



Government-wide position on the use of generative AI

Government-wide position on the use of generative AI

The government wants to harness the power of generative AI to benefit its citizens and employees. The technology offers plenty of opportunities to address societal issues and improve services, such as simplifying administrative processes, strengthening data analytics and improving communications. Generative AI technology offers opportunities for the government in numerous areas. At the same time, government use of generative AI is not without risk. To maximise the opportunities of AI technology as a uniformly operating government, it is essential to adopt a clear position on the use thereof. This position offers clarity on how government organisations can use generative AI effectively, responsibly and ethically, providing a framework within which government organisations can act, with room for innovation and attention to risk. We have taken the current state of technology into account, and this position can be updated as necessary to reflect new insights.

This position has been created through active cooperation between all levels of the government, coordinated by the Ministry of the Interior and Kingdom Relations (BZK). A guide has been prepared to support government organisations in the practical implementation of this position.¹

What is the position?

The government promotes innovation for improved service delivery using generative AI. In addition, government organisations are encouraged to include this technology in the search for solutions to problems.

Generative AI can be procured, developed or used when:

- The use of generative AI complies with existing laws and regulations (e.g., the GDPR, the AI Act, and the Constitution).
- The purpose served by the use of generative AI is sufficiently clarified.
- Risk analyses have been conducted to determine whether a specific generative AI application is also acceptable and safe. These must be conducted with input from as diverse a group of internal or external experts as possible. The Algorithm Framework can be valuable in selecting an appropriate review tool for the risk analysis (e.g., a DPIA or pre-DPIA scan when processing personal data).² Important questions to consider in a risk analysis include³:
 - What important problem does the use of this technology solve that cannot be solved by other means (necessity and proportionality)?
 - Are the right prerequisites in place (e.g., AI literacy, clear governance or information management)?
 - Have lower-impact alternatives been explored (subsidiarity)?

1 For more information see the Government-wide Guide to the Use of Generative AI.

2 The algorithm framework contains laws and regulations, tips and tools for the responsible use of algorithms and AI, providing a good, though not exhaustive, overview of several common tools. See: <https://minbzk.github.io/Algoritmekader/>.

3 These considerations were inspired in part by the opinion issued by the Association of Provincial Authorities (IPO) on the use of ChatGPT by provincial authorities: <https://www.ipo.nl/nieuws/advies-over-het-gebruik-van-chatgpt-in-provincies/>.

- Is there a plan in place for dealing with the limitations of the technology (diligence)?
- Is there a clear exit strategy that makes it possible to stop processes quickly if problems arise (safety)?
- The use of generative AI and its associated risk assessments fits within the relevant policy context. Among other things, this means that organisations include the appropriate roles in decision-making regarding the use of generative AI, and that it aligns with existing policy frameworks.
- Specific agreements have been made with suppliers during the procurement of generative AI models and/or applications. Important topics for these agreements (and/or business terms) include data sharing, data retention, and whether or not to allow model retraining with user input (e.g., prompts).

For whom is this position intended?

This position is intended for all government organisations that procure, develop or use generative AI. In practice, many government employees—data scientists, procurement consultants, policy officers, administrators, and more—will be affected by it.

When does this position apply?

- This position applies to both individual use of generative AI by government employees and project-based use. It applies to all types of work and all stages of a project. This includes, research on improving internal processes using generative AI, experimental use in pilots or living labs, and performing business management tasks.
- The position also applies to external parties providing services to government organisations using generative AI.
- In the event of unavoidable use, such as generative AI built into search engine search results or automatic machine translations of websites, individual employees are expected to use the technology with integrity.

What organisational and societal considerations come into play?

The use of generative AI by a government actor is always part of a broader organisational and societal context. It is important to take this into account, both directly in the organisation and with regard to the broader societal effects that can be associated with the use of generative AI. Specifically, this includes:

- When using generative AI, employees should be adequately informed about the effective and responsible use of the technology. This includes fostering AI literacy by offering training and guidelines for the responsible use of AI.⁴ In addition to responsible use of AI, training is necessary to unlock the full potential of a generative AI application.
- The use of generative AI applications by government employees under consumer conditions, including for individual use, is not allowed because agreements cannot be made with the suppliers.⁵ New (generative) AI systems may be explored in a controlled, non-production environment, provided no sensitive or confidential data is used. Government-wide exploration is underway to identify promising ongoing experiments that can be scaled up. Efforts are also underway to improve agreements with suppliers for the procurement of generative AI applications.
- Because not all generative AI applications are equally adept at processing Dutch and other European languages, and to avoid undesirable dependencies on countries outside the EU, the use of applications developed and managed in Europe is recommended. These may include explicit focus on Dutch-language resources during the training process.
- The preferred approach is to use open-source software and models, or to work towards using these.⁶ Open source models offer more insight and transparency. However, transparency should not come at the expense of the security of generative AI models.

Technical measures and developments

- Generative AI applications can be extremely helpful but are sometimes based on outdated or incorrect information and can make mistakes. Because they are designed to provide plausible answers rather than completely accurate ones, these applications can also exhibit hallucinations.⁷ This problem is particularly prevalent when they are used in unfamiliar or

4 The European AI Act states that organisations using AI must ensure an adequate level of AI literacy among employees.

5 For more information see the Data Protection Authority's [recommendation](#) on the use of generative AI in government, which emphasises that the use of online generative AI involves numerous risks involving personal data.

6 For a definition of open source AI, see the Open Source Initiative definition: <https://opensource.org/ai/open-source-ai-definition>.

7 Factually incorrect information generated by a large language model (LLM). The answers generated by the

highly specialised situations. In such cases, it can often be difficult to find out exactly where the information comes from. It is important to take this into account in development and procurement by ensuring that outputs are based on up-to-date, reliable and context-specific information. There are also techniques to reduce the probability of error and better tailor the answers to the specific context in which generative AI is used.

- It is prudent to examine which AI model(s) will deliver the best results on a per-application basis. Smaller models may be better suited to certain use cases, especially for specialised or complex topics, because they are simpler and more adaptable to the context. For government organisations and others operating within the same field, we recommend collaboration in model customisation, as this contributes to improved traceability, consistency, efficiency and reliability of the technology.

For a more detailed explanation of organisational and technical measures that can be concretely taken to ensure the responsible use of generative AI within government organisations, see the guide to the responsible use of generative AI.

Relationship between this position and the European AI Act

The European AI Act includes requirements for AI models that are capable of performing many different tasks and can be integrated into many different AI systems (often using generative AI). The AI Act refers to these as 'general-purpose AI models' (GPAI models). Providers of such models are required to disclose information on how the model works with developers integrating the model in their systems. It should also be clear how copyright is respected, and all AI-generated content must explicitly indicate that this is the case. These requirements help to better understand which content was created with AI and may be questionable, incorrect or unverified. In many cases, government parties will be providers and/or users of an AI system based on a GPAI model. This means that when using generative AI, government organisations will frequently face requirements (e.g., regarding transparency or high-risk) from the AI Act. In addition to the AI Act, other legal frameworks may apply to the use of generative AI. The government-wide guide provides more information on this.

model are not based on the given inputs or actual information from the training data. Hallucinations can have various causes, including a lack of specific information in the training data, a lack of context, or erroneous and inconsistent information in the training data.

Ministerie van Binnenlandse Zaken en Koninkrijksrelaties

P.O. Box 20011

2500 EA The Hague

Date

January 2025