Google have expanded their grip over journalism beyond tools and distribution channels the infrastructure on which journalism broadly depends for its viability and, subsequently, its capacity to fulfil its role in democratic societies.

These tech firms have systemically leveraged their access to, and sway over, capital, cloud infrastructure, data, and talent to shape the development of GenAI in their favour (Autorité de la concurrence, 2024; van der Vlist et al., 2024). In other words, they have used their pre-existing power over digital infrastructures to shape a new playing field that, in turn, allows them to deepen their control over those large technical systems. Common practices include vertical integration of key technology for the training and deployment of AI models and complementary innovation strategies allowing dominant players to corner specific markets, (Luitse, 2024). For

example, hyper-scaler cloud providers often bundle together different software as an unified service, thereby forcing their clients to engage with a pre-determined production environment while (often) locking them into their infrastructure. These examples illustrate an infrastructural environment characterised by market concentration, lack of meaningful regulation, and astronomical investments from a small number of players whose broader strategic objectives are mostly profit-oriented (Vipra & Myers West, 2023; Widder & Kim, 2025).

When GenAI-fuelled systems become embedded in our information ecosystem, the question of “who can receive or share what kinds of information and in what contexts” (Mansell et al., 2025, p. 75) is greatly affected by the infrastructural paradigm and political economy laid out above. In fact, this entanglement of infrastructures with the political economy of journalism is not abstract: it is already reconfiguring the broader media environment (Sevignani et al., 2025). Licensing agreements with AI providers, for instance, reduce journalism to modular snippets of information ready to be recombined at the prompt of a user in a service like ChatGPT, while more direct integrations reorganise newsroom operations around third-party systems. As such, knowledge in the form of data – especially as produced by journalists – is a crucial resource for GenAI providers, not simply for training their AI models but also to differentiate their offerings by providing services like up-to-date news summarisations (e.g., accessing news via ChatGPT or Google’s AI Overview). Therefore, search for and distribution of news is increasingly mediated by GenAI-fuelled digital infrastructure. This will very likely have consequences for the visibility, reachability and revenue-making of news media organisations. For example, at time of writing, OpenAI has announced seven deals with news media publishers headquartered in Europe (Italy, France, Germany, Nordics, and Spain, and two in the UK). This approach to deal-making has implications for the reliability and diversity of media sources on which AI models can draw; media organisations’ decision to move their content behind paywalls to prevent it from AI companies from using it can also restrict the broader public’s access to information. Developing and enforcing licensing agreements that ensure news content can be accessible to humans without being used to train AI models may mitigate this issue.

Infrastructural dependencies may also undermine news media’s capacity to exercise journalism outside the technology value chains of GenAI providers (see again Section 2.1).

By way of illustration, *The Washington Post* (owned by Nash Holdings, Jeff Bezos’s private investment company) developed ArcXP, a CMS for news media organisations hosted within AWS. The adoption of ArcXP by news media organisations like *Le Parisien* (Arc XP, 2024) encapsulates the dynamic described above: a newsroom task as simple as summarising or tagging a story now passes through AWS’s cloud stack, embedding external governance, values and organisational logics into the conditions of journalistic labour. In doing so, these infrastructures may transform newsrooms into tenants of their proprietary systems, modularising editorial work, dictating how value is created and circulated, and broadly, reshaping journalism’s critical role for a functioning information ecosystem. As succinctly put by Schneiders and Stark: “the value and quality of information stand and fall with the quality of the production process. If independence, diversity, and relevance are not guaranteed, journalism loses its functions – and thus its *epistemic value* for the audience” (Schneiders & Stark, 2025, p. 16)(emphasis added).

Let us consider this in the context of an example discussed in Section 2.1: the agreements between press publishers and AI tech firms like OpenAI are based on the burgeoning technique of “Retrieval-Augmented Generation” (RAG), wherein—basically—a GenAI text generator like ChatGPT will try to answer to a user prompt using not only its training data but also another (potentially more up-to-date) dataset. This is what ChatGPT does when answering to a prompt asking for news on a topic and pulls up information – not always accurately, as shown in section 2.1(d) – from the *Financial Times*, with which OpenAI has signed a multi-million-dollar deal. Similarly, the Associated Press has been reconfiguring its archive to be ‘AI-ready’ to ensure that GenAI systems can use RAG techniques for their own consumer or enterprise offerings (Helberger et al., 2022). The locus of value creation of journalism, then, is expanded and broadened from circulation and advertising to the value chains of AI tech firms, including conditions of access to both these firms’ APIs, and the press publishers’ AI-ready archives.

Crucially, in this context, we do not simply talk about economic value, notwithstanding its importance for journalism’s viability. Rather, we are also concerned with epistemic and democratic value. That is how citizens make sense of the world in order to partake in democratic deliberations and political life more broadly. Therefore, the example of monetising publishers’ archives, and by extension people’s capacity for meaning-making, underscores the transactional value that the current (infra)structural paradigm assigns to journalism. This process on proprietary and paywalled ecosystems supported by, predominantly, the cloud computing infrastructure owned by the same tech firms that are purchasing access to these information troves.

At the same time, we see emerging initiatives in the EU, seeking to experiment with alternative governance arrangements for licensing access to news content. For instance, a consortium of Dutch organisations led by the Netherlands Organisation for Applied Scientific Research (TNO) is developing what they call a “sovereign” Dutch-language LLM, GPT-NL (GPT-NL, 2025). The project is based on a licensing model for publishers that is framed around consent, fair compensation, and regulatory compliance, with some major publishers like DPG Media, Mediahuis, and the press agency ANP having already agreed to license their archive material. Such initiative stands in contrast to the industry-led licensing marketplaces embedded within proprietary cloud ecosystems mentioned earlier. Building and maintaining competitive models, however, requires substantial financial resources, long-term political commitment, and access to compute infrastructure that is itself highly controlled by a few dominant companies. Better licensing terms may, thus, rebalance certain aspects of value extraction and enhance transparency for publishers, potentially improving conditions for democratic access to information, but they do not necessarily lead to a decoupling of European AI ecosystems from globally concentrated hardware and cloud dependencies. In that sense, the persistence of these underlying dependencies indicate broader systemic and structural issues discussed in the next section.

2.3.3 Infrastructural dependency and democratic backsliding

As technologies are developed and operated behind a veil of opacity, and citizens are left without meaningful tools to contest the choices and logics that make up their data-driven environment, the impact these technologies might have on their digital environment and

livelihoods is often unilaterally determined by others without much transparency or insight (Naudts, 2024). This position of technological domination has also been identified as a strategic, geo-political asset.

The disproportionate focus on scale that characterised the development of GenAI has resulted in a profound emmeshing of corporate and state interests, as explicitly shown by the tech leaders' rush to court the new US administration in January 2025. Despite being privately owned, the computational infrastructures underlying the development of AI tools were developed in partnership with U.S. state actors and they still exert considerable sway over them (Holt, 2024; Hu, 2015). Incidents such as the suspension of AWS services in Iran (Motamedi, 2019) or Microsoft's account of ICC prosecutor Karim Khan (Quell, 2025) have shown that the American Department of State is willing to leverage the infrastructural power of companies under their jurisdiction to achieve strategic goals. As the Dutch government concluded in the aftermath of Khan's suspensions, the reliance of public institutions on these infrastructures could negatively impact their autonomy (Veldkamp & van Marum, 2025). GenAI is further entrenching these dynamics. On the one hand, dominant actors are using the development of AI as an excuse to further extend their control over computational infrastructures (Chander & Gürses, 2024). On the other hand, users—including public institutions—are becoming further dependent on those infrastructures through the integration of GenAI in their daily activities (Brennan et al., 2025).

More specifically, one should be mindful of the bundling and reinforcement of powers between (authoritarian) regimes (States) and platform owners based and operating on their territories. According to Rolf and Schindler, global technological business power has become concentrated in the United States and China, whereby digital actors are strategic partners in response to economic and geopolitical tensions (Rolf & Schindler, 2023). In turn, these preferred partners gain economic and regulatory benefits, which also help explain the pushback the European Union has faced. Importantly, digital platforms, including owners of general-purpose models, can be used to push and hide certain political narratives. For example, if words like "biases", "disability", and "diversity" draw scrutiny in scientific research (Johnson et al., 2025), it is not unthinkable that in the future instructions are given to private actors to neutralise, depoliticise, or shadowban these notions in GenAI. In turn, these interventions could have a downstream effect where these models are relied upon to generate journalistic, electoral, or political content. For instance, a procured model used to streamline journalistic articles could be trained to purposefully prioritise language that reflects the gender binary in an effort to render non-binary people invisible. Systems might also rationalise these decisions (Kotek et al., 2023), obscuring the discriminatory or harmful nature of such value-laden choices.

2.3.4 Conclusion

Infrastructural concentration adds another layer of vulnerability for democracy and its institutions. This vulnerability can be hard to detect. While the negative implications of a viral deepfake of a political candidate are readily apparent, infrastructure operates in the background and impacts democracy in subtle ways. Control over infrastructure does not need to translate into censorship; by controlling the technologies and other resources that facilitate communication, technology companies can shape what kind of communication is possible in the first place. This control manifests, for example, through what kind of AI tools are easy to

develop, what kind of language AI tools use by default, and how AI tools can be deployed on cloud infrastructure. When control over infrastructure is exerted in a more abusive manner (for example by raising the prices or cutting off access to critical cloud infrastructure), it is often in a situation where users are already heavily dependent on infrastructure.

The concept of infrastructure is useful to understand how technology companies have come to influence many of the technologies and resources needed to communicate in the digital era, from minerals to chips, and data centres to software. It also elucidates how much of this infrastructure is operated in a manner that is at odds with democratic principles. Namely, by a few concentrated actors, in an opaque manner, and with little regard for values that underpin communication in a healthy information environment, such as diversity, media freedom, and non-discrimination. However, as is clear from the section above, the concept of infrastructure also captures a wide array of different technologies and resources. By zooming in, then, on the infrastructural layer of the information environment, it becomes clear that what we are witnessing is not merely a concentration of power in the hands of a few corporate actors – who are also often in cahoots with powerful state actors – but an unprecedented convergence of political, social, economic, and cultural domains. The same infrastructures that undergird our increasingly digital lives also underpin our information environment, from digital journalism to political campaigns. This profound enmeshment has equally profound implications for the way information and the value assigned to it are produced. In that sense, considering the infrastructural layer means also considering the means and relations of production, including the value of said production and quality. Said otherwise, this convergence collapses – in the clearest way possible – boundaries between economic and democratic value, rendering the infrastructures of our information ecosystem – rather than just its content – a primary arena for the contestation of power.

3. Implications for human rights, political participation and democratic principles

In the previous sections, we demonstrated how GenAI has been deployed, experimented with, or is considered for future use in journalism and political communication, and creates risks as well as opportunities for democracy and democratic processes. We also discussed some framework conditions, including the prominent role of, and dependency on a concentrated number of technology companies that shape market dynamics. The goal of this section is to place these dynamics into the context of existing Council of Europe normative frameworks and instruments. To this end, we will first discuss the use of GenAI in journalism, political communication and the public information infrastructure through the lens of human rights law, with a special focus on the rights to non-discrimination and freedom of expression. The core question to answer for this section is the extent to which the dynamics and implications identified in the previous section bear human rights relevance, and where we see some risks, but also opportunities for realising human rights. We will then turn to the body of recommendations of the Council to identify to what extent the Council already has addressed some of the issues arising and forwarded concepts and solutions to build on. In a last section, we will then turn towards the Council of Europe's Framework Convention with the question to what extent the Convention could be a tool to address some of the risks and opportunities identified.

3.1 Adding the human rights dimension: GenAI from the perspective of the rights to equality, non-discrimination and political participation

The ability to participate in public decision-making, and do so on equal and non-discriminatory terms is at the heart of the democratic idea, as well as the very foundations of the Council of Europe: “Its central pillars of human rights, democracy and the rule of law are grounded on the fundamental belief in the equality and dignity of all humans as independent moral agents” (CoE, 2019, para. 9). Inclusivity refers to the belief that policies or decisions that significantly impact people’s living sphere are normatively legitimate only when those affected have been meaningfully included, through participation or representation, in the decision-making process (Young, 2002, p. 23).

Generative technologies mediate how people are seen or heard within social and cultural life. In so doing, they can also unduly limit the access to, and exercise of, civil and political rights *and* social and economic freedoms, of some individuals and groups to the benefit of others. In this context, the right to non-discrimination – enshrined within Article 14 of the European Convention on Human Rights (hereinafter “ECHR” or “the Convention”) and Protocol No. 12 to the Convention – can play a pivotal role as it helps define under what conditions decisions and policies that (in)directly result in people’s differential enjoyment of human rights, and associated freedoms, can be objectively and reasonably justified. However, to challenge the histories of disadvantage underlying the AI-pipeline *and* combat the discriminatory impact generative technologies may have upon deployment, states, governments, and private actors, must move beyond the passive or negative duty not to discriminate or prevent unfair differential treatment. Rather, as Bartoletti and Xenidis note, stakeholders must both revisit and adapt existing measures in light of technological change and challenges, and embrace positive action and obligations as a means to transform patterns of discrimination and accommodate differences within society and alongside the AI value chain (Bartoletti & Xenidis, 2023, p. 67).

3.1.1 Inequalities and structural implications

In this regard, the ideal of realising socio-economic and political equality shares a mutually reinforcing relationship with the right to non-discrimination and the promotion of human rights. Whereas people need political agency over their environments to ensure their human rights are safeguarded, those same rights ensure everybody has the necessary civil, political and socio-economic opportunities to participate in the public sphere. In this context, the use of GenAI in journalism and political campaigning, combined with growing infrastructural dependencies, triggers at least two core concerns. First, due to their widespread availability, generative technologies risk entrenching and deepening structural inequalities, like sexism and racism, in areas of life crucial to democratic processes of deliberation and participation at scale. Second, they risk flattening diversity within spaces of political debate.

On a structural level, the discriminatory effects of AI-systems have been well-documented, and GenAI systems too have been shown to reflect (latent) societal prejudice, prevailing forms of discrimination and negative stereotypes, including racist and gender biases (see for example: Abid et al., 2021; Birhane et al., 2021, 2024; Kotek et al., 2023; Tao et al., 2024; Gillespie 2024). These biases might emerge throughout the entire AI lifecycle, from data

collection, to labelling; from design to deployment (Bartoletti & Xenidis, 2023; Chen et al., 2021).

The widespread availability and ease-of-use of these systems, combined with the infrastructural centralisation of power in terms of their design, has the potential to further institutionalise these social forms of injustice. If left unchecked, each downstream application in journalism or political campaigning built on top of a model is liable to reinforce the inequalities the latter reflects. Systemic harms, like the exclusion and discrimination of historically disadvantaged and marginalised groups, might therefore be introduced on a larger scale, with a greater immediate impact than non-general purpose AI systems would. At the same time, those downstream applications too can introduce (structural) inequalities, if models are enriched with biased custom data or used to re-enforce biases and stereotypes, as examples of the use of GenAI in political campaigning have demonstrated (see section 2.2.3).

As cautioned by the European Court of Human Rights (hereinafter “ECtHR” or “the Court”), stereotypes can undermine “group’s sense of identity and the feelings of self-worth and self-confidence of members of the group (*Aksu v. Turkey*, 2012, para. 58).” According to Timmer, stereotypes should be conceived as “covert structures that underlie equality and discrimination” (Timmer, 2014, p. 221). Stereotypes constitute preconceived beliefs about particular social groups and unduly limit people’s social and economic mobility, including their enjoyment of human rights. Ethnic profiling strategies, for example, are discriminatory, because they extrapolate from the stereotypical perception an automatic connection between anti-social behaviour and a person’s ethnicity (see also *Lingurar v. Romania*, 2019, para. 76). In this context, stereotypes, social inequality, and discrimination, constitute an interrelated invidious cycle: “stereotypes can form a manifestation of discrimination, as well as a rationalisation and a cause of discrimination” (Timmer, 2014, p. 221). In doing so, they help maintain existing social and power structures. Given this reinforcing dynamic, traditions, generalised assumptions or prevailing attitudes, such as racial stereotypes or gendered roles, are insufficient to justify differential treatment (see also *Konstantin Markin v. Russia*, 2012). Yet, the latter risks happening where GenAI systems help mediate people’s knowledge environments, like news outlets or political campaigns, and (indirectly) inform democratic decision-making processes.

In this context, particular attention must also be paid to the position of vulnerabilised persons or groups. Vulnerability as we use it here, is a relational concept in the sense that it is informed and shaped by people’s embedding and position within wider social, economic and institutional relationships and dependencies, including structural inequalities and hierarchies (Fineman, 2008). In this context, stereotypes and vulnerability are reinforcing social status harms. On the one hand, vulnerable groups are often the subject of pervasive stereotypes. On the other hand, in recognising a group’s vulnerable status, we can identify those instances where the application of generalised assumptions is more likely to cause harm (Timmer, 2014, pp. 246–248; Case of *Carvalho Pinto De Sousa Morais v. Portugal*, 2017; Naudts, 2023). At the same time, the use of GenAI to enable new forms of (hyper)personalisation (see sections 2.1.1 and 2.2.4), which too brings with it the danger to further exploit, or create new vulnerabilities (Helberger et al., 2022). Already in its Declaration on the manipulative capabilities of algorithmic processes the Council of Europe’s Committee of Ministers warned that the use of computational means can result in “the sorting of individuals into categories, thereby reinforcing different forms of social, cultural, religious, legal and economic segregation and

discrimination. It also facilitates the micro-targeting of individuals based on profiles in ways that may profoundly affect their lives.” The border to unacceptable manipulation is reached, so the Council, where technology is used to “exploit existing vulnerabilities or cognitive biases” (CoE, 2019, para. 9e).

The second major concern from a human rights perspective that we would like to highlight are the potential implications for inclusion and the danger of flattening diversity as another core principle to the human rights to non-discrimination, political participation but also freedom of expression (see section 3.2).

Generative models might not handle diversity well. Large Language Models for instance risk functioning as flawed crowd-sourced consensus-machines (Waldo & Boussard, 2024). When an LLM is prompted on topics of which there is limited data online, they are more likely to produce answers that are inaccurate while doing so in ways that are framed as convincing or realistic. Conversely, LLMs perform better on topics that are more popular or have reached consensus. Section 2.2.2 also showed how LLMs, in interaction with algorithmic content moderation, can result in a prevalence of certain types of content, like emotional content, over others. Likewise, generative systems’ output may reflect majoritarian or dominant cultural standards. For instance, and unless prompted to do otherwise, contemporary LLMs still tend to promote and reproduce prevailing cultural identities and narratives, “rarely representing less common arrangements and perspectives” (Gillespie, 2024). Put differently, representation is not a given but must be queried. This flattening of diversity might also occur as generative models are asked to perform tasks on new input data. Certain AI-powered apps promise efficient document analysis, like summarisation, of citizen feedback received during public consultations (Small et al., 2023; Lazar & Manuali, 2026). However, in performing these tasks too, generative technologies have been shown to fabricate information, introduce (subtle) changes in meaning, reflect political biases, and hide outliers or minority voices (Lazar & Manuali, 2026).

As Young notes: “Too strong a commitment to consensus as a common good can incline some or all to advocate removing difficult issues from discussion for the sake of agreement and preservation of the common good (Young, 2002, p. 45).” Yet, she continues: “Sometimes those difficult issues matter deeply to one group because they perceive themselves as suffering a basic injustice, but they are the sources of deep disagreement because others in the society perceive rectifying this alleged injustice as coming at too great a cost to them (Young, 2002, pp. 44–45). While dissent can offer insight, those perspectives are unfortunately drowned out when prioritising consensus-based or majoritarian viewpoints. Similarly, where generative technologies are used by bad-faith actors to falsely inflate stakeholders’ preferences regarding certain policy positions, it reduces the visibility of genuine political expressions, including those co-produced through AI (Kreps & Kriner, 2023, pp. 123–125; see also Section 2.2.5). Whether LLMs are used to summarise, aggregate or represent people’s viewpoints, under current conditions, certain perspectives are more likely to be rendered invisible, rather than made an integral component of the processes of deliberation and contestation (Mouffe, 1999; Young, 2002). The risk of generative systems flattening diversity also counters ideals of pluralism embraced by the ECtHR, which has proclaimed that diversity should never be considered a threat, but rather an enrichment to the democratic vision (*D.H. and Others v. The Czech Republic*, 2007, para. 176) (see section 3.2).

3.1.2 The potential of non-discrimination law to address these concerns

The aforementioned socio-technical dynamics influence the output that will be produced by GenAI. The output produced by GenAI can be non-inclusive, outright demeaning or abusive, or otherwise inadequately represent or misrepresent certain social groups (see also Hacker et al., 2025). Scholars have long warned however how the grammar and vocabulary of European non-discrimination law risks becoming outdated for the digital era (Gerards & Xenidis, 2021; Gerards & Zuiderveen Borgesius, 2022; Naudts, 2023; Wachter, 2023; Xenidis, 2025). Unsurprisingly then, current non-discrimination frameworks might fail to address both the deepening of structural inequalities and the danger of invisibilising and excluding diverse voices. In the following, we highlight some areas of policy where the democratic impact of generative systems vis-à-vis the right to equality and non-discrimination can be improved.

To adequately capture the negative impact generative models and systems might have on people's ability to participate as equals within society, any evaluation of these technologies must be informed by a broader perspective on how societal and technological elements dynamically interact with each other to limit people's enjoyment of human rights. Indeed, legal interventions tend to target specific applications and instances of discrimination alongside the AI value-chain, rather than its societal roots. When focusing on a particular piece of content or decision however, the interconnected nature and unique amplification properties of generative models, and more so when they are general-purpose, might become obscured, ignoring how risks to rights accumulate and intensify over time (Liu et al., 2018; Naudts, 2023; Vogiatzoglou, 2025). Yet, in certain cases, when viewed in isolation, an output might be biased, or even discriminatory, but in itself, pose too little risk to people's overall human rights. Rather, it is the accumulation of, or the tendency of the model to reproduce, these types of injustice at scale, which reinforces the larger societal narrative that ultimately does impact people's overall freedoms and opportunities. In a similar manner, the individualised harms of generative systems produce, emerge from group-level data operations (Viljoen, 2021). For instance, the output and impact of GenAI might be the result of the many actions of many individuals, including decisions from both designers and end-users. The latter might be the case where citizens' interactions are used as data to inform future output.

Of course, this does not preclude the possibility that certain outputs constitute a human rights violation, for instance, when they are directly insulting, demeaning, harmful, hateful, or abusive to individuals or groups (see also Hacker et al., 2025). These observations illustrate how technologies must be placed within and evaluated as part of an intricate and interconnected network of socio-technical elements (Naudts et al., 2024, 2025). Importantly then, where third-party models are used to inform both private and public decision-making, the logics and interactions that inform their functioning and the normative assumptions and consequences those choices entail for individuals and society at large, cannot be hidden from decision-makers and citizens. Under conditions of such opacity, stakeholders cannot assess the proportionality of their decisions, as benefits might be projected, but potential disadvantages affected users might face remain hidden. Doing so would not only endanger the (democratic) legitimacy of AI-driven policies, but also their meaningful contestation by the broader public, including citizens, equality bodies and human rights institutions.

The data that inform the decisions and interactions of GenAI are varied, and consequently, the (socio-economic) injustice they introduce cannot always be delineated alongside salient or historically disadvantaged social groups, but may affect individuals and aggregates across novel dimensions, like their preferences or other behaviours solely monitorable through technologies (Gerards & Zuiderveen Borgesius, 2022; Naudts, 2023, 2024; Wachter, 2023). As part of these discussions, the single axis focus of European non-discrimination laws is often highlighted as an obstacle. As people's positioning within data structures is typically not shaped by one, but rather multiple, ever-evolving data points, in certain domains, the disadvantage people experienced might not be traced back to one characteristic, but multiple ones. An intersectional perspective, which recognises how discriminatory harms are shaped at the confluence of multiple (social) axes, might be more befitting the intricacies and complexities of the digital era (Atrey & Dunne, 2020; Bartoletti & Xenidis, 2023; Crenshaw, 1989). Though the ambit of non-discrimination law could be extended to cover novel forms of discrimination (Article 14 ECHR is an open-ended principle), the question has been raised as to the normative limits of non-discrimination law and the types of injustice it is meant to address (Bartoletti & Xenidis, 2023). However, if the overall goal is to ensure for each person an equal opportunity to participate as social and political equals within democratic processes, the law must provide the intervening mechanisms when that overall ambition is being threatened. In this regard, human rights could be positively harnessed to abolish data-driven structures of oppression and domination. In this context, positive actions must target historical patterns of discrimination, which for too long have operated "to exclude or marginalise the voice and influence of some groups while magnifying the influence of others" (Young, 2002, p. 34).

3.1.3 Conclusion

Though much more can be said about the limitations and challenges faced by non-discrimination law, the abovementioned observations indicate that decision-making structures informed by the output produced by GenAI should reflect on the capacity of the associated digital infrastructures to further institutionalise and entrench stereotypes and vulnerabilities, as these add to the political and socio-economic disadvantage experienced by people, and hence, the overall impact of the policy decisions under consideration (Malgieri & Niklas, 2020; Helberger et al., 2022; Naudts, 2024). In addition, we emphasise the importance of State's positive obligations to promote (public) oversight and accountability, whereby systemic risks to democracy are also assessed based on the cumulative harms GenAI can produce, and whereby private actors, under certain circumstances, are viewed as critical actors in shaping society's democratic structures.

3.2 GenAI & freedom of expression

In a functioning democracy, the rights to political participation and non-discrimination are invariably linked and interconnected with the right to freedom of expression, which, according to the ECtHR "constitutes one of the essential foundations of any democratic society, one of the basic conditions for its progress and for the development of every man" (*Handyside v. the United Kingdom*, 1976, para. 61). Or to speak with Dworkin, a fair democracy requires "that each citizen has not just a vote but a voice: a majority decision is not fair unless everyone has had a fair opportunity to express his or her attitudes or opinions or fears or tastes or presuppositions or prejudices or ideals, not just in the hope of influencing others (though that

hope is crucially important), but also just to confirm his or her standing as a responsible agent in, rather than a passive victim, of collective action.” (Dworkin, 1994). In the following, we will analyse some of the trends and practices described in section 2 from a freedom of expression perspective.

3.2.1 Adding the human rights dimension: GenAI from the perspective of the right to freedom of expression

The right to freedom of expression under Art. 10 ECHR includes the right to receive, hold and impart information and ideas, the protection of speech, including political speech, from interference and censorship, as well as institutional guarantees of the press and the conditions for its functioning, such as media pluralism. The types of speech Article 10 ECHR protects include text, but also images, videos and other forms of expression of ideas or information, irrespective of the form in which it is expressed (Bychawska-Siniarska, 2017). Importantly for the given context, freedom of expression can also extend to the (technical) means of producing and distributing information (Helberger et al., 2020) with the consequence that freedom of expression can also extend to the use of technical means, including AI-driven tools for the production and distribution of content. More recently, Bouché et al. have argued that GenAI, too, should be considered a communication technology as GenAI applications can “further narrow the boundaries between the editorial sphere and the material backend of media” (Bouché et al., 2025).

Opportunities and risks

As section 2 demonstrated very clearly, GenAI can give citizens as well as other societal actors a powerful tool to exercise their voice, and with that their right to freedom of expression. In a situation in which more and more citizens but also political parties or civil society representatives have access to powerful GenAI models, the barriers to create content to impart are significantly lowered. Translation and spell-checking services are an evident example how GenAI can help citizens to exercise their voice and equalise the opportunities to do so. Then there is the use of GenAI in citizen reporting or reporting in a way to circumvent algorithmic speech control on social media platforms (see section 2.1). GenAI has been successfully used as a digital strategy to avoid silencing and reclaim political presence. For example, in Venezuela, journalists have turned to GenAI to create virtual journalistic anchors. In so doing, journalists can ensure plurality in media reporting, while protecting their own identities from prosecution (Phillips & Torres, 2024). And new emerging forms of journalism, and of providing the audience access to journalism (see section 2.1) but also the use of GenAI in political campaigning (section 2.1 and 2.2) can also facilitate and even democratise access to information and political participation. If new, more engaging forms of offering media content invite parts of society who so far tended to avoid news or were even excluded (e.g. due to language barriers or disability), this is another win for freedom of expression. This presupposes, of course, that there is sufficient room for, and incentives to experiment with, and promote social innovation. This is because the opposite can also be true, and GenAI can be equally used to limit freedom of expression, for example through censorship, crowding out voices in a cacophony of synthetic content and disinformation, reducing trust in the public sphere and speech institutions. From a freedom of expression perspective it is therefore also problematic to note the apparent dissonance between, on the one hand, the growing professional use of GenAI in journalism and political practices and, on the other hand, the

many concerns that citizens still have in terms of the protection of their privacy, their inability to discern synthetic from non-synthetic content and concerns about abuse (section 2.1.3). AI empowerment goes beyond teaching citizens how to use GenAI, but includes involving them in the development of the solutions that must help them to advance their rights in the first place. This the more since GenAI is structurally reshaping communicative practices and the public sphere, from individual communicative practices to the shape of the underlying communication infrastructure.

The other dimension of free speech, freedom of thought, has so far received lesser attention (O’Callaghan et al., 2024) but will gain in importance to the extent that GenAI models become more persuasive, and are being used to persuade and manipulate. The use of GenAI in political campaigning is a point in case, but also in commercial advertising (as freedom of speech encompasses commercial communication). The Council of Europe has warned in its Declaration on the Manipulative Capabilities of algorithmic processes way before GenAI was widely available that “fine grained, sub-conscious and personalised levels of algorithmic persuasion may have significant effects on the cognitive autonomy of individual and their right form opinions and take independent decisions” (CoE, 2019). Remarkably the Council continued by observing: “Not only may they weaken the exercise and enjoyment of individual human rights, but they may lead to the corrosion of the very foundation of the Council of Europe. Its central pillars of human rights, democracy and the rule of law are grounded on the fundamental belief in the equality and dignity of all humans as independent moral agents.” As Bayer point out, forms of AI-assisted (micro)targeting can cause double harm to citizens and voters, through violating the rights to privacy, individual autonomy and the right not to receive information of some, as well as excluding others from access to media content or political communication (Bayer, 2020).

Institutional dimension of freedom of expression

The Council of Europe’s quote above is a powerful reminder that the right to freedom of expression is not only an individual right but also has an institutional dimension to it, as it also protects the institutions that enable citizens to exercise their right to freedom of expression in the first place. Most prominently that is the role of the free press and independent journalism, but also the democratic function of political parties. Under human rights doctrine, the press has traditionally held a special position in a democracy as it is upon the press to impart the information that the public has a right to receive (*Sunday Times vs The United Kingdom*, 1979). The resilience of a democracy is thereby invariably linked to the health of a society’s media ecosystem. And it is because of this democratic role that the freedom of the media, and individual journalists, to impart information is constitutionally protected (Oster, 2015, p. 22). GenAI, as sections 2.1.2 and 2.1.5 demonstrated, can affect not only concrete journalistic processes, such as researching an article, writing headlines or even production. It can also implicate the press as an institution, for example if its proliferation affects trust and democratic viability of the press, or if the practices of technology developers affect its economic sustainability. Seen through the lens of freedom of expression theory, it becomes clear that the challenges to the economic viability of journalism, and in particular smaller and independent outlets, is not only a problem of industrial policy but warning signal of (yet another) impending democratic crisis.

This is also because next to imparting information that the audience has a right to receive, providing a public sphere, and scrutinising those in power are often described as other democratic core functions (Wetenschappelijke Raad voor Regeringsbeleid (WRR), 2024). Both functions are under pressure, and even more so under the impact of GenAI. As far as the ability to scrutinise power is concerned, GenAI and its capacity to process unfathomable amounts of data gives investigative journalist a potent tool for new forms of data and AI-assisted journalism. Examples abound, like recent Pulitzer prizes for journalists that used GenAI in war reporting. And yet, the use of GenAI in journalism is a double-edged sword as it also comes, as we demonstrated in section 2.1.5, with deep structural dependencies on a handful of technology companies and the computational infrastructures under their control. The role of journalism to hold power to account extends to the companies that exercise more and more economic and political power in our democracy. And yet, there is a real danger that technical dependency translates into reduced autonomy and independence of the press (Schibstedt, 2024; van Drunen, 2025).

The various ways in which technology companies enable but also threaten such a public sphere have been extensively discussed in a number of Council of Europe recommendations (see also section 4.1). It is here that the dynamics of the concentration of economic, technological, political, and cultural power are most visibly at play – dynamics that will potentially be further re-enforced by the rise of GenAI, and resulting processes of vertical integration. As section 2.1.5 has shown, the very players that control some of the most influential social media platforms on earth are also the most powerful driver behind the development of leading GenAI systems and LLMs.

Policy and legal efforts to address these power imbalances and dependencies have so far offered little meaningful remedy. The largest media organisations have pursued their own protective measures, either by striking licensing deals with AI companies, as seen with Springer and OpenAI, or by filing lawsuits to defend their copyrights, such as the *New York Times* case against OpenAI and Microsoft.² Yet these strategies remain accessible only to well-resourced players, leaving smaller media without comparable means of protection. This re-emphasises the deeper structural issues explained before. Moreover, outdated media concentration regulations and the approval of media mergers, frequently justified as measures to prevent the disappearance of smaller outlets, may ultimately prove detrimental to the preservation of media pluralism in the long-term (Seipp et al., 2024). In this context, standard-setting and human rights bodies like the Council of Europe play a crucial role in developing frameworks that sustain journalism amid these technological and economic power shifts. Moving forward, policy should focus on fostering collaboration rather than consolidation, for instance, by publicly funding shared digital infrastructures, enabling collective bargaining among news organisations, especially local, establishing inclusive standards for technology procurement, and supporting public-interest tech start-ups that offer innovative, journalism-aligned solutions.

² https://intellectual-property-helpdesk.ec.europa.eu/news-events/news/openai-partners-news-publisher-axel-springer-meta-faces-claims-authors-infringing-copyright-2023-12-22_en ; <https://harvardlawreview.org/blog/2024/04/nyt-v-openai-the-timess-about-face/>

Duties and responsibilities

It is important to acknowledge that the media's lack of digital autonomy is in parts also the result of very strategic choices of the media to team up with Big Tech. From a human rights theory point of view, the decision of the media to use extends to the freedom to decide which technologies to use when the media exercises its democratic function (Helberger et al., 2020; Oster, 2015, p. 77). Insofar, the decision to use Facebook or YouTube as central distribution platform or build future information processing pipelines on the products and infrastructures of OpenAI or Microsoft are part of the media's freedom to decide how to inform and impart information. On the other hand, the choices that the media make can have implications that are not only relevant for the individual outlet, but for the shape and openness of our common public information infrastructure. This raises the question of whether when deciding in favour or against entering deals with Open AI & Co, the constitutional protection of the media also requires re-thinking existing duties and responsibilities of the press. Traditionally, those responsibilities would focus on the output of journalistic activity, e.g. that content is produced in line with the ethics of journalism (Bychawska-Siniarska, 2017; Wetenschappelijke Raad voor Regeringsbeleid (WRR), 2024). Moving forward, an argument could be made that the media's duties and responsibilities should extend beyond the responsible production of content and must also include the responsible use of technology and responsible technology choices. Elsewhere, we argued that duties and responsibilities can also include the obligation to develop journalistic ethics for the use of AI (Bastian et al., 2021), which is increasingly also happening (Becker et al., 2023; Cools & Diakopoulos, 2024). But duties and responsibilities can also be thought further to entail that the technology choices the media make contribute to creating a public information infrastructure that respects core democratic values (such as respect for human rights) but also wider societal interests, for instances, taking into account the information needs of a diverse citizenry, the protection of the environment or refraining from unfair exploitation of authors, media, workers and planet.

Pluralism and diversity

Pluralism is a central concept in freedom of expression theory and of democracy ("There can be no democracy without pluralism" (*Centro Europa 7 S.R.L. and Di Stefano v Italy*, 2012, p. 129). It is also a central concept to consider the implications from GenAI for democracy. On the level of GenAI models, the aspect of flattening diversity, and failure to represent cultural or linguistic diversity in both training data and the model, affects not only the quality of the output of GenAI models, but also risks embedding structural biases along the entire GenAI value chain.

On the level of GenAI systems and applications, like hyper-personalisation of news or targeted political messages, diversity is potentially at risk through reducing diverse exposure, particularly if citizens are presented with the one and convincingly-crafted answer and no need to seek further (see section 2.1), the invisibilisation of other perspectives, or biased answers (see section 2.2.5 demonstrating how chatbots can favour only a small selection of political parties).

Finally, on the structural level, the emergence, respectively re-enforcement of gatekeeper roles and concentration of economic, political and technological power in the hands of few are in direct conflict with basis principles of democracy. Or to speak in the words of Edwin Baker:

“[A] country is democratic only to the extent that the media, as well as elections, are structurally egalitarian and politically salient. The best institutional interpretation of this democratic vision of the public sphere is, I suggest, an egalitarian distribution of control... This is a democratic distribution principle for communicative power – a claim that democracy implies as wide as practical a dispersal of power within public discourse” (Baker, 2006).

3.2.2. The potential of freedom of expression doctrine to contribute to the democratic governance of AI

Creating a conceptual risk assessment framework

The right to freedom of expression can serve as an important benchmark to identify in which instances GenAI applications and systems are beneficial, or problematic, namely when their development and deployment undermine the conditions that are needed to be able to exercise freedom of expression rights. The right to freedom of expression also can help to explain why certain applications of GenAI deserve more regulatory attention than others, namely to the extent that they are closer to the core of the realisation of free speech rights. For example, constitutional speech protection also extends to speech that “shocks or offends” (*Handyside v. the United Kingdom*, 1976, para. 49). In that sense, even the use of GenAI to create polarising or even shocking content (see section 2.2) can still fall under the protection of free speech rights as it can be part of the democratic discourse. The border to impermissible speech is crossed where it risks tolerance and respect for the equal dignity of all human beings, which are, according to the Court, a foundation of a democratic and pluralistic society. According to the ECHR, “it may be necessary in democratic societies to sanction, or even prevent, all forms of expression that spread, incite, promote, or justify hatred based on intolerance” (*Féret vs Belgium*, 2009, para. 64). This is an example of how Art. 10 ECHR can help to identify forms of GenAI use that are societally not permissible and require political action exactly because they run contrary to fundamental principles of democracy.

The Council of Europe has already taken an important step towards concretising and systematising the freedom of expression implications from GenAI in its Guidance Note on Generative AI and Freedom of Expression (Council of Europe, 2025). The recommendation scrutinises the freedom of expression implications of GenAI for each level of the GenAI stack. This includes the foundational layer, and the way biased but also inaccurate information can affect the right to receive information or endanger linguistic diversity; the tool layer and the potential of content moderation and filtering to suppress and censor free speech, as well as the product layer where applications like chatbots or AI Agents can create access to information, but also result in manipulation, re-enforcement of stereotypes or inaccuracies, but also structural implications by, for example, failing to display sources and appropriating media content to develop competing gateways of information. Attempts to systematise and exemplify substantive challenges to the right to freedom of expression, both on an individual as well as institutional level, can inform risk assessment practices but also empower academics and civil society when auditing existing LLMs (see e.g. the interesting example of (Calvet-Bademunt et al., 2025, scrutinising national laws and corporate practices that governing AI upon their compatibility with freedom of expression). It could also enable the creation of effective monitoring and signalling framework either within the Council or through empowering and setting up monitoring networks of/by academics, civil society, fact checkers and regulators.

In a similar way, freedom of speech theory can inform us what responsible/public interest uses of GenAI are, for example when GenAI is being used to empower citizens to impart information, to reach groups in society that are typically excluded or less likely to be able to encounter independent journalism or political information, or to promote the accessibility of information as these are all actions that can create an enabling environment for the realisation of freedom of expression, and thereby might also fall under the positive protection duties of member states.

Foregrounding the perspective of public interest innovation

For governments, media freedom so far entailed both, an obligation to refrain from interference and censorship (GenAI creates new tools to do so) as well as positive obligations to create the conditions so that the media can exercise its media freedom and citizens benefit from freedom of expression (*Dink v. Turkey*, 2010). Arguably, this latter dimension – positive obligations – becomes increasingly more important because of the immense asymmetries of power between citizens, and some platforms and technology companies, as well as the transformative potential of GenAI. In the past decade and on various occasions the Council of Europe has emphasised and explained the role of states in creating a favourable information environment.

Traditionally, such “positive obligations” of states to protect and promote media freedom ranged from enabling a resilient and plural media landscape, establishing journalistic privileges such as access to publicly-held information or the protection of sources, provide them with financial and non-financial support as well as safeguarding the safety of journalists and of citizens to be able to express their ideas and opinions without fear of repression (Pentney, 2024). Considering the potential of GenAI to also contribute to the realisation of freedom of expression rights or the societal transformation of the press, such positive obligations could potentially also include measures to encourage and enable digital autonomy and alternative infrastructures, or incentivise public interest oriented media innovation, e.g. through funding program, tax incentives, sandboxes, extended public service mandates and concretising the role and mission of public service media, or other forms of incentive-based regulation. As Bouché et al argue, “[a]cknowledging the role of media (technology) infrastructure and law in shaping public formation, ideal expectations regarding the state’s role in guarding pluralism should account for the creation of both legal and infrastructural conditions that support the multiplicity of ‘publics’ and their different aspirations” (Bouché et al., forthcoming).

Next to states, platforms and technological companies themselves come into the picture as addressees of human rights obligations. While traditionally human rights would be effective in the relationship government citizens, more recent Council of Europe recommendations increasingly also include large technology providers, and social media platforms in particular, in the circle of those that are liable to respect and protect citizens’ human rights, including freedom of expression. The risk-based approach in the Framework Convention is a point in case, but as we will see later (section 4.2), the Framework Convention is more ambitious and explicitly also emphasises the need to realise the potential of AI technologies to create value for society through promoting and advancing the realisation of human rights and democratic principles.

4. Council of Europe legal framework

4.1 Situating GenAI in Council of Europe Recommendations

Council of Europe recommendations have a long history of adapting human right standards to account for technological change. Some instruments focus on specific technologies, such as the recommendations on intermediaries or algorithmic systems (CoE, 2018, 2020). More commonly, recommendations analyse the role of technology as part of a comprehensive analysis of a specific issue, such as hate speech or elections (CoE, 2022a, 2022b).

Council of Europe recommendations have not yet dealt with the impact of GenAI on democracy directly, although a Guidance note was developed in 2025 on the implications of GenAI for freedom of expression. Instead, recommendations have focused on technology used to distribute information. For example, recommendations have analysed the responsibilities of intermediaries to take down hate speech distributed through their systems (CoE, 2022b, para. 30), ensure their targeted advertising technologies respect users' privacy (CoE, 2022a, para. 5), and promote more diverse information and high-quality journalism (CoE, 2018, para. 2.5, 2022a, 2022b).

Recommendations have analysed precursors and technologies closely related to GenAI. For example, in 2017 the Committee of Ministers adopted a recommendation on Big Data for culture, literacy and democracy (CoE, 2017). Large volumes of data are, of course, a key resource enabling the current advances in GenAI, and the recommendation already called on Member States to improve user literacy and set up multistakeholder processes to guide the future of Big Data in line with human rights. However, the recommendation focused on archiving and improving the accessibility of cultural data, which is of limited relevance to the challenges and opportunities currently created by GenAI. The recommendation on the human rights impact of algorithmic systems arguably comes closest to engaging with GenAI directly. Its broad definition of algorithmic systems (which is also visible in other recommendations (CoE, 2022a para. 4, 2022c para. 7)) understands them as “applications that, often using mathematical optimisation techniques, perform one or more tasks such as gathering, combining, cleaning, sorting, classifying and inferring data, as well as selection, prioritisation, the making of recommendations and decision making” (CoE, 2020 para. 2). While this definition does not refer directly to algorithms that generate data, it does capture many of the steps (gathering, combining, inferring) in GenAI. More importantly, the substance of the recommendation itself emphasises the importance of the underlying infrastructure and datasets used in algorithmic systems (CoE, 2020 para. 2.2). As the above sections have shown, these technologies continue to play a key role in the context of GenAI.

In short, while GenAI is a new technology to Council of Europe recommendations, recommendations have a long history of adapting human rights standards to technological change, and have begun to show how to do so with technologies related to GenAI. The following section will draw on these lessons and the preceding analysis risks GenAI's risks and opportunities to see what we can learn from Council of Europe recommendations about what GenAI's role in democratic society should be.

4.1.1 CoE recommendations and the role of GenAI in democratic society

Harmful content

One key challenge posed by GenAI is that it can be used to create *false* messages that discriminate against specific societal groups (see section 3.1). As the sections on the use of GenAI in politics and journalism explained in depth, GenAI significantly lowers the barriers to produce content, including discriminatory or hateful content. For example, section 2.2 showed how political actors (particularly those that already have fewer normative concerns about using AI) can use the technology to quickly and easily produce more extreme organic content. In some cases, such content may amount to illegal hate speech. While less extreme content might not reach the required level of toxicity to be considered illegal hate speech or harassment, it can still negatively impact the right to non-discrimination and undermine tolerance (see also Hacker et al. 2024).

The Recommendation on hate speech already acknowledged the danger of more widespread exposure to hate speech distributed through intermediaries, noting “that the persistent availability of hate speech online exacerbates its impact” (CoE, 2022b). To address this issue, the recommendation emphasised the importance of ensuring swift take down procedures and due diligence requirements for intermediaries. Such procedures are arguably even more important as GenAI facilitates the large-scale production of hate speech, which can be distributed through intermediaries and amplified by algorithms (see political campaigning section). However, the argument that technology companies have a responsibility to ensure their services are not abused also applies to the companies whose AI systems are used to generate hate speech. Arguably, the argument applies with greater force in this context, given that genAI companies do not simply distribute hate speech uploaded by their users, but control how this content is produced through the design of their models. Combined, the recommendation on hate speech and recommendations that engage with technology companies’ human rights obligations more broadly provide many starting points to operationalise such a responsibility (CoE, 2018b, 2020, 2022a, 2022e). For example, GenAI companies could integrate human rights into their content policies (akin to intermediaries’ content moderation policies), hire staff with sufficient expertise in the relevant human rights standards, account for legal, local, cultural, socio-political and historical context, and put in place due diligence processes that ensure the safety of their systems. Throughout, the response to hate speech must be calibrated to its severity, and account for freedom of expression.

GenAI can also inadvertently amplify societal biases when it is used by actors that have no intention of producing harmful content. For example, as section 2.1 showed, journalists using GenAI to gather data or draft texts may unintentionally adopt stereotypes embedded in its training data. In that context, the recommendation on hate speech emphasises public officials, political parties and the media should avoid disseminating hate speech or expressions that foster intolerance (CoE, 2022b para. 28, 41). Regulation of these expressions is sensitive, as the freedom of media and political actors to express themselves, including in offensive ways, is protected strongly under art. 10 ECHR (see section 3.2). However, as noted, technology companies do have an obligation to ensure the safety of their systems - and this arguably includes a responsibility to ensure that their systems do not steer their users to unintentionally produce harmful content. In addition, literacy initiatives that alert users to the danger that

GenAI recreates harmful stereotypes may prevent unintentional misuse of GenAI. Finally, the recommendation on hate speech emphasises the importance of counter speech. Political actors can openly condemn the use of AI to produce hate speech, while the media can provide a counterweight by providing reliable and diverse information.

Manipulation

GenAI also has the potential to undermine individuals' ability to make informed decisions. As preceding sections have shown, GenAI can be used to manipulate individuals' political choices, for example by the creation of messages tailored to a specific individual's vulnerabilities and characteristics. Existing Council of Europe recommendations, particularly the Recommendation on electoral communication, Declaration on the manipulative capabilities of algorithmic processes, and the guidelines on the protection of personal data in political campaigns for Convention 108 already provide useful guidance on when such personalised political communication should be seen as impermissible (CoE, 2019, 2021a, 2022c). One key lesson for regulators in this context is to base any adaptations of these standards on the way GenAI is used for hyperpersonalised manipulation in practice. In the previous iteration of this debate, fears over the potential for highly effective microtargeted political advertisements were not borne out by the evidence on their actual use and impact. Taking a similar approach to GenAI would draw valuable regulatory attention away from risks that are arguably more pressing.

GenAI may also undermine individuals' ability to inform themselves. As section 2.2 showed, by lowering the barriers to content creation, GenAI enables actors to produce large volumes of low-quality information (also referred to as 'slop') as well as inaccurate information (also referred to as 'hallucinations'). From the perspective of Council of Europe recommendations, it is useful to distinguish between two different implications of the lowered barriers to content production. First, GenAI can be intentionally misused to undermine the democratic debate by allowing actors to more easily produce large volumes of disinformation. The recommendation on electoral communication already highlighted the danger of bots and fake accounts being used to spread disinformation during campaigns. To address these issues, the recommendation emphasised the importance of monitoring disinformation campaigns and called on intermediaries to ensure the integrity of their services while accounting for freedom of expression. GenAI increases the importance of such measures (CoE, 2023a, para. 21); in line with the analysis of harmful content above, it arguably also requires the providers of GenAI models to ensure the safety of their technology.

Second, GenAI facilitates the production of large volumes of low-quality information. The problem in that context is not so much that individuals rely on false information. It is that by making it easy to produce information that looks good but has little informative value, GenAI makes it much harder for individuals to find (and consequently for media organisations to produce and fund) high quality information. Concretely: investigative reporting costs a lot of time and money. Creating a good-looking video with GenAI does not. Going further, genAI allows AI companies to produce summaries of news reporting themselves. Even where such summaries link to the underlying news stories, there is no guarantee that (and indeed little incentive for) users to follow these links to the original source. And while AI companies have begun to license news media's content to use in their AI models, these deals have only benefited a few publishers, and their terms are not open to the public.

Previous Council of Europe recommendations have already warned extensively of the negative implications of “[t]he unrelenting pace of information sharing online” that is facilitated by intermediaries, noting it challenges the funding of quality journalism, the depth and accuracy of reporting, and individuals’ access to reliable information (CoE, 2018a, 2022d). Unfortunately, there is no clear-cut solution. While there is widespread recognition that technology companies’ control over attention (and associated revenues) for journalism is problematic, the question of how attention and revenue for the media should be organised instead remains controversial. This is particularly true for mechanisms that would require technology companies to ensure the visibility of media organisations’ content on their services (CoE, 2021b; OSCE, 2025). Nevertheless, the recommendation on a favourable environment on quality journalism does surface several policy tools that may be more easily adapted to improve the relationship between AI companies and media, including the need to update media concentration law (for example, to allow media organisations to negotiate collectively with AI companies), and ensure data sharing (so media organisations know how their content is used and monetised by AI companies) (CoE, 2022d). In addition to redressing the relationship between media and technology companies, states may use existing mechanisms to fund and produce journalism, such as tax benefits, direct funding of investigative journalism, and support for public service media (CoE, 2022d, para. 1.3). In many countries, such mechanisms are already available, and thus can potentially be scaled up relatively quickly. In addition, these mechanisms have the benefit that they strengthen the media’s direct relationship to its audiences, rather than increase their visibility on services controlled by AI companies.

Infrastructure

Infrastructure is a key but often hidden aspect of GenAI, as explored in detail in section 2.c. Control over the technological infrastructure needed to deploy and develop GenAI is centralised in the hands of a few technology companies. This allows these actors to shape the way GenAI can be used and developed, and potentially to cut off access to critical infrastructure. Technological infrastructure has thus far been a secondary topic in Council of Europe recommendations, being dealt with at in subsections (CoE, 2022e, para. 2.1-2.7) and in passing (CoE, 2020 para. 2.3, 2022a para. 11). However, from these recommendations and earlier work on media ownership and network neutrality (CoE, 2016, 2018a), several principles that should inform any future regulation of infrastructure can already be identified.

Council of Europe recommendations emphasise the importance of preventing the centralised control over infrastructure we are currently witnessing in the context of GenAI. Such centralisation can limit states’ ability to fulfil their human rights obligations (CoE, 2020 para. 2.3) and enable problematic concentrations of power over public opinion (CoE, 2022e para. 2.4). To limit such centralised control, standards note the importance of effective (ex-ante) application and amendments to competition and media concentration law (CoE, 2018a, 2022e, para. 2.4, 2022a para. 11). In addition, recommendations note the importance of stimulating innovation that leads to alternative infrastructures (CoE, 2020 para. 2.3), for example through state support for the public service media or through investment by the media itself (CoE, 2022d para. 2.4.2, 2022e para. 2.6).

Furthermore, states can require that infrastructure is designed to be open, interoperable, and allow for data portability. Such measures can prevent user lock-in, which not only hampers competition but also leaves users vulnerable to infrastructure providers raising their prices or abusing their powerful position (CoE, 2022e para. 2.3). Finally, where users are dependent on an infrastructure provider, states can intervene to ensure fair, reasonable, and non-discriminatory access conditions (CoE, 2022a para. 11, 2022e para. 2.2). Such measures can prevent (or at least redress) retaliation by infrastructure providers against specific actors based on their political viewpoints; more broadly, they also help ensure that the infrastructure needed to use GenAI to contribute to the public debate is available to everyone (CoE, 2016).

Opportunities of GenAI

Council of Europe recommendations also consistently emphasise the importance of going beyond addressing the risks posed by new technologies, and actively create the conditions for technology to support democracy. Recommendations have for example highlighted the potential for digital technologies to increase the range and diversity of information to which individuals are exposed and the creation of new networks of solidarity (CoE, 2018a, 2022e), and emphasised the need for states to create a “safe and enabling environment for public debate and communication” as well as a favourable environment for quality journalism (CoE, 2022b, 2022d).

One keyway for states to operationalise such positive obligations is to facilitate the opportunities brought by GenAI. Such an approach will not prevent actors that intentionally use GenAI to undermine democratic values from doing so. However, it can prevent unintentional irresponsible use of genAI, and make it easier for actors (such as media organisations and political parties) who want to use GenAI in a way that supports democracy to do so. One example that has been noted in different contexts throughout this report, is that GenAI can be used to make information more accessible by adapting it to an individual’s specific needs or cognitive/physical impairments.

States could facilitate the responsible use of GenAI in a number of ways. One is to issue concrete guidance on how genAI impacts the media’s responsibilities, and how media organisations can go about implementing genAI responsibly (see extensively CoE, 2023b). Another is to stimulate innovation that seizes such opportunities (CoE, 2023b). In line with the preceding section, such a strategy should also account for the underlying infrastructure on which GenAI relies, which can steer the type of innovation which is possible. Of course, creating the conditions for a healthy democracy is not limited to incentivising better use of genAI. On the contrary, given the difficulties genAI presents for individuals to distinguish authentic and inauthentic content it is arguably all the more important to invest in traditional measures to safeguard democracy. Particularly noteworthy in this context is the role of the media to provide reliable, high-quality information. Such information can provide an important counterweight to the large volumes of low-quality or inaccurate information that can be generated with AI.

4.1.2 Conclusion

Council of Europe standards are arguably a particularly useful policy tool to address the challenges posed by AI. As explored above, genAI’s main challenge to democracy arguably

does not lie in new, highly impactful content such as deepfakes and synthetic illegal hate speech. Rather, it is (1) the large volumes of low-quality, polarising, and discriminatory content that genAI makes it possible to produce, and (2) the expansion of technology companies' power to the production of content and underlying infrastructure, and corresponding loss of control by the media over its technological innovation, funding, and access to audiences.

Council of Europe recommendations can help address these issues by laying out a more positive vision, which goes beyond trying to mitigate genAI's excesses and clarifies what AI that supports democracy may look like, and what actions the different stakeholders (including states, media, and AI companies themselves) can take to realise this vision. Providing such a vision is useful in and of itself, to push back against the framing of AI as the vehicle for innovation and economic growth that must not be weighed down by regulations that address its impact on democratic society. But it can also ground ongoing efforts to make AI more responsible in European human rights standards. Without such guidance, there is considerable risk that such efforts prioritise immediate economic interests over abstract notions of human rights, or fail to do justice to the nuance of European human rights captured in Council of Europe standards.

4.2 Framework Convention on Artificial Intelligence

The Council of Europe's Framework Convention on Artificial Intelligence is the first-ever international for Parties legally binding treaty to ensure that AI systems respect human rights, democracy and the rule of law. Similar to the European Union AI Act and other comparable frameworks (like the pending Brazilian AI law), the goal of the Framework Convention is to create a horizontal regulatory framework that can apply across a range of sectors to ensure responsible development and deployment of AI. There are, however, a number of important differences that can potentially make the Framework Convention a powerful instrument to advance human rights, democracy and the rule of law in an age of GenAI.

4.2.1 Human rights and democracy as point of departure

The advancement of democracy is the central goal of the Framework Convention. In so doing it is grounded in a long tradition of the Council of Europe as a human rights organisation. Elsewhere, the Council of Europe writes: "The connection between human rights and democracy is deep, and goes both ways: each is in some way dependent on the other, and incomplete without the other."³ This distinguishes the Framework Convention from other AI governance frameworks that foreground economic and market regulatory aspects (see also para 5 Explanatory Memorandum to the Framework Convention). The European AI Act is a point in case: grounded in product liability regulation, the AI Act also aims to protect democracy and human rights, but this is only one out of several objectives. In addition, the AI Act must balance the protection of human rights and democracy against other goals, such as supporting the internal market and innovation. Effectively, this makes the protection of human rights and democracy one out of many objectives to optimise for, and vulnerable to economic and geo-political pressure. In addition, product liability law does not have a tradition in the protection of human rights protection or of democracy.

³ <https://www.coe.int/en/web/compass/democracy>, visited on 12 November 2025.

4.2.2 Risks to human rights and democracy

Like the European AI Act and comparable governance approaches elsewhere, the Framework Convention follows a risk-based approach. According to Article 16 of the Framework Convention, member states must handle, adopt, or maintain a risk management framework, and do so with an explicit focus on human rights, democracy and the rule of law. Possible examples of risk can be the use of GenAI to instil hate, or maintain and reinforce stereotypes and structural forms of discrimination. In the case law of the ECtHR, which is developing on the issues related to AI and remains limited with regard to GenAI, and in the Council of Europe's recommendations, some general guidance can be found when AI systems and/or GenAI conflict with human rights and democracy. An example of more specific, AI-related guidance is the Guidance Note on the Implications of Generative Artificial Intelligence for Freedom of Expression elaborated by the Steering Committee on Media and Information Society in December 2025 (Council of Europe, 2025). This guidance note systematises specific challenges that AI systems and GenAI can trigger to freedom of expression. Therefore, building on and expanding AI specific guidance for other human rights and democratic principles (including situations where the application of GenAI is incompatible with human rights or democratic principles) would be a worthwhile path forward, and one in which the Council of Europe can build on its considerable expertise and specific human rights mandate. The development of the HUDERIA methodology as a tool that has been designed to specifically identify and mitigate risks to the protection and promotion of human rights, democracy and the rule of law, and the HUDERIA Model with concrete steps to operationalise the Methodology is an important step in this direction.⁴

4.2.3 Positive obligations to protect the realisation of human rights and democracy

Unlike other governance instruments, like the AI Act, the Framework Convention does not stop at addressing risks to human rights and democracy. Article 10 of the Framework Convention calls on parties to not only ensure that AI is respectful of equality but also to adopt measures to overcome inequalities to achieve fair, just and equitable outcomes. As section 3.2 of this study argued, tackling the underlying roots of discriminatory, stereotyping or exclusionary practices is essential to address the discriminatory effects of GenAI, exactly because of the way the technology influences and is influenced by underlying societal dynamics and data. In a similar vein, Article 7 of the Framework Convention requires measures to respect and promote human dignity and individual autonomy, which is critical for addressing the manipulative effects of (hyper)personalised services. After all, these effects cannot be seen in isolation from underlying business models that are optimised to exploit individual vulnerabilities and reduce users to profitable sources of data and advertisement revenue. This is in line with the approach of the Framework Convention to not only focus on AI systems in a narrow, technical sense, but to situate these systems in their broader socio-economic context (see Preambles).

As this report demonstrates, the development and deployment of GenAI can benefit democratic institutions such as journalism and an independent press, but it can also undermine its integrity and independence. The concentration of control over key elements of

⁴ More information on the methodology and the model can be found here: <https://rm.coe.int/huderia-methodology-and-model-adopted-provisional-version-2026/48802ac001>

the GenAI infrastructure, the dependency of journalistic institutions on a handful of commercially operating, profit-driven companies and the significant power asymmetries can not only challenge journalism or political communication as democratic practices and institutions, but also the balance of power in a democracy. The importance of preventing centralised control over critical communication infrastructures has already been an aspect that the Council of Europe emphasised in its recommendations (see section 4.1). Accordingly, developing strategies to reduce this dependency, for example by incentivising alternatives, but also calling upon the responsibility of all parties (including technology companies and journalistic institutions) must be an important point of attention in the realisation of the obligations from this Convention.

4.2.4 Public interest innovation

Another important difference is the Framework Convention's approach to concentrate not only on the potential negative implications and risks from AI, but also emphasise the importance of actively promoting a positive contribution of innovation to the realisation of human rights, democracy and the rule of law (Article 13 and Explanatory Memorandum, para 90).

In other words, underlying the philosophy of the Convention is not only the wish to protect citizens and society from harmful or irresponsible AI, but to promote the use of AI to advance human rights and democracy. For example, the Explanatory Memorandum to the Convention explains that the drafters explicitly included references to 'human dignity and individual autonomy as foundational values and principles that are essential for the full realisation of human rights, democracy and the rule of law' as well as undertaking measures aimed at overcoming structural and historical inequalities that can feed into, and cause algorithmic bias (Explanatory Memorandum, para 77). While the goal of risk-based approaches is to protect society from unsafe AI, a value-based approach like the Framework Convention can create powerful incentives to drive the development of public interest AI. Through the focus on human rights, concrete guidance (on what those rights entail in relation to (gen) AI and more generally, the experience and expertise of the Council of Europe in analysing human rights in the light of technological developments, the Convention could also form an important basis for human rights standard setting and substantive guidance (also, for example, in the context of the application of the AI Act that requires a more procedural approach).

Positive measures to advance and promote public interest innovation can take various forms, ranging from the promotion of inclusive co-design strategies, the adoption of fairness metrics, and performance of equality impact assessments, to strengthening the role of civil society to better enable political mobilisation and knowledge dissemination or public procurement guidelines. In this context, assessments made during the ideation, development or procurement of technologies, including potential risk assessments, should be made publicly accessible to foster public accountability and democratic contestability. Defining what exactly "public interest" innovation entails could be an important means to support member states and private parties in driving demand and adopting AI systems that not only 'do no harm' but are dedicated to advancing human rights and democracy. In more practical terms, enabling the findability of, and access to technical solutions that are in line with human rights standards, like the UNESCO tools radar,⁵ can be important means to help in particularly smaller providers

⁵ <https://radar.ircai.org/en/about/>

and independent journalists to contribute to public interest innovation. Championing public interest innovation would also be an important, and much needed counter-narrative to other more technological and economic understandings of innovation that currently dominate much of the global AI debate.

4.2.5 Democratising AI Governance

AI has been referred to as a ‘systems technology’ (Wetenschappelijke Raad voor Regeringsbeleid (WRR), 2021) to highlight the potentially far-reaching and thorough systemic impact of AI (and GenAI) on large parts of society and societal processes, be that in the area of labour, commerce, manufacturing, politics, health, justice, journalism, education and many *other* sectors. Seeing the transformative potential of AI, it is not only important that AI respects and advances human rights, democracy and the rule of law; but also the governance of AI has to be more democratic and inclusive. This is important both on its own merits (to ensure those affected by a societally transformative technology have a say in how it is governed) *and* to improve democratic governance and accountability of the development and deployment of GenAI (by bringing in the insights of those affected by the technology in practice).

The making of the European Union’s AI Act is a case in point of how exactly because of the transformative potential of the technology, the legislative process can be characterised by immense power asymmetries, industry lobby and difficulty of reaching and including a broader group of societal stakeholders and, importantly, affected citizens in the decision-making process - despite all good intentions. Developing governance approaches that are inclusive, participatory and guaranteeing meaningful involvement of groups with less negotiation power or that are traditionally underrepresented in these debates will be paramount to democratic AI governance. The Framework Convention could be an important starting point to develop and implement principles of more democratic AI governance.

According to Article 19 of the Framework Convention, “[e]ach Party shall seek to ensure that important questions raised in relation to artificial intelligence systems are, as appropriate, duly considered through public discussion and multistakeholder consultation in the light of social, economic, legal, ethical, environmental and other relevant implications”. This provision is an important building block of a more democratic approach to AI governance as it requires member states to create fora and processes for including citizens and other societal stakeholders not only on the governance of AI technology strict sense but also the broader societal implications of governance choices. Moving forward, there is a significant challenge to make these processes inclusive, meaningful and effective in practice.

In this context, Member States should promote meaningful citizen engagement, in the form of participation, representation and contestation mechanisms, throughout the GenAI life cycle to ensure truly human-centric driven AI. The latter might also involve the resourcing of equality bodies, discrimination ombudspersons or national human rights institutions with the mandate and powers to identify, monitor and measure the impact of GenAI on society, among others by granting them access to the necessary documents related to the data, design choices and organisational structures surrounding tech development and deployment, a mandate currently missing in their toolbox (Bartoletti & Xenidis, 2023, p. 70; On the role of equality bodies, national human rights structures, and the legal protection against algorithmic discrimination,

see also: Shrishak & Pénicaud, 2025; Xenidis, 2025). Beyond technology-oriented measures, digital literacy campaigns can play a critical role too.

Making AI governance more democratic therefore also means addressing the existing power asymmetries that amplify some voices or narratives over others. In that context, academics and civil society act as an important source of counterpower as they help individuals and members of communities articulate experiences and perceptions that are not part of dominant political and economic discourse, and through collective efforts, can help expose injustice in both state and economic power (Young, 2002, p. 155). To fulfil this role, it is key that academia and civil society also critically reflect on whether and how they can sufficiently account for the perspectives of those that are affected by GenAI but have little say in how it is governed.

Next to creating processes and procedures to include societal stakeholders in AI governance upfront, creating an effective responsibility framework and concrete rights for affected parties to contest particular instances of AI development or deployment that conflict with their human rights or the rule of law create an important societal feedback loop, and are indispensable for democratic oversight. Again, and unlike other regulatory frameworks, the Framework Convention provides critical building blocks, including the requirement for parties to ensure accountability and responsibility for the adverse impacts on human rights, democracy and the rule of law (Art. 9 Framework Convention) as well as the obligation for states to ensure transparency, access to data and the effective possibility for persons concerned to lodge a complaint (Art. 14 Framework Convention) and offer concrete rights and procedural safeguards to persons affected (Art. 15 Framework Convention).

Finally, whether or not governments and supra-national regulators will be able to govern AI (and the economic and political interests behind it), depends on their political will as well as their ability to govern. Both are under pressure. Over the past decade, the relationship between regulators and technology companies has changed substantially due to the growth in political and economic power of some of the leading technology companies, the extent to which society has become reliant on their technologies, the level of *infrastructural investment* and loss of control over the production environment; and the gradual loss of digital autonomy. How much this constellation is affecting national governments' capacity to govern is exemplified, again, by current debates in the EU where technology companies exercise pressure on the European Commission to pause or even repeal existing regulation. Therefore, requiring governments to protect their own digital autonomy and ability to govern and reducing strategic dependency on dominant technology companies is another critical element of democratic AI governance.

4.2.6 Conclusion

There could be an important role for the Council of Europe's Steering Committee on Democracy (CDDem) to advance a human-rights and democracy-oriented counter narrative to dominant narratives of innovation, speed and scale, and the Framework Convention could be an important instrument for member states to work with the Committee towards structural and individual safeguards to advance a democratic and societal perspective on GenAI and GenAI governance. This also means that the Framework Convention allocates a more prominent role for states to proactively address risks and impacts to human rights, democracy and the rule of law, including in relation to private actors. Doing so requires a framework to

operationalise the democratic governance of AI, grounded in human rights law and recalling that human rights not only include the obligation to refrain from interference but also create the conditions for individuals' ability to exercise their human rights.

5. Overall conclusions and recommendations

This report has outlined some of the primary ways in which GenAI transforms not only democratic practices, such as citizens informing themselves, political campaigning or the involvement of civil society and societal actors in decision making, but also democratic institutions, such as journalism as well as the broader public information infrastructure. It has highlighted some of the opportunities, and given examples of practices that are at odds or even harmful for democracy and the realisation of human rights, such as the right to non-discrimination or freedom of expression. Finally, it argued that GenAI and its implications for human rights and democracy must be seen in the context of the configurations of infrastructural and political power, and who is able and entitled to decide the course and character of innovation.

From an analysis of some of the most salient human rights implications of GenAI and the existing body of recommendations as well as the Framework Convention it becomes clear that the challenge moving forward is two-fold: addressing the risks, *and* realising the opportunities from GenAI for the advancement of human rights and democratic practices and institutions. Seeing the transformative potential of AI as well as the fact that this potential cannot be realised without drawing substantially on societal resources, be they labour, content, environmental or personal data, the notion of "innovation" with GenAI needs to be interpreted in the sense of public interest innovation where the interests of society as a whole outweigh purely economic interests of a handful of technology companies. With its focus on human rights, democracy and the rule of law, the Council of Europe's Framework Convention could be an important starting point to realise such a more societal perspective on AI innovation, and ensuring that it benefits society as a whole.

In this context, the Council of Europe, and within the scope of its mandate, the CDDEM could have an important role in helping to concretise the role and operationalisation of the Framework Convention for democracy, and defining the conditions for realising democratic opportunities from GenAI, while strengthening and protecting democratic institutions, processes and governance frameworks. In the following, and based on our analysis, we make three recommendations for possible policy priorities for the Organisation and its bodies.

Substantive guidance

Unlike other international governance instruments, the Council of Europe's Framework Convention on AI focuses specifically on addressing the risks and opportunities from AI for democracy, human rights and the rule of law. This also insulates the Convention at least to some extent from the pressure to lower its standards for the sake of purely technical and profit-driven innovation. Seeing the mandate of the Council of Europe as a human rights organisation, the extensive expertise and rich body of work (e.g. in the form of various relevant recommendations and the case law of the ECtHR), the Council of Europe is particularly well prepared to develop more substantive guidance on when GenAI is posing risks to human rights and democracy, and which uses of GenAI are unacceptable from a human rights

standpoint. In this report we highlighted in particular the danger of flattening diversity, structural inequality, manipulation, undermining trust and the practical and institutional conditions for the realisation of freedom of expression. These risks are further aggravated by infrastructural dependencies and considerable power asymmetries as well as the prominent role of private, profit-driven actors that are not subject to fundamental rights obligations and rule of law safeguards in deciding which risks are (un)acceptable for society. With the HUDERIA methodology and model and recommendations like the Guidance Note on Generative AI and Freedom of Expression, the Council of Europe has already made important inroads into providing more substantive guidance. The Council of Europe could have an important role in identifying, concretising and monitoring risks to democracy and political equality, be that in the way GenAI is used in democratic practices like elections, the more structural implications for democratic institutions (such as the free press) as well as the democratic governance of AI.

In addition to substantive guidance on when GenAI is at odds with human rights and democratic principles, guidance on when and how AI can advance democracy and human rights is much needed. In particular, there is currently little guidance on what “public interest” innovation exactly means, or which values to prioritise in public procurement policies. Also this question would benefit from substantive guidance, there may be scope for work on standardisation and helping human rights bodies.

Next to substantive guidance, the Council of Europe could also play a role in standardisation and helping human rights bodies and governments across the member states to agree on standards of public interest innovation and public procurement.

Incentivise and enable public interest innovation

In recent policy and regulatory debates around GenAI, the focus has been strongly on risk management and making these technologies safer and more trustworthy while still economically viable. Comparatively less attention has been dedicated to governance approaches that enable and incentivise public interest innovation, in the sense of GenAI development and deployment to advance human rights and democratic principles. An important contribution of the Framework Convention that distinguishes it from other regulatory frameworks is that it goes beyond a merely risk-based approach, and requires member states to promote the use of AI to advance human rights and democracy. This is in line with existing human rights doctrine and the positive obligations of states to create the conditions so that citizens can benefit from human rights.

The Council of Europe could play an important role in making more visible the democratic and human rights implications of procurement decisions, and helping to identify public interest innovators, the obstacles to public interest innovation and what is needed to overcome them. It could investigate mechanisms and regulatory solutions to incentivise and support the realisation of resilient public interest alternatives and investment in high quality trust institutions, like local media, public service media and media outlets that focus on high quality output, openness and standards of integrity - whether they have the form of traditional media, or are new forms of media, enabled by AI. And it could also bring in an important alternative narrative counterpoint to dominant technology narratives, by prioritising value-oriented

innovation, open source datasets, decentralisation and interoperability above the promotion of scale and innovation driven by a concentrated number of hyper-scalers.

Strengthen democratic resilience and ability to govern

Democracy is also about the ability of a society to self-determine the role of technology in society, and the values it adheres to. With the growing integration of a handful of proprietary, profit-driven technology solution in democratic institutions and the resulting dependencies, and the growing political influence of technology companies, we distinguished five dimensions of democratic self-determination and governance of technology:

- The use of GenAI in *elections and processes of democratic decision making* and the implications for inclusion, equal opportunities to participate and benefit from political freedoms
- The integration into *democratic infrastructures*, and the extent to which GenAI can contribute or risk the quality of the public sphere and the resilience of democratic institutions
- Susceptibility and responsiveness of technology companies and downstream deployers of GenAI to *democratic accountability, observability and control*
- The ability of governments and regulatory authorities to *autonomous decision making*
- The *inclusivity and representativeness* of democratic decision making and ability to mitigate power asymmetries and digital vulnerability.

Monitoring the impact of GenAI on these five dimensions of democratic decision making and helping governments to develop a strategy to preserve their digital autonomy and promote democratic resilience could develop into an important focal point of the Council of Europe and the CDDEM.

Concretely, this could include monitoring and developing guidance for the responsible use of GenAI in elections, mapping existing dependencies, and exploring mechanisms and regulatory solutions to encourage the development of European-based, infrastructural solutions that offer an alternative to the ecosystems of hyper-scalers. Other similar measures could be a re-consideration of the remit of public interest media, to also include the development of mission-critical elements of the European digital ecosystem, and working with member states towards favourable provisions for the development of local and community-led alternatives to Big Tech products.

A different strand could include helping member states to develop ways to empower public accountability watchdog actors, like academics, civil society, investigative reporters. Examples could be through fostering collaborations, creating feedback loops between public accountability actors, companies and regulators, funding and legal frameworks that enable strategic litigation. In this context, the important and urgent need for transparency needs to be re-iterated, on the individual, institutional and societal level. As the first part of the report has demonstrated, the lack of transparency and observability is a major obstacle to democratic accountability, and runs counter to the very explicit expectations and needs of citizens that must navigate the digital environment. This leads us to our final recommendation:

So far, much of the AI governance debate has been led above the heads of citizens, both in terms of developing GenAI based solutions and practices, as well as governance frameworks. For the governance of GenAI to become more democratic, it is necessary to develop ways to promote meaningful citizen engagement (participation, representation and contestation) throughout the GenAI life cycle, paying particular attention to the structural inequalities faced by affected persons and (historically) disadvantaged groups. The Framework Convention already includes a set of important obligations in this context, ranging from the obligation to consider the perspectives of those whose rights may be impacted in risk and impact management frameworks (Art. 16 (2)(c) Framework Convention), to the demand for organising public consultations (Art. 19 Framework Convention) and the importance of digital literacy (Art. 20 Framework Convention), transparency (Art. 14 (2) Framework Convention) and concrete rights and procedural safeguards for citizens to ensure contestability (Art. 15 Framework Convention). Considering the concrete interests and rights of citizens in the development of European digital infrastructure, the use of GenAI in democratic decision making processes, and the quality of democratic decision making about GenAI could be another important contribution of the CDDEM.

References

- Abid, A., Farooqi, M., & Zou, J. (2021). Persistent Anti-Muslim Bias in Large Language Models. *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society, AIES '21*, 298–306. <https://doi.org/10.1145/3461702.3462624>
- Amoore, L., Campolo, A., Jacobsen, B., & Rella, L. (2024). A world model: On the political logics of generative AI. *Political Geography*, 113, 103134. <https://doi.org/10.1016/j.polgeo.2024.103134>
- Arc XP. (2024). Incorporating AI in the Newsroom: Le Parisien's Journey to Innovation and Efficiency. *Arc XP*. <https://arcmarketing-arc-marketing-prod-staging.web.arc-cdn.net/2024/07/22/incorporating-ai-in-the-newsroom-le-parisiens-journey-to-innovation-and-efficiency/>
- Argyle, L. P., Bail, C. A., Busby, E. C., Gubler, J. R., Howe, T., Rytting, C., Sorensen, T., & Wingate, D. (2023). Leveraging AI for democratic discourse: Chat interventions can improve online political conversations at scale. *Proceedings of the National Academy of Sciences*, 120(41), e2311627120. <https://doi.org/10.1073/pnas.2311627120>
- Atrey, S., & Dunne, P. (Eds). (2020). *Intersectionality and Human Rights Law*. Hart Publishing. <https://doi.org/10.5040/9781509935321>
- Autoriteit Persoonsgegevens. (2025, October 21). *AP waarschuwt: Chatbots geven vertekend stemadvies | Autoriteit Persoonsgegevens*. <https://www.autoriteitpersoonsgegevens.nl/actueel/ap-waarschuwt-chatbots-geven-vertekend-stemadvies>
- Autorité de la concurrence. (2024). *Avis relatif au fonctionnement concurrentiel du secteur de l'intelligence artificielle générative* (Nos 24-A-05). <https://www.autoritedelaconcurrence.fr/fr/avis/relatif-au-fonctionnement-concurrentiel-du-secteur-de-lintelligence-artificielle-generative>
- Baker, E. (2006). *Media Concentration and Democracy: Why Ownership Matters*. Cambridge University Press.
- Balayn, A., & Gürses, S. (2024, September 30). Misguided: AI regulation needs a shift in focus | Internet Policy Review. *Internet Policy Review*. <https://policyreview.info/articles/news/misguided-ai-regulation-needs-shift/1796>
- Bartoletti, I., & Xenidis, R. (2023). *Study on the impact of artificial intelligence systems, their potential for promoting equality, including gender equality, and the risks they may cause in relation to non-discrimination* (p. 94). Council of Europe. <https://edoc.coe.int/fr/intelligence-artificielle/11649-study-on-the-impact-of-artificial-intelligence-systems-their-potential-for-promoting-equality-including-gender-equality-and-the-risks-they-may-cause-in-relation-to-non-discrimination.html>
- Bastian, M., Helberger, N., & Makhortykh, M. (2021). Safeguarding the Journalistic DNA: Attitudes towards the Role of Professional Values in Algorithmic News Recommender Designs. *Digital Journalism*, 9(6), 835–863. <https://doi.org/10.1080/21670811.2021.1912622>
- Bayer, J. (2020). Double harm to voters: Data-driven micro-targeting and democratic public discourse. *Internet Policy Review*, 9(1). <https://policyreview.info/articles/analysis/double-harm-voters-data-driven-micro-targeting-and-democratic-public-discourse>

- Becker, K. B., Simon, F. M., & Crum, C. (2023). *Policies in Parallel? A Comparative Study of Journalistic AI Policies in 52 Global News Organisations* [SocArXiv].
<https://osf.io/preprints/socarxiv/c4af9>
- Beckett, C., & Yaseen, M. (2023). *Generating Change. A Global Survey of what news organisations are doing with AI*. Journalism AI.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623.
<https://doi.org/10.1145/3442188.3445922>
- Bhattacharya, A. (2024, September 19). Political campaigns embrace AI to reach voters across language barriers. *Rest of World*. <https://restofworld.org/2024/aapi-victory-alliance-ai-voter-outreach/>
- Birhane, A., Dehdashtian, S., Prabhu, V., & Boddeti, V. (2024). The Dark Side of Dataset Scaling: Evaluating Racial Classification in Multimodal Models. *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1229–1244.
<https://doi.org/10.1145/3630106.3658968>
- Birhane, A., Prabhu, V. U., & Kahembwe, E. (2021). *Multimodal datasets: Misogyny, pornography, and malignant stereotypes* (arXiv:2110.01963). arXiv.
<https://doi.org/10.48550/arXiv.2110.01963>
- Blagojev, T., Da Costa Leite Borges, D., Brogi, E., Kermer, J. E., Trevisan, M., & Verza, S. (2023). *News Deserts in Europe: Assessing Risks for Local and Community Media in the 27 EU Member States* (EUI, RSC, Preliminary Report). Centre for Media Pluralism and Media Freedom (CMPF).
- Bonin, A.-K., Pecher, M., Peter, T., Schönacher, N., & Wettinger, A. (2025). Between Authenticity and Automation: Nonprofit-Organizations Navigating AI-driven Communication. In A. Godulla, F. Decker, L. Hartmann, C. Pankoke, M. Pecher, S. Pütter, & C. Santangelo (Eds), *Strategic Communication in Disruptive Times: How Sociopolitical Polarization, Virtual Media and AI Reshape Organizational Communication* (pp. 93–113).
- Borchardt, A. (2022). Go, Robots, Go! The Value and Challenges of Artificial Intelligence for Local Journalism. *Digital Journalism*, 10(10), 1919–1924.
<https://doi.org/10.1080/21670811.2022.2149584>
- Bouché, G., Helberger, N., de Rijke, M., & van Drunen, M. Z. (2025). LLMs as media technology: Implications under article 10 ECHR. *International Journal of Law and Information Technology*, 33, eaaf027. <https://doi.org/10.1093/ijlit/eaaf027>
- Braun, A. F. (2025, April 8). Independent Media has an Infrastructural Problem. *Tech Policy Press*. <https://techpolicy.press/independent-media-has-an-infrastructural-problem>
- Brennan, K., Kak, A., & Myers West, S. (2025). *Artificial Power: 2025 Landscape Report*. AI Now Institute. <https://ainowinstitute.org/publications/research/ai-now-2025-landscape-report>
- Bychawska-Siniarska, D. (2017). *Protecting the right to Freedom of Expression under the European Convention on Human Rights—A handbook for legal practioners*. Council of Europe.
- Calvet-Bademunt, J., Mchangama, J., & Anzabi, I. (2025). *That Violates My Policies> AI Laws, Chatbots, and The Future of Expression*. The Future of Speech.
<https://futurefreespeech.org/wp-content/uploads/2025/10/AI-Report-2025-Full-Report.pdf>

- Caswell, D. (2024). Audiences, automation, and AI: From structured news to language models. *AI Magazine*, 45(2), 174–186. <https://doi.org/10.1002/aaai.12168>
- Chander, S., & Gürses, S. (2024, October 15). From Infrastructural Power to Redistribution: How the EU's Digital Agenda Cements Securitization and Computational Infrastructures (and How We Build Otherwise). *AI Now Institute*. <https://ainowinstitute.org/publications/from-infrastructural-power-to-redistribution-how-the-eus-digital-agenda-cements-securitization-and-computational-infrastructures-and-how-we-build-otherwise>
- Chen, J., Dong, H., Wang, X., Feng, F., Wang, M., & He, X. (2021). *Bias and Debias in Recommender System: A Survey and Future Directions* (arXiv:2010.03240). arXiv. <https://doi.org/10.48550/arXiv.2010.03240>
- Cohen, J. E. (2019). *Between truth and power: The legal constructions of informational capitalism*. Oxford University Press.
- Cools, H., & de Vreese, C. H. (2025). From Automation to Transformation with AI-Tools: Exploring the Professional Norms and the Perceptions of Responsible AI in a News Organization. *Digital Journalism*, 1–20. <https://doi.org/10.1080/21670811.2025.2505982>
- Cools, H., & Diakopoulos, N. (2024). Uses of Generative AI in the Newsroom: Mapping Journalists' Perceptions of Perils and Possibilities. *Journalism Practice*, 1–19. <https://doi.org/10.1080/17512786.2024.2394558>
- Costello, T. H., Pennycook, G., & Rand, D. G. (2024). Durably reducing conspiracy beliefs through dialogues with AI. *Science*, 385(6714), eadq1814. <https://doi.org/10.1126/science.adq1814>
- Crenshaw, K. (1989). *Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics*. 31.
- Danaher, J. (2024). Generative AI and the future of equality norms. *Cognition*, 251, 105906. <https://doi.org/10.1016/j.cognition.2024.105906>
- de Brún, P., & Sawyeddollah, M. (2025, February 17). *Meta's New Content Policies Risk Fueling More Mass Violence And Genocide*. Tech Policy Press. <https://techpolicy.press/metass-new-content-policies-risk-fueling-more-mass-violence-and-genocide>
- de León, E., Votta, F., Araujo, T., & de Vreese, C. (2025, October 29). 1 in 10 Dutch citizens are likely to ask AI for election advice. This is why they shouldn't. *University of Amsterdam / AlgoSoc*. <https://www.uva.nl/en/shared-content/faculteiten/en/faculteit-der-maatschappij-en-gedragswetenschappen/news/2025/10/dont-ask-ai-for-election-advice.html>
- Deuze, M., & Beckett, C. (2022). Imagination, Algorithms and News: Developing AI Literacy for Journalism. *Digital Journalism*, 10(10), 1913–1918. <https://doi.org/10.1080/21670811.2022.2119152>
- Diakopoulos, N., Helberger, N., Cools, H., Li, C., Kung, E., & Rinehart, A. (2024a). *Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem*.
- Diakopoulos, N., Helberger, N., Cools, H., Li, C., Kung, E., & Rinehart, A. (2024b). *Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem*.
- Dodds, T., Vandendaele, A., Simon, F. M., Helberger, N., Resendez, V., & Yeung, W. N. (2025b). Knowledge Silos as a Barrier to Responsible AI Practices in Journalism?

- Exploratory Evidence from Four Dutch News Organisations. *Journalism Studies*, 26(6), 740–758. <https://doi.org/10.1080/1461670X.2025.2463589>
- Dodds, T., Zamith, R., & Lewis, S. C. (2025a). The AI turn in journalism: Disruption, adaptation, and democratic futures. *Journalism*, 14648849251343518. <https://doi.org/10.1177/14648849251343518>
- Dommett, K., Kefford, G., Kruschinski, S., Dommett, K., Kefford, G., & Kruschinski, S. (2024). *Data-Driven Campaigning and Political Parties: Five Advanced Democracies Compared*. Oxford University Press.
- Dommett, K., Kefford, G., & Power, S. (2021). The digital ecosystem: The new politics of party organization in parliamentary democracies. *Party Politics*, 27(5), 847–857. <https://doi.org/10.1177/1354068820907667>
- Dworkin, R. (1994). A new map of censorship. *Index on Censorship*, 23(1–2), 9–15. <https://doi.org/10.1080/03064229408535633>
- EBU. (2025). *News Integrity in AI Assistants. An international PSM study*.
- Eder, M., & Sjøvaag, H. (2025). Falling behind the adoption curve: Local journalism's struggle for innovation in the AI transformation. *Journal of Media Business Studies*, 22(4), 325–343. <https://doi.org/10.1080/16522354.2025.2473301>
- Elliott, V. (2025). The WIRED AI Elections Project. *Wired*. <https://www.wired.com/story/generative-ai-global-elections/>
- Feenstra, W., & Sabel, P. (2025, October 27). Twee PVV-Kamerleden vallen met nepbeelden anoniem Timmermans aan, GroenLinks-PvdA doet aangifte. *de Volkskrant*. <https://www.volkskrant.nl/politiek/twee-pvv-kamerleden-vallen-met-nepbeelden-anoniem-timmermans-aan-groenlinks-pvda-doet-aangifte~b6958ea1/>
- Fineman, M. A. (Ed.). (2008). The vulnerable subject: Anchoring equality in the human condition. In *Transcending the Boundaries of Law* (0 edn, pp. 177–191). Routledge-Cavendish. <https://doi.org/10.4324/9780203848531-26>
- Foos, F. (2024). The Use of AI by Election Campaigns. *LSE Public Policy Review*, 3(3), 8. <https://doi.org/10.31389/lseppr.112>
- Garton Ash, T. (2016). *Free Speech: Ten Principles for a Connected World*. Yale University Press. (1227501). <https://research.ebsco.com/linkprocessor/plink?id=5a1596fe-8094-3709-af3b-f82edf3fbb49>
- Gerards, J., & Xenidis, R. (2021). *Algorithmic discrimination in Europe: Challenges and opportunities for gender equality and non discrimination law*. Publications Office of the European Union. <https://data.europa.eu/doi/10.2838/544956>
- Gerards, J., & Zuiderveen Borgesius, F. (2022). Protected Grounds and the System of Non-Discrimination Law in the Context of Algorithmic Decision-Making and Artificial Intelligence. *Colorado Technology Law Journal*, 20, 56.
- Gillespie, T. (2024). Generative AI and the politics of visibility. *Big Data & Society*, 11(2), 20539517241252131. <https://doi.org/10.1177/20539517241252131>
- Gondwe, G. (2025). Perceptions of AI-driven news among contemporary audiences: A study of trust, engagement, and impact. *AI & SOCIETY*, 1–12. <https://doi.org/10.1007/s00146-025-02294-x>
- Guler, G., Remfry, E., Kherroubi Garcia, I., Barrow, N., Duarte, T., & Ibison, Y. (2024). *Grassroots and non-profit perspectives on generative AI*. Joseph Rowntree Foundation. <https://www.jrf.org.uk/ai-for-public-good/grassroots-and-non-profit-perspectives-on-generative-ai>
- Guzman, A. L., & Lewis, S. C. (2024). What Generative AI Means for the Media Industries, and Why it Matters to Study the Collective Consequences for Advertising,

- Journalism, and Public Relations. *Emerging Media*, 2(3), 347–355.
<https://doi.org/10.1177/27523543241289239>
- Hacker, P., Zuiderveen Borgesius, F., Mittelstadt, B. D., & Wachter, S. (2025). *Generative discrimination: What happens when generative AI exhibits bias, and what can be done about it* (1st edn). Oxford University Press.
<https://ora.ox.ac.uk/objects/uuid:cf87e259-368f-440a-a2df-5285433ffcb1>
- Helberger, N., Sax, M., Strycharz, J., & Micklitz, H.-W. (2022). Choice Architectures in the Digital Economy: Towards a New Understanding of Digital Vulnerability. *Journal of Consumer Policy*, 45(2), 175–200. <https://doi.org/10.1007/s10603-021-09500-5>
- Helberger, N., van Drunen, M., Eskens, S., Bastian, M., & Möller, J. (2020). A freedom of expression perspective on AI in the media: With a special focus on editorial decision making on social media platforms and in the news media. *European Journal of Law and Technology*, 11(30).
- Hofman, E., & Veerbeek, J. (2024, December 11). Verantwoording bij het data-onderzoek naar de AI-beelden van radicaal-rechts. *De Groene Amsterdammer*.
<https://www.groene.nl/artikel/data-onderzoek-naar-de-ai-beelden-van-radicaal-rechts>
- Holt, J. (2024). *Cloud Policy: A History of Regulating Pipelines, Platforms, and Data*. The MIT Press. <https://doi.org/10.7551/mitpress/14683.001.0001>
- Hu, T.-H. (2015). *A Prehistory of the Cloud*. The MIT Press.
<https://doi.org/10.7551/mitpress/9780262029513.001.0001>
- Jaźwińska, K., & Chandrasekar, A. (2025, March). AI Search Has A Citation Problem. *Columbia Journalism Review*. https://www.cjr.org/tow_center/we-compared-eight-ai-search-engines-theyre-all-bad-at-citing-news.php
- Jones, R., & Jones, B. (2019). Atomising the News: The (In)Flexibility of Structured Journalism. *Digital Journalism*, 7(8), 1157–1179.
<https://doi.org/10.1080/21670811.2019.1609372>
- Jungherr, A. (2023). Artificial Intelligence and Democracy: A Conceptual Framework. *Social Media + Society*, 9(3), 20563051231186353.
<https://doi.org/10.1177/20563051231186353>
- Jungherr, A., & Rauchfleisch, A. (2025). Artificial Intelligence in deliberation: The AI penalty and the emergence of a new deliberative divide. *Government Information Quarterly*, 42(4), 102079. <https://doi.org/10.1016/j.giq.2025.102079>
- Kapoor, S., & Narayanan, A. (2025, April 15). We Looked at 78 Election Deepfakes. Political Misinformation is not an AI Problem. *AI as Normal Technology*.
<https://www.normaltech.ai/p/we-looked-at-78-election-deepfakes>
- Kieslich, K., Diakopoulos, N., & Helberger, N. (2024a). Anticipating impacts: Using large-scale scenario-writing to explore diverse implications of generative AI in the news environment. *AI and Ethics*. <https://doi.org/10.1007/s43681-024-00497-4>
- Kieslich, K., Diakopoulos, N., & Helberger, N. (2024b). Anticipating impacts: Using large-scale scenario-writing to explore diverse implications of generative AI in the news environment. *AI and Ethics*. <https://doi.org/10.1007/s43681-024-00497-4>
- Kotek, H., Dockum, R., & Sun, D. Q. (2023). Gender bias and stereotypes in Large Language Models. *Proceedings of The ACM Collective Intelligence Conference*, 12–24. <https://doi.org/10.1145/3582269.3615599>
- Kreps, S., & Kriner, D. (2023). How AI Threatens Democracy. *Journal of Democracy*, 34(4), 122–131.
- Kruschinski, S., Jost, P., Fecher, H., & Scherer, T. (2025). *Künstliche Intelligenz in politischen Kampagnen*. Otto Brenner Stiftung.

- Kruschinski, S., & Votta, F. (2025). *CampAlignTracker.de: Detection and Analysis of Visual GenAI in the 2025 German Federal Election*. <https://doi.org/10.17605/OSF.IO/CZEB6>
- Kruschinski, S., & Votta, F. (forthcoming). Efficiency Tool or Manipulation Machine? The Use, Transparency, and Communication Strategies of Visual Generative AI on Social Media in the 2025 German Federal Election. In S. Geise & M.-M. Leidecker-Sandmann (Eds.), *Die Medien im Wahlkampf—Analysen zur Bundestagswahl 2025*. Wiesbaden: Springer VS.
- Larkin, B. (2013). The Politics and Poetics of Infrastructure. *Annual Review of Anthropology*, 42(Volume 42, 2013), 327–343. <https://doi.org/10.1146/annurev-anthro-092412-155522>
- Laude, L., & Daum, A. (2024). *KI als neues Wahlkampfinstrument*. <https://doi.org/10.59704/9e96c5f69bbfa94c>
- Lazar, S., & Manuali, L. (2026). Using LLMs to Enhance Democracy. *Minds and Machines*, 36(1), 12. <https://doi.org/10.1007/s11023-026-09767-y>
- Leerssen, P. (2025). From Murdoch to Musk: Platform ownership and the political economy of online content governance. *Platforms & Society*, 2, 29768624251386260. <https://doi.org/10.1177/29768624251386260>
- Lewis, S. C., Guzman, A. L., Schmidt, T. R., & Lin, B. (2025). Generative AI and its disruptive challenge to journalism: An institutional analysis. *Communication and Change*, 1(1), 9. <https://doi.org/10.1007/s44382-025-00008-x>
- Lin, B., & Lewis, S. C. (2022). The One Thing Journalistic AI Just Might Do for Democracy. *Digital Journalism*, 10(10), 1627–1649. <https://doi.org/10.1080/21670811.2022.2084131>
- Liu, L. T., Dean, S., Rolf, E., Simchowitz, M., & Hardt, M. (2018). Delayed Impact of Fair Machine Learning. *arXiv:1803.04383 [Cs, Stat]*. <http://arxiv.org/abs/1803.04383>
- Luitse, D. (2024). Platform power in AI: The evolution of cloud infrastructures in the political economy of artificial intelligence. *Internet Policy Review*, 13(2). <https://policyreview.info/articles/analysis/platform-power-ai-evolution-cloud-infrastructures>
- Maher, B. (2025, February 21). ChatGPT referrals to top publishers up eight times in six months but still negligible. *Press Gazette*. <https://pressgazette.co.uk/platforms/chatgpt-perplexity-news-ai-referral-search/>
- Malgieri, G., & Niklas, J. (2020). Vulnerable data subjects. *Computer Law & Security Review*, 37, 105415. <https://doi.org/10.1016/j.clsr.2020.105415>
- Mansell, R., Durach, F., Kettemann, M., Lenoir, T., Procter, R., Tripathi, G., & Tucker, E. (2025). Information ecosystem and troubled democracy. *OID*. <https://observatory.informationdemocracy.org/report/information-ecosystem-and-troubled-democracy/>
- Mattis, N., & de Vreese, C. (2025). Breaking the News? Generative AI's Impact on Journalism and Its Implications for Disinformation. *International Journal of Communication*, 19, 1–25.
- Mayer, V. (2023). When do we go from here? Data center infrastructure labor, jobs, and work in economic development time and temporalities. *New Media & Society*, 25(2), 307–323. <https://doi.org/10.1177/14614448221149947>
- McKinney, S. (2024). Integrating Artificial Intelligence into Citizens' Assemblies: Benefits, Concerns and Future Pathways. *Journal of Deliberative Democracy*, 20(1). <https://doi.org/10.16997/jdd.1556>

- Meese, J., & Seipp, T. (2025). The Press as Platform: Institutional Isomorphism and the Strategic Adoption of Platform Logics. *Digital Journalism*, 1–17. <https://doi.org/10.1080/21670811.2025.2573074>
- Mitova, E., Blassnig, S., Strikovic, E., Urman, A., De Vreese, C., & Esser, F. (2024a). Exploring users' desire for transparency and control in news recommender systems: A five-nation study. *Journalism*, 25(10), 2001–2021. <https://doi.org/10.1177/14648849231222099>
- Mitova, E., Blassnig, S., Strikovic, E., Urman, A., De Vreese, C., & Esser, F. (2024b). Exploring users' desire for transparency and control in news recommender systems: A five-nation study. *Journalism*, 25(10), 2001–2021. <https://doi.org/10.1177/14648849231222099>
- Monzer, C., Moeller, J., Helberger, N., & Eskens, S. (2020). User Perspectives on the News Personalisation Process: Agency, Trust and Utility as Building Blocks. *Digital Journalism*, 8(9), 1142–1162. <https://doi.org/10.1080/21670811.2020.1773291>
- Morosoli, S., Naudts, L., Cools, H., Venkatraj, K., Helberger, N., & De Vreese, C. (2025). Public accountability and regulatory expectations for AI in journalism: Qualitative evidence from focus groups with Dutch citizens. *AI & SOCIETY*. <https://doi.org/10.1007/s00146-025-02591-5>
- Motamedi, M. (2019, October 2). Locked out: Why is Amazon blocking Iranians from its services? *Al Jazeera*. <https://www.aljazeera.com/economy/2019/10/2/locked-out-why-is-amazon-blocking-iranians-from-its-services>
- Mouffe, C. (1999). Deliberative Democracy or Agonistic Pluralism? *Social Research*, 66(3), 745–758.
- Naudts, L. (2023). *Fair or Unfair Differentiation? Reconsidering the Concept of Equality for the Regulation of Algorithmically Guided Decision-Making (Doctoral Dissertation)*. KU Leuven.
- Naudts, L. (2024). The Digital Faces of Oppression and Domination: A Relational and Egalitarian Perspective on the Data-driven Society and its Regulation. *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT '24*, 701–712. <https://doi.org/10.1145/3630106.3658934>
- Naudts, L., Helberger, N., Sax, M., & Veale, M. (2024). Toward Constructive Optimisation: A new perspective on the regulation of recommender systems and the rights of users and society. In N. Helberger, B. Kas, H.-W. Micklitz, M. Namysłowska, L. Naudts, M. Sax, & M. Veale (Eds), *Digital Fairness for Consumers* (pp. 25–68). BEUC - The European Consumer Organisation.
- Naudts, L., Helberger, N., Veale, M., & Sax, M. (2025). A Right to Constructive Optimization: A Public Interest Approach to Recommender Systems in the Digital Services Act. *Journal of Consumer Policy*. <https://doi.org/10.1007/s10603-025-09586-1>
- Nechushtai, E. (2018). Could digital platforms capture the media through infrastructure? *Journalism*, 19(8), 1043–1058. <https://doi.org/10.1177/1464884917725163>
- Newman, N., & Cherubini, F. (2025). *Journalism, media, and technology trends and predictions 2025*. Reuters Institute.
- Nieuwsuur. (2024, August 5). *Chatbots mogen niets zeggen over verkiezingen, maar dat lukt niet altijd*. <https://nos.nl/nieuwsuur/artikel/2531896-chatbots-mogen-niets-zeggen-over-verkiezingen-maar-dat-lukt-niet-altijd>
- O'Callaghan, P., Cronin, O., Kelly, B. D., Shiner, B., Walmsley, J., & McCarthy-Jones, S. (2024). The right to freedom of thought: An interdisciplinary analysis of the UN

- special rapporteur's report on freedom of thought. *The International Journal of Human Rights*, 28(1), 1–23. <https://doi.org/10.1080/13642987.2023.2227100>
- OECD. (2025). *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*. OECD Publishing. <https://doi.org/10.1787/795de142-en>
- OpenAI. (2024, March 13). Disrupting malicious uses of AI by state-affiliated threat actors. *OpenAI*. <https://openai.com/index/disrupting-malicious-uses-of-ai-by-state-affiliated-threat-actors/>
- OSCE. (2025). *Safeguarding Media Freedom in the Age of Big Tech Platforms and AI*. <https://www.osce.org/representative-on-freedom-of-media/598525>
- Oster, J. (2015). *Media Freedom as a Fundamental Right*. Cambridge University Press.
- Pentney, K. (2024). States' positive obligation to create a favourable environment for participation in public debate: A principle in search of a practical effect? *Journal of Media Law*. (world). <https://www.tandfonline.com/doi/abs/10.1080/17577632.2024.2326278>
- Phillips, T., & Torres, P. (2024, August 27). 'Being on camera is no longer sensible': Persecuted Venezuelan journalists turn to AI. *The Guardian*. <https://www.theguardian.com/world/article/2024/aug/27/venezuela-journalists-nicolas-maduro-artificial-intelligence-media-election>
- Piasecki, S., Morosoli, S., Helberger, N., & Naudts, L. (2024). AI-generated journalism: Do the transparency provisions in the AI Act give news readers what they hope for? *Internet Policy Review*, 13(4). <https://doi.org/10.14763/2024.4.1810>
- Pybus, J., & Côté, M. (2024). Super SDKs: Tracking personal data and platform monopolies in the mobile. *Big Data & Society*, 11(1), 20539517241231270. <https://doi.org/10.1177/20539517241231270>
- Quell, M. (2025, May 15). Trump's sanctions on ICC prosecutor have halted tribunal's work. *AP News*. <https://apnews.com/article/icc-trump-sanctions-karim-khan-court-a4b4c02751ab84c09718b1b95cbd5db3>
- Reiss, M. V., Knor, E. L., Stöwing, E., Merten, L., & Möller, J. (2025). Zwischen Neugier und Skepsis: Nutzung und Wahrnehmung generativer KI zur Informationssuche in Deutschland. *Arbeitspapiere des Hans-Bredow-Instituts*, 76, 39 S. <https://doi.org/10.21241/SSOAR.100907>
- Rikap, C. (2021). *Capitalism, Power and Innovation: Intellectual Monopoly Capitalism Uncovered* (1st edn). Routledge. <https://doi.org/10.4324/9780429341489>
- Rolf, S., & Schindler, S. (2023). The US–China rivalry and the emergence of state platform capitalism. *Environment and Planning A: Economy and Space*, 55(5), 1255–1280. <https://doi.org/10.1177/0308518X221146545>
- Ryan, Ó. (2025, October 22). Deepfake AI video depicting Catherine Connolly quitting presidential race removed by Meta. *The Irish Times*. <https://www.irishtimes.com/politics/2025/10/22/meta-removes-ai-video-purporting-to-show-catherine-connolly-quitting-presidential-race/>
- Sanseviero, O., Cuenca, P., Passos, A., & Whitaker, J. (2025). *Hands-On Generative AI with Transformers and Diffusion Models*. <https://www.oreilly.com/library/view/hands-on-generative-ai/9781098149239/>
- Schibstedt. (2024). *Editorial media as defenders of democracies*. <https://cdn.schibstedmedia.com/wp-content/uploads/2024/11/25125955/Editorial-media-as-defenders-of-democracies.pdf>

- Schiele, M., Gittmann, Y., Ilchmann, S., Gojsalić, A., Jurinčić, D., & Klempt, P. (2024). *Voting Advice Applications: Implementation of RAG-supported LLMs*. Preprints. <https://doi.org/10.36227/techrxiv.172115156.64500701/v1>
- Schiff, K. J., Schiff, D. S., & Bueno, N. S. (2025). The Liar's Dividend: Can Politicians Claim Misinformation to Evade Accountability? *American Political Science Review*, 119(1), 71–90. <https://doi.org/10.1017/S0003055423001454>
- Schneiders, P., & Stark, B. (2025). Ensuring News Quality in Platformized News Ecosystems: Shortcomings and Recommendations for an Epistemic Governance. *Media and Communication*, 13, 10042. <https://doi.org/10.17645/mac.10042>
- Seipp, T. J., Helberger, N., De Vreese, C., & Ausloos, J. (2024). Between the cracks: Blind spots in regulating media concentration and platform dependence in the EU. *Internet Policy Review*, 13(4). <https://doi.org/10.14763/2024.4.1813>
- Sevignani, S., Theine, H., & Tröger, M. (2025). Unpacking Property: Media, Ownership, and Power in Transformation| Toward Media Environment Capture: A Theoretical Contribution on the Influence of Big Tech on News Media. *International Journal of Communication*, 19, 21–21.
- Shrishak, K., & Pénicaud, S. (2025). *European Policy Guidelines on AI and Algorithm-Driven Discrimination for Equality Bodies and Other National Human Rights Structures* (p. 14). Council of Europe. <https://rm.coe.int/policy-guidelines-eng-web/48802a38a5>
- Simchon, A., Edwards, M., & Lewandowsky, S. (2024). The persuasive effects of political microtargeting in the age of generative artificial intelligence. *PNAS Nexus*, 3(2), pgae035. <https://doi.org/10.1093/pnasnexus/pgae035>
- Simon, F. M. (2024a). Escape Me If You Can: How AI Reshapes News Organisations' Dependency on Platform Companies. *Digital Journalism*, 12(2), 149–170. <https://doi.org/10.1080/21670811.2023.2287464>
- Simon, F. M. (2024b). Escape Me If You Can: How AI Reshapes News Organisations' Dependency on Platform Companies. *Digital Journalism*, 12(2), 149–170. <https://doi.org/10.1080/21670811.2023.2287464>
- Simon, F. M. (2024b). *Artificial Intelligence in the News: How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena*. <https://doi.org/10.7916/ncm5-3v06>
- Simon, F. M., & Altay, S. (2025). Don't panic (yet): Assessing the evidence and discourse around generative AI and elections. *Knight First Amendment Institute at Columbia University*. <https://ora.ox.ac.uk/objects/uuid:ea305c74-27f2-41cd-a6f5-010dd90e60eb>
- Simon, F. M., Nielsen, R. K., & Fletcher, R. (2025). *Generative AI and news report 2025: How people think about AI's role in journalism and society*. Reuters Institute for the Study of Journalism. <https://doi.org/10.60625/RISJ-5BJV-YT69>
- Sjøvaag, H., Ferrer-Conill, Raul, & Olsen, R. Kr. (2024). Capture Beyond the Platforms: The Material and Infrastructural Conditions for Digital Journalism. *Digital Journalism*, 0(0), 1–20. <https://doi.org/10.1080/21670811.2024.2377078>
- Small, C. T., Vendrov, I., Durmus, E., Homaei, H., Barry, E., Cornebise, J., Suzman, T., Ganguli, D., & Megill, C. (2023). *Opportunities and Risks of LLMs for Scalable Deliberation with Polis* (arXiv:2306.11932). arXiv. <https://doi.org/10.48550/arXiv.2306.11932>
- Stamboliev, E. (2023). Proposing a Postcritical AI Literacy: Why We Should Worry Less about Algorithmic Transparency and More about Citizen Empowerment. *Media Theory*, 7(1), 201–232. <https://doi.org/https://doi.org/10.70064/mt.v7i1.890>
- Stryker, Cole, C., & Scapicchio, M. (2025). What is generative AI. *IBM.Com*. <https://www.ibm.com/think/topics/generative-ai>

- Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus*, 3(9), pgae346. <https://doi.org/10.1093/pnasnexus/pgae346>
- Timmer, A. (2014). *Strengthening the Equality Analysis of the European Court of Human Rights: The Potential of the Concepts of Stereotyping and Vulnerability*. Ghent University.
- Tobitt, C. (2025, November 4). Who's suing AI and who's signing: Getty signs with Perplexity, Penske Media sues Google over AI Overviews. *Press Gazette*. <https://pressgazette.co.uk/platforms/news-publisher-ai-deals-lawsuits-openai-google/>
- Tsai, L. L., Pentland, A. S., Braley, A., Chen, N. L., Ramón Enríquez, J., & Reuel, A. (2024). Generative AI for Pro-Democracy Platforms. *An MIT Exploration of Generative AI*. <https://mitgovlab.org/research/generative-ai-for-pro-democracy-platforms/>
- van der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society*, 11(1). <https://doi.org/10.1177/20539517241232630>
- van Drunen, M. Z. (2025). Safeguarding media freedom from infrastructural reliance on AI companies: The role of EU law. *Telecommunications Policy*, 49(7), 102990. <https://doi.org/10.1016/j.telpol.2025.102990>
- Veldkamp, C., & van Marum, E. (2025). *Beantwoording vragen over het bericht 'Wake-up call: Microsoft sluit e-mail ICC zonder pardon af.'* [Kamerstuk]. Ministerie van Algemene Zaken. <https://www.rijksoverheid.nl/documenten/kamerstukken/2025/07/01/beantwoording-vragen-over-het-bericht-wake-up-call-microsoft-sluit-e-mail-icc-zonder-pardon-af>
- Verma, P., & Zakrzewski, C. (2024, April 23). AI deepfakes threaten to upend global elections. No one can stop them. *The Washington Post*. <https://www.washingtonpost.com/technology/2024/04/23/ai-deepfake-election-2024-us-india/>
- Viljoen, S. (2021). A Relational Theory of Data Governance. *The Yale Law Journal*, 82.
- Vipra, J., & Myers West, S. (2023, September 27). Computational Power and AI. *AI Now Institute*. <https://ainowinstitute.org/publications/compute-and-ai>
- Votta, F., & Kruschinski, S. (2025). CampAlignTracker.nl: Detection and Analysis of Visual GenAI in the 2025 Dutch Parliamentary Election. *OSF*. <https://osf.io/sjc2b>
- Wachter, S. (2023). The Theory of Artificial Immutability: Protecting Algorithmic Groups Under Anti-Discrimination Law. *Tulane Law Review*, 97(2), 149.
- Waldo, J., & Boussard, S. (2024). GPTs and Hallucination: Why do large language models hallucinate? *Queue*, 22(4), 19–33. <https://doi.org/10.1145/3688007>
- Weikmann, T., Greber, H., & Nikolaou, A. (2025). After Deception: How Falling for a Deepfake Affects the Way We See, Hear, and Experience Media. *The International Journal of Press/Politics*, 30(1), 187–210. <https://doi.org/10.1177/19401612241233539>
- Weikmann, T., & Lecheler, S. (2023). Visual disinformation in a digital age: A literature synthesis and research agenda. *New Media & Society*, 25(12), 3696–3713. <https://doi.org/10.1177/14614448221141648>
- Weikmann, T., Wouters, F., Tulin, M., Hameleers, M., de Vreese, C. H., Zarouali, B., & Opgenhaffen, M. (under review). *On the same page? Experts are mostly, but not always aligned in their perceptions of self-competence, risks and responsibilities surrounding disinformation in times of generative AI*.

- Wetenschappelijke Raad voor Regeringsbeleid (WRR). (2021). *Opgave AI. De nieuwe systeemtechnologie*. <https://www.wrr.nl/publicaties/rapporten/2021/11/11/opgave-ai-de-nieuwe-systeemtechnologie>
- Wetenschappelijke Raad voor Regeringsbeleid (WRR). (2024). *Aandacht voor media. Naar nieuwe waarborgen voor hun democratische functies*. WRR.
- Whittaker, M. (2021). The steep cost of capture. *Interactions*, 28(6), 50–55. <https://doi.org/10.1145/3488666>
- Widder, D. G., & Kim, N. (2025). *How Big Cloud becomes Bigger: Scrutinizing Google, Microsoft, and Amazon's investments* (SSRN Scholarly Paper No. 5377426). Social Science Research Network. <https://doi.org/10.2139/ssrn.5377426>
- Xenidis, R. (2025). *Legal Protection Against Algorithmic Discrimination in Europe: Current Frameworks and Remaining Gaps* (p. 73). Council of Europe. <https://rm.coe.int/legal-protection-eng-web/48802a38d7>
- Young, I. M. (2002). *Inclusion and Democracy* (online edition). Oxford University Press. <https://doi.org/10.1093/0198297556.001.0001>
- Zuiderveen Borgesius, F. J., Möller, J., Kruikemeier, S., Ó Fathaigh, R., Irion, K., Dobber, T., Bodo, B., & De Vreese, C. (2018). Online Political Microtargeting: Promises and Threats for Democracy. *Utrecht Law Review*, 14(1), 82. <https://doi.org/10.18352/ulr.420>

Council of Europe instruments

Treaties

- CoE (1950). Convention for the Protection of Human Rights and Fundamental Freedoms (ETS No. 005) <https://www.coe.int/en/web/conventions/full-list?module=treaty-detail&treatynum=005>
- CoE (1981 and 2018). Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data (ETS No. 108) <https://www.coe.int/en/web/conventions/full-list?module=treaty-detail&treatynum=108> and Protocol amending the Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data (CETS No. 223) <https://www.coe.int/en/web/conventions/full-list?module=treaty-detail&treatynum=223>
- CoE (2024). Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225) <https://www.coe.int/en/web/conventions/full-list?module=treaty-detail&treatynum=225>

Other instruments

- CoE. (2016). Recommendation CM/Rec(2016)1 of the Committee of Ministers to member States on protecting and promoting the right to freedom of expression and the right to private life with regard to network neutrality <https://search.coe.int/cm?i=09000016805c1e59>
- CoE. (2017). Recommendation CM/Rec(2017)8 of the Committee of Ministers to member States on Big Data for culture, literacy and democracy <https://search.coe.int/cm?i=0900001680750d68>
- CoE. (2018a). Recommendation CM/Rec(2018)1 of the Committee of Ministers to member States on media pluralism and transparency of media ownership <https://search.coe.int/cm?i=0900001680790e13>

- CoE. (2018b). Recommendation CM/Rec(2018)2 of the Committee of Ministers to member States on the roles and responsibilities of internet intermediaries
<https://search.coe.int/cm?i=0900001680790e14>
- CoE. (2019). Declaration by the Committee of Ministers on the manipulative capabilities of algorithmic processes (adopted by the Ministers' Deputies on 13 February 2019)
<https://search.coe.int/cm?i=090000168092dd4b>
- CoE. (2020). Recommendation CM/Rec(2020)1 of the Committee of Ministers to member States on human rights impacts of algorithmic systems
<https://search.coe.int/cm?i=09000016809e1154>
- CoE. (2021a). Guidelines on the Protection of Individuals with regard to the Processing of Personal Data by and for Political Campaigns. Consultative Committee of the Convention for the protection of individuals with regard to automatic processing of personal data.
https://edoc.coe.int/fr/module/ec_addformat/download?cle=8c3c27ac7d298331a1bdfd0a5e8703d3&k=0b8b7ab709bee8cea2965e31e20bec12
- CoE. (2021b). Guidance Note on the Prioritisation of Public Interest Content Online. Steering Committee on Media and Information Society.
<https://www.google.si/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&ved=2ahUKEwiq4sG57vGQAxXS3gIHHTQBGiQQFnoECAsQAQ&url=https%3A%2F%2Frm.coe.int%2Fcdmsi-2021-009-guidance-note-on-the-prioritisation-of-pi-content-e-ado%2F1680a524c4&usq=AOvVaw18AuthqoSpK6RQgchJ44Dp&opi=89978449>
- CoE. (2022a). Recommendation CM/Rec(2022)11 of the Committee of Ministers to member States on principles for media and communication governance
<https://search.coe.int/cm?i=0900001680a61712>
- CoE. (2022b). Recommendation CM/Rec(2022)16 of the Committee of Ministers to member States on combating hate speech
<https://search.coe.int/cm?i=0900001680a67955>
- CoE. (2022c). Recommendation CM/Rec(2022)12 of the Committee of Ministers to member States on electoral communication and media coverage of election campaigns
<https://search.coe.int/cm?i=0900001680a6172e>
- CoE. (2022d). Recommendation CM/Rec(2022)4 of the Committee of Ministers to member States on promoting a favourable environment for quality journalism in the digital age
<https://search.coe.int/cm?i=0912594880267f95>
- CoE. (2022e). Recommendation CM/Rec(2022)13 of the Committee of Ministers to member States on the impacts of digital technologies on freedom of expression
<https://search.coe.int/cm?i=0900001680a61729>
- CoE. (2023a). Guidance Note on countering the spread of online mis- and disinformation through fact-checking and platform design solutions in a human rights compliant manner. Steering Committee on Media and Information Society.
https://www.google.si/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&cad=rja&uact=8&ved=2ahUKEwiLuLXx7_GQAxVL9gIHHfFtMWsQFnoECAsQAQ&url=https%3A%2F%2Fedoc.coe.int%2Fen%2Finternet%2F11885-guidance-note-on-countering-the-spread-of-online-mis-and-disinformation-through-fact-checking-and-platform-design-solutions-in-a-human-rights-compliant-manner.html&usq=AOvVaw1oZUxdL_VuOIImYTfQ9B-F1&opi=89978449
- CoE. (2023b). Guidelines on the Responsible Implementation of Artificial Intelligence Systems in Journalism. Steering Committee on Media and Information Society.

<https://rm.coe.int/cdmsi-2023-014-guidelines-on-the-responsible-implementation-of-artific/1680adb4c6>

CoE. (2025). *Guidance Note on the Implications of Generative Artificial Intelligence for Freedom of Expression* (CDMSI(2025)15rev). <https://rm.coe.int/cdmsi-2025-15rev-guidance-note-on-the-implications-of-generative-artif/488029df80>

Case law of the European Court of Human Rights

Aksu v. Turkey, 4149/04, 41029/04 (ECtHR [GC] 15 March 2012).

<https://hudoc.echr.coe.int/eng?i=001-109577>

Carvalho Pinto De Sousa Morais v. Portugal, No. 17484/15 (ECtHR 25 July 2017)

<https://hudoc.echr.coe.int/fre?i=001-175659>

Centro Europa 7 S.R.L. and Di Stefano v. Italy, No. 38433/09 (ECtHR [GC] 7 June 2012)

<https://hudoc.echr.coe.int/eng?i=001-111399>

D.H. and Others v. the Czech Republic, No. 57325/00 (ECtHR [GC] 13 November 2007).

<https://hudoc.echr.coe.int/fre?i=001-83256>

Dink v. Turkey, No. 2668/07 (ECtHR 14 September 2010).

<https://hudoc.echr.coe.int/eng?i=002-808>

Féret v. Belgium, No. 15615/07 (ECtHR 16 July 2009)

<https://hudoc.echr.coe.int/eng?i=001-93626>

Handyside v. the United Kingdom, No. 5493/72 (ECtHR 7 December 1976)

<https://hudoc.echr.coe.int/eng?i=001-57499>

Konstantin Markin v. Russia, No. 30078/06 (ECtHR [GC] 22 March 2012)

<https://hudoc.echr.coe.int/fre?i=001-109868>

Lingurar v. Romania, No. 48474/14 (ECtHR 16 April

2019)<https://hudoc.echr.coe.int/eng?i=001-192466>

Sunday Times v. The United Kingdom, No. 6538/74 (ECtHR 26 April 1979)

<https://hudoc.echr.coe.int/eng?i=001-57584>