

Using generative AI in science communications: A guide for communications professionals

Introduction

Generative AI tools, such as ChatGPT, Google Gemini, and Microsoft Copilot, can create text and multimedia from just a few simple written instructions, known as prompts. For science communications teams, this can be a game changer: the tools can speed up the researching of complex scientific topics to aid the writing of blog posts or press releases, help to edit written or multimedia content, and spark new ideas for how to make complex scientific topics accessible to diverse audiences.

However, use of these tools in communications work also comes with risks. They can produce inaccurate or biased content, causing public confusion and misinformation; mishandle sensitive or embargoed data; and blur the line between human and machine voices. It is therefore essential that science communications professionals carefully and critically consider their

use of generative AI in communications work. This guide provides some key considerations and guidelines to help communications professionals use generative AI in a responsible, transparent way that strengthens, rather than undermines, good science communication. Hopefully, it serves as a starting point for the development or refinement of your own organisational guidelines.

Contents

Considerations

- Transparency and explainability
- Accuracy
- Data use and confidentiality
- Fairness, inclusion, and bias
- Copyright and attribution considerations
- Environmental sustainability
- Impact on creative professionals
- Misuse and risk of harm
- Training needs

Guidelines on using generative AI in science communications work

- Additional guidelines on using generative AI in content production
- Additional guidelines on using generative AI with multimedia

Use cases of generative AI in science communications

Considerations

The following points are some key factors to consider when making any decisions around:

- **Whether** you should use generative AI tools at all
- **Which tasks** you will use generative AI tools for
- **Which generative AI model** or tool you are going to use
- What **internal guidelines** you and your team are going to follow on the use of generative AI.

Transparency and explainability

It is best practice to declare the use of generative AI in content creation to both internal and external stakeholders. Within your team, you should standardise the degree of generative AI use which needs to be declared and the format your declarations will take. Such declarations are critical when generative AI tools have played a substantial role in content creation (rather than a more minor editing role, such as proofreading human-generated work).

You may also wish to use explainable generative AI tools (see the info box). These are specifically designed to allow users to better understand and trust the results and output created by the generative AI algorithms. The use of explainable tools can help to counteract the “black box” problem.

This transparency, both in your use of generative AI tools and in the functioning of the tools you use, helps build and maintain trust in your organisation and content.

Explainability in any form of AI is a design principle that aims to ensure that the process leading to the final output of a model can be fully understood, and explained, by a human user. This can be achieved in several ways, from developing models using publicly available training datasets, to showing the logic used after an output has been generated. It is a particularly important concept when AI is used in fields such as finance or healthcare.

The “**black box**” problem refers to any system which produces a specific output from a specific input without a clear explanation. It is a particular issue in the use of generative AI as it can make it significantly harder to make corrections to how the models are performing.

Accuracy

Generative AI tools can hallucinate – include unverified or inaccurate statements or data in the content generated. While you may be able to reduce the probability of hallucination by carefully designing and developing your prompts to increase their success rate – a process known as prompt engineering – AI-generated content can never be assumed to be correct.

It is therefore essential to fact check all AI-generated content. Pay particular attention to numbers and references, as these are commonly affected by hallucinations. An error could seriously undermine your organisation’s reputation as a trusted source, or even damage public trust in science.

Data use and confidentiality

Generative AI tools should have clearly stated policies around the storage, use, and reuse of uploaded data. These should be used to inform decisions about what data you choose to upload. “Free” tools typically use uploaded data for future training purposes, while internal or enterprise-licensed tools may retain data only for the organisation’s use (see the info box). Regardless of licence, the data retention policy for any tool should be reviewed before use. IT or legal teams may be able to assist with this.

Responsible use of personal and restricted data with generative AI tools both minimises the risk to your organisation and ensures that no harm can come to the subjects of any data. This type of data must only ever be uploaded to generative AI tools in accordance with GDPR and other applicable legislation.

You should also consider standardising an approach for other types of confidential data (embargoed information or business-sensitive data, for example) to ensure that any risk to your organisation is minimised.

There are several categories of generative AI tools you may encounter.

- **Free** tools do not require payment to use, and may not even require you to create an account before use. Many tools are available as a free version, but do not guarantee data security during free use.
- **Enterprise-licensed** tools are typically provided for commercial use by large technology companies. These will often be contracted and paid for at an organisational level, and usually have significantly more robust data policies. Many tools available for free are also available as an enterprise-licensed version, including ChatGPT, Google Gemini, and Microsoft Copilot.
- **Internal** tools are those specifically developed by an organisation for use only by their employees. These are typically held on organisational servers, and often retain data solely for company use.

If your organisation provides generative AI tools for use and you are unsure into which category they fall, check with your IT team.

You should also be aware that many well-known and longstanding software brands and programs are introducing generative-AI-powered features and tools. You may need to check whether any new feature you use is based on generative AI before distributing content which has relied on the feature's use.

Fairness, inclusion, and bias

Generative AI models are trained on large datasets gathered from the internet.

This method of training raises several concerns around how generative AI tools perform for different demographics. The datasets used may reflect societal biases and stereotypes (around gender or race, for example). Biases in the types of content available within those datasets may also mean generative AI tools perform less well for certain demographics (non-English-speaking audiences or people with visual impairments, for example).

AI-generated content should therefore always be carefully reviewed and edited, with any inherent biases in the data actively considered, to ensure it is inclusive and accessible.

Copyright and attribution considerations

Copyright and generative AI is an evolving legal area, and to minimise your organisation's risk you should be cautious. There are two key aspects to consider:

- The large datasets used to train generative AI models may include copyrighted content from individuals who did not consent to the use of their content for such a purpose.
- Copyright restrictions on the use of AI-generated content can be unclear and vary by country. In the UK, how existing copyright laws apply to AI-generated content is currently ambiguous. If you create any such content, you should be prepared to potentially have no control over its further use.

Broadly, be aware that the legal status of AI-generated content and generative AI models is likely to change in the future. Keeping up to date with major developments on this topic can help minimise the risk to your organisation.



Environmental sustainability

Generative AI tools have a significant environmental impact — primarily due to the high requirements for energy and water, which run and cool the servers that generative AI systems are housed on. Different query types have different levels of environmental impact — multimedia creation, for example, has a significantly higher impact than a simple text-only query.

With every query having an environmental impact, it is always important to consider alternative ways to answer queries, prioritise the use of generative AI tools for tasks where such use adds the most value, and identify tasks where the environmental trade-off means using generative AI may not be a sensible choice.

The known environmental impacts of a tool may be an important factor in your choice of generative AI tool. It can also be worth considering if, given its environmental impact, generative AI use is appropriate for your organisation at all — organisations with a sustainability focus, for instance, may feel that any use of generative AI could damage their public image.



Impact on creative professionals

Generative AI has made it much easier for many communications teams to produce the content they need — particularly multimedia — for a lower cost. However, this has led to concerns within the creative and writing industries that there will be a dramatic drop in the number of creative professionals able to successfully maintain a career. In this broader context, the use of AI-generated content may be viewed negatively by other stakeholders.

It is therefore worth considering whether the use of AI-generated content in external contexts — particularly in high-profile instances — may harm your and your organisation's reputation, both with the public and with creative professionals you may wish to work with in the future.



Misuse and risk of harm

Generative AI can be used for harmful purposes as well as for neutral and beneficial ones. While this may not directly affect your work, it should be a consideration when encountering AI-generated content, and when considering how any AI-generated content you produce may be viewed by other stakeholders.

Generative AI is already being used maliciously to help spread misinformation and to produce fake content intended to be interpreted as real. Such content can also be used for phishing or social engineering on a large scale. In this broader context — and coupled with the increasing prevalence of AI-generated content and the increasing difficulty in distinguishing it from human-authored content — poor and unreviewed use of generative AI tools risks eroding trust and confidence in the organisations using them.



Training needs

Generative AI tools are still tools — and a tool is only as good as its user. Teams that both appreciate the key considerations of using generative AI tools and understand how best to operate them will benefit most from their use.

It is important to make sure that you and your team are suitably trained in using generative AI tools effectively. There are many free and paid-for resources available for these purposes.

You should also be aware of the risks of deskilling — the decrease in the ability of people to perform tasks when they rely heavily on generative AI tools for assistance with those tasks. Deskilling can have negative effects on your job performance and the content you produce, and should always be considered when making decisions on long-term generative AI usage.

You may work with other stakeholders who you believe have not followed best practice for using generative AI tools in communications work. This guide could be a useful tool to aid you when discussing such issues.

Guidelines on using generative AI in science communications work

- Develop your own **organisational guidelines** on the **responsible use of generative AI** systems in communications work.
- Be aware that you and your colleagues are **fully accountable for any work produced with the assistance of generative AI tools**. For those in a management position, it is critical that this is **communicated to your team**.
- Standardise what **degree of generative AI use you will consistently declare**. As a minimum, you should transparently declare when generative AI has played a **substantial role in content creation**.
- In any declaration, make clear **how you have used generative AI tools** in producing the relevant content, and **how a human professional has subsequently verified or checked the content**. If possible, declare which specific tools you have used, making clear that they use generative AI.
- **Do not** use any generative AI tools with **missing or unclear data retention policies**.
- **Do not** upload **copyrighted data** into any generative AI tool. This includes scientific papers which have not been published under an **Open Access** licence.
- **Do not** upload **personal or private data** into a free generative AI tool. When using an enterprise-licensed tool, check the data storage policies carefully before upload to ensure that you are not breaking any organisational or legal policies.
- **Exercise caution** when considering uploading **confidential or embargoed data** into a free generative AI tool. Ensure that you are not breaking any organisational or legal policies by doing so. This type of data typically includes scientific datasets.
- **Check the conditions and restrictions** regarding AI-generated content when producing content for **distribution or use by third parties** (including platforms like EurekaAlert! and news outlets).
- If you work with **researchers who intend to publish** their work in an academic context, advise them to **check the conditions and restrictions of academic publishers** concerning the use of generative AI tools in the writing of academic papers.

Additional guidelines on using generative AI in content production

- **Fact-check** all content produced by generative AI tools against a reliable, non-AI-generated source. **Check numbers and references** particularly carefully.
- Edit all text copy produced by generative AI tools with a critical eye to ensure that it fits your organisation's **house style and tone of voice**. As part of this, make sure you **check for "tortured phrases"**, where a generative AI tool has inappropriately attempted to reword an accepted scientific term.
- Use terms such as "generative AI system" or "generative AI tool" in your content rather than referring to just "an AI" to help frame generative AI as software and **reduce the anthropomorphisation** of generative AI systems.
- Use the **correct terminology** in content which **describes an AI system** or model. Some basic definitions are available from The Alan Turing Institute – see [here](#) for their glossary of terms.

Additional guidelines on using generative AI with multimedia

The use of generative AI for creating and editing multimedia is a rapidly evolving area with many different viewpoints. Some publishers and platforms have policies banning all AI-generated figures, images, audio, and video, but others allow AI-generated multimedia provided the use of generative AI is explicitly declared.

Before choosing to use generative AI for multimedia-related purposes, it is worth considering whether you should instead engage a creative professional for this work. You may also want to consider whether there are specific topics for which using AI-generated multimedia may be acceptable (such as abstract concepts, or generative AI itself).

- **Do not remove or obscure any watermarks** that have been added by generative AI tools to indicate the content was AI-generated or edited.
- Use generative AI-powered multimedia editing tools **ethically** – you should **not change people's key or protected characteristics** (such as their ethnicity, gender, or facial expression).
- When working with existing multimedia, try to **work with the original creative professionals**, to gain their consent and involve them in the process.
- If you are provided with multimedia that you think may have been AI-generated or modified, **ask the original creator to clarify any use of AI**.
- When **editing existing multimedia**, use generative AI tools **only to correct and make minor edits**. This could include changing image dimensions or orientations, combining existing images or figures, enhancing the contrast or saturation of existing images or figures, or reducing background noise or visuals in videos. Ensure that **any modifications you make are permitted by the multimedia's licence**.
- **Do not make substantive edits** that change existing multimedia to such an extent that it is essentially new content.
- **Do not create or use AI-generated multimedia of a real person** without their express written permission and their approval of any script or prompt you use. You will therefore likely only be able to generate such multimedia of colleagues in your organisation. This should be clearly **declared in the caption** when the multimedia is published.
- **Do not create or use photorealistic AI-generated multimedia** that could be mistaken for an actual photograph.
- **Strongly consider avoiding** the use of multimedia that has been **created solely by generative AI**. If you do need to use such content, **edit it as necessary** to ensure it is appropriate for use, **check that it is appropriate** with a subject matter expert if relevant, and **clearly declare in the caption** that the multimedia is AI-generated.

Using AI-generated images from stock image libraries is a grey area. You may consider addressing it as a distinct aspect of your own organisational guidelines, to clarify whether you want to avoid using externally sourced AI-generated images entirely or instead have clear rules around their use. If you are choosing not to use AI-generated content at all, most stock image sites have filters which remove AI-generated content when searching. However, you should still exercise caution – if in doubt, it is more sensible to find a different image.

Use cases of generative AI in science communications

These typical use cases for generative AI in communications are provided to help encourage you to explore the possibilities of the various available tools.

1

Drafting and editing materials

Generative AI tools can help write first drafts or summaries of press releases, blog posts, web content, or other materials. They can be particularly useful when turning technical research papers into plain-language summaries. They can also assist with editing and proofreading for clarity, tone, or structure, as well as sense-checking explanations of unfamiliar topics.



Use responsibly: Only upload data and content within the guidelines above. Always fact-check, review, and edit the text before publication, and be transparent about generative AI assistance where appropriate.

2

Summarising research and background materials

When working with large documents (academic papers or policy briefings, for example), generative AI tools can produce concise summaries or key point lists to speed up internal understanding.



Use responsibly: Only upload data and content within the guidelines above. Check the produced summaries for accuracy and nuance.

3

Draft compelling pitches

Generate tailored email pitches for different media outlets, highlighting the most newsworthy angles of your research.



Use responsibly: Treat outputs from generative AI tools as a starting point, not a finished pitch. Ensure accuracy, avoid exaggerated claims, and tailor messages using human judgement and media knowledge. Only upload data and content within the guidelines above.

4

Create media Q&A briefs

Prepare anticipated questions and suggested answers for researchers before interviews, based on the topic and outlet. Or role-play how different types of journalists might respond to a pitch or press release, helping you refine your approach.



Use responsibly: Only upload data and content within the guidelines above.

5

Repurpose content across formats and audiences

Turn a research announcement into a blog post, tweet thread, LinkedIn update, or short video script. Generative AI tools can help brainstorm story angles, event titles, or social media captions. It is especially useful for exploring different tones, or different phrasing options for complex scientific topics suitable for different audiences.



Use responsibly: Use generative AI tools to explore options, not to finalise messaging. Ensure scientific accuracy is preserved when adapting content for different formats or audiences, and avoid oversimplification or sensationalism. Review all outputs for tone, accessibility, platform appropriateness, and institutional voice, with a communications professional having the final editorial judgement.

6

Identify emerging science topics

Use generative AI tools to scan global media and preprint servers for rising areas of interest that align with your institution's work.



Use responsibly: Validate trends using reliable external sources and human expertise. Be alert to algorithmic bias, regional blind spots, and over-representation of well-resourced research communities.

7

Drafting routine communications

Generative AI tools can speed up the creation of routine messages, such as internal updates, event invitations, or email templates.



Use responsibly: Keep sensitive or personal data out of prompts, ensure GDPR compliance, and review all outputs for tone, clarity, and appropriateness before sending.

8

Supporting multimedia and accessibility

Generative AI tools can create transcripts, captions, or summaries for videos and podcasts, or suggest edits to improve pacing and clarity.



Use responsibly: Check outputs carefully for accuracy and inclusivity. Avoid generating synthetic voices, images, or deepfakes that could mislead audiences or misrepresent real people, and follow platform or publisher rules on AI-generated media.

This guide has been produced for **Stempra**, a UK-based network for science, technology, engineering, and mathematics PR and communications professionals. We care about the quality of communication in science and want to support our members in negotiating the balance between encouraging public interest in STEM subjects, and retaining balance and accuracy. Our aim is for everyone involved in science public relations and press work to have the requisite skills, professional knowledge and recognition, and ethical guidance to do this.

Stempra is run by a voluntary committee of science communication professionals from around the UK, operating under the constitution that has been approved by the membership.

For more information, please visit our website or social media pages.



stempra.org.uk



LinkedIn



Bluesky

This guide was written and developed by:

Alex Durk, Springer Nature
Stempra committee member

Beth Baker
PLOS

Clare Green, The Francis Crick Institute
Stempra committee member

Ed Day, Science Media Centre
Stempra committee member

Edition 01 | FEBRUARY 2026

Thank you to
EurekAlert! for making
this resource possible

stempra

 **EurekAlert!**
 AAAS