



Ofcom

Future Technology and Media Literacy: Applications of Generative AI

Published 13 November 2024

Making Sense
of Media



Contents

Section

1. Overview.....	3
2. The news sector and personalisation.....	7
3. Personalisation and adaptation	10
4. Content creation and education	13
5. Data protection concerns.....	16
6. Conclusion	18

1. Overview

Generative artificial intelligence (“generative AI”) has rapidly gone from being a relatively unknown technology to a topic that has attracted significant mainstream attention from users, media and investors. As this attention has grown, so too has availability of the technology for public use.

Generative AI is being integrated into services such as social media, search, gaming, dating, and mainstream productivity applications (e.g., word-processing and spreadsheets). Examples include social media platform Snap, which has integrated into its app a conversational chatbot called MyAI that engages with users to provide information, suggestions, and recommendations¹ and standalone companion apps such as Replika, which provide AI companions². In 2024 Microsoft updated its search service Bing, so that generative AI can provide a summary of live search results from across the web,³ and the gaming platform Roblox has developed a tool that can generate new virtual worlds.⁴ Ofcom’s Online Nation report has found that 79% of 12 to 17-year-olds are using generative AI tools and services.⁵

In our [previous discussion paper](#) on generative AI, we examined the accuracy of generative AI, mis and disinformation, and the amplification of bias in generative AI. In this paper we will consider the applications and impacts of generative AI across users’ lives. This document explores where shifts may occur, what the opportunities and risks could be, and how platforms, the media literacy sector and users could respond.

As discussed at the UK government AI Safety Summit 2023 there are immediate threats that AI, including generative AI, could pose to international and national security as well as individual rights and safety.⁶ Indeed, Ofcom has identified and assessed the risks of harm from Generative AI on online services in our Illegal Harms and Protection of Children consultations.⁷ We have also published research exploring the risks posed by harmful deepfakes, and interventions that could help services to address them.⁸ This paper has focussed on the media literacy implications of generative AI, particularly in relation to understanding and trust in online content.

Key areas this paper explores in relation to media literacy:

News and personalisation

¹ https://techcrunch.com/2023/04/19/snapchat-opens-its-ai-chatbot-to-global-users-says-the-ai-will-later-snap-you-back/?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlMmNvbS8&guce_referrer_sig=AQAAADZV-8q7ODLm75-edQjrVDGwnDKHdu9FbTAHEulv7_ssYuDwFuE_UmUa8EtIrfDzZrveaRa19sfb1Xg26c5M-zK0asG8qlgArDA8Mccai0AMFL_dLHq9rp_B2Dj4LqD8rGxXCMoUntb_NutcuKGPPuoQZZxkPh2_LOT_J_gBee0

² <https://www.theverge.com/24216748/replika-ceo-eugenia-kuyda-ai-companion-chatbots-dating-friendship-decoder-podcast-interview>

³ <https://www.reuters.com/technology/microsoft-infuse-software-with-more-ai-google-rivalry-heats-up-2023-02-07/>

⁴ <https://aibusiness.com/ml/roblox-gets-generative-ai-users-can-build-virtual-worlds-from-text>

⁵ <https://www.ofcom.org.uk/news-centre/2023/gen-z-driving-early-adoption-of-gen-ai>

⁶ <https://www.aisafetysummit.gov.uk/>

⁷ As part of our work as regulator of the Online Safety Act 2023, Ofcom has published a number of consultations including [Consultation: Protecting People from Illegal Harms Online – Ofcom](#); and [Consultation: Protecting children from harms online - Ofcom](#).

⁸ [A deep dive into deepfakes that demean, defraud and disinform - Ofcom](#)

Personalisation and adaptation
Content creation and education
Data protection concerns

Introduction to the future technology trends project

This discussion paper forms part of a series Ofcom is producing on future technology trends and their related potential media literacy implications, in order to support those working on media literacy to better understand what opportunities and challenges could arise.

To outline Ofcom's role in this area, where this remit originates and how Ofcom defines media literacy, we have created an [anchor document](#).⁹ It describes how we will select the future technology trends that will be looked at in this discussion series and the lenses through which we will assess the media literacy implications of those trends.

A Definition of Generative AI

Generative AI refers broadly to machine learning models that can create new content. The most popular models generate text and images from text prompts, but some use other inputs such as images to create audio, video, and images. This contrasts with other types of artificial intelligence systems, which primarily process and analyse existing data.

Generative AI is based on foundation models, which include large language models (text), image generators, audio generators and code generators. They are trained on significant quantities of data. The last time OpenAI shared ChatGPT's training size was with GPT-3, when the company said it was trained on 300 billion tokens (strings of words) at the time¹⁰. They are then fine-tuned for purposes such as answering questions with conversational data and human feedback. This results in a model that can produce a plausible response written in language that can be understood, although there is no guarantee that the response is factually correct.

Generative AI is already being used in a wide range of applications, from creative tasks such as art and content generation to practical applications such as data augmentation¹¹, summarising text, and language translation. It's an area of AI that continues to evolve and has the potential to support human creation and learning, as well as supporting technologies to improve inclusivity through its adaptive and language capabilities.

Opportunities and risks

⁹ https://www.ofcom.org.uk/__data/assets/pdf_file/0025/263374/Anchor-Document.pdf

¹⁰ <https://www.cnn.com/2023/05/16/googles-palm-2-uses-nearly-five-times-more-text-data-than-predecessor.html>

¹¹ <https://research.aimultiple.com/data-augmentation/>

Like all new technologies, generative AI also comes with challenges including concerns about the accuracy of the generated content, as well as issues related to the spread of mis and disinformation, bias, privacy, and security. Since the launch of ChatGPT in November 2022, generative AI has gone from being a relatively unknown technology to one that millions of us interact with every day. It is now being used to power features in a range of online services, including gaming platforms, dating apps, search engines and social media sites. While these generative AI applications are creating significant benefits for users, they also pose risks. For example, we know that bad actors have used generative AI to create child sexual abuse material, low-cost deepfake adverts, and synthetic terrorist content.¹² Although many of the risks associated with generative AI are similar to those posed by longstanding and less sophisticated forms of AI, they may be amplified given the volume of content generative AI models can create and the increase in their use by both companies and individuals.

As generative AI technologies advance and become more widely used, addressing these challenges will become increasingly important. In our paper exploring harmful deepfakes, we discuss four routes for actors to mitigate deepfakes.

Prevention involves efforts to limit the creation of harmful deepfakes. This can include adopting prompt filters to prevent models being instructed to create certain types of content (e.g., nude content); removing harmful content from model training datasets; and blocking outputs before they are presented to users.

Embedding involves attaching contextual information to synthetic media through tools including watermarks, provenance metadata and labels. Many organisations have signed up to the Coalition for Content Provenance and Authenticity (C2PA) scheme, which has been described as a ‘nutritional’ label for content.¹³

Detection encompasses efforts to distinguish real from fake content, even where no contextual data has been attached to that content. One means of detecting deepfakes is by using machine learning classifiers that have been trained on known deepfake content.

Enforcement involves online services setting clear rules about the types of synthetic content that can be created and shared on them. It also involves acting against users that breach those rules, for example through suspending or removing user accounts.¹⁴

Media Literacy has an important role to play, and can help to address some of these risks, but it will likely not be a silver-bullet, particularly given the potential for malicious actors to exploit the features of generative AI to perpetrate harm. It may be that the level of media literacy skill and consistency of application needed to mitigate the risks of generative AI may not be realistic for users to achieve.

As outlined in the anchor document for this series, we have used Ofcom’s media literacy definition (*‘the ability to use, understand and create media and communications across multiple formats and services’*)¹⁵ and the European Commission’s Digital Competencies framework¹⁶ (the “Framework”) as

¹² <https://www.ofcom.org.uk/siteassets/resources/documents/consultations/discussion-papers/red-teaming/red-teaming-for-gen-ai-harms.pdf?v=370762>

¹³ <https://www.bbc.co.uk/rd/blog/2024-03-c2pa-verification-news-journalism-credentials>

¹⁴ <https://www.ofcom.org.uk/siteassets/resources/documents/consultations/discussion-papers/deepfake-defences/deepfake-defences.pdf?v=370754>

¹⁵ <https://www.ofcom.org.uk/research-and-data/media-literacy-research>

¹⁶ Vuorikari, R., Kluzer, S. and Punie, Y., DigComp 2.2: The Digital Competence Framework for Citizens - With new examples of knowledge, skills and attitudes, Publications Office of the European Union, Luxembourg, 2022

a structure for examining the media literacy opportunities and risks that may arise in relation to generative AI. In this paper we have explored some of the areas where there is the potential for a significant shift in either the opportunity or risk (and sometimes both) that could arise from increased use and prevalence of generative AI.

As outlined in the Making Sense of Media Annual Plan 2023-24,¹⁷ media literacy is about both people and platforms. The considerations below are designed to form part of a sector-wide discussion on how media literacy could be supported through improving people's skills and the actions of platforms.

¹⁷ Ofcom, Making Sense of Media annual plan 2023-24, <https://www.ofcom.org.uk/research-and-data/media-literacy-research/approach>

2. The news sector and personalisation

Generative AI will pose specific challenges and opportunities for the news sector. While there could be numerous benefits to integrating generative AI into their businesses, this could impact users' trust in the information they see online. How news organisations choose to use generative AI, alongside the media literacy of their audience, are both factors which could impact the future of trust in news.

The issue of trust for news organisations is a significant one. This could have far-reaching implications for news organisation sustainability, but also for users individually and for society as a whole. The use of generative AI in news is already a matter of concern for users. A Reuters study found that only 12% of UK users trusted news organisations to use generative AI responsibly.¹⁸

It is already the case that AI is being used across the news sector.

Some publishers, in the UK and around the world, are already experimenting with AI-generated content.¹⁹ Reach plc used the technology to create listicles, pulling together existing online content and repurposing it into articles published on their platform. Technology news organisation CNET have also tried to integrate generative AI into their news production but had to delete and correct articles after the stories were found to contain substantial factual errors. CNET have since pulled back on their plans to further implement generative AI-created content.²⁰

Aside from generating the news articles themselves, generative AI can be used to create charts and graphs that may otherwise require significant time and a level of expertise to create. This is an area where generative AI may be able to support the production of news without being directly involved in writing the stories. Generative AI could also be used to automate other news production tasks, such as editing. It is not only within news production that generative AI could be integrated into the work of news organisations.

The integration of generative AI in newsrooms could prompt a reassessment of financial strategies. While there may be initial investments required for technology adoption and employee training, the long-term implications for financial viability could be varied. The potential for efficiency gains through automation of routine tasks can lead to cost reductions, allowing news companies to allocate resources more strategically leading to increased high-quality journalism and increased financial sustainability. However, staying competitive in the evolving media landscape may necessitate ongoing investments in technology upgrades and employee skill development. Balancing these financial considerations could require news companies to strategically balance innovation and fiscal responsibility.

Generative AI can also be employed to repurpose existing content at speed and without significant resource. This could allow news organisations to create the same news story in various formats tailored for different news, search, and social platforms. This efficiency may contribute to a more versatile and dynamic presentation of news across a multitude of mediums, driving more users to view news sites.

¹⁸ <https://reutersinstitute.politics.ox.ac.uk/what-does-public-six-countries-think-generative-ai-news>

¹⁹ <https://www.theguardian.com/business/2023/mar/07/mirror-and-express-owner-publishes-first-articles-written-using-ai>

²⁰ <https://www.cnet.com/tech/cnet-is-testing-an-ai-engine-heres-what-weve-learned-mistakes-and-all/>

The use of generative AI in news also has the potential to present hyper personalised content tailored to individual preferences. News content for those who consume news via social media is already personalised by algorithms that choose which stories get served up in feeds²¹. With the advent of generative AI as a content creator, it is possible that in future users on social media will see news that not only matches their existing viewpoints but has been created to specifically appeal to them.

While this provides users with content aligned with their interests, it could also create information bubbles, limiting exposure to diverse perspectives, reinforcing users' views. This could create an environment where users are presented with either views they agree with, or views presented in a way to provoke them- as ways to keep users engaging with the content they see.

As well as creating personalised news for users, generative AI could also enable news delivery direct to users, bypassing traditional news platforms. This disintermediation poses challenges to established media outlets while offering opportunities for more direct and immediate engagement with the audience. A sense of how this might work can be seen with models such as Google's Search Gemini and AI Overviews. AI overviews provides audiences with a summary of search results.

The tension between amalgamating news from various sources and maintaining trust in the authenticity of the content is a critical challenge. News outlets are responding to this challenge in different ways. While some news organisations are pro-actively working with generative AI developers, selling their archives to support the training of the developer's large language model²², others have chosen to work to prevent their work being scraped to supply training data.²³ If users are reliant on amalgamated news summaries, there is a risk that the good information from trusted sources become amalgamated with bad information from untrustworthy sources from other areas of the internet. If trusted news organisations remove the ability for their information to be used by generative AI developers, there is the risk that developers are left with less reliable information on which to train their models, increasing the risk of mis and disinformation reaching users. Striking the right balance is essential to prevent the spread of misinformation while ensuring users have access to diverse perspectives.

While legitimate news organisations are able to utilise generative AI in these ways, there are also risks that bad faith actors can utilise many of these same tools to create news that appears real to users, or feeds into the news production of legitimate news organisations.

As generative AI becomes more sophisticated, the risk of creating content that mimics the style and tone of legitimate news articles is a significant concern for news companies. It will become much harder to distinguish between genuine journalism and AI-generated content, and harder for audiences to know where to look in order to verify information. News companies may need to develop advanced authentication measures to safeguard their brand value. This could involve developing and implementing technologies that can detect subtle nuances indicative of human versus AI-generated writing. Additionally, proactive disclosure and clear labelling of AI-generated content can contribute to transparency, helping users differentiate between AI-assisted and human-produced news.

The reliance of generative AI on extensive training data provides news companies with unprecedented opportunities to refine and enhance their algorithms. However, this dependence also raises critical challenges. Ensuring the quality, diversity, and credibility of the training data is

²¹ <https://www.ofcom.org.uk/media-use-and-attitudes/media-plurality/social-media-online-news/>

²² <https://apnews.com/article/openai-chatgpt-associated-press-ap-f86f84c5bcc2f3b98074b38521f5f75a>

²³ <https://www.bbc.co.uk/mediacentre/articles/2023/generative-ai-at-the-bbc>

paramount to prevent the perpetuation of biases or the propagation of misinformation. If news organisations intend to use generative AI, they need to establish rigorous standards for data collection and verification, considering the ethical implications of using certain datasets. Moreover, transparency about data sources is crucial for building and maintaining trust with the audience. Collaborations with reputable data providers and adherence to ethical guidelines are essential to mitigate these challenges.

News and Personalisation: media literacy considerations

- Radicalisation and personalisation

The current landscape of news consumption reveals a significant proportion of users getting their news from non-traditional sources²⁴. This may suggest an appetite for content tailored to cater to the specific preferences of each user, presenting information in a way that resonates with their individual tastes. However, amid this shift towards personalisation, news organisations will need to strike a delicate balance between catering to individual preferences and preserving the broader spectrum of information. While personalisation enhances user satisfaction, ensuring a plurality of content is essential in ensuring users can engage with a variety of viewpoints.

- Trust in news

Generative AI introduces a significant challenge to the credibility of news sources. Firstly, the technology increases the ease with which almost any actor can create a convincingly realistic deepfake, raising concerns about the authenticity of the content. Secondly, as AI-generated content becomes more prevalent, there is a risk of eroding confidence in traditionally trusted sources. Users may struggle to distinguish between AI generated content and human-produced journalism, leading to scepticism about the reliability of information. This erosion of trust in authoritative sources is a critical reason why users rely on credible news outlets. Addressing this challenge requires a concerted effort from both technology developers and news organisations to implement robust authentication measures and transparent disclosure practices.

- Mis and disinformation

As generative AI becomes more involved in news creation, there is increased urgency to develop and deploy provenance-promoting technologies. Provenance technologies could prove essential for verifying the origin and authenticity of news content, helping to distinguish between genuine reporting and potentially manipulated or fabricated information. These technologies could act as safeguards, providing users with the tools to assess the trustworthiness of the content they encounter. The development of standards and protocols for provenance is not only a technological necessity but a foundational step towards maintaining the integrity of journalistic practices in the face of evolving technologies. As generative AI becomes more sophisticated, the implementation of provenance-promoting technologies could become necessary to preserve the public's trust in the news ecosystem.

²⁴ https://www.ofcom.org.uk/__data/assets/pdf_file/0024/264651/news-consumption-2023.pdf

3. Personalisation and adaptation

Generative AI excels at hyper personalisation in ways that other AI technologies struggle to match. This hyper personalisation has uses across the online world in advertising, in user adaptability, and in producing content that resonates with users across the globe, from different backgrounds and with a large variety of interests.

Generative AI also enhances personalisation in content delivery. Used consistently, this capability creates the impression of content that is highly tailored to an individual user's preferences.

It is also possible for generative AI to use personalisation across a variety of mediums to tailor user experiences by considering multiple forms of communication simultaneously, such as text, images, and voice. This can ensure that content and interactions are highly customised, engaging, and effective. This capability allows for holistic experiences that resonate with diverse audiences.

There are also uses for hyper personalisation that go beyond simply engaging users with the content created. Adaptive user interfaces allow generative AI to adapt interfaces based on individual preferences and behaviours. For example, it can rearrange app layouts, suggest relevant features, or even change visual elements to better suit each user's needs and preferences,²⁵ making technology more user-friendly.

Accessibility and generative AI

Generative AI could play a significant role in advancing accessibility by creating tools and solutions that cater to the specific needs of disabled people and people with impairments.

A significant development has been the ability of generative AI models to convert text into natural-sounding speech. This is invaluable for individuals with visual impairments or with impairments that can make reading difficult, as it makes written content accessible through audio.

For users who have impairments or illnesses that impact the ease or ability to speak, text to speech conversion can allow users to communicate verbally using a voice that is either chosen by them, or for users to record their own voice, creating the possibility that even as speech is lost, users could continue to communicate in their own voice.

Generative AI, when combined with other tools, can also translate spoken language into sign language through avatars or animations. This could enable hearing individuals to communicate with people from the d/Deaf community and those who are hard of hearing, and vice versa, aiding communication and understanding between these communities. Sign language avatars could also be customised to represent different signing languages and regional variations, ensuring that individuals can access content in their preferred sign language. This could help ensure d/Deaf users are able to participate with services and settings where a lack of sign language on the part of the service provider has previously been a barrier to access.

²⁵

<https://citeseerx.ist.psu.edu/viewdoc/download;jsessionid=4A94E07F693F56CDA83544114DDDD47E?doi=10.1.1.35.631&rep=rep1&type=pdf>

These developments, alongside the potential for real-time translation and subtitles and customised content as discussed above, could mark the opportunity for both content and delivery of information and communication to be more accessible and inclusive than has previously been the case.

Personalisation and adaptation: media literacy considerations

- Can personalised content create echo chambers?

The ability to use generative AI to create large amounts of personalised content could enable organisations to generate content that speaks directly to individual users across a number of formats including images, text, audio and video. This allows users to interact with content that is engaging and speaks directly to their wants and needs.

There is a risk that users who interact with a lot of hyper-personalised content and content created to meet their needs are not seeing a wide variety of the content available. This stretches beyond current concerns that algorithms are shaping users' online content²⁶ and could mean that users largely interact with content that was designed specifically for them. This content can be positive but can also include chatbots responding to users with highly sexualised content, self-harm, and pro-anorexia content.^{27 28} This level of personalised content will necessitate users being aware that content has been specifically designed for them and that they may therefore not be seeing the breadth of content available, meaning they are missing other perspectives or ideas.

- Does advertising become more effective with hyper personalisation?

The ability to personalise content to a greater extent is beginning to be incorporated into how brands and companies advertise their product to potential consumers. The volume of advertising that can be created with generative AI and the speed with which it can be generated will allow advertisers to create content that is highly personalised. Advertising containing content that is personalised or created for individual users will increase the likelihood that the consumer engages with the advertising.²⁹

This has the potential to make advertising incredibly effective. Users will therefore need to be aware that advertising can be specifically targeted to them and that this is designed to increase the likelihood of them purchasing a product. Companies could signal to users that content has been designed in this way, to support users to make informed decisions. Specific consideration should be given on how to protect children and young people, alongside other vulnerable users, from advertising that is highly personalised in order to prevent financial harm.

- Personalised scams

Scams can be highly effective when a user feels like they are being specifically spoken to, or when scammers prey on their interests, needs, or vulnerabilities. The level of personalisation available through AI generated content means that targeted scam content can be produced at high volume and speed, making it more likely that consumers will fall victim to the scam.

²⁶ <https://www.gov.uk/government/publications/findings-from-the-drcf-algorithmic-processing-workstream-spring-2022/the-benefits-and-harms-of-algorithms-a-shared-perspective-from-the-four-digital-regulators>

²⁷ <https://voicebox.site/sites/default/files/2023-10/Coded%20Companions%20VoiceBox%20Report.pdf>

²⁸ <https://www.wsj.com/articles/eating-disorder-chatbot-ai-2aeb179>

²⁹ <https://techcrunch.com/2023/09/25/snap-partners-with-microsoft-on-ads-in-its-my-ai-chatbot-feature/>

Further, generative AI can be used for voice cloning that can mimic intonation, speech patterns and even emotional nuances. This imitation of not only sound, but delivery, could make it difficult for users to differentiate between the voice of a human, even one who is known well to them, and an AI-generated voice clone. This raises the possibility of advanced scams that are both hyper-personalised in content and delivered through a voice that is familiar to the victim. In this scenario, users will need to be resilient in the face of these scams that seem to be speaking directly to their interests and will need to utilise their media literacy skills to check the validity of potential scams.

This is also likely to be a significant challenge for the media literacy sector as it works to support users to spot signs of scams and be vigilant against increasingly targeted and convincing scams. The sector could consider what is currently effective in supporting users to spot scams, and how what currently works will need to be developed to meet this challenge. This challenge is likely to be particularly acute when considering users who may be more vulnerable to scams, either because of age, skills or other compounding vulnerabilities.

- What could adaptive user interfaces mean for users?

There is the potential for generative AI to be a key part of adaptive user interfaces, which could allow users with different needs and abilities to access technology that they have previously been unable to engage with. Adaptive user interfaces will need to be developed alongside users who have previously been marginalised in their use of technologies to meet their potential to open up access to technologies.

- Access to technology

The potential for generative AI to support more inclusive and accessible online spaces is a promising development. For this potential to be realised, it will be key that users have access to the appropriate technologies. Without this, there could be an increasing divide between those users that have access to technology, and the accessibility benefits that could come from this, and those who do not.

- Skills for non-disabled users

It will also be important that non-disabled users understand what generative AI can do to support accessibility and the creation of inclusive online spaces. Non-disabled users will also need to be able to utilise these capabilities to ensure that the potential of this technology is realised. Without this, the onus will remain on disabled users and those from the d/Deaf community to work out how to access online spaces and services, as opposed to ensuring the responsibility of removing barriers to full inclusion is taken on by all users.

4. Content creation and education

Generative AI is a technology that holds huge potential for content creation. The technology leverages machine learning algorithms to generate content ranging from text to images, music and more. Generative AI's impact on content creation has the potential to radically change the way we produce, innovate, and express ourselves online, with outputs that can be shared offline.

Generative AI impacts creation through content generation. It can be prompted to produce written articles, poetry, code, and stories. This is achieved by training AI models on vast datasets of human-created content and then using those models to generate new, coherent content. Generative AI can generate digital art, illustrations, and designs. Some generative AI models³⁰ can produce lifelike images of non-existent people or merge artistic styles to create entirely new visual aesthetics. In music composition, AI-driven algorithms can compose original music that imitates the style of famous composers and can generate music that adapts to a specific mood or setting. Generative AI has streamlined the process of creating characters and scenes in filmmaking and animation. This can significantly reduce production times and costs while enabling filmmakers to explore creative ideas that were previously challenging to execute.

While generative AI offers numerous opportunities for creators, it also raises important ethical, societal, and legal considerations. There are questions about the ethical implications of generative AI models being trained on the creative work of humans and that work being used to generate new content. This also creates legal and financial questions of whether creators should be remunerated for the content that is used to train generative AI models and for work that uses the likeness or recognisable patterns from that work. This has already been a key driver of industrial disputes in the USA with the SAG-AFTRA strike.³¹

More broadly, there are questions about whether creativity driven by generative AI is a positive or negative development because it is currently not known how generative AI will be taken up, by users' or industry. Some are concerned that the work and thought of artists, authors, musicians, and creators will be replaced by generative AI content³², whilst others believe that generative AI will become a tool that empowers and augments human creativity³³.

Education and generative AI

As well as influencing the expression of human creativity, generative AI has the potential to significantly impact education in various ways. While the extent of its impact depends on its development and adoption, generative AI could change the way education is delivered.

Generative AI can generate personalised educational content, adapting to the specific needs and abilities of each student. It can create customised lesson plans, practice exercises, and learning materials tailored to the individual's learning style, pace, and strengths. This may be particularly relevant when it comes to tutoring and homework help, where generative AI-powered chatbots and

³⁰ Such as StyleGAN

³¹ <https://www.nbcnews.com/news/sag-aftra-strike-what-know-actors-writers-wga-rcna94075>

³² <https://www.theguardian.com/technology/2022/nov/12/when-ai-can-make-art-what-does-it-mean-for-creativity-dall-e-midjourney>

³³ <https://hbr.org/2023/07/how-generative-ai-can-augment-human-creativity>

virtual assistants can provide instant homework help and tutoring support to students. They can answer questions, explain complex concepts, and provide guidance, while challenging students on topics where it has identified specific needs.

Generative AI could influence teaching, as well as learning. For example, it is possible that generative AI could be used in curriculum development, helping to design a curriculum based on specific requirements, the latest research on a topic and best practices in education. It could then develop learning plans for the year or term and put together lesson plans. The use of generative AI in this way could remove many of the human elements of curriculum planning, which could be either positive or negative and most likely both.

This use of generative AI for teaching could also be used to create educational materials, such as textbooks, quizzes, and instructional materials that meet the aims of the curriculum.

It is possible that generative AI, if used in these ways, could save significant time and resource for schools and teachers, freeing up teachers' time to provide additional in-person connection and support for their students. However, it is also possible that AI generated curriculum and materials do not meet the needs of the students. If the generative AI is used to create a general curriculum or learning materials, rather than to create personalised materials based on learners it already holds data on, it will not understand the learners' educational needs or pre-existing knowledge and will not understand how they learn both individually and as a class or unit. It will be key that teachers understand where generative AI can be a valuable resource and where their expertise as educators is vital in giving their students the best possible education.

The UK government ran a call for evidence on the use of generative AI in education which closed on 23 August 2023.³⁴

Creation and education: media literacy considerations

- Access to creation and digital inclusion

Generative AI offers the potential for users to have access to tools that enable creativity across several mediums. This could make it easier for users to create digital content to support their own ideas, opinions, and creative desires. It could also enable users to express their personal opinions in ways that meet the desired purpose and in the appropriate format. It could also give users access to mediums of creativity and self-expression that they have not previously had access to.

However, for this technology to be truly transformative in supporting human self-expression and creation, it has to be accessible to all. If not, it risks deepening the digital divide that already exists in the UK.³⁵

The rollout of affordable technologies that allow users access to generative AI tools is not simply an issue for platforms and manufacturers but may also be of relevance for governments and societies around the world. Ensuring that devices are accessible across society regardless of social or economic status would go some way to preventing the compounding of the current digital divide.

- Rights to content

³⁴ <https://www.gov.uk/government/calls-for-evidence/generative-artificial-intelligence-in-education-call-for-evidence>

³⁵ <https://www.ofcom.org.uk/research-and-data/data/opendata>

As more generative AI models are developed, and as the volume of data they are trained on increases, there are likely to be more questions asked about the rights users have to content that may be used as part of the training data.

Informing users of their ownership rights over content they have created is likely to be complex, as there is not currently a consensus as to what ownership rights mean within generative AI models. Platforms informing users of their rights in an accessible and comprehensive manner, although challenging, will be key in shaping the user experience. Users who do not understand their ownership rights could find that their content is used in the training of generative AI models or reproduced by generative AI in response to user prompts without appropriate credit or remuneration given to them as the creator.

- Device divide and generative AI

Device inequality already drives the digital divide and can result in worse outcomes for those without access to the technology that would assist them in performing the task they are undertaking. This is particularly the case if generative AI leads to additional learning opportunities in tutoring and homework, where those without access to devices could miss out on improved educational understanding and achievement. Ensuring that devices are accessible to all those in education regardless of social or economic status would go some way towards preventing the reproduction of existing educational inequalities.

- Generative AI knowledge as neutral

If not carefully balanced, the information that generative AI conveys could be biased or inaccurate. Understanding that the information and content produced by generative AI is not neutral will be key to utilising generative AI to support education. Without this understanding it could be that inaccuracies and biases are taught and reproduced in formal and informal education systems. To this end it may be beneficial to consider how teachers are themselves educated on new and emerging technologies such as generative AI, through ongoing professional development or through ‘train the trainer’-type programmes, which could support education professionals. Ensuring those working with learners understand the opportunities, technological capabilities, and limits of generative AI will prove key in ensuring they can obtain the most amount of benefit from generative AI, while avoiding some of the risks outlined.

5. Data protection concerns

While generative AI provides new means of creation, creativity, and efficiency, it also raises significant data protection concerns. Some key issues surrounding generative AI have been addressed by the Global Privacy Assembly³⁶ who are clear that as *'generative AI must be designed, developed, and deployed in a manner that is responsible and trustworthy, based on the principles of data protection, privacy, human control, transparency, and democratic values'*³⁷

One of the foremost data protection concerns with generative AI is the processing of personal data used for model training.³⁸ Generative AI models, particularly large language models, are trained on extremely large datasets containing a wide variety of information, including publicly available content scraped from the internet, which may include personal information.³⁹

When training data includes personal information, data protection concerns will arise⁴⁰ and subsequently, there are many questions that developers and deployers will need to consider.⁴¹

For example, there is a risk that this information could be unintentionally re-produced in generated content, or malicious actors could query models to reveal sensitive information.⁴² This could result in a personal data protection breach, especially if the generated content is publicly accessible.⁴³ Even if personal information does not immediately identify an individual, generative AI can inadvertently generate content that, when combined with other publicly available information, can identify individuals. This raises concerns about the potential loss of anonymity and the violation of personal privacy.

The potential for re-identification undermines online anonymity and pseudonymity, which many individuals rely on to protect their privacy and, in some cases, their safety.

Organisations developing or using generative AI must consider their data protection obligations from the outset. The UK GDPR requires a data protection by design and default approach, to correctly implement data protection principles and safeguard individual rights.⁴⁴ Organisations should consider key data protection principles such as whether there is a lawful basis for processing personal data⁴⁵, and if a data protection impact assessment is required.⁴⁶ The Information Commissioner's Office provides guidance and tools to help organisations fulfil their privacy obligations.^{47,48,49}

³⁶ The Global Privacy Assembly has brought together data protection and privacy authorities from local, national, and international levels since 1979, first as the International Conference of Data Protection and Privacy Commissioners then as the Assembly from the 41st conference in 2019.

³⁷ https://edps.europa.eu/system/files/2023-10/edps-gpa-resolution-on-generative-ai-systems_en.pdf

³⁸ [How should we assess security and data minimisation in AI? | ICO](#)

³⁹ [What is personal information: a guide | ICO](#)

⁴⁰ [edps-gpa-resolution-on-generative-ai-systems_en.pdf \(europa.eu\)](#)

⁴¹ [Generative AI: eight questions that developers and users need to ask | ICO](#)

⁴² For example in a prompt injection attack <https://www.theguardian.com/technology/2023/aug/30/uk-cybersecurity-agency-warns-of-chatbot-prompt-injection-attacks>

⁴³ [Personal data breaches: a guide | ICO](#)

⁴⁴ [Data protection by design and default | ICO](#)

⁴⁵ [How do we ensure lawfulness in AI? | ICO](#)

⁴⁶ [What are the accountability and governance implications of AI? | ICO](#)

⁴⁷ [Guidance on AI and data protection | ICO](#)

⁴⁸ [AI and data protection risk toolkit | ICO](#)

⁴⁹ [Generative AI: eight questions that developers and users need to ask | ICO](#)

Data protection concerns: media literacy considerations

- Training generative AI models

Generative AI models are trained on large data sets that may contain sensitive or personal information. This information may have been obtained by the developer of the model without the permission or knowledge of the individual(s) whose personal information is included in this data set. The training of generative AI should be fair and transparent. This requires users to be fully informed that their data is being used to train generative AI models, and generative AI developers being clear about the lawful basis on which personal data is being used. Failing to do so risks both users' safety and their trust in using generative AI models.

Organisations and developers should be alive to the risk of generating content that includes personal information, or other identifiable information and consider the risk in the context the model is to be used and the personal data that may be being processed, for example the risk associated with customer retail data and biometric health data. It is important that when training generative AI models, significant efforts have been made to identify and remove any personal information and that developers are only collecting data that is adequate to fulfil their stated purpose.

- Collection of data at the user interface

Many developers of generative AI models utilise users' inputs and outputs that have been collected at the user interface to continue to train their models. Users may not be aware that the information from their interactions with the model could be used to further train the model. It is important that users are aware that generative AI models may use the information from their interactions with the model to further train models, understand the potential implications for their privacy, and consider whether they are comfortable engaging with the generative AI model on that basis. Organisations developing generative AI models could support users in this by making clear how users' content could be used, having clear data protection measures, privacy notices and have user friendly option controls around data sharing.

6. Conclusion

Generative AI provides many opportunities for access to information, exploration of ideas, creativity and the possibility of a world that is more accessible and inclusive than is currently the case.

However, the realisation of these opportunities will require awareness and mitigation of the risks posed by generative AI, supporting users to understand how the information landscape is changing and how to protect themselves against malicious actors.

A new way of generating and engaging with online content does not necessarily mean completely new media literacy skills are needed. Many of the skills referred to throughout this discussion paper are skills that are already needed to navigate the internet today, but with different applications and skill levels needed in increasingly complex environments.

These questions may begin a useful conversation as users, industry, and those working in media literacy seek to grapple with what generative AI may mean:

1. If generative AI opens up the possibility of more accessible and inclusive interactions, how do we ensure this technology can be accessed by those who may benefit the most?
2. What will generative AI mean for human-to-human interaction and how will this impact how we engage with the online and offline world?
3. What are users' attitudes towards generative AI? To what extent are users critically engaging with generative AI, and does that differ across age, gender, and socio-economic backgrounds?
4. What is it reasonable to expect of users and where are the boundaries of media literacy compared to technological innovation and regulation.