



Generative AI in Grant Applications

White Paper

Dr Steven Wooding
Sam Gilbert

VILLUM FONDEN

Generative AI in Grant Applications
March 2025

Published by:
Villum Foundation
Tobaksvejen 10
2860 Søborg
Denmark

Authors:
Dr Steven Wooding and Sam Gilbert
Bennett Institute for Public Policy, University of Cambridge

Uses of Generative AI in this White Paper:
The right-hand image in Figure 1 and the image in Figure 2 were generated using Bing Image Creator. A first draft of the Conclusion was generated by giving Google Gemini access to the main text of the document. Features of Google Docs were used for spelling and grammar checking.

Image Credits:
The left-hand image in Figure 1 is an internet meme attributed to Karen Zack.
Figure 3 is inspired by the Gartner Hype Cycle.

Sharing and Attribution:
Sharing and reproduction of this white paper is permitted with proper acknowledgment of the source and the authors.

Preface

AI – and in particular generative AI based on large language models – is transforming the conduct of science and, perhaps even more, the reporting of science. This holds true for the writing of research papers, popular science articles, and, indeed, of grant applications.

Villum Foundation has a long history of supporting excellence in science in Denmark and of working closely with the research community to ensure alignment between the foundation's aims and the reality faced by researchers in their daily work.

It has therefore been an obvious step for us to reach out to our grantees to discuss the impact of generative AI on the writing of research proposals. On 29 October 2024 almost 100 of our Villum Experiment grant holders gathered for a day of inspiration and networking. As part of the programme we asked Sam Gilbert and Dr Steven Wooding to facilitate a workshop on the use of generative AI in the preparation of research proposals. The outcome is the present white paper. We wish to thank all the participants in the workshop for their valuable input, as well as Sam and Steve for expertly conducting the workshop and writing up the white paper.

We believe that the 'synthetic policy' outlined below, condensing the workshop outcome, can contribute to the important dialogue on the use of AI that we look forward to continue with the research community and with other stakeholders.

Anders Smith
Head of Programme

Thomas Bjørnholm
Chief Scientific Officer

Lars Bo Nielsen
Executive Director

VILLUM FONDEN

1. Executive Summary

Drawing on a workshop at the Villum Experiment Annual Meeting on 29 October 2024, this white paper discusses the implications of generative AI for the grant-making process, and proposes a draft policy for the use of generative AI in making grant applications. Recommendations have been co-created with almost 100 scientists from the Villum Experiment community.

We begin by briefly taking stock of current public discourse around generative AI, arguing that it is too consequential a development for foundations to ignore. We summarize the results of surveys of Villum Experiment applicants and participants, highlighting best practices for using generative AI in making grant applications.

We then describe the workshop, in which Villum Experiment participants used Large Language Models (LLMs) to develop and refine the outputs of group brainstorming, and analyse the results.

Finally, we combine the suggestions of the workshop participants, which they also prioritised, into the following synthetic draft policy for the use of generative AI, which foundations can use as a starting point for developing their own policies.



Synthetic policy on the use of generative AI for proposal preparation¹

- Each applicant remains fully responsible for the content of their application, regardless of AI assistance. Applicants should be aware of the shortcomings of the generative AI tools they use and they must thoroughly review all content and references.
- Applicants are encouraged to use generative AI tools to improve the flow of language (editing, spelling, grammar, phrasing, translation²) and review the logical flow of their applications.
- Applicants can use generative AI tools as idea storming partners.
- Applicants can use generative AI to aid exploration of previously identified literature.
- Applicants must disclose which categories of use, and which tools, they have used generative AI for in the preparation of their applications. Application forms should include checkboxes for common tools and use cases to allow easy disclosure.

¹ This policy only includes the perspectives of the scientists who participated in the workshop

² Although not covered in the workshop, the Villum Foundation secretariat's experience is that generative AI translation into Danish is poor at producing idiomatic translations, and generally benefits from revision by a fluent Danish speaker.

2. Background: should foundations believe the AI hype?

Generative AI represents an important step in the evolution of machine learning systems. Previously, the pattern-matching capabilities of such systems were used for classification and ranking tasks; now they can be used to “generate” new examples of a pattern – in text, image, audio or video format.

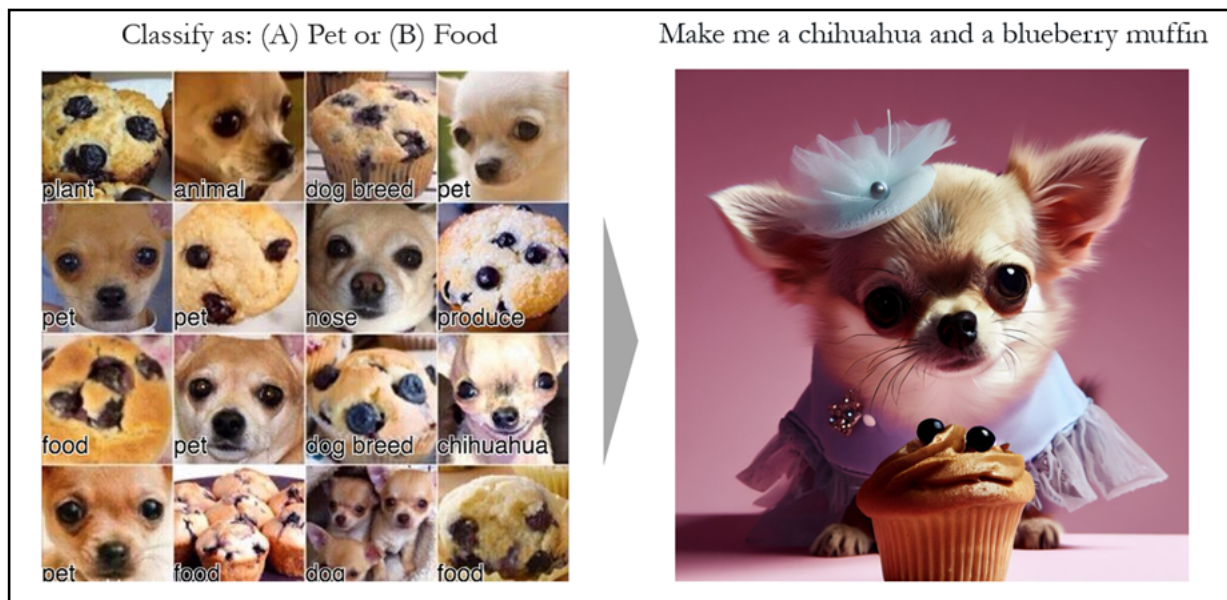


Figure 1: From image recognition to image generation.

Adoption of generative AI technologies has been rapid, and OpenAI’s consumer app ChatGPT is now used by more than [200m people each week](#).¹ Many find these new capabilities miraculous, raising expectations that generative AI will have a transformative impact on human society. The film producer Tyler Perry, for example, [paused the \\$800m expansion of his studio complex](#) after seeing a demo of the text-to-video model Sora, commenting:

I no longer would have to travel to locations. If I wanted to be in the snow in Colorado, it’s text. If I wanted to write a scene on the moon, it’s text...If I wanted to have two people in the living room in the mountains, I don’t have to build a set in the mountains, I don’t have to put a set on my lot. I can sit in an office and do this with a computer.²

However, recent months have brought a change in sentiment, as the limitations and drawbacks of generative AI have come into focus. The reputation of LLMs has suffered as a result of their tendency to “hallucinate” – that is, to fabricate plausible-sounding facts, hyperlinks, academic references, and [court decisions](#) (see Figure 2).³ In addition to longstanding fears about job displacement and existential risk from AI, the public seem increasingly concerned about AI companies’ [appropriation of data](#)⁴ for model training, about the effect of “[slop](#)”⁵ on the information environment, and about the climate impact of models’ water and [electricity consumption](#).^{6 7}

1 The Verge, “ChatGPT’s weekly users have doubled in less than a year”, 29 August 2024

2 The Hollywood Reporter, “Tyler Perry Puts \$800M Studio Expansion on Hold After Seeing OpenAI’s Sora: “Jobs Are Going to Be Lost”, 22 February 2024

3 The New York Times, “The ChatGPT Lawyer Explains Himself”, 8 June 2023

4 NMW, “Radiohead’s Thom Yorke, Abba’s Björn Ulvaeus and the BPI sign statement against using creatives’ work to train AI”, 22 October 2024

5 Wired, “AI Slop is Flooding Medium”; 28 October 2024

6 Goron Noble and Fiona Barry, “Power-hungry AI is driving a surge in tech giant carbon emissions. Nobody knows what to do about it”, The Conversation, 8 July 2024

7 We acknowledge that there are serious ethical and legal critiques of generative AI. The workshop and this white paper are predicated on the assumption that settled norms will emerge from the contested process of legal challenge, regulatory scrutiny and public debate that is currently underway.



- Steven A. Schwartz, a lawyer with 30 years' experience, made headlines by using ChatGPT for legal research
- He asked it to find relevant court decisions he could cite in a personal injury lawsuit (Mata v. Avianca, Inc.)
- ChatGPT “hallucinated” plausible-sounding cases that did not exist, which Schwartz included in court filings
- Schwartz and his firm were sanctioned and fined for acting in bad faith and making false and misleading statements in court
- At the hearing Schwartz revealed he had heard about ChatGPT from his teenage children and “falsely assumed [it] was, like, a super search engine”

Note: This image does not depict Steven A. Schwartz but was generated using Bing Image Creator.

Figure 2: “The ChatGPT Layer”: A Cautionary Tale

Are these technologies just over-hyped novelties, with negative externalities? Experience of the software industry suggests not, with a growing body of evidence demonstrating that they can materially increase productivity. The developer platform Github claims that its Copilot product enables code to be written [55% faster](#)⁸ than previously, while the CEOs of BP and Amazon have reported a [70% increase in developer productivity](#)⁹ and [savings of \\$260m](#)¹⁰ respectively. Admittedly these are self-serving claims that may exaggerate the benefits of generative AI in software development – but even relatively modest gains such as the [26% improvement in developer productivity](#) demonstrated by a recent preprint academic study would be highly significant when scaled across the global economy.¹¹

In our view, then, the economic potential of generative AI and its widespread adoption – both [in Denmark](#)¹² and [in academia](#)¹³ – mean it cannot be ignored by Danish foundations. It follows that foundations should develop policies on how generative AI technologies may be used in making and reviewing applications for grant funding.

⁸ Github, “GitHub Copilot for Business is now available”, 14 February 2023

⁹ bp, “1Q 2024 Results: Webcast Q&A Transcript”, 7 May 2024

¹⁰ CNBC, “Amazon CEO Andy Jassy says Gen AI saved \$260 million and 4,500 developer years”, 26 August 2024

¹¹ Cui, Zheyuan and Demirer, Mert and Jaffe, Sonia and Musolf, Leon and Peng, Sida and Salz, Tobias, The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers (September 03, 2024).

¹² Anders Humlum and Emilie Vestergaard, “The Adoption of ChatGPT”, The University of Chicago: Becker Friedman Institute for Economics

¹³ McDonald, Paula, Hay, Stephen, Cathcart, Abby, & Feldman, Alicia (2024) Apostles, Agnostics and Atheists: Engagement with Generative AI by Australian University Staff. QUT Centre for Decent Work and Industry, Brisbane, Qld.

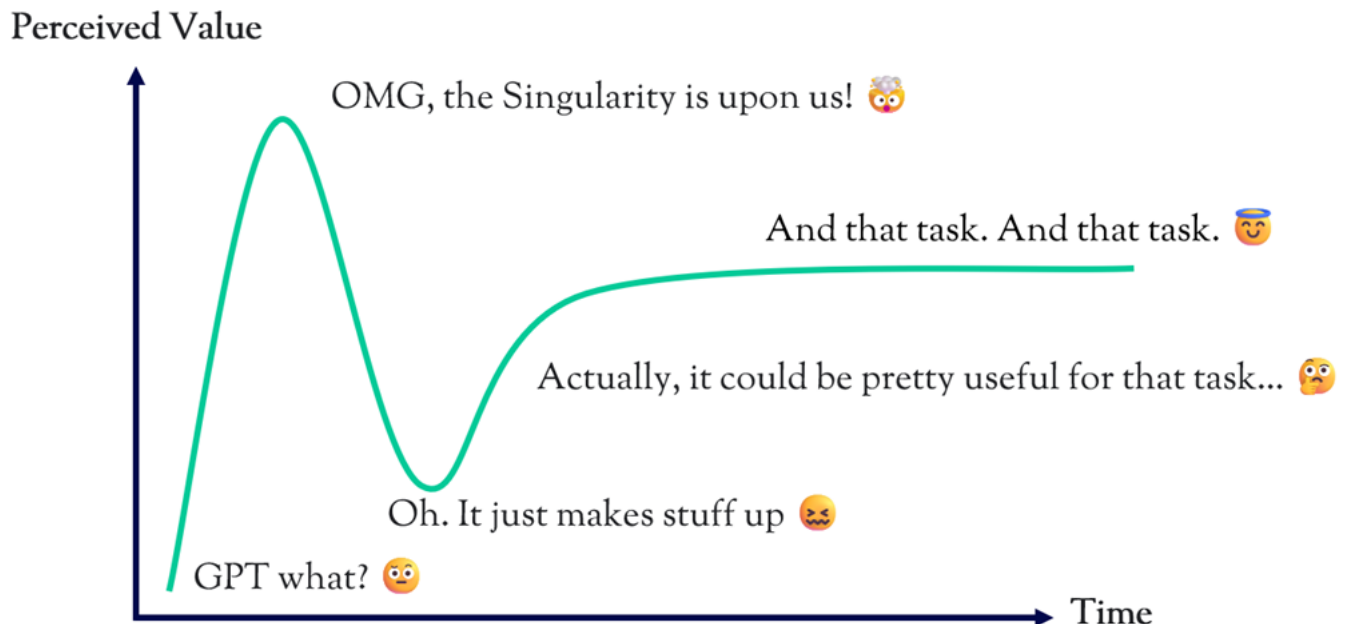


Figure 3: The Generative AI Learning Curve

3. How are scientists using generative AI today?

Prior to the Villum Experiment Annual workshop, invitees were surveyed about their familiarity with generative AI technologies and how they had used the technologies in preparing their grant application (101 invitees surveyed, 88 responded, 87% response rate). The majority of participants – all research scientists at Danish universities – reported that they had some experience with generative AI, and around a third (30 of 88) had used generative AI tools in preparing grant applications. Within this subgroup, the most popular use cases were editorial tasks including phrasing suggestions (21 of 30) and grammar checking (11 of 30).

Other use cases included writing early drafts, generating graphics, and reviewing relevant literature. A complete list is provided in Figure 5.

The use of AI tools for translation or “to correct typos, improve grammar, and shorten the text” is likely to be uncontroversial. However, foundations might be concerned that grant application text produced using LLMs will inevitably contain so-called hallucinations. In fact, the risk of hallucination can be mitigated with detailed “prompts” (that is, the instructions given to the LLM by the user). Best practice is to write a “system prompt”, which clearly explains the persona the LLM is expected to adopt, and a “user prompt”, which provides the specific

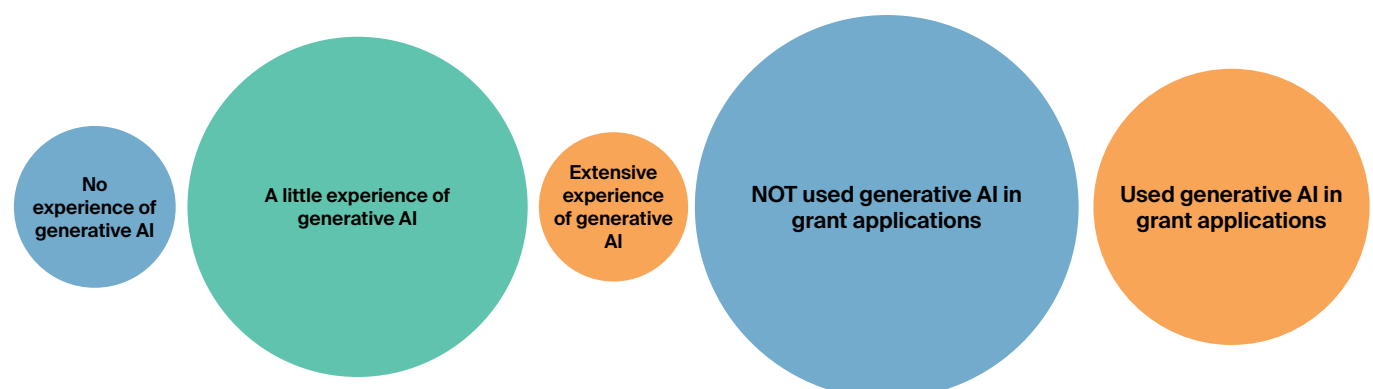


Figure 4: Workshop participants' experience with and use of generative AI

Use Case	Illustrative Verbatim Quote
Writing	“to generate a very first draft of the proposal”
Brainstorming	“to brainstorm words and acronyms for the project name”
Translation	“for translating the abstract to Danish”
Editing	“to correct typos, improve grammar, and shorten the text”
Graphics	“Figure 1b was created with the aid of DALL-E 3”
Literature Review	“to gather first information, which I then did fact checking on/ based additional search on.”

Figure 5: Use cases and illustrative verbatim quotes

instructions for the task the LLM is expected to complete, and the real data that is needed for it to complete the task effectively.¹⁴ (Examples can be found in the Appendix.)

Concerns over hallucinated academic references might also give foundations pause over the use of LLMs in literature reviews. There are, however, several tools researchers can use to assist with the process of literature review while mitigating the risk of hallucination. AI features embedded in Microsoft’s [Edge browser](#)¹⁵ and [Adobe Acrobat](#)¹⁶ allow users to “chat” with webpages and pdfs, with easily auditable referencing.




Figure 6: Auditable references provided by Adobe Acrobat AI Assistant

¹⁴ These techniques are less effective when applied to diffusion models and we would advise against using tools such as DALL-E3 and Midjourney to produce scientific or technical illustrations. For examples of pitfalls see Retraction Watch, “[Giant rat penis redux: AI-generated diagram, errors lead to retraction](#)”, 22 July 2024

¹⁵ See <https://www.microsoft.com/en-us/edge/features/copilot>

¹⁶ See <https://www.adobe.com/acrobat/generative-ai-pdf.html>

Google's [NotebookLM](https://notebooklm.google.com/)¹⁷ extends the scope of this functionality to 50 documents, which can include YouTube videos, audio files, or Google Slides. Meanwhile, AI applications like [Consensus](https://consensus.app/)¹⁸ and [ScholarAI](https://scholarai.io/)¹⁹ help with the discovery of new literature around a research question by referring to a [database of over 220 million peer-reviewed scientific papers](https://www.semanticscholar.org/).²⁰ In all these cases, the LLM is drawing on real data to generate a response, rather than relying on its training corpus, leading to significantly more accurate results. While these tools are not a substitute for close reading, they enable researchers to explore a much wider range of literature than was previously possible.

4. Policies for the use of generative artificial intelligence in grant applications

What do researchers see as the benefits and risks?

At the start of the workshop session we provided an introduction to recent developments in generative AI, its promise, perils and some pragmatic guidance on its use. The material from this presentation is summarised in the previous sections of the white paper.

The participants in the workshop were drawn from across the natural and technical sciences and included a range of seniorities from postdoc to heads of department.

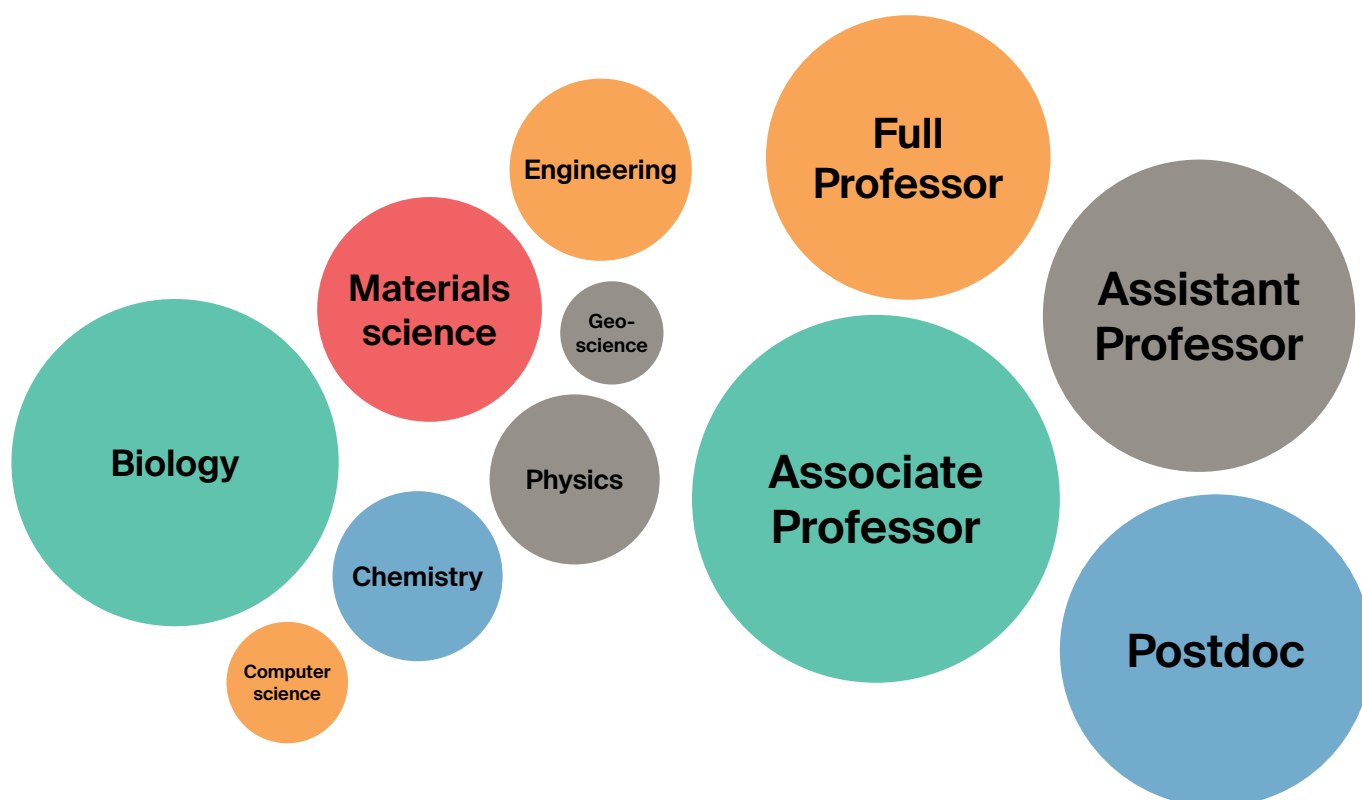


Figure 7: Variety of disciplines represented at the workshop

Figure 8: Distribution of career stages represented at the workshop

¹⁷ See <https://notebooklm.google.com/>

¹⁸ See <https://consensus.app/>

¹⁹ See <https://scholarai.io/>

²⁰ See <https://www.semanticscholar.org/>

After the introductory presentation we asked the participants to idea-storm ‘Good’ and ‘Bad’ issues around the use of generative AI for grant writing. We used a technique called ‘Crazy Eights’ which aims to broaden thinking by applying time pressure to idea-storming. Through this each participant produced eight issues – four ‘Good’ and four ‘Bad’ – each on individual Post-it notes. They then worked in table groups, comparing the issues they had identified and grouping them by similarity into named clusters on A1 posters. An example of the clustering is shown in Figure 9.

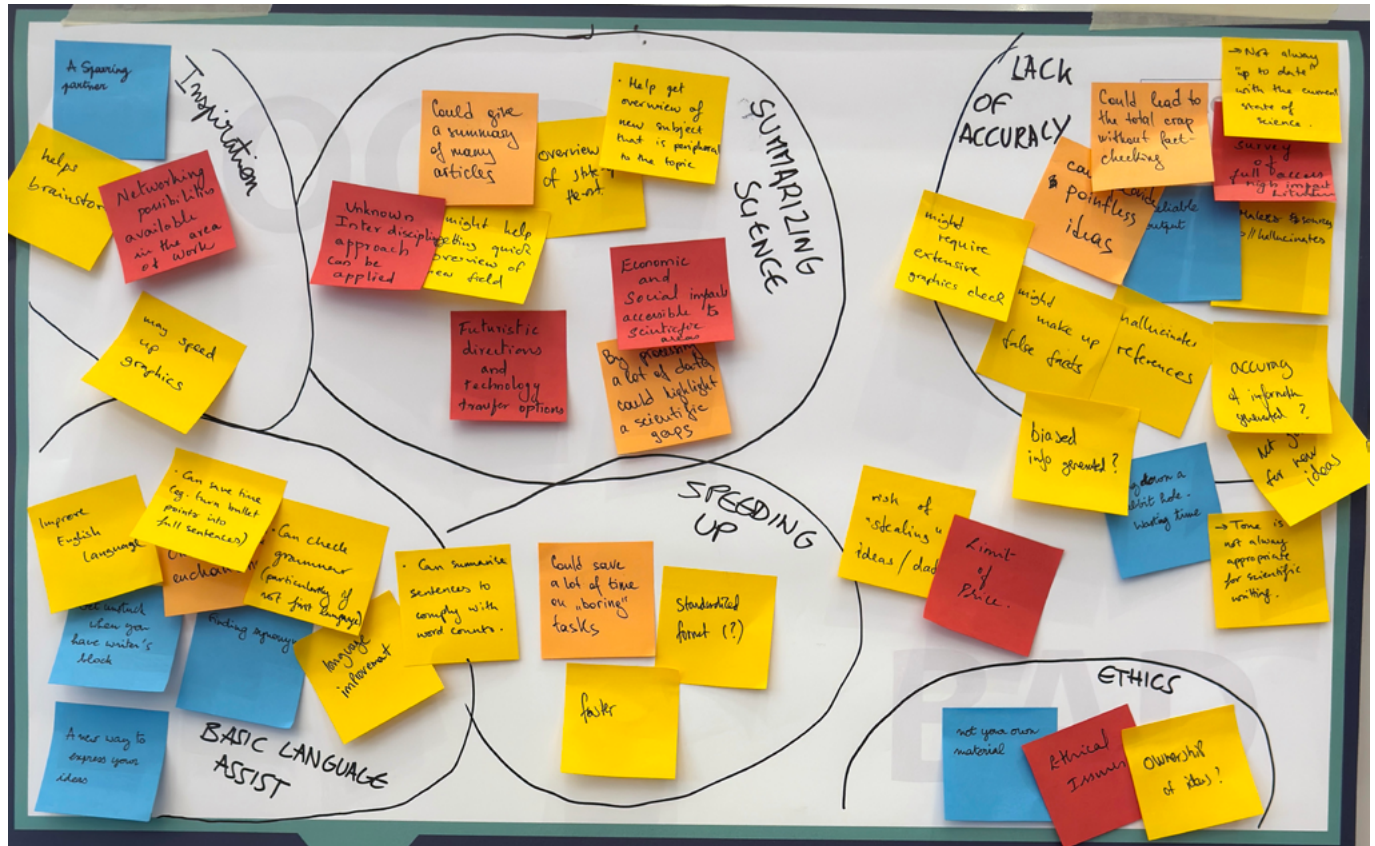


Figure 9: Clustering of ‘Good’ and ‘Bad’ issues around generative AI use

To synthesize these idea-storms we grouped similar concepts across tables. The visualisation in Figure 10 illustrates the key issues sized according to the number of tables who mentioned them. We have also tried to co-locate related concepts.

The three commonest positive aspects of using generative AI were seen as improving writing style; summarising literature; and as a sparring partner for brainstorming. It was also seen as useful for creating arresting images, acronyms and titles. The concerns were more diverse, but the dominating one was that generative AIs lie, hallucinate and make reasoning errors. There were also concerns that applicants might lose their distinctive ‘voice’ by using generative AI and it would lead to future dependence if writing skills were lost.

There was a split of opinion on whether generative AI would save time – a similar number of groups suggesting it would, to those suggesting that the additional checking and review needed would remove these benefits.



Figure 10: Areas emphasised in policy bullets

Reflections on the use of generative AI

At the end of the workshop we asked the researchers to reflect on the benefits, problems and their concerns about the use of generative AI after their experiences during the workshop. Probably because this activity came at the end of a busy day, the level of engagement was much lower than in the other activities of the afternoon, and only seven (of ten) groups made suggestions.

The suggestions largely mirrored those provided in the initial exercise described above. Benefits highlighted included improving language and saving time. Researchers also suggested generative AI can provide a foil for ideation, but could erode the social aspects of science and undermine confidence of researchers in developing their own writing skills. Just under half of the groups that completed this section (three of seven) used this opportunity to express moral and ethical concerns with the use of generative AI. Aspects that had not previously been mentioned were concerns about intellectual property and the sensitivity of generative AI to the system prompts that are provided.

One group asked the question ‘Where is the originality line?’ highlighting the challenge of any collaboration involving assigning and judging ideas.

What sorts of policy on generative AI use do researchers recommend?

We then asked the participants to devise a policy for the use of generative AI in the preparation of grant applications. We asked that the policy be expressed as a series of 3-5 bullet points and draw on their previous discussions about the ‘Good’ and ‘Bad’ issues identified in the first exercise.²¹

As hands-on guidance in using generative AI we provided example ‘system prompts’ the participants could use to ask for improved wording and critique their draft policies.

The initial bullet points the participants wrote tended to be short and informal. The policies that emerged from the process of review and comment by the AI Assistants were much more formally phrased, were always longer, and provided less room for interpretation. Note 4.1 shows examples of the revision of initial policy bullets.

Note 4.1: Examples of revision of policy bullets

- You are responsible for content → Researchers are responsible for ensuring the accuracy, integrity and inclusivity of all AI-generated content in grant applications
- Explicitly allow basic language assistance → Explicitly permit the use of generative AI for basic language assistance such as grammar and style corrections
- Generative AI is a useful tool for brainstorming → Brainstorming: Encouraged use! Leverage generative AI as a brainstorming tool to generate ideas, explore different angles and inspire creative thinking. This can help in developing innovative approaches and solutions.

²¹ One table could not reach consensus on a single approach and wrote two separate policies – one suggesting no restrictions on AI use and one emphasising transparency.

What aspects of generative AI policy do researchers think are most important?

To identify the policy suggestions with most support across the group we held a vote. We gave each participant 3 green stars (to imply endorsement and support) and 1 red star (to indicate a policy point they disagreed with). We asked participants to vote only on other groups' posters.

As many groups suggested closely related policy points, to summarise them we have grouped them by theme (and subtheme), this time visualising them according to the number (and type) of votes they received. In Figure 11 the number of supportive votes is shown in green overlapped with the number of negative votes in red. The 'Implementation approaches' cluster was a series of disparate ideas for implementation of a generative AI policy so we have differentiated it in blue.

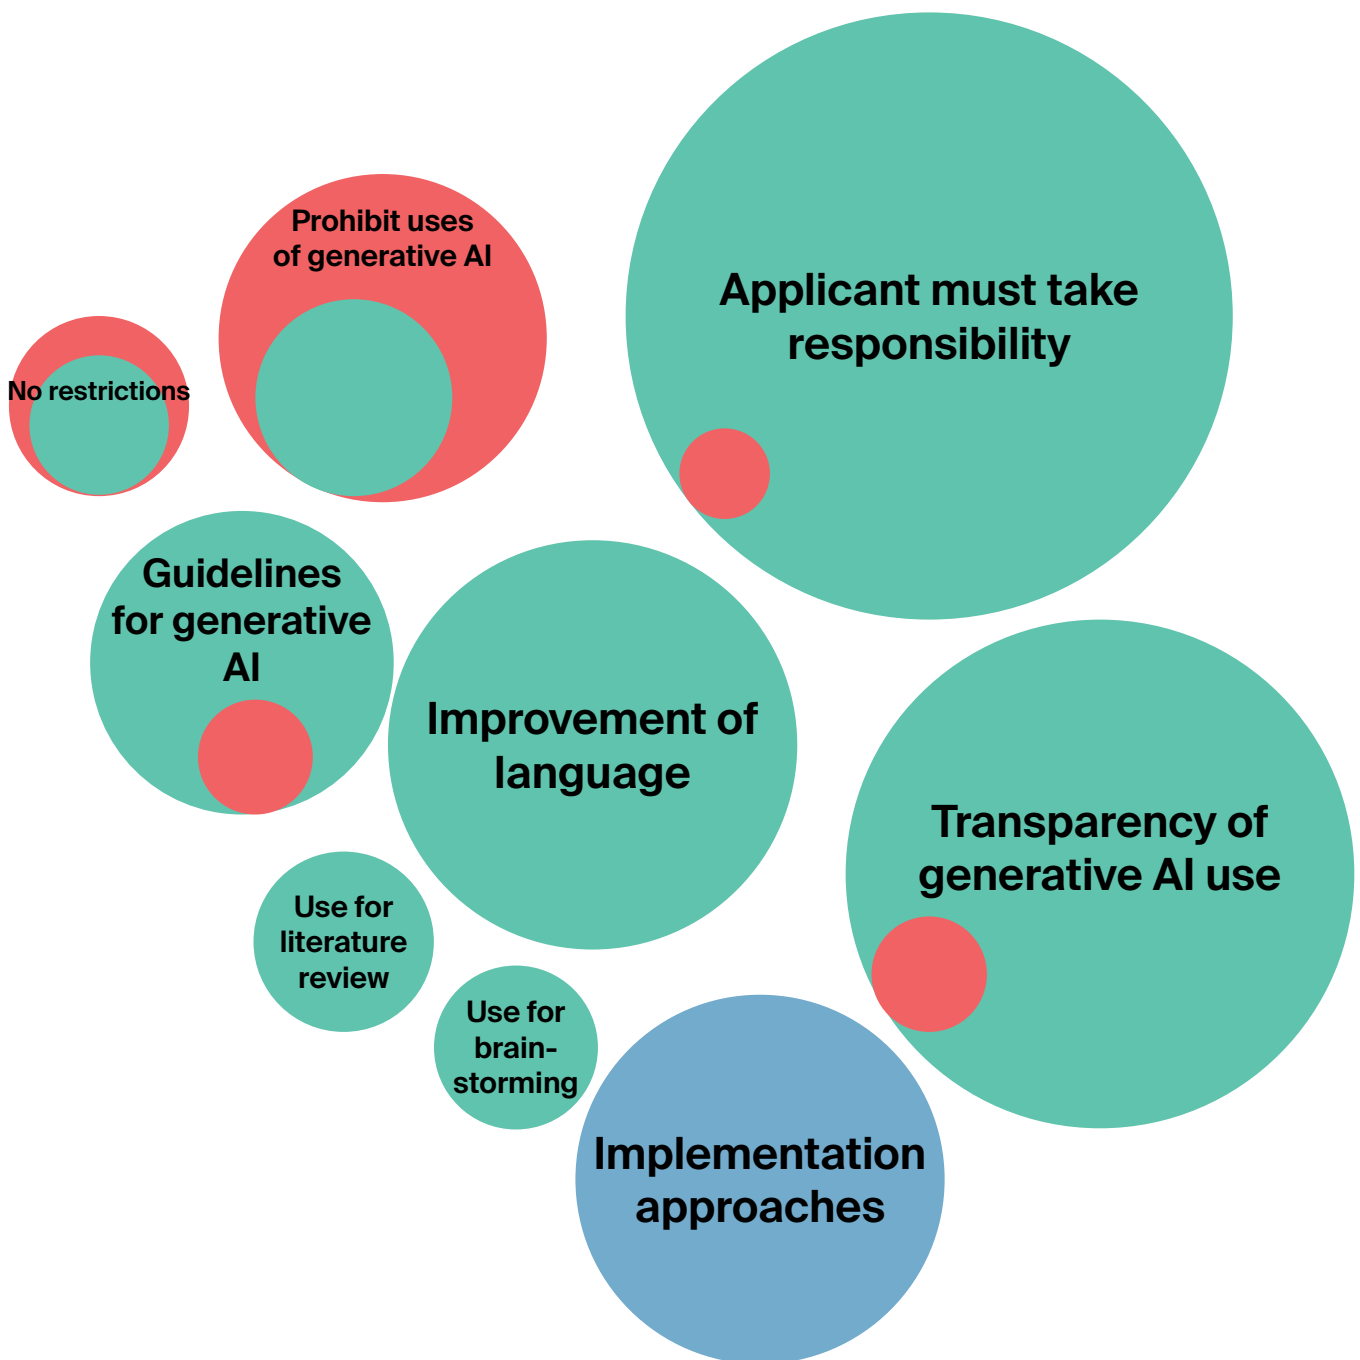


Figure 11: Clusters of suggested policy bullets

What might a policy look like?

Is a policy for use of generative AI valuable?

Only two clusters had more negative votes than positive ones – and both concerned whether use of AI should be controlled and whether it should be explicitly limited (see Note 4.2).

The idea of having ‘no restrictions’ was only suggested once and condemned by the majority of votes it received, suggesting general support for the idea of a policy on the use of generative AI. The second cluster with a majority of negative votes concerned the idea of prohibiting specific uses of generative AI.

Note 4.2: Policy bullets

- No specific policy on the use of generative AI in grant applications (5-, 3+)
- Do not use generative AI to generate the initial draft of your research proposal. The foundational ideas and strategy should be developed by the researchers themselves to maintain originality and intellectual integrity (11-, 0+)
- Prohibited AI uses: generating research ideas, drafting text from scratch, creating pictures or graphs (5-, 2+)

Applicant Responsibility

The strongest support was given for the idea that applicants must take responsibility for everything in their application – while allowing them to use generative AI tools; alongside advice from colleagues and other support as they saw fit. The three policy bullets, of six, with most support in this cluster are included in Note 4.3.

Note 4.3: Policy bullets

- Each applicant remains fully responsible for the content of their application, regardless of AI assistance (20+)
- Researchers are responsible for ensuring the accuracy, integrity and inclusivity of all AI generated content in grant applications (15+)
- All references and statements generated by AI must undergo thorough review and validation by human applicant to ensure scientific accuracy (14+)

Transparency in use of generative AI

Transparency in generative AI use was the second most strongly supported cluster, with researchers feeling that it was important to acknowledge how generative AI tools had been used. This might come from a trade-off between the complexity of cataloguing and verifying the detail of how generative AI has been used, and the desire to establish accountability for generative AI use. The four bullets, of eight, with the most support are shown in Note 4.4.

Note 4.4: Policy bullets

- Disclosure standards: applicants must provide detailed documentation of AI usage, including metadata and the specific AI tools employed (9+)
- Declare the AI tool and its purpose for transparency and ethics (9+)
- Disclosure: applicant must disclose AI usage in proposal (6+)
- Verification and training: University may verify AI use, training on ethical AI is encouraged (5+)

Improvement of language

Many of the groups had policy bullets that aimed to make clear that particular uses of generative AI were encouraged, where they judged these uses could improve the clarity of applications, making their assessment easier. These suggestions tended to be around improving language and grammar – with two suggestions of other uses: as a brainstorming aid and to support literature review. The most strongly supported policy bullets, of seven, in these clusters are shown in Note 4.5.

Note 4.5: Policy bullets

- Text editing and logical assessment: Encouraged use: Utilize generative AI tools to edit and evaluate the logical flow and structure of proposal, this helps ensure clarity and coherence (7+)
- Permitted AI uses: editing, spelling, grammar, phrasing, translation, enhancing text structure and clarity (7+)

Guidelines for generative AI

This cluster focussed on careful use and harm reduction rather than outright prohibition, it was a less draconian version of the prohibition cluster, and garnered much more support. It suggested use of generative AI should be done in a way that would minimise the potential problems and in line with wider ethical standards, but left detailed judgements about use in the hands of the researchers preparing applications. The three bullets, of four, with the most support are shown in Note 4.6.

Note 4.6: Policy bullets

- Originality: primary intellectual contributions must originate from the applicant (7+, 1-)
- Ethical standards: AI must align with university and academia standards (4+)
- Applicants using AI in grant writing must conduct AI-assisted literature searches within approved database-limited resources (2+, 1-)

Implementation approaches

A final cluster of suggestions concerned how a generative AI policy should be implemented. There was strong support for providing a simple mechanism to report on generative AI use in proposal writing – focussing on a set of tick boxes covering a standard set of uses, alongside a list of AI assistants/services. There was ambivalence about, tipping into resistance to, the idea that all funders, universities and other stakeholders in the research system should adopt the same rules for generative AI use.

Note 4.7: Policy bullets

- Include multiple check-boxes addressing specific ethical considerations, transparency in AI usage and the role of AI in idea generation (13+, 0-)
- Same rules everywhere (school, university, Villum) (4+, 7-)
- Ensure all rules are enforceable to prevent exploitation by unethical applicants (3+, 1-)
- Affirm that the use of generative AI will not impact the evaluation process (2+)

5. Synthetic Policy

Looking across the views expressed during the workshop we have drafted an outline policy for generative AI based on that suggested by those at the workshop. We aimed to hybridise the language used by different table groups, so it could act as a starting point for further discussion.

There is one notable omission from the policy: the use of generative AI for image generation. This is because the use of generative AI to produce images was only mentioned in one policy bullet – this bullet received five negative and only two positive votes – but was part of a wider set of prohibited uses. This absence is slightly surprising given the problematic examples we highlighted in our introduction to the topic – and could reflect that participants felt we had put image generation out of scope.

We judge it would be beneficial to explore whether it could be helpful to frame guidance about using generative AI tools in terms of the purposes they have been designed/tested for. For example – standard LLMs are notoriously bad at identifying literature, and many of the researchers had stories of phantom references, but those same tools can provide reasonable summaries of literature that has already been identified. This situation probably underlies some of the recommendations of the workshop.

However, as discussed in Section 2, newer generative AI tools have been designed to identify literature from a delineated set of sources. Indeed one group suggested careful use in identification of literature, but the idea only gained 2 positive votes against one negative one ('Applicants using AI in grant writing must conduct AI-assisted literature searches within approved database-limited resources' (2+, 1-)). Experimenting with this use case was not included as part of the workshop. This type of consideration highlights the challenge of drafting guidelines that are specific enough to be useful and flexible enough to accommodate the evolution of generative AI tools.

Synthetic policy on the use of generative AI for proposal preparation

- Each applicant remains fully responsible for the content of their application, regardless of AI assistance. Applicants should be aware of the shortcomings of the generative AI tools they use and they must thoroughly review all content and references.
- Applicants are encouraged to use generative AI tools to improve the flow of language (editing, spelling, grammar, phrasing, translation¹) and review the logical flow of their applications.
- Applicants can use generative AI tools as idea storming partners.
- Applicants can use generative AI to aid exploration of previously identified literature.
- Applicants must disclose which categories of use, and which tools, they have used generative AI for in the preparation of their applications. Application forms should include checkboxes for common tools and use cases to allow easy disclosure.

¹ Although not covered in the workshop, the Villum Foundation secretariat's experience is that generative AI translation into Danish is poor at producing idiomatic translations, and generally benefits from revision by a fluent Danish speaker.

6. Conclusion

This white paper has explored the implications of generative AI for the grant-making process. We have argued that generative AI represents a potentially transformative technology that foundations cannot afford to ignore. Drawing on a survey of research scientists, we have highlighted best practices for using generative AI in grant applications, focusing on the importance of smart prompts and real data in mitigating the risk of hallucination. We have also described the workshop in which Villum Experiment participants used LLMs to help co-create a draft policy for the use of generative AI in grant applications.

Reflecting the views of the scientists involved in the workshop, the draft policy emphasizes accountability and transparency, underlining the importance of researchers taking responsibility for the accuracy and integrity of all AI-generated content in grant applications. The policy also encourages the use of generative AI for brainstorming and improving the clarity of applications, and for assisting with literature review. We believe this policy can serve as a useful starting point for further discussion and development by foundations.

About the Authors

Dr Steven Wooding is Head of Research on Research in the Research Strategy Office at the University of Cambridge, an affiliated researcher at the Bennett Institute for Public Policy and a Visiting Research Fellow at the Centre for Science Policy. He leads the [Action Research on Research Culture](#) project. His central interest is how research itself works: from funding processes; to the social machinery of research and the wider impact of research. He has recently looked at the evidence behind using peer review for funding allocation and the value of block grant funding. He works with foundations to develop improved methods of selecting research projects to boost innovation. Steve has broader interests in the use of evidence in policy; and approaches to visualising concepts and data. Previously, Steve was Director of the Innovation Health and Science team at RAND Europe. He advises on various projects on the funding and use of research.

Sam Gilbert is an affiliated researcher at the Bennett Institute for Public Policy at the University of Cambridge where he has written influential reports on online safety, data ethics, crypto, web3, the metaverse, and AI. He is the author of [Good Data: An Optimist's Guide to Our Digital Future](#) (Hachette, 2021) and co-author (with Ankur Shah) of the e-book [Generative AI: An Executive Guide](#). Previously, he was Employee No. 1 and CMO at the insurtech “unicorn” ManyPets, and held senior roles at Experian and Santander.

About Villum Experiment

Villum Experiment is an innovative programme from Villum Foundation which targets out-of-the-ordinary research proposals – proposals that challenge the norm and have the potential to change fundamentally the way we approach important topics. Applications are judged anonymously to sharpen the focus on the research idea and to allow researchers to think freely in relation to their past merits. Members of the review panel are asked to emphasise ideas perceived as potential real breakthroughs with high-risk/high-reward. Each panel member has the option of awarding a single decisive vote for a proposal they find to be really out of the ordinary.

About Villum Foundation

Villum Foundation is a Danish charitable foundation supporting research and education in the technical and natural sciences, as well as environmental, social and cultural purposes in Denmark and internationally. The foundation was established in 1971 by civil engineer Villum Kann Rasmussen – the inventor of the VELUX window and founder of the company group VKR Holding A/S.

Appendix: Workshop materials

Timetable

Start	Duration	Topic
11:40	00:30	Presentation on generative AI assistants
12:10	00:15	Questions
12:25	01:00	Lunch
13:25	00:08	Overview of workshop session
13:33	00:07	Brainstorm good and bad
13:40	00:15	Cluster good and bad on to top half of poster
13:55	00:25	Draft bullets using clusters of ideas as starting point
14:20	00:10	Revise with LLMs
14:30	00:10	Supplement with LLMs
14:40	00:10	Critique with LLMs
14:50	00:05	Write final bullets on to poster
14:55	00:05	Demonstration of changing perspective
15:00	00:05	Discuss and add benefits/problems/concerns
15:05	00:05	Demonstration of changing perspective
15:10	00:05	Vote on bullets – build a portfolio

Materials

We arranged the attendees into table groups of 7-8 researchers and gave each attendee a pen, a small pad of post-it notes for individual brainstorming of ideas. Each table had an A1 copy of the poster to record their discussions. Each attendee was also given four sticky stars – 3 green and 1 red, to vote on their preferred, and least favourite, policy suggestions. The poster is shown to the right.

GROUP:

GOOD

BAD

DRAFT BULLETS

REVISED BULLETS

BENEFITS

PROBLEMS

CONCERNS

Workshop presentation

We guided the workshop through a presentation explaining each stage and using the example of a collaboration policy for human collaborators as an example.

AI Prompts

We used the document below uploaded as a Google Doc to provide attendees with starting prompts for working with the generative AI assistants. We duplicated the instructions and prompts so each group could have their own copy to modify as they saw fit.

Introduction

This document provides example prompts that you can take and modify to improve the Generative AI policy bullet points that you write in the workshop. You can also use it to store drafts of bullet points or other notes as you work in your groups.

There is a 'template' section at the end of the document in case you accidentally delete something from your table's section.

Popular LLM-based applications:

- <https://claude.ai/>
- <https://copilot.microsoft.com/>
- <https://chatgpt.com/>
- <https://gemini.google.com/>

Example Generative AI policies from academic publishers

- <https://us.sagepub.com/en-us/nam/artificial-intelligence-policy>
- <https://www.elsevier.com/en-gb/about/policies-and-standards/the-use-of-generative-ai-and-ai-assisted-technologies-in-writing-for-elsevier>
- <https://taylorandfrancis.com/our-policies/ai-policy/>

1. Revise the brainstorm output with LLMs

Give bullets to LLMs and ask for rewording, revise accordingly

Example Prompt

You are an experienced researcher whose responsibilities include writing policies and guidelines at an elite university.

The following bullet points are the output of a brainstorming session with other researchers, and list the key elements of a new policy covering the ethical use of generative AI technologies in writing grant applications.

Please edit these bullet points for clarity, succinctness, and academic tone.

[insert bullet points]

2. Ask for improvements to the revised output with LLMs

Give set of bullets to LLMs and ask for suggestions of improvements (eg what has been missed), revise accordingly

Example Prompt

You are an experienced researcher whose responsibilities include writing policies and guidelines at an elite university.

The following bullet points are a draft policy covering the ethical use of generative AI technologies in writing grant applications.

Please suggest substantive additions and improvements.

[insert bullet points]

3. Critique from different perspectives with LLMs

Ask for the likely objections from different groups of perspectives, or the likely problems that such a policy would cause, revise as necessary

Example Prompt

You are the head of the research standards office at an elite university.

The bullet points below are a draft policy covering the ethical use of generative AI technologies in writing grant applications. The policy has been prepared by your team and shared with you for comments.

Please critique it from the perspective of [Diversity, Equality and Inclusion / data ethics / managerial efficiency] and recommend improvements.

[insert bullet points]

VILLUM FONDEN

Villum Fonden
Tobaksvejen 10
2860 Søborg
Denmark

info@villumfonden.dk
www.villumfonden.dk