Ministerie van Binnenlandse Zaken en Koninkrijksrelaties

www.rijksoverheid.nl
www.facebook.com/minbzk
www.twitter.com/minbzk
www.linkedin.com/company/ministerie-van-bzk

Reference

Your reference

Appendix(es)
1

Date	17 January 2024
Subject	Cover letter for the (government-wide) vision on generative AI

We hereby present to you, on behalf of the Minister of the Interior and Kingdom Relations and the Minister of Social Affairs and Employment, the government-wide vision on generative artificial intelligence (AI).[1] The presentation of this vision responds to the motion by the members Dekker-Abdulaziz and Rajkowski, which was adopted by a large majority of the House of Representatives of the Netherlands in April 2023.[2]

Through this government-wide vision, we stress the importance of taking action in view of the opportunities and challenges of this disruptive yet promising technology. Generative AI should be used in the service of enhancing human welfare, prosperity, sustainability, justice and security. In doing so, we start from the assumption that generative AI is human-centered and meets our public values at all times. This vision aligns with the ambitions of the 'Value-Driven Digitalisation Work Agenda': everyone can participate in the digital age, everyone can trust the digital world, and everyone is in control of their digital lives.[3] The government focuses on the impact of this new digital technology on society in particular. This is closely related to the Strategy with regard to the Digital Economy[4], particularly in terms of creating the right conditions for well-functioning digital markets and services, stimulating digital innovation and strengthening cybersecurity.

This letter summarises the key points of the vision. The government-wide vision is attached in its entirety as Appendix 1 to this letter. This vision also addresses of a number of motions and commitments.

[1] Generative AI is a type of AI that uses complex algorithms to generate new content such as text, images, computer code or videos. ChatGPT is a chatbot that has gained recognition in the field.
[2] Amended motion by members of the House of Representatives Dekker-Abdulaziz and Rajkowski of 4 April 2023
'Integrale visie op nieuwe AI-producten' (Parliamentary papers 2022/23 26 643, no 1003).
[3] https://www.rijksoverheid.nl/documenten/rapporten/2022/11/04/bijlage-1-werkagenda-waardengedreven-digitaliseren
[4] https://www.rijksoverheid.nl/documenten/kamerstukken/2022/11/18/strategie-digitale-economie

The vision has been developed in collaboration with public sector bodies, implementing organisations, civil society organisations, businesses, knowledge and higher education institutions, AI developers and citizens. Considering the importance of this rapidly evolving topic, these discussions will continue even after the publication of this vision.

Introduction

AI is a system technology that has a major impact on all domains and sectors of our society and will affect all areas of government policiy.[5] Generative AI is a specific form of AI. It assists people in generating text, computer code, images and audio and has already proven to be a powerful extension of human analytical and creative abilities. Especially since the launch of chatbot ChatGPT in late 2022, and the introduction of other generative AI tools such as Google Bard, Midjourney and DALL-E, the use of AI has become part of the daily lives of many people in the Netherlands.

When combined with related technologies, generative AI offers promising possibilities for addressing societal and scientific challenges. For example, generative AI can contribute to making accurate and efficient medical diagnoses, improving healthcare and contributing to medical research. Economists predict productivity growth not only for large companies, but also for small and medium-sized enterprises (SMEs), as its deployment makes performing tasks such as financial analysis and legal processes cost-efficient. In the cultural sector, generative AI is used to support the creative process, for instance in creating marketing content or generating descriptions of artworks.

At the same time, there are significant risks, and the realisation of public values and fundamental rights such as transparency, privacy, autonomy and non-discrimination may come under pressure from irresponsible development and the use or abuse of generative AI. For instance, generative AI can be used to create and disseminate misinformation; it can reinforce discriminatory dynamics; and it can reinforce socio-economic inequalities. The use of generic AI can also result in systemic risks such as disruption of vital processes.

Developments in generative AI continue unabated. Major powers and tech companies recognise the potential of such advanced AI systems and are investing significant funds in their further development.

The wide availability of generative AI, the scale and the pace at which the technology is currently developing, require a future-proof vision. Your House of Representatives of the Netherlands has stressed the importance of this on several occasions through various motions. The actions resulting from this vision will help Dutch society reap the benefits of generative AI and cope with the risks associated with it.

[5] We follow the OECD's recently revised definition of "AI system" (2023): a machine-based system that derives, for explicit or implicit purposes, from the inputs it receives, how to generate outputs such as predictions, content, recommendations or decisions that can affect physical or virtual environments. Different AI systems vary in their degree of autonomy and adaptability after their implementation/deployment.

Impact and underlying issues

Generative AI will have a (major) impact on our economy, society and on both central and decentralised public bodies. It is important to note that the possible applications of generative AI cannot always be identified beforehand as either opportunity or risk. Whether it ultimately manifests as an opportunity or a risk depends on the specific development, the application of the technology and the intentions or expertise of the user. This emphasises the importance of adequate control over development and use of generative AI by the government.

i. Opportunities and possibilities

Generative AI offers several possibilities to optimise processes and support users in a wide range of tasks. It can contribute to improving efficiency, cost savings, better decision-making, better services and numerous innovative solutions. Wide applicability creates opportunities for citizens, businesses, the government and society as a whole. Given the potential impact of generative AI models and systems on society and the economy, the government is committed to promoting responsible experimentation and use across various sectors and domains.

Generative AI can bring productivity growth. New tasks and responsibilities are also expected to emerge within organisations and in the economy as a whole as a result of the deployment of generative AI. This can lead to the creation of new employment opportunities and the emergence of new professions, necessitating an ongoing development of skills and expertise within society. Furthermore, the content of many work-related tasks will undergo changes. There lies an opportunity to enhance work by automating repetitive tasks and creating more room, for instance, for social interaction or creative thinking. In various industrial processes, generative AI can quickly produce a large number of design alternatives. As a result, the design process in the manufacturing industry could potentially be significantly accelerated. The impact of (generative) AI on work is expected to be widely felt.

In addition, the use of generative AI within organisations can contribute to the improvement of internal organisational procedures and services. Provided that legal and ethical considerations are considered, its deployment can benefit clear and inclusive communication with users, increased efficiency in legal and administrative procedures, and improved decision-making based on data-driven policy making and evaluation. The government can also benefit from this deployment of generative AI, thereby making its services proactive and optimised.

In the context of learning and development, generative AI can, for instance, assist in analysing (a large volume of) documents, explain complex concepts, and generate practice materials. This contributes to a more personalised approach to learning and offers users the opportunity to expand their knowledge in various fields.

Furthermore, within the context of scientific research, generative AI can contribute to devising, improving, and accelerating research methods, for instance in discovering or refining new concepts for drug development. The impact of generative AI is not limited to specific fields and has the potential to support a wide range of scientific disciplines.

Generative AI can thus contribute as a problem solver in solving major societal problems. Despite concerns about the energy use of (generative) AI technology, this is counterbalanced by the potential contribution that (generative) AI can make to the sustainability transition. Generative AI, for instance, can be used for the analysis of natural ecosystems or predicting climate trends.

Finally, the deployment of generative AI can be used to support and expedite administrative processes, assist customer services, write computer code, and automate industrial processes. This provides users with more room for substantive tasks and in this way contributes to an increased overall quality of performed tasks.

ii. Risks and challenges

The use of generative AI may affect public values such as non-discrimination, privacy and transparency. If it leads to the deterioration of our information ecosystem, it thereby affects democracy and our rule of law. For a well-functioning democracy, well-informed citizens and a common view of reality are important prerequisites, as is broad support for democratic institutions.

A not to be underestimated challenge with generative AI is the potential bias and selectivity that can be embedded in the systems as a result of the used training data and model parameters. This puts equal treatment and non-discrimination under pressure. Additionally, the risk of disinformation grows as generative AI facilitates its creation and dissemination.[6] An example of this is the ease of producing deepfakes. Concerns also exist regarding possible violations of rights related to privacy, data protection, and copyrights.

Regarding work, there are concerns about the impact of generative AI on employment, the quality of work, and income distribution. Both in the short and long term, generative AI can lead to unemployment. Although new tasks and jobs are emerging and the job losses are not the sole outcome, these new opportunities may not always be readily accessible to those workers displaced by automation. Here lies a challenge to support people in continuously developing their skills for the labour market. Moreover, the interaction between generative AI and labour also influences the quality of work. As mentioned, AI can enrich work content. However, there is also a risk that it narrows down tasks, limits workers' autonomy and puts pressure on the humanity of work relationships.[7] In doing so, generative AI can have an impact on the income and job security distribution in

[6] Dutch Data Protection Authority (2023). Report AI & algorithm risks Netherlands (RAN) - autumn 2023.
[7] https://osha.europa.eu/sites/default/files/artificial-intelligence-osh-summary_en.pdf

the longer term.[8] Thus, generative AI is a potential driver of polarisation in the labour market and can exacerbate income inequality. This also underscores the challenge of digital inclusion. Providing everyone with the tools to understand and use innovative technology can empower them to better adapt to the changes it brings.[9]

Another challenge concerns concentration of power. The emergence of generative AI technology amplifies the already existing dynamics in digital markets, increasing the risk of abuse of market power. The most widely used models and services in the Netherlands come from a small number of tech companies that have access to large amounts of data, computational resources and development capacity. As a result, these companies have an advantage and can integrate generative AI into their (existing) services (ecosystem formation) or wield market power to strengthen their position. Despite intense competition among these companies, this can lead to increased entry barriers, which could potentially hinder productive market dynamics. This could, for instance, result in unfair trade practices, higher prices for access and use of AI applications and infrastructure, and reduced choices for consumers and innovation.

At present, the Netherlands largely relies on language models from non-European countries, as outlined in the 'Agenda Digitale Open Strategische Autonomie' (*Agenda Digital Open Strategic Autonomy*).[10] This can pose a risk when these models do not align with our European standards and values. The challenge is to create a market where generative AI applications are offered that comply with all Dutch and European values and laws. Furthermore, we are fully committed to the development and use of public and open alternative models that respect these values and norms.

Concerning society as a whole, the systemic risks associated with the potentially disruptive nature of generative AI pose a significant challenge that should not be underestimated. As a consequence of the application of generative AI, social and economic inequality could potentially increase. A final, yet crucial aspect, concerns the growing energy demands for training and using generative AI, which contradicts our goal of conserving energy and reducing CO2 emissions.

In conclusion, it can be said that generative AI deeply permeates society. Dealing with the opportunities and challenges of generative AI is, therefore, a task that should be tackled collectively, intergovernmentally, and intersectorally, with a learning and experimental approach. This entails continuously engaging in a broad societal dialogue in the Netherlands and based on that, seeking international cooperation within the EU and globally. To achieve this, we need to stay well-informed about the latest developments in generative AI and be aware of the consequences on socio-cultural, socio-economic, and sustainability fronts. Only

[8] https://www.theguardian.com/technology/2023/feb/08/ai-chatgpt-jobs-economy-inequality
[9] https://www.technologyreview.com/2023/03/25/1070275/chatgpt-revolutionize-economy-decide-what-looks-like
[10] https://www.rijksoverheid.nl/documenten/rapporten/2023/10/17/bijlage-agenda-dosa-tgpdfa

then can we anticipate and prepare for the socio-economic impact and further digital transitions. In doing so, as the government, in collaboration with business community and other stakeholders, we must continue to learn and experiment ourselves, for instance, through sandboxes where we deploy (generative) AI in a controlled (testing) environment to seize opportunities.

The impact of generative AI, including both opportunities and possibilities as well as risks and challenges, is discussed in more detail in the vision (Appendix 1).

Connection with current policies, laws and regulations
The Dutch government has been working on policies related to AI for some time now. It is given priority because AI is a key technology, and the associated public values must be safeguarded. The Netherlands aspires to be a front-runner within Europe in the application and regulation of safe and just generative AI.

i. Existing policy
To arrive at a future-proof vision, it is essential to briefly explain existing policies and current laws and regulations. For the most part, the current policy in place is appropriate for the challenges and opportunities that generative AI presents. Key initiatives undertaken in this context include the following:

- Commitment to responsible AI applications. Through the Dutch AI Coalition (NL-AIC), the government, business community, educational and research institutions, and civil society organisations collaborate on socially responsible AI applications. This includes labs, such as the ELSA-labs, where scientists, entrepreneurs, and public institutions conduct research on the ethical, legal, and social aspects of AI.
- The AiNed programme is a public-private multi-year initiative within the National Growth Fund. The programme aims to position the Netherlands permanently in the forefront of AI nations and seeks to contribute to economic recovery and growth, the structural strengthening of the economic foundation in the Netherlands, and a human-centered and responsible application of AI.
- It is important that Dutch government facilitates a support structure that manages the development of AI for education. With funding from the National Growth Fund, the Ministry of Education, Culture and Science and the Ministry of Economic Affairs and Climate Policy are investing substantially in the National Education Lab AI (NOLAI) for a period of ten years.[11] In addition, the National Growth Fund programme Npuls is developing a nationwide AI centre and AI vision for vocational education (mbo), higher education (hbo), and academic education (wo) to prepare these sectors for the transformation of education with AI and to collaborate with partners and institutions in shaping these changes.
- Major investments have already been made in the field of secure AI. Through the Innovation Centre for Artificial Intelligence (ICAI), extensive experimentation and research are conducted in collaboration with the

[11] https://www.rijksoverheid.nl/documenten/kamerstukken/2023/07/06/visiebrief-digitalisering-in-het-funderend-onderwijs

business community, public sector, and the knowledge sector.

ii. (Current) applicable laws and regulations

Generative AI has potentially negative impacts on the realisation of fundamental rights, including the prohibition of discrimination and privacy and data protection law, protection of copyright and trade secrets.[12]

Applicable legal frameworks include the prohibition of discrimination and the principle of equality as contained in the Constitution, the European Convention on Human Rights (ECHR) and the International Covenant on Civil and Political Rights (ICCPR). Privacy and data protection law is protected in the Charter of Fundamental Rights of the European Union, the Treaty on the Functioning of the European Union, the Constitution and the GDPR and, additionally for the Netherlands, the GDPR Implementation Act (UAVG) and the Police Data Act (Wpg) and the Judicial Data and Criminal Records Act (Wjsg).

When it comes to supervising compliance with the GDPR, the independent supervisory authority, in the Netherlands being the Dutch Data Protection Authority (Autoriteit Persoonsgegevens, AP), is responsible. The AP has a broad mandate and extensive powers to investigate whether parties comply with their obligations under the GDPR and, based on that, to take necessary corrective and sanctioning measures. It is therefore not the task of the government to monitor the legality of data processing in the private sector in specific cases.

The European Data Protection Board (EDPB) is a collaborative and independent body consisting of all national privacy regulators from the EU and the European Data Protection Supervisor (EDPS). Because generative AI is a cross-border phenomenon that calls for a harmonised approach, the EDPB has established a task force ChatGPT to promote cooperation and information exchange regarding possible enforcement measures.[6] The Netherlands Authority for Consumers and Markets (ACM) is the regulator responsible for the supervision of compliance with the Telecommunications Act. Copyright is protected under the Copyright Act, and trade secrets are safeguarded under the Trade Secrets Protection Act.

iii. European AI Act

The AI Act is considered to be the primary legislative framework for the development and use of AI in the European Union. Requirements have been established for generative AI, and a system for monitoring these requirements has been put in place. On 8 December 2023, a preliminary political agreement was reached in Brussels concerning the AI Act. The Act will come into force once it has been approved by all EU Member States and the full European Parliament. The aim of this European act is to facilitate the integration of secure AI systems into the internal market, while also protecting public health and fundamental rights. In order to achieve this goal, certain criteria will be established for AI systems, taking into account the level of risk they may pose. While some AI practices are prohibited, other AI systems are subject to high requirements due to

[12] The vision (Appendix 1) considers in more detail the current legal framework around generative AI and how this technology relates to various (fundamental) rights.

their high-risk scope, such as for recruitment and selection or law enforcement. The act will have direct effect in the Netherlands and the regulation will automatically become part of Dutch law. Part of this, such as the supervision of prohibited and high-risk AI applications, will be further established by a Dutch law. The AI Act covers Generative AI and the powerful AI models that often underpin it, which can be used for a wide range of applications. These models are also known as general-purpose AI (GPAI) models. For instance, it is mandated that all GPAI models adhere to transparency standards. This guarantees that businesses utilising these models for particular AI applications are provided with the essential technical documentation to conform with the AI Act. For the most powerful GPAI models with systemic risks, additional obligations will be in place, including monitoring of severe incidents and conducting model evaluations. These obligations will be implemented through codes of practice that the European Commission will develop in collaboration with industry, science, civil society, and other stakeholders. The European Commission will establish a regulatory body, known as the AI Office, to enforce the new regulations for GPAI models.

For generative AI systems, such as chatbots and systems that generate images and videos, the AI Act includes additional transparency requirements. Providers of generative AI systems must ensure that it is clear to people when they are interacting with AI or when content has been generated by AI.

This approach aligns with the Netherlands' commitment to the AI Act. The government considers it proportionate to impose obligations on businesses and governments to protect people against the risks that AI applications may pose. This is crucial for the safe development and use of (generative) AI, and thus, for building trust from both society and the market to harness the opportunities this technology offers.[13]

The European AI Act has been under negotiation between EU member states and with the European Parliament since 21 April 2021. After approval by EU member states and the European Parliament, the law will come into effect. Based on the political agreement reached in December 2023, the Netherlands' authorities and businesses will have between half a year and two years to ensure that AI systems being developed, purchased, and used comply with the requirements of the AI Act. This timeline depends on the level of risk involved; for instance, some AI practices may be prohibited as early as 6 months. For high-risk AI applications, there are timelines of 24 and 36 months, during which the implementation law will be developed in consultation with stakeholders and presented to parliament. For GPAI models, which include most large AI models generating content, the requirements apply after 12 months, and within that period, European supervision will be established.

Government-wide vision for generative AI

Generative AI is part of broader developments in the field of digitalisation and (traditional) AI, albeit with a more disruptive nature. This form of AI can have both strongly positive and strongly negative effects. This underscores the

[13] Parliamentary paper 2023-2024, 21501-33, no 1034.

importance of governance. In line with broader digitalisation policies, such as the 'Value-Driven Digitalisation Work Agenda', the attached vision, therefore, adopts a values-driven approach.

Generative AI is a promising technology. It is essential that this approach serves to enhance human wellbeing and autonomy, sustainability, prosperity, justice, and security. These principles also align seamlessly with the Sustainable Development Goals (SDGs).[14] With our distinctive values-driven approach, the Netherlands has the opportunity to take the lead in Europe and, together with Europe, on the global stage. Our aim is to create a strong AI ecosystem in the Netherlands and the EU, in which many responsible generative AIs can be innovated. We do this by establishing conditions for their development and use, while preserving our digital open strategic autonomy. Therefore, the government focuses on four key principles.

Generative AI in the Netherlands:

I. Is developed and applied in a safe manner.
II. Is developed and applied in an equitable manner.
III. Promotes human wellbeing and autonomy.
IV. Contributes to sustainability and our prosperity.

With its values-driven approach, the Netherlands has the opportunity to play a leading role in Europe and globally. By focusing specifically on responsible applications of generative AI, we are seizing the opportunities presented by this technology. We do this for all sectors. By taking a responsible and open approach to the application of generative AI, society as a whole can benefit.

Therefore, there is a government-wide effort to:

- Develop and apply new generations of generative AI in collaboration with Dutch and European partners, as openly as possible, and grounded in public values and fundamental rights.
- Contributing to the regulation and promotion of AI systems, including generative AI, at the national, European, and global levels. Develop new policies and, if necessary, modernise laws and regulations for the long-term development of generative AI and related technologies.
- Increasing awareness of and knowledge about generative AI in society, so that everyone can benefit from it while also being resilient when it comes to the risks.

The proposed approach requires both a proactive and learning attitude as well as vision and courage. In this regard, open collaboration between the government, execution, business, citizens, and science is crucial. This enables the early detection of signals and the ability to make necessary adjustments. Furthermore, the Netherlands plays an active role both within the EU and globally in stimulating international collaboration and regulation of generative AI.

From vision to action

[14] https://www.rijksoverheid.nl/onderwerpen/ontwikkelingssamenwerking/internationale-afspraken-ontwikkelingssamenwerking

To ensure responsible use, development, and maintain control over the impact of generative AI in the Netherlands, action is required. The government-wide vision on generative AI leads to six action lines. The potential effects of generative AI highlight the importance of examining any shortcomings in policy, as they can have significant consequences for individuals, the economy, and society. This may prompt the assessment of potential additional policy actions and the establishment of conditions for values-driven development and use of generative AI. Some actions are already underway, while others are new. Additionally, a few potential new actions are being explored or researched.[15] The vision document (Appendix 1) provides an extensive overview of the complete set of actions, whether already initiated or new.

1. Collaborate

Generative AI poses some fundamental challenges. Encouraging social dialogue and cooperation at various levels is crucial. Therefore, the vision has been developed in close collaboration with all levels of government, various implementation organisations, businesses, domains and sectors. This cooperation remains essential in the future as well. This means actively involving residents, implementers, employees, employers, entrepreneurs, businesses and professionals. It is essential that they have a voice in shaping the future where generative AI plays a role. The Dutch AI Coalition plays a crucial role in this, including through the AI-Parade. Furthermore, we will continue the conversation with residents about the opportunities and risks of generative AI and digitalisation for society in 2024 through various channels (such as through the vision trajectory generative AI and dialogue sessions organised by the Rathenau Institute).

The safe and responsible development, use, and integration of generative AI require intergovernmental collaboration. Maximising the opportunities that generative AI offers requires a coordinated effort from all levels of government, the private sector, and civil society. Together with local governments and implementation organisations, we are exploring ways in which generative AI can support in making communication with citizens clearer and more inclusive. We are also experimenting with pilots to assess the added value of applying generative AI in improving government services.

Generative AI is a cross-border phenomenon and therefore requires international collaboration. In the EU context, the Netherlands actively participates in initiatives to strengthen the knowledge and innovation position in generative AI, such as the Alliance for Language Technologies EDIC (ALT-EDIC).[16] At the same time, ongoing discussions on international AI governance are being intensified, with the government actively contributing to them. The governance of generative AI has geopolitical aspects. As a result, it can have implications for the international

[15] A number of possible new policy actions are being explored or researched. The outcomes of these actions will also be presented to the next government. Decision-making regarding the potential implementation of these new policy actions, as well as the financial coverage for them, will be (potentially) up to the next government.
[16] An EDIC is a legal framework in the EU that assists member states in setting up and implementing multi-country projects.

order. This underscores the necessity of intensive collaboration with both EU countries and countries outside the EU. Collaboration with like-minded nations on aspects such as human rights (the right to privacy, information, and non-discrimination) and personal and international security concerning the development and use of generative AI is essential. Bodies that play a significant role in this include the Committee on Artificial Intelligence (CAI) of the Council of Europe, the G7 Hiroshima AI Process, and the United Nations General Assembly. Through these channels, the government aims to facilitate the establishment of regulations for generative AI models that prioritise security and interpretation, while also promoting research into responsible and transparent AI models. We do this from a strongly values-driven approach, focusing on safe AI and protection against risks, so that the opportunities of this technology can be optimally utilised.

2. Closely monitoring all developments

Given the expected exponential growth that generative AI will undergo in the coming years, it is of great importance for the government to ensure that employees possess the right knowledge and skills (training) and to closely monitor relevant developments. This means that the government will remain in close contact with local governments, implementaton organisations, knowledge institutions, commercial entities, public interest organisations, employers, employees, and citizens on this subject.

An AI advisory body (or Rapid Response Team) at the highest (Dutch) political level could therefore (potentially) advise the government on pressing issues regarding significant developments in the field of (generative) AI. The government is currently exploring the establishment of such an advisory body,[17] in accordance with the Dekker-Abdulaziz and Rajkowski motion. To keep track of the development and use of generative AI within the government, we are collaborating with local governmental bodies on a monitoring system.[18]

Additionally, specific attention should be given to the consequences for the labour market. This includes both employment and the quality of work. For instance, the Social and Economic Council (Sociaal-Economische Raad, SER) will be mapping the impact of AI on labour productivity, quantity, and quality of work.

The government also pays special attention to the impact of this technology on sustainability and climate. Generative AI can contribute to solving climate challenges, but so far, it has only led to an increase in the ecological footprint. Therefore, it is essential to make these technologies more sustainable through more energy-efficient methods and promoting responsible and sustainable development and use. As a government, sustainability is a key principle in the deployment of generative AI, meaning that as a government, we do not employ this technology if it causes significant harm to people and the planet.

[17] Parliamentary paper 2023-2024, 26 643, no 1075.
[18] To achieve this, meetings are being organised with subnational governments, among others.

3. Shaping and applying laws and regulations

Citizens and businesses must be able to trust that the government establishes appropriate frameworks and rules for generative AI. In this regard, attention should be paid to how existing and upcoming legislation is capable of developing and using generative AI in a safe and fair manner. Clear national, European (or international) frameworks can contribute to accelerating innovation. The European AI Act, which will also apply to generative AI, provides a solid foundation in this respect. The government will start the government-wide implementation of the AI Act in 2024, where guidance and education towards stakeholders is of significant importance.

With a rapidly developing technology like generative AI, it is important that regulation is as technology-neutral and adaptive as possible. Therefore, the government is committed to continuously monitoring – both at the national and European level – whether the current policy and legal framework are sufficient.[19]

4. Increase knowledge and skills

The government is proactively working to increase knowledge and skills to fully capitalise on the opportunities offered by generative AI. In doing so, we, as the government, aim to set the right example by providing guidance to national government organisations for the deployment of generative AI and by tightening procurement conditions with regard to generative AI. A knowledge centre for generative AI within the (national) government could potentially bring such developments together. The government is committed to developing and retaining AI talent, so that we can develop generative AI that meets European norms and values. This is also of added value for the digital open strategic autonomy of Europe. We are also exploring the (necessary) investments in large-scale scientific and technological infrastructure (including supercomputers and computational power) at the national and EU level to be competitive in the field of LLMs (Large Language Models) and other forms of generative AI.

The use of generative AI requires skills in assessing the content generated by this technology. This necessitates further efforts to increase media literacy among various target groups. Therefore, in the entire educational system, skills such as critical and structured thinking are becoming increasingly important. In this context, the National Education Lab AI (NOLAI) is researching the pedagogical, societal, and social consequences of generative AI.

5. Innovating with generative AI

A government that wants to be in control should be a government that, within the existing legal frameworks, experiments with safe and responsible generative AI to discover where risks and opportunities lie in specific applications. It is important that in the Netherlands, an environment is fostered where there is room for experimenting, testing, and scaling up reliable and transparent (generative) AI models and tools.

Therefore, the government is investigating the establishment of a public national

[19] See also: Rathenau Institute (2023), Rathenau scan on generative AI.

AI testing facility for responsible generative AI in the Netherlands. Moreover, the importance of reliable datasets in the Dutch language is increasing as a basis for generative AI models. A National AI Validation Team facilitates publicly available benchmarking and tools (such as bias detection) to responsibly deploy generative AI in the Netherlands.

With the development of its own open language model, GPT-NL, the starting signal has been given for promoting the development of (open) Dutch and European *large language models*, in line with public values. GPT-NL will be a virtual facility open to partners who wish to contribute data and knowledge to a language model and develop applications, for instance in the areas of safety, health, education, and government services. In the further development of the implementation framework for algorithms (Implementatiekader voor Algoritmen, IKA) for government use, ethical frameworks for the responsible use of generative AI will also be developed.

6. *Strong and clear supervision and enforcement*

It is important for developers, policymakers, and regulators at both the European and national levels to remain vigilant for any potential adverse effects that may arise from the use of generative AI in the future. To promote the development of desirable generative AI, it is important to take a proactive approach. It may be beneficial for the government to establish frameworks from the outset to guide its development in a positive direction and mitigate any potential negative outcomes.

Strong and clear supervision requires that regulators acquire knowledge during the development and use of these technologies to keep up with developments and make necessary adjustments. Therefore, good collaboration between regulators, (judicial) authorities, politicians, and society is of great importance.
In this context, it is important that regulators use their practical experience to advise in legislative processes, and play a key role in evaluating whether current policy and laws and regulations are still sufficient in light of the latest developments. In the coming years, the number of generative AI applications will increase significantly. At the same time, the responsibilities of regulators within the digital domain are only increasing. Therefore, it remains important to continually evaluate whether regulators have the knowledge, skills, capacity, and resources to effectively perform their duties now and in the future.

In conclusion

It is widely anticipated that generative AI will have an increasingly significant impact on society in the future, potentially affecting a wide range of areas in our lives. The government thus emphasises the importance of continuing to monitor and analyse the developments and consequences of generative AI.

For the entire Dutch society to benefit from generative AI, it is crucial that the Netherlands, as part of the EU, remains in control. The Netherlands, therefore, aims to be the front-runner within the EU in the field of safe and responsible generative AI.

The government has an important responsibility to steer the development, application and embedding of generative AI in the right direction. This government-wide vision serves as a starting point. Your House of Representatives of the Netherlands will be informed of progress by the end of 2024. In this process, we will also explore the possibility of implementing new actions or policies, if deemed necessary. Other tiers of government and implementation organisations will remain actively involved, especially with regard to feasibility. The government intends to engage in further discussions with the business community on the topic of generative AI. The aim is to explore how this technology can be leveraged to address societal challenges. The rapid developments justify an iterative and learning approach. The (expected) impact of generative AI on many sectors and domains makes the active involvement of the public sector essential, so that all individuals in the Netherlands can truly reap the benefits of this technology and feel protected.

A.C. van Huffelen
State Secretary for Kingdom Relations and Digitalisation

M.A.M. Adriaansens
Minister of Economic Affairs and Climate Policy

R.H. Dijkgraaf
Minister of Education, Culture, and Science