

The Norwegian Board of Technology

Date
31.08.2026

Artificial intelligence in knowledge work – experiences with the use of generative AI in technology assessment

English summary of the report

Generative artificial intelligence (AI) is expected to significantly transform knowledge work. For the Norwegian Board of Technology, this raises questions about how AI can strengthen or challenge our work in providing technology policy advice to the Norwegian Parliament.

The Norwegian Board of Technology is an independent advisory body to the Norwegian Parliament. We are a typical knowledge-based organisation, where language, textual analysis and communication are central to our work. These are areas in which generative AI has considerable potential. To assess what generative AI could mean for our organisation, we conducted an experience-based study focusing on our core tasks. During two periods in 2025, all nine employees at the Norwegian Board of Technology tested AI tools in a total of 35 cases. All participants documented how they used the tools and their perceived usefulness in structured logs. Based on these logs, we systematised our experiences and analysed usefulness, quality and time spent.

Our study is set against a backdrop of numerous analyses and studies suggesting that generative AI will have a major impact on knowledge-based occupations, either by replacing human work or by complementing it, freeing up time and stimulating creativity. Despite studies pointing to considerable potential for generative AI in knowledge work, the large-scale effects have so far failed to materialise. Surveys show that many employees experience substantial individual benefits, but few organisations can document measurable gains.

We chose a low-threshold approach, using off-the-shelf AI tools without customisation or integration with internal systems. AI was used as an assistant for individual tasks, with project manager asking questions and the AI providing responses. Our experiences can be summarised in the following main findings:

- **The benefits vary depending on the type of task.** For *structured and clearly defined tasks*, such as summarising, translating and producing meeting notes, we obtained good results quickly. For *unstructured and analytical tasks*, the main value lies in the working process itself. In these cases, AI serves as a sparring partner in an extended dialogue rather than as a tool for quickly producing a finished result.

- **AI contributed little or no time savings.** We saved some time on highly structured tasks with clearly defined objectives. However, these account for only a limited share of our overall work. For creative and less structured tasks, it is uncertain whether we saved any time at all. Providing AI with sufficient context and steering the dialogue towards relevant results was time-consuming. Generated responses may also contain errors, lack precision or reflect priorities that we do not share. In clearly defined tasks such as summarising and translating, errors were generally relatively easy to identify and correct. The situation is different for more complex and analytical tasks: texts may appear well structured and well written while containing substantive weaknesses that are not immediately apparent. Identifying these weaknesses can require considerable review and revision.
- **AI can affect our work in ways that both strengthen and challenge our quality criteria of reliability, relevance and legitimacy.**

Reliability means that our results are of high professional quality and can be verified. This requires precise language, transparent sources and a balanced presentation of different perspectives. We found that AI can help improve wording and thereby strengthen communication. In several cases, AI also helped identify a broader range of arguments and perspectives than we had initially considered. At the same time, there is a risk that AI frames an issue in a particular way, giving the tool the power to define how an issue is understood that can be difficult to detect. Another challenge is that AI-generated text often appears fluent and convincing even when the content is incorrect. This can create a false sense of confidence and make weaknesses more difficult to identify.

Relevance concerns how well the knowledge produced through technology assessment responds to political and societal needs. It must address the right issues, be understandable and available at the right time, and provide insights that can inform policymaking. Ensuring that topics and content are relevant to political decision-makers requires a sound understanding of the political landscape, current issues and the available policy options. Our experience suggests that AI can only replace this contextual understanding to a limited extent. When it comes to making texts accessible to their intended audience, however, AI can contribute to clearer structure and effective wording. At the same time, we found that AI-generated text can become generic or overly sensational, or subtly alter the intended meaning.

Legitimacy concerns transparency, independence and trust in the interaction between knowledge production and policymaking. Technology assessment must remain free from party-political or interest-based influence. In some cases, the use of AI may strengthen independence by providing rapid access to a broad range of information without direct input from individual stakeholders. At the same time, AI models may contain inherent professional or political biases because they are trained on large datasets that are not neutral. In some cases, we found that AI gave greater prominence to topics, individuals or policy measures that had received substantial media attention, while other relevant perspectives were less visible.

- **The expertise and judgement of the individual project manager remain essential.** The use of AI does not change the fundamental standards governing our work. Responsibility for assessments and recommendations cannot be delegated to a tool. The project manager must still steer the process, critically evaluate AI-generated suggestions and take responsibility for the final result.

The relatively modest gains we have observed so far are not surprising. One reason may be that knowledge work is complex and context-dependent, relying heavily on judgement, contextual understanding and continuous assessment, whereas AI models perform best when tasks are clearly defined and bounded. In addition, our quality criteria require us to retain full control over both the process and the final result, limiting how much can be delegated to an AI model. Finally, although the benefits are currently limited, this may change. AI models are improving rapidly, and more domain-specific solutions and new forms of human–AI collaboration may change working practices over time.

We believe that increasingly capable AI models and new functionality is likely to become more relevant to knowledge work in the years ahead. We will therefore continue to test the technology and develop our practices as new opportunities and challenges emerge.